Tell us about your PDF experience.

Specifications

Read the detailed specification for the C# language and the detailed specifications for the latest features.

ECMA specifications and latest features

OVERVIEW

Specification process

Detailed ECMA specification contents

REFERENCE

Latest feature specifications

Draft C# ECMA specification - Introductory material

REFERENCE

Foreword

Introduction

REFERENCE

Scope

Normative references

Terms and definitions

General description

Conformance

Draft C# ECMA specification - Language specification

REFERENCE

Lexical structure

Basic concepts

Types

Variables

Conversions

Patterns

Expressions

Statements

REFERENCE

Namespaces

Classes

Structs

Arrays

Interfaces

Enums

Delegates

Exceptions

Attributes

Unsafe code

Draft C# ECMA specification - Annexes

REFERENCE

Grammar

Portability issues

Standard library

Documentation comments

Bibliography

C# standard specification

Article • 05/08/2024

The C# language specification is the definitive source for the C# language. The ECMA C# standard committee (TC49-TG2) produces the specification. The committee is currently working on version 8 of the standard. The draft published here includes some, but not all, of the C# 8 features. The committee uses Microsoft specifications and language design meeting (LDM) notes to produce the specification.

This section contains the latest working draft of the C# language specification. The latest working draft is published here before being submitted to ECMA for approval. The committee works in the dotnet/csharpstandard repository. You can track the committee's progress and participate in the standard work there.

> **Note**
>
> The committee works to update the C# language specification, but it has fallen behind the current implementation. For specifications about current features, see Feature specifications.

See also

- Feature specifications

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

.NET feedback

.NET is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Detailed table of contents

Article • 02/07/2024

- Foreword
- Introduction
- §1 Scope
- §2 Normative references
- §3 Terms and definitions
- §4 General description
- §5 Conformance
- §6 Lexical structure
  - §6.1 Programs
  - §6.2 Grammars
    - §6.2.1 General
    - §6.2.2 Grammar notation
    - §6.2.3 Lexical grammar
    - §6.2.4 Syntactic grammar
    - §6.2.5 Grammar ambiguities
  - §6.3 Lexical analysis
    - §6.3.1 General
    - §6.3.2 Line terminators
    - §6.3.3 Comments
    - §6.3.4 White space
  - §6.4 Tokens
    - §6.4.1 General
    - §6.4.2 Unicode character escape sequences
    - §6.4.3 Identifiers
    - §6.4.4 Keywords
    - §6.4.5 Literals
      - §6.4.5.1 General
      - §6.4.5.2 Boolean literals
      - §6.4.5.3 Integer literals
      - §6.4.5.4 Real literals
      - §6.4.5.5 Character literals
      - §6.4.5.6 String literals
      - §6.4.5.7 The null literal
    - §6.4.6 Operators and punctuators
  - §6.5 Pre-processing directives
    - §6.5.1 General
    - §6.5.2 Conditional compilation symbols

        - §6.5.3 Pre-processing expressions
        - §6.5.4 Definition directives
        - §6.5.5 Conditional compilation directives
        - §6.5.6 Diagnostic directives
        - §6.5.7 Region directives
        - §6.5.8 Line directives
        - §6.5.9 Pragma directives
- §7 Basic concepts
    - §7.1 Application startup
    - §7.2 Application termination
    - §7.3 Declarations
    - §7.4 Members
        - §7.4.1 General
        - §7.4.2 Namespace members
        - §7.4.3 Struct members
        - §7.4.4 Enumeration members
        - §7.4.5 Class members
        - §7.4.6 Interface members
        - §7.4.7 Array members
        - §7.4.8 Delegate members
    - §7.5 Member access
        - §7.5.1 General
        - §7.5.2 Declared accessibility
        - §7.5.3 Accessibility domains
        - §7.5.4 Protected access
        - §7.5.5 Accessibility constraints
    - §7.6 Signatures and overloading
    - §7.7 Scopes
        - §7.7.1 General
        - §7.7.2 Name hiding
            - §7.7.2.1 General
            - §7.7.2.2 Hiding through nesting
            - §7.7.2.3 Hiding through inheritance
    - §7.8 Namespace and type names
        - §7.8.1 General
        - §7.8.2 Unqualified names
        - §7.8.3 Fully qualified names
    - §7.9 Automatic memory management
    - §7.10 Execution order
- §8 Types
    - §8.1 General

- §8.2 Reference types
  - §8.2.1 General
  - §8.2.2 Class types
  - §8.2.3 The object type
  - §8.2.4 The dynamic type
  - §8.2.5 The string type
  - §8.2.6 Interface types
  - §8.2.7 Array types
  - §8.2.8 Delegate types
- §8.3 Value types
  - §8.3.1 General
  - §8.3.2 The System.ValueType type
  - §8.3.3 Default constructors
  - §8.3.4 Struct types
  - §8.3.5 Simple types
  - §8.3.6 Integral types
  - §8.3.7 Floating-point types
  - §8.3.8 The Decimal type
  - §8.3.9 The Bool type
  - §8.3.10 Enumeration types
  - §8.3.11 Tuple types
  - §8.3.12 Nullable value types
  - §8.3.13 Boxing and unboxing
- §8.4 Constructed types
  - §8.4.1 General
  - §8.4.2 Type arguments
  - §8.4.3 Open and closed types
  - §8.4.4 Bound and unbound types
  - §8.4.5 Satisfying constraints
- §8.5 Type parameters
- §8.6 Expression tree types
- §8.7 The dynamic type
- §8.8 Unmanaged types

- §9 Variables
  - §9.1 General
  - §9.2 Variable categories
    - §9.2.1 General
    - §9.2.2 Static variables
    - §9.2.3 Instance variables
      - §9.2.3.1 General
      - §9.2.3.2 Instance variables in classes

- §9.2.3.3 Instance variables in structs
- §9.2.4 Array elements
- §9.2.5 Value parameters
- §9.2.6 Reference parameters
- §9.2.7 Output parameters
- §9.2.8 Input parameters
- §9.2.9 Local variables
  - §9.2.9.1 Discards
- §9.3 Default values
- §9.4 Definite assignment
  - §9.4.1 General
  - §9.4.2 Initially assigned variables
  - §9.4.3 Initially unassigned variables
  - §9.4.4 Precise rules for determining definite assignment
    - §9.4.4.1 General
    - §9.4.4.2 General rules for statements
    - §9.4.4.3 Block statements, checked, and unchecked statements
    - §9.4.4.4 Expression statements
    - §9.4.4.5 Declaration statements
    - §9.4.4.6 If statements
    - §9.4.4.7 Switch statements
    - §9.4.4.8 While statements
    - §9.4.4.9 Do statements
    - §9.4.4.10 For statements
    - §9.4.4.11 Break, continue, and goto statements
    - §9.4.4.12 Throw statements
    - §9.4.4.13 Return statements
    - §9.4.4.14 Try-catch statements
    - §9.4.4.15 Try-finally statements
    - §9.4.4.16 Try-catch-finally statements
    - §9.4.4.17 Foreach statements
    - §9.4.4.18 Using statements
    - §9.4.4.19 Lock statements
    - §9.4.4.20 Yield statements
    - §9.4.4.21 General rules for constant expressions
    - §9.4.4.22 General rules for simple expressions
    - §9.4.4.23 General rules for expressions with embedded expressions
    - §9.4.4.24 Invocation expressions and object creation expressions
    - §9.4.4.25 Simple assignment expressions
    - §9.4.4.26 && expressions
    - §9.4.4.27 || expressions

        - §9.4.4.28 ! expressions
        - §9.4.4.29 ?? expressions
        - §9.4.4.30 ?: expressions
        - §9.4.4.31 Anonymous functions
        - §9.4.4.32 Throw expressions
        - §9.4.4.33 Rules for variables in local functions
        - §9.4.4.34 is-pattern expressions
    - §9.5 Variable references
    - §9.6 Atomicity of variable references
    - §9.7 Reference variables and returns
        - §9.7.1 General
        - §9.7.2 Ref safe contexts
            - §9.7.2.1 General
            - §9.7.2.2 Local variable ref safe context
            - §9.7.2.3 Parameter ref safe context
            - §9.7.2.4 Field ref safe context
            - §9.7.2.5 Operators
            - §9.7.2.6 Function invocation
            - §9.7.2.7 Values
            - §9.7.2.8 Constructor invocations
            - §9.7.2.9 Limitations on reference variables
- §10 Conversions
    - §10.1 General
    - §10.2 Implicit conversions
        - §10.2.1 General
        - §10.2.2 Identity conversion
        - §10.2.3 Implicit numeric conversions
        - §10.2.4 Implicit enumeration conversions
        - §10.2.5 Implicit interpolated string conversions
        - §10.2.6 Implicit nullable conversions
        - §10.2.7 Null literal conversions
        - §10.2.8 Implicit reference conversions
        - §10.2.9 Boxing conversions
        - §10.2.10 Implicit dynamic conversions
        - §10.2.11 Implicit constant expression conversions
        - §10.2.12 Implicit conversions involving type parameters
        - §10.2.13 Implicit tuple conversions
        - §10.2.14 User-defined implicit conversions
        - §10.2.15 Anonymous function conversions and method group conversions
        - §10.2.16 Default literal conversions
        - §10.2.17 Implicit throw conversions

- §10.3 Explicit conversions
  - §10.3.1 General
  - §10.3.2 Explicit numeric conversions
  - §10.3.3 Explicit enumeration conversions
  - §10.3.4 Explicit nullable conversions
  - §10.3.5 Explicit reference conversions
  - §10.3.6 Explicit tuple conversions
  - §10.3.7 Unboxing conversions
  - §10.3.8 Explicit dynamic conversions
  - §10.3.9 Explicit conversions involving type parameters
  - §10.3.10 User-defined explicit conversions
- §10.4 Standard conversions
  - §10.4.1 General
  - §10.4.2 Standard implicit conversions
  - §10.4.3 Standard explicit conversions
- §10.5 User-defined conversions
  - §10.5.1 General
  - §10.5.2 Permitted user-defined conversions
  - §10.5.3 Evaluation of user-defined conversions
  - §10.5.4 User-defined implicit conversions
  - §10.5.5 User-defined explicit conversions
- §10.6 Conversions involving nullable types
  - §10.6.1 Nullable Conversions
  - §10.6.2 Lifted conversions
- §10.7 Anonymous function conversions
  - §10.7.1 General
  - §10.7.2 Evaluation of anonymous function conversions to delegate types
  - §10.7.3 Evaluation of lambda expression conversions to expression tree types
- §10.8 Method group conversions

- §11 Patterns and pattern matching
  - §11.1 General
  - §11.2 Pattern forms
    - §11.2.1 General
    - §11.2.2 Declaration pattern
    - §11.2.3 Constant pattern
    - §11.2.4 Var pattern
  - §11.3 Pattern subsumption
  - §11.4 Pattern exhaustiveness
- §12 Expressions
  - §12.1 General
  - §12.2 Expression classifications

- §12.2.1 General
- §12.2.2 Values of expressions
- §12.3 Static and Dynamic Binding
  - §12.3.1 General
  - §12.3.2 Binding-time
  - §12.3.3 Dynamic binding
  - §12.3.4 Types of subexpressions
- §12.4 Operators
  - §12.4.1 General
  - §12.4.2 Operator precedence and associativity
  - §12.4.3 Operator overloading
  - §12.4.4 Unary operator overload resolution
  - §12.4.5 Binary operator overload resolution
  - §12.4.6 Candidate user-defined operators
  - §12.4.7 Numeric promotions
    - §12.4.7.1 General
    - §12.4.7.2 Unary numeric promotions
    - §12.4.7.3 Binary numeric promotions
  - §12.4.8 Lifted operators
- §12.5 Member lookup
  - §12.5.1 General
  - §12.5.2 Base types
- §12.6 Function members
  - §12.6.1 General
  - §12.6.2 Argument lists
    - §12.6.2.1 General
    - §12.6.2.2 Corresponding parameters
    - §12.6.2.3 Run-time evaluation of argument lists
  - §12.6.3 Type inference
    - §12.6.3.1 General
    - §12.6.3.2 The first phase
    - §12.6.3.3 The second phase
    - §12.6.3.4 Input types
    - §12.6.3.5 Output types
    - §12.6.3.6 Dependence
    - §12.6.3.7 Output type inferences
    - §12.6.3.8 Explicit parameter type inferences
    - §12.6.3.9 Exact inferences
    - §12.6.3.10 Lower-bound inferences
    - §12.6.3.11 Upper-bound inferences
    - §12.6.3.12 Fixing

- §12.6.3.13 Inferred return type
- §12.6.3.14 Type inference for conversion of method groups
- §12.6.3.15 Finding the best common type of a set of expressions
- §12.6.4 Overload resolution
  - §12.6.4.1 General
  - §12.6.4.2 Applicable function member
  - §12.6.4.3 Better function member
  - §12.6.4.4 Better parameter-passing mode
  - §12.6.4.5 Better conversion from expression
  - §12.6.4.6 Exactly matching expression
  - §12.6.4.7 Better conversion target
  - §12.6.4.8 Overloading in generic classes
- §12.6.5 Compile-time checking of dynamic member invocation
- §12.6.6 Function member invocation
  - §12.6.6.1 General
  - §12.6.6.2 Invocations on boxed instances
- §12.7 Deconstruction
- §12.8 Primary expressions
  - §12.8.1 General
  - §12.8.2 Literals
  - §12.8.3 Interpolated string expressions
  - §12.8.4 Simple names
  - §12.8.5 Parenthesized expressions
  - §12.8.6 Tuple expressions
  - §12.8.7 Member access
    - §12.8.7.1 General
    - §12.8.7.2 Identical simple names and type names
  - §12.8.8 Null Conditional Member Access
  - §12.8.9 Invocation expressions
    - §12.8.9.1 General
    - §12.8.9.2 Method invocations
    - §12.8.9.3 Extension method invocations
    - §12.8.9.4 Delegate invocations
  - §12.8.10 Null Conditional Invocation Expression
  - §12.8.11 Element access
    - §12.8.11.1 General
    - §12.8.11.2 Array access
    - §12.8.11.3 Indexer access
  - §12.8.12 Null Conditional Element Access
  - §12.8.13 This access
  - §12.8.14 Base access

- §12.8.15 Postfix increment and decrement operators
- §12.8.16 The new operator
  - §12.8.16.1 General
  - §12.8.16.2 Object creation expressions
  - §12.8.16.3 Object initializers
  - §12.8.16.4 Collection initializers
  - §12.8.16.5 Array creation expressions
  - §12.8.16.6 Delegate creation expressions
  - §12.8.16.7 Anonymous object creation expressions
- §12.8.17 The typeof operator
- §12.8.18 The sizeof operator
- §12.8.19 The checked and unchecked operators
- §12.8.20 Default value expressions
- §12.8.21 Stack allocation
- §12.8.22 The nameof operator
- §12.8.23 Anonymous method expressions

- §12.9 Unary operators
  - §12.9.1 General
  - §12.9.2 Unary plus operator
  - §12.9.3 Unary minus operator
  - §12.9.4 Logical negation operator
  - §12.9.5 Bitwise complement operator
  - §12.9.6 Prefix increment and decrement operators
  - §12.9.7 Cast expressions
  - §12.9.8 Await expressions
    - §12.9.8.1 General
    - §12.9.8.2 Awaitable expressions
    - §12.9.8.3 Classification of await expressions
    - §12.9.8.4 Run-time evaluation of await expressions
- §12.10 Arithmetic operators
  - §12.10.1 General
  - §12.10.2 Multiplication operator
  - §12.10.3 Division operator
  - §12.10.4 Remainder operator
  - §12.10.5 Addition operator
  - §12.10.6 Subtraction operator
- §12.11 Shift operators
- §12.12 Relational and type-testing operators
  - §12.12.1 General
  - §12.12.2 Integer comparison operators
  - §12.12.3 Floating-point comparison operators

- §12.12.4 Decimal comparison operators
- §12.12.5 Boolean equality operators
- §12.12.6 Enumeration comparison operators
- §12.12.7 Reference type equality operators
- §12.12.8 String equality operators
- §12.12.9 Delegate equality operators
- §12.12.10 Equality operators between nullable value types and the null literal
- §12.12.11 Tuple equality operators
- §12.12.12 The is operator
  - §12.12.12.1 The is-type operator
  - §12.12.12.2 The is-pattern operator
- §12.12.13 The as operator
- §12.13 Logical operators
  - §12.13.1 General
  - §12.13.2 Integer logical operators
  - §12.13.3 Enumeration logical operators
  - §12.13.4 Boolean logical operators
  - §12.13.5 Nullable Boolean & and | operators
- §12.14 Conditional logical operators
  - §12.14.1 General
  - §12.14.2 Boolean conditional logical operators
  - §12.14.3 User-defined conditional logical operators
- §12.15 The null coalescing operator
- §12.16 The throw expression operator
- §12.17 Declaration expressions
- §12.18 Conditional operator
- §12.19 Anonymous function expressions
  - §12.19.1 General
  - §12.19.2 Anonymous function signatures
  - §12.19.3 Anonymous function bodies
  - §12.19.4 Overload resolution
  - §12.19.5 Anonymous functions and dynamic binding
  - §12.19.6 Outer variables
    - §12.19.6.1 General
    - §12.19.6.2 Captured outer variables
    - §12.19.6.3 Instantiation of local variables
  - §12.19.7 Evaluation of anonymous function expressions
  - §12.19.8 Implementation Example
- §12.20 Query expressions
  - §12.20.1 General
  - §12.20.2 Ambiguities in query expressions

        - §12.20.3 Query expression translation
            - §12.20.3.1 General
            - §12.20.3.2 Query expressions with continuations
            - §12.20.3.3 Explicit range variable types
            - §12.20.3.4 Degenerate query expressions
            - §12.20.3.5 From, let, where, join and orderby clauses
            - §12.20.3.6 Select clauses
            - §12.20.3.7 Group clauses
            - §12.20.3.8 Transparent identifiers
        - §12.20.4 The query-expression pattern
    - §12.21 Assignment operators
        - §12.21.1 General
        - §12.21.2 Simple assignment
        - §12.21.3 Ref assignment
        - §12.21.4 Compound assignment
        - §12.21.5 Event assignment
    - §12.22 Expression
    - §12.23 Constant expressions
    - §12.24 Boolean expressions
- §13 Statements
    - §13.1 General
    - §13.2 End points and reachability
    - §13.3 Blocks
        - §13.3.1 General
        - §13.3.2 Statement lists
    - §13.4 The empty statement
    - §13.5 Labeled statements
    - §13.6 Declaration statements
        - §13.6.1 General
        - §13.6.2 Local variable declarations
            - §13.6.2.1 General
            - §13.6.2.2 Implicitly typed local variable declarations
            - §13.6.2.3 Explicitly typed local variable declarations
            - §13.6.2.4 Ref local variable declarations
        - §13.6.3 Local constant declarations
        - §13.6.4 Local function declarations
    - §13.7 Expression statements
    - §13.8 Selection statements
        - §13.8.1 General
        - §13.8.2 The if statement
        - §13.8.3 The switch statement

- §13.9 Iteration statements
  - §13.9.1 General
  - §13.9.2 The while statement
  - §13.9.3 The do statement
  - §13.9.4 The for statement
  - §13.9.5 The foreach statement
- §13.10 Jump statements
  - §13.10.1 General
  - §13.10.2 The break statement
  - §13.10.3 The continue statement
  - §13.10.4 The goto statement
  - §13.10.5 The return statement
  - §13.10.6 The throw statement
- §13.11 The try statement
- §13.12 The checked and unchecked statements
- §13.13 The lock statement
- §13.14 The using statement
- §13.15 The yield statement

- §14 Namespaces
  - §14.1 General
  - §14.2 Compilation units
  - §14.3 Namespace declarations
  - §14.4 Extern alias directives
  - §14.5 Using directives
    - §14.5.1 General
    - §14.5.2 Using alias directives
    - §14.5.3 Using namespace directives
    - §14.5.4 Using static directives
  - §14.6 Namespace member declarations
  - §14.7 Type declarations
  - §14.8 Qualified alias member
    - §14.8.1 General
    - §14.8.2 Uniqueness of aliases

- §15 Classes
  - §15.1 General
  - §15.2 Class declarations
    - §15.2.1 General
    - §15.2.2 Class modifiers
      - §15.2.2.1 General
      - §15.2.2.2 Abstract classes
      - §15.2.2.3 Sealed classes

- §15.2.2.4 Static classes
  - §15.2.2.4.1 General
  - §15.2.2.4.2 Referencing static class types
- §15.2.3 Type parameters
- §15.2.4 Class base specification
  - §15.2.4.1 General
  - §15.2.4.2 Base classes
  - §15.2.4.3 Interface implementations
- §15.2.5 Type parameter constraints
- §15.2.6 Class body
- §15.2.7 Partial declarations
- §15.3 Class members
  - §15.3.1 General
  - §15.3.2 The instance type
  - §15.3.3 Members of constructed types
  - §15.3.4 Inheritance
  - §15.3.5 The new modifier
  - §15.3.6 Access modifiers
  - §15.3.7 Constituent types
  - §15.3.8 Static and instance members
  - §15.3.9 Nested types
    - §15.3.9.1 General
    - §15.3.9.2 Fully qualified name
    - §15.3.9.3 Declared accessibility
    - §15.3.9.4 Hiding
    - §15.3.9.5 this access
    - §15.3.9.6 Access to private and protected members of the containing type
    - §15.3.9.7 Nested types in generic classes
  - §15.3.10 Reserved member names
    - §15.3.10.1 General
    - §15.3.10.2 Member names reserved for properties
    - §15.3.10.3 Member names reserved for events
    - §15.3.10.4 Member names reserved for indexers
    - §15.3.10.5 Member names reserved for finalizers
- §15.4 Constants
- §15.5 Fields
  - §15.5.1 General
  - §15.5.2 Static and instance fields
  - §15.5.3 Readonly fields
    - §15.5.3.1 General
    - §15.5.3.2 Using static readonly fields for constants

- §15.5.3.3 Versioning of constants and static readonly fields
- §15.5.4 Volatile fields
- §15.5.5 Field initialization
- §15.5.6 Variable initializers
  - §15.5.6.1 General
  - §15.5.6.2 Static field initialization
  - §15.5.6.3 Instance field initialization
- §15.6 Methods
  - §15.6.1 General
  - §15.6.2 Method parameters
    - §15.6.2.1 General
    - §15.6.2.2 Value parameters
    - §15.6.2.3 By-reference parameters
      - §15.6.2.3.1 General
      - §15.6.2.3.2 Input parameters
      - §15.6.2.3.3 Reference parameters
      - §15.6.2.3.4 Output parameters
    - §15.6.2.4 Parameter arrays
  - §15.6.3 Static and instance methods
  - §15.6.4 Virtual methods
  - §15.6.5 Override methods
  - §15.6.6 Sealed methods
  - §15.6.7 Abstract methods
  - §15.6.8 External methods
  - §15.6.9 Partial methods
  - §15.6.10 Extension methods
  - §15.6.11 Method body
- §15.7 Properties
  - §15.7.1 General
  - §15.7.2 Static and instance properties
  - §15.7.3 Accessors
  - §15.7.4 Automatically implemented properties
  - §15.7.5 Accessibility
  - §15.7.6 Virtual, sealed, override, and abstract accessors
- §15.8 Events
  - §15.8.1 General
  - §15.8.2 Field-like events
  - §15.8.3 Event accessors
  - §15.8.4 Static and instance events
  - §15.8.5 Virtual, sealed, override, and abstract accessors
- §15.9 Indexers

    - §15.9.1 General
    - §15.9.2 Indexer and Property Differences
  - §15.10 Operators
    - §15.10.1 General
    - §15.10.2 Unary operators
    - §15.10.3 Binary operators
    - §15.10.4 Conversion operators
  - §15.11 Instance constructors
    - §15.11.1 General
    - §15.11.2 Constructor initializers
    - §15.11.3 Instance variable initializers
    - §15.11.4 Constructor execution
    - §15.11.5 Default constructors
  - §15.12 Static constructors
  - §15.13 Finalizers
  - §15.14 Iterators
    - §15.14.1 General
    - §15.14.2 Enumerator interfaces
    - §15.14.3 Enumerable interfaces
    - §15.14.4 Yield type
    - §15.14.5 Enumerator objects
      - §15.14.5.1 General
      - §15.14.5.2 The MoveNext method
      - §15.14.5.3 The Current property
      - §15.14.5.4 The Dispose method
    - §15.14.6 Enumerable objects
      - §15.14.6.1 General
      - §15.14.6.2 The GetEnumerator method
  - §15.15 Async Functions
    - §15.15.1 General
    - §15.15.2 Task-type builder pattern
    - §15.15.3 Evaluation of a task-returning async function
    - §15.15.4 Evaluation of a void-returning async function
- §16 Structs
  - §16.1 General
  - §16.2 Struct declarations
    - §16.2.1 General
    - §16.2.2 Struct modifiers
    - §16.2.3 Ref modifier
    - §16.2.4 Partial modifier
    - §16.2.5 Struct interfaces

        - §16.2.6 Struct body
    - §16.3 Struct members
    - §16.4 Class and struct differences
        - §16.4.1 General
        - §16.4.2 Value semantics
        - §16.4.3 Inheritance
        - §16.4.4 Assignment
        - §16.4.5 Default values
        - §16.4.6 Boxing and unboxing
        - §16.4.7 Meaning of this
        - §16.4.8 Field initializers
        - §16.4.9 Constructors
        - §16.4.10 Static constructors
        - §16.4.11 Automatically implemented properties
        - §16.4.12 Safe context constraint
            - §16.4.12.1 General
            - §16.4.12.2 Parameter safe context
            - §16.4.12.3 Local variable safe context
            - §16.4.12.4 Field safe context
            - §16.4.12.5 Operators
            - §16.4.12.6 Method and property invocation
            - §16.4.12.7 stackalloc
            - §16.4.12.8 Constructor invocations
- §17 Arrays
    - §17.1 General
    - §17.2 Array types
        - §17.2.1 General
        - §17.2.2 The System.Array type
        - §17.2.3 Arrays and the generic collection interfaces
    - §17.3 Array creation
    - §17.4 Array element access
    - §17.5 Array members
    - §17.6 Array covariance
    - §17.7 Array initializers
- §18 Interfaces
    - §18.1 General
    - §18.2 Interface declarations
        - §18.2.1 General
        - §18.2.2 Interface modifiers
        - §18.2.3 Variant type parameter lists
            - §18.2.3.1 General

        - §18.2.3.2 Variance safety
        - §18.2.3.3 Variance conversion
      - §18.2.4 Base interfaces
    - §18.3 Interface body
    - §18.4 Interface members
      - §18.4.1 General
      - §18.4.2 Interface methods
      - §18.4.3 Interface properties
      - §18.4.4 Interface events
      - §18.4.5 Interface indexers
      - §18.4.6 Interface member access
    - §18.5 Qualified interface member names
    - §18.6 Interface implementations
      - §18.6.1 General
      - §18.6.2 Explicit interface member implementations
      - §18.6.3 Uniqueness of implemented interfaces
      - §18.6.4 Implementation of generic methods
      - §18.6.5 Interface mapping
      - §18.6.6 Interface implementation inheritance
      - §18.6.7 Interface re-implementation
      - §18.6.8 Abstract classes and interfaces
- §19 Enums
    - §19.1 General
    - §19.2 Enum declarations
    - §19.3 Enum modifiers
    - §19.4 Enum members
    - §19.5 The System.Enum type
    - §19.6 Enum values and operations
- §20 Delegates
    - §20.1 General
    - §20.2 Delegate declarations
    - §20.3 Delegate members
    - §20.4 Delegate compatibility
    - §20.5 Delegate instantiation
    - §20.6 Delegate invocation
- §21 Exceptions
    - §21.1 General
    - §21.2 Causes of exceptions
    - §21.3 The System.Exception class
    - §21.4 How exceptions are handled
    - §21.5 Common exception classes

- §22 Attributes
  - §22.1 General
  - §22.2 Attribute classes
    - §22.2.1 General
    - §22.2.2 Attribute usage
    - §22.2.3 Positional and named parameters
    - §22.2.4 Attribute parameter types
  - §22.3 Attribute specification
  - §22.4 Attribute instances
    - §22.4.1 General
    - §22.4.2 Compilation of an attribute
    - §22.4.3 Run-time retrieval of an attribute instance
  - §22.5 Reserved attributes
    - §22.5.1 General
    - §22.5.2 The AttributeUsage attribute
    - §22.5.3 The Conditional attribute
      - §22.5.3.1 General
      - §22.5.3.2 Conditional methods
      - §22.5.3.3 Conditional attribute classes
    - §22.5.4 The Obsolete attribute
    - §22.5.5 The AsyncMethodBuilder attribute
    - §22.5.6 Caller-info attributes
      - §22.5.6.1 General
      - §22.5.6.2 The CallerLineNumber attribute
      - §22.5.6.3 The CallerFilePath attribute
      - §22.5.6.4 The CallerMemberName attribute
  - §22.6 Attributes for interoperation
- §23 Unsafe code
  - §23.1 General
  - §23.2 Unsafe contexts
  - §23.3 Pointer types
  - §23.4 Fixed and moveable variables
  - §23.5 Pointer conversions
    - §23.5.1 General
    - §23.5.2 Pointer arrays
  - §23.6 Pointers in expressions
    - §23.6.1 General
    - §23.6.2 Pointer indirection
    - §23.6.3 Pointer member access
    - §23.6.4 Pointer element access
    - §23.6.5 The address-of operator

        - §23.6.6 Pointer increment and decrement
        - §23.6.7 Pointer arithmetic
        - §23.6.8 Pointer comparison
        - §23.6.9 The sizeof operator
    - §23.7 The fixed statement
    - §23.8 Fixed-size buffers
        - §23.8.1 General
        - §23.8.2 Fixed-size buffer declarations
        - §23.8.3 Fixed-size buffers in expressions
        - §23.8.4 Definite assignment checking
    - §23.9 Stack allocation
- §A Grammar
    - §A.1 General
    - §A.2 Lexical grammar
    - §A.3 Syntactic grammar
    - §A.4 Grammar extensions for unsafe code
- §B Portability issues
    - §B.1 General
    - §B.2 Undefined behavior
    - §B.3 Implementation-defined behavior
    - §B.4 Unspecified behavior
    - §B.5 Other issues
- §C Standard library
    - §C.1 General
    - §C.2 Standard Library Types defined in ISO/IEC 23271
    - §C.3 Standard Library Types not defined in ISO/IEC 23271
    - §C.4 Format Specifications
    - §C.5 Library Type Abbreviations
- §D Documentation comments
    - §D.1 General
    - §D.2 Introduction
    - §D.3 Recommended tags
        - §D.3.1 General
        - §D.3.2 <c>
        - §D.3.3 <code>
        - §D.3.4 <example>
        - §D.3.5 <exception>
        - §D.3.6 <include>
        - §D.3.7 <list>
        - §D.3.8 <para>
        - §D.3.9 <param>

        - §D.3.10 <paramref>
        - §D.3.11 <permission>
        - §D.3.12 <remarks>
        - §D.3.13 <returns>
        - §D.3.14 <see>
        - §D.3.15 <seealso>
        - §D.3.16 <summary>
        - §D.3.17 <typeparam>
        - §D.3.18 <typeparamref>
        - §D.3.19 <value>
    - §D.4 Processing the documentation file
        - §D.4.1 General
        - §D.4.2 ID string format
        - §D.4.3 ID string examples
    - §D.5 An example
        - §D.5.1 C# source code
        - §D.5.2 Resulting XML
- §E Bibliography

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Foreword

Article • 02/07/2024

This specification replaces ECMA-334:2022. Changes from the previous edition include the addition of the following:

- Binary integer literals
- Embedded digit separators in numeric literals
- Leading-digit separators in binary and hexadecimal integer literals
- `out` variables
- Discards
- Tuple types
- Pattern Matching
- `ref` locals and returns, conditional `ref` expressions, `ref` with `this` in extension methods, and reassignment of `ref` local variables
- Local Functions
- More expression-bodied members
- `throw` Expressions
- Generalized `async` return types
- `async Main` method
- `default` literal expressions
- Non-trailing named arguments
- `private protected` access modifier
- `in` parameter modifier
- `readonly` structs
- `ref` structs
- Indexing movable fixed buffer without pinning
- Initializers on `stackalloc` arrays
- Pattern-based `fixed` statements
- `System.Delegate` and `System.Enum` as *class_type* constraints.
- Additional generic constraints
- Allow expression variables in more locations
- Attach attributes to the backing field of auto-implemented properties
- Reduce ambiguity of overload resolution

Collaborate with us on GitHub

.NET

ECMA C# draft specification feedback

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Introduction

Article • 03/30/2024

This specification is based on a submission from Hewlett-Packard, Intel, and Microsoft, that described a language called C#, which was developed within Microsoft. The principal inventors of this language were Anders Hejlsberg, Scott Wiltamuth, and Peter Golde. The first widely distributed implementation of C# was released by Microsoft in July 2000, as part of its .NET Framework initiative.

Ecma Technical Committee 39 (TC39) [later renamed to TC49] Task Group 2 (TG2) was formed in September 2000, to produce a standard for C#. Another Task Group, TG3, was also formed at that time to produce a standard for a library and execution environment called Common Language Infrastructure (CLI). (CLI is based on a subset of the .NET Framework.) Although Microsoft's implementation of C# relies on CLI for library and run-time support, other implementations of C# need not, provided they support the features and API required by this C# Standard (see Annex C).

As the definition of C# evolved, the goals used in its design were as follows:

- C# is intended to be a simple, modern, general-purpose, object-oriented programming language.
- The language, and implementations thereof, should provide support for software engineering principles such as strong type checking, array bounds checking, detection of attempts to use uninitialized variables, and automatic garbage collection. Software robustness, durability, and programmer productivity are important.
- The language is intended for use in developing software components suitable for deployment in distributed environments.
- Source code portability is very important, as is programmer portability, especially for those programmers already familiar with C and C++.
- Support for internationalization is very important.
- C# is intended to be suitable for writing applications for both hosted and embedded systems, ranging from the very large that use sophisticated operating systems, down to the very small having dedicated functions.
- Although C# applications are intended to be economical with regard to memory and processing power requirements, the language was not intended to compete directly on performance and size with C or assembly language.

The name C# is pronounced "C Sharp".

The name C# is written as the LATIN CAPITAL LETTER C (U+0043) followed by the NUMBER SIGN # (U+0023).

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

- Open a documentation issue
- Provide product feedback

1 Scope

Article • 02/07/2024

This specification describes the form and establishes the interpretation of programs written in the C# programming language. It describes

- The representation of C# programs;
- The syntax and constraints of the C# language;
- The semantic rules for interpreting C# programs;
- The restrictions and limits imposed by a conforming implementation of C#.

This specification does not describe

- The mechanism by which C# programs are transformed for use by a data-processing system;
- The mechanism by which C# applications are invoked for use by a data-processing system;
- The mechanism by which input data are transformed for use by a C# application;
- The mechanism by which output data are transformed after being produced by a C# application;
- The size or complexity of a program and its data that will exceed the capacity of any specific data-processing system or the capacity of a particular processor;
- All minimal requirements of a data-processing system that is capable of supporting a conforming implementation.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

2 Normative references

Article • 02/07/2024

The following normative documents contain provisions, which, through reference in this text, constitute provisions of this specification. For dated references, subsequent amendments to, or revisions of, any of these publications do not apply. However, parties to agreements based on this specification are encouraged to investigate the possibility of applying the most recent editions of the normative documents indicated below. For undated references, the latest edition of the normative document referred to applies. Members of ISO and IEC maintain registers of currently valid specifications.

ISO/IEC 23271:2012, *Common Language Infrastructure (CLI), Partition IV: Base Class Library (BCL), Extended Numerics Library, and Extended Array Library.*

ISO 80000-2, *Quantities and units — Part 2: Mathematical signs and symbols to be used in the natural sciences and technology*.

ISO/IEC 2382, *Information technology — Vocabulary*.

ISO/IEC 60559:2020, *Information technology — Microprocessor Systems — Floating-Point arithmetic*

The Unicode Consortium. The Unicode Standard, https://www.unicode.org/standard/standard.html

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

3 Terms and definitions

Article • 02/07/2024

For the purposes of this specification, the following definitions apply. Other terms are defined where they appear in ***italic*** type or on the left side of a syntax rule. Terms explicitly defined in this specification are not to be presumed to refer implicitly to similar terms defined elsewhere. Terms not defined in this specification are to be interpreted according to ISO/IEC 2382.1. Mathematical symbols not defined in this specification are to be interpreted according to ISO 80000-2.

- **application** – assembly with an entry point
- **application domain** – entity that enables application isolation by acting as a container for application state
- **argument** – expression in the comma-separated list bounded by the parentheses in a method or instance constructor call expression or bounded by the square brackets in an element access expression
- **assembly** – one or more files output by the compiler as a result of program compilation
- **behavior** – external appearance or action
- **behavior, implementation-defined** – unspecified behavior where each implementation documents how the choice is made
- **behavior, undefined** – behavior, upon use of a non-portable or erroneous construct or of erroneous data, for which this specification imposes no requirements
- **behavior, unspecified** – behavior where this specification provides two or more possibilities and imposes no further requirements on which is chosen in any instance
- **character** (when used without a qualifier)
  - In the context of a non-Unicode encoding, the meaning of character in that encoding; or
  - In the context of a character literal or a value of type char, a Unicode code point in the range U+0000 to U+FFFF (including surrogate code points), that is a UTF-16 code unit; or
  - Otherwise, a Unicode code point
- **class library** – assembly that can be used by other assemblies
- **compilation unit** – ordered sequence of Unicode characters that is input to a compiler
- **diagnostic message** – message belonging to an implementation-defined subset of the implementation's output messages
- **error, compile-time** – error reported during program translation

- **exception** – exceptional condition reported during program execution
- **implementation** – particular set of software (running in a particular translation environment under particular control options) that performs translation of programs for, and supports execution of methods in, a particular execution environment
- **module** – the contents of an assembly produced by a compiler. Some implementations may have facilities to produce assemblies that contain more than one module. The behavior in such situations is outside the scope of this specification
- **namespace** – logical organizational system grouping related program elements
- **parameter** – variable declared as part of a method, instance constructor, operator, or indexer definition, which acquires a value on entry to that function member
- **program** – one or more compilation units that are presented to the compiler and are run or executed by an execution environment
- **unsafe code** – code that is permitted to perform such lower-level operations as declaring and operating on pointers, performing conversions between pointers and integral types, and taking the address of variables
- **warning, compile-time** – informational message reported during program translation, which is intended to identify a potentially questionable usage of a program element

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

4 General description

Article • 02/07/2024

This text is informative.

This specification is intended to be used by implementers, academics, and application programmers. As such, it contains a considerable amount of explanatory material that, strictly speaking, is not necessary in a formal language specification.

This specification is divided into the following subdivisions: front matter; language syntax, constraints, and semantics; and annexes.

Examples are provided to illustrate possible forms of the constructions described. References are used to refer to related clauses. Notes are provided to give advice or guidance to implementers or programmers. Annexes provide additional information and summarize the information contained in this specification.

End of informative text.

Informative text is indicated in the following ways:

1. Whole or partial clauses or annexes delimited by "**This clause/text is informative**" and "**End of informative text**".
2. *Example*: The following example ... code fragment, possibly with some narrative ... *end example* The *Example:* and *end example* markers are in the same paragraph for single paragraph examples. If an example spans multiple paragraphs, the end example marker should be its own paragraph.
3. *Note*: narrative ... *end note* The *Note*: and *end note* markers are in the same paragraph for single paragraph notes. If a note spans multiple paragraphs, the *end note* marker should be its own paragraph.

All text not marked as being informative is normative.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

5 Conformance

Article • 02/07/2024

Conformance is of interest to the following audiences:

- Those designing, implementing, or maintaining C# implementations.
- Governmental or commercial entities wishing to procure C# implementations.
- Testing organizations wishing to provide a C# conformance test suite.
- Programmers wishing to port code from one C# implementation to another.
- Educators wishing to teach Standard C#.
- Authors wanting to write about Standard C#.

As such, conformance is most important, and the bulk of this specification is aimed at specifying the characteristics that make C# implementations and C# programs conforming ones.

The text in this specification that specifies requirements is considered ***normative***. All other text in this specification is ***informative***; that is, for information purposes only. Unless stated otherwise, all text is normative. Normative text is further broken into ***required*** and ***conditional*** categories. ***Conditionally normative*** text specifies a feature and its requirements where the feature is optional. However, if that feature is provided, its syntax and semantics shall be exactly as specified.

Undefined behavior is indicated in this specification only by the words 'undefined behavior.'

A ***strictly conforming program*** shall use only those features of the language specified in this specification as being required. (This means that a strictly conforming program cannot use any conditionally normative feature.) It shall not produce output dependent on any unspecified, undefined, or implementation-defined behavior.

A ***conforming implementation*** of C# shall accept any strictly conforming program.

A conforming implementation of C# shall provide and support all the types, values, objects, properties, methods, and program syntax and semantics described in the normative (but not the conditionally normative) parts in this specification.

A conforming implementation of C# shall interpret characters in conformance with the Unicode Standard. Conforming implementations shall accept compilation units encoded with the UTF-8 encoding form.

A conforming implementation of C# shall not successfully translate source containing a #error preprocessing directive unless it is part of a group skipped by conditional

compilation.

A conforming implementation of C# shall produce at least one diagnostic message if the source program violates any rule of syntax, or any negative requirement (defined as a “shall” or “shall not” or “error” or “warning” requirement), unless that requirement is marked with the words “no diagnostic is required”.

A conforming implementation of C# is permitted to provide additional types, values, objects, properties, and methods beyond those described in this specification, provided they do not alter the behavior of any strictly conforming program. Conforming implementations are required to diagnose programs that use extensions that are ill formed according to this specification. Having done so, however, they can compile and execute such programs. (The ability to have extensions implies that a conforming implementation reserves no identifiers other than those explicitly reserved in this specification.)

A conforming implementation of C# shall be accompanied by a document that defines all implementation-defined characteristics, and all extensions.

A conforming implementation of C# shall support the class library documented in Annex C. This library is included by reference in this specification.

A ***conforming program*** is one that is acceptable to a conforming implementation. (Such a program is permitted to contain extensions or conditionally normative features.)

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

6 Lexical structure

Article • 02/07/2024

6.1 Programs

A C# ***program*** consists of one or more source files, known formally as ***compilation units*** (§14.2). Although a compilation unit might have a one-to-one correspondence with a file in a file system, such correspondence is not required.

Conceptually speaking, a program is compiled using three steps:

1. Transformation, which converts a file from a particular character repertoire and encoding scheme into a sequence of Unicode characters.
2. Lexical analysis, which translates a stream of Unicode input characters into a stream of tokens.
3. Syntactic analysis, which translates the stream of tokens into executable code.

Conforming implementations shall accept Unicode compilation units encoded with the UTF-8 encoding form (as defined by the Unicode standard), and transform them into a sequence of Unicode characters. Implementations can choose to accept and transform additional character encoding schemes (such as UTF-16, UTF-32, or non-Unicode character mappings).

> *Note*: The handling of the Unicode NULL character (U+0000) is implementation-specific. It is strongly recommended that developers avoid using this character in their source code, for the sake of both portability and readability. When the character is required within a character or string literal, the escape sequences `\0` or `\u0000` may be used instead. *end note*

> *Note*: It is beyond the scope of this specification to define how a file using a character representation other than Unicode might be transformed into a sequence of Unicode characters. During such transformation, however, it is recommended that the usual line-separating character (or sequence) in the other character set be translated to the two-character sequence consisting of the Unicode carriage-return character (U+000D) followed by Unicode line-feed character (U+000A). For the most part this transformation will have no visible effects; however, it will affect the interpretation of verbatim string literal tokens (§6.4.5.6). The purpose of this recommendation is to allow a verbatim string literal to produce the same character sequence when its compilation unit is moved between systems that support

> differing non-Unicode character sets, in particular, those using differing character sequences for line-separation. *end note*

6.2 Grammars

6.2.1 General

This specification presents the syntax of the C# programming language using two grammars. The ***lexical grammar*** (§6.2.3) defines how Unicode characters are combined to form line terminators, white space, comments, tokens, and pre-processing directives. The ***syntactic grammar*** (§6.2.4) defines how the tokens resulting from the lexical grammar are combined to form C# programs.

All terminal characters are to be understood as the appropriate Unicode character from the range U+0020 to U+007F, as opposed to any similar-looking characters from other Unicode character ranges.

6.2.2 Grammar notation

The lexical and syntactic grammars are presented in the ANTLR grammar tool's Extended Backus-Naur form.

While the ANTLR notation is used, this specification does not present a complete ANTLR-ready "reference grammar" for C#; writing a lexer and parser, either by hand or using a tool such as ANTLR, is outside the scope of a language specification. With that qualification, this specification attempts to minimize the gap between the specified grammar and that required to build a lexer and parser in ANTLR.

ANTLR distinguishes between lexical and syntactic, termed parser by ANTLR, grammars in its notation by starting lexical rules with an uppercase letter and parser rules with a lowercase letter.

> *Note*: The C# lexical grammar (§6.2.3) and syntactic grammar (§6.2.4) are not in exact correspondence with the ANTLR division into lexical and parser grammers. This small mismatch means that some ANTLR parser rules are used when specifying the C# lexical grammar. *end note*

6.2.3 Lexical grammar

The lexical grammar of C# is presented in §6.3, §6.4, and §6.5. The terminal symbols of the lexical grammar are the characters of the Unicode character set, and the lexical

grammar specifies how characters are combined to form tokens (§6.4), white space (§6.3.4), comments (§6.3.3), and pre-processing directives (§6.5).

Many of the terminal symbols of the syntactic grammar are not defined explicitly as tokens in the lexical grammar. Rather, advantage is taken of the ANTLR behavior that literal strings in the grammar are extracted as implicit lexical tokens; this allows keywords, operators, etc. to be represented in the grammar by their literal representation rather than a token name.

Every compilation unit in a C# program shall conform to the *input* production of the lexical grammar (§6.3.1).

6.2.4 Syntactic grammar

The syntactic grammar of C# is presented in the clauses, subclauses, and annexes that follow this subclause. The terminal symbols of the syntactic grammar are the tokens defined explicitly by the lexical grammar and implicitly by literal strings in the grammar itself (§6.2.3). The syntactic grammar specifies how tokens are combined to form C# programs.

Every compilation unit in a C# program shall conform to the *compilation_unit* production (§14.2) of the syntactic grammar.

6.2.5 Grammar ambiguities

The productions for *simple_name* (§12.8.4) and *member_access* (§12.8.7) can give rise to ambiguities in the grammar for expressions.

> *Example*: The statement:
>
> ```
> C#
> F(G<A, B>(7));
> ```
>
> could be interpreted as a call to `F` with two arguments, `G < A` and `B > (7)`. Alternatively, it could be interpreted as a call to `F` with one argument, which is a call to a generic method `G` with two type arguments and one regular argument.
>
> *end example*

If a sequence of tokens can be parsed (in context) as a *simple_name* (§12.8.4), *member_access* (§12.8.7), or *pointer_member_access* (§23.6.3) ending with a

type_argument_list (§8.4.2), the token immediately following the closing `>` token is examined, to see if it is

- One of `( ) ] } : ; , . ? == != | ^ && || & [`; or
- One of the relational operators `< > <= >= is as`; or
- A contextual query keyword appearing inside a query expression; or
- In certain contexts, *identifier* is treated as a disambiguating token. Those contexts are where the sequence of tokens being disambiguated is immediately preceded by one of the keywords `is`, `case` or `out`, or arises while parsing the first element of a tuple literal (in which case the tokens are preceded by `(` or `:` and the identifier is followed by a `,`) or a subsequent element of a tuple literal.

If the following token is among this list, or an identifier in such a context, then the *type_argument_list* is retained as part of the *simple_name*, *member_access* or *pointer_member-access* and any other possible parse of the sequence of tokens is discarded. Otherwise, the *type_argument_list* is not considered to be part of the *simple_name*, *member_access* or *pointer_member_access*, even if there is no other possible parse of the sequence of tokens.

> *Note*: These rules are not applied when parsing a *type_argument_list* in a *namespace_or_type_name* (§7.8). *end note*

> *Example*: The statement:
>
> ```
> C#
> F(G<A, B>(7));
> ```
>
> will, according to this rule, be interpreted as a call to `F` with one argument, which is a call to a generic method `G` with two type arguments and one regular argument. The statements
>
> ```
> C#
> F(G<A, B>7);
> F(G<A, B>>7);
> ```
>
> will each be interpreted as a call to `F` with two arguments. The statement
>
> ```
> C#
> x = F<A> + y;
> ```

will be interpreted as a less-than operator, greater-than operator and unary-plus operator, as if the statement had been written `x = (F < A) > (+y)`, instead of as a *simple_name* with a *type_argument_list* followed by a binary-plus operator. In the statement

```
C#
x = y is C<T> && z;
```

the tokens `C<T>` are interpreted as a *namespace_or_type_name* with a *type_argument_list* due to the presence of the disambiguating token `&&` after the *type_argument_list*.

The expression `(A < B, C > D)` is a tuple with two elements, each a comparison.

The expression `(A<B,C> D, E)` is a tuple with two elements, the first of which is a declaration expression.

The invocation `M(A < B, C > D, E)` has three arguments.

The invocation `M(out A<B,C> D, E)` has two arguments, the first of which is an `out` declaration.

The expression `e is A<B> C` uses a declaration pattern.

The case label `case A<B> C:` uses a declaration pattern.

end example

When recognising a *relational_expression* (§12.12.1) if both the "*relational_expression* `is` *type*" and "*relational_expression* `is` *constant_pattern*" alternatives are applicable, and *type* resolves to an accessible type, then the "*relational_expression* `is` *type*" alternative shall be chosen.

6.3 Lexical analysis

6.3.1 General

For convenience, the lexical grammar defines and references the following named lexer tokens:

```
ANTLR
```

```
DEFAULT  : 'default' ;
NULL     : 'null' ;
TRUE     : 'true' ;
FALSE    : 'false' ;
ASTERISK : '*' ;
SLASH    : '/' ;
```

Although these are lexer rules, these names are spelled in all-uppercase letters to distinguish them from ordinary lexer rule names.

> *Note*: These convenience rules are exceptions to the usual practice of not providing explicit token names for tokens defined by literal strings. *end note*

The *input* production defines the lexical structure of a C# compilation unit.

ANTLR

```
input
    : input_section?
    ;

input_section
    : input_section_part+
    ;

input_section_part
    : input_element* New_Line
    | PP_Directive
    ;

input_element
    : Whitespace
    | Comment
    | token
    ;
```

> *Note*: The above grammar is described by ANTLR parsing rules, it defines the lexical structure of a C# compilation unit and not lexical tokens. *end note*

Five basic elements make up the lexical structure of a C# compilation unit: Line terminators (§6.3.2), white space (§6.3.4), comments (§6.3.3), tokens (§6.4), and pre-processing directives (§6.5). Of these basic elements, only tokens are significant in the syntactic grammar of a C# program (§6.2.4).

The lexical processing of a C# compilation unit consists of reducing the file into a sequence of tokens that becomes the input to the syntactic analysis. Line terminators,

white space, and comments can serve to separate tokens, and pre-processing directives can cause sections of the compilation unit to be skipped, but otherwise these lexical elements have no impact on the syntactic structure of a C# program.

When several lexical grammar productions match a sequence of characters in a compilation unit, the lexical processing always forms the longest possible lexical element.

> *Example*: The character sequence `//` is processed as the beginning of a single-line comment because that lexical element is longer than a single `/` token. *end example*

Some tokens are defined by a set of lexical rules; a main rule and one or more sub-rules. The latter are marked in the grammar by `fragment` to indicate the rule defines part of another token. Fragment rules are not considered in the top-to-bottom ordering of lexical rules.

> *Note*: In ANTLR `fragment` is a keyword which produces the same behavior defined here. *end note*

6.3.2 Line terminators

Line terminators divide the characters of a C# compilation unit into lines.

ANTLR

```
New_Line
    : New_Line_Character
    | '\u000D\u000A'    // carriage return, line feed
    ;
```

For compatibility with source code editing tools that add end-of-file markers, and to enable a compilation unit to be viewed as a sequence of properly terminated lines, the following transformations are applied, in order, to every compilation unit in a C# program:

- If the last character of the compilation unit is a Control-Z character (U+001A), this character is deleted.
- A carriage-return character (U+000D) is added to the end of the compilation unit if that compilation unit is non-empty and if the last character of the compilation unit is not a carriage return (U+000D), a line feed (U+000A), a next line character (U+0085), a line separator (U+2028), or a paragraph separator (U+2029).

> *Note*: The additional carriage-return allows a program to end in a *PP_Directive* (§6.5) that does not have a terminating *New_Line*. *end note*

6.3.3 Comments

Two forms of comments are supported: delimited comments and single-line comments.

A ***delimited comment*** begins with the characters `/*` and ends with the characters `*/`. Delimited comments can occupy a portion of a line, a single line, or multiple lines.

> *Example*: The example
>
> C#
>
> ```
> /* Hello, world program
>    This program writes "hello, world" to the console
> */
> class Hello
> {
>     static void Main()
>     {
>         System.Console.WriteLine("hello, world");
>     }
> }
> ```
>
> includes a delimited comment.
>
> *end example*

A ***single-line comment*** begins with the characters `//` and extends to the end of the line.

> *Example*: The example
>
> C#
>
> ```
> // Hello, world program
> // This program writes "hello, world" to the console
> //
> class Hello // any name will do for this class
> {
>     static void Main() // this method must be named "Main"
>     {
>         System.Console.WriteLine("hello, world");
>     }
> }
> ```
>
> shows several single-line comments.

> *end example*

ANTLR

```
Comment
    : Single_Line_Comment
    | Delimited_Comment
    ;

fragment Single_Line_Comment
    : '//' Input_Character*
    ;

fragment Input_Character
    // anything but New_Line_Character
    : ~('\u000D' | '\u000A'   | '\u0085' | '\u2028' | '\u2029')
    ;

fragment New_Line_Character
    : '\u000D'  // carriage return
    | '\u000A'  // line feed
    | '\u0085'  // next line
    | '\u2028'  // line separator
    | '\u2029'  // paragraph separator
    ;

fragment Delimited_Comment
    : '/*' Delimited_Comment_Section* ASTERISK+ '/'
    ;

fragment Delimited_Comment_Section
    : SLASH
    | ASTERISK* Not_Slash_Or_Asterisk
    ;

fragment Not_Slash_Or_Asterisk
    : ~('/' | '*')    // Any except SLASH or ASTERISK
    ;
```

Comments do not nest. The character sequences `/*` and `*/` have no special meaning within a single-line comment, and the character sequences `//` and `/*` have no special meaning within a delimited comment.

Comments are not processed within character and string literals.

> *Note*: These rules must be interpreted carefully. For instance, in the example below, the delimited comment that begins before `A` ends between `B` and `C()`. The reason is that

```
```

C#

```
// B */ C();
```

is not actually a single-line comment, since `//` has no special meaning within a delimited comment, and so `*/` does have its usual special meaning in that line.

Likewise, the delimited comment starting before `D` ends before `E`. The reason is that `"D */ "` is not actually a string literal, since the initial double quote character appears inside a delimited comment.

A useful consequence of `/*` and `*/` having no special meaning within a single-line comment is that a block of source code lines can be commented out by putting `//` at the beginning of each line. In general, it does not work to put `/*` before those lines and `*/` after them, as this does not properly encapsulate delimited comments in the block, and in general may completely change the structure of such delimited comments.

Example code:

C#

```
static void Main()
{
    /* A
    // B */ C();
    Console.WriteLine(/* "D */ "E");
}
```

end note

*Single_Line_Comment*s and *Delimited_Comment*s having particular formats can be used as *documentation comments*, as described in §D.

6.3.4 White space

White space is defined as any character with Unicode class Zs (which includes the space character) as well as the horizontal tab character, the vertical tab character, and the form feed character.

ANTLR

```
Whitespace
    : [\p{Zs}]  // any character with Unicode class Zs
    | '\u0009'  // horizontal tab
    | '\u000B'  // vertical tab
    | '\u000C'  // form feed
    ;
```

6.4 Tokens

6.4.1 General

There are several kinds of ***tokens***: identifiers, keywords, literals, operators, and punctuators. White space and comments are not tokens, though they act as separators for tokens.

ANTLR

```
token
    : identifier
    | keyword
    | Integer_Literal
    | Real_Literal
    | Character_Literal
    | String_Literal
    | operator_or_punctuator
    ;
```

> *Note*: This is an ANTLR parser rule, it does not define a lexical token but rather the collection of token kinds. *end note*

6.4.2 Unicode character escape sequences

A Unicode escape sequence represents a Unicode code point. Unicode escape sequences are processed in identifiers (§6.4.3), character literals (§6.4.5.5), regular string literals (§6.4.5.6), and interpolated regular string expressions (§12.8.3). A Unicode escape sequence is not processed in any other location (for example, to form an operator, punctuator, or keyword).

ANTLR

```
fragment Unicode_Escape_Sequence
    : '\\u' Hex_Digit Hex_Digit Hex_Digit Hex_Digit
    | '\\U' Hex_Digit Hex_Digit Hex_Digit Hex_Digit
```

```
            Hex_Digit Hex_Digit Hex_Digit Hex_Digit
    ;
```

A Unicode character escape sequence represents the single Unicode code point formed by the hexadecimal number following the "\u" or "\U" characters. Since C# uses a 16-bit encoding of Unicode code points in character and string values, a Unicode code point in the range `U+10000` to `U+10FFFF` is represented using two Unicode surrogate code units. Unicode code points above `U+FFFF` are not permitted in character literals. Unicode code points above `U+10FFFF` are invalid and are not supported.

Multiple translations are not performed. For instance, the string literal `"\u005Cu005C"` is equivalent to `"\u005C"` rather than `"\"`.

> *Note*: The Unicode value `\u005C` is the character "`\`". *end note*

> *Example*: The example
>
> C#
>
> ```
> class Class1
> {
>     static void Test(bool \u0066)
>     {
>         char c = '\u0066';
>         if (\u0066)
>         {
>             System.Console.WriteLine(c.ToString());
>         }
>     }
> }
> ```
>
> shows several uses of `\u0066`, which is the escape sequence for the letter "`f`". The program is equivalent to
>
> C#
>
> ```
> class Class1
> {
>     static void Test(bool f)
>     {
>         char c = 'f';
>         if (f)
>         {
>             System.Console.WriteLine(c.ToString());
>         }
> ```

```
    }
}
```

end example

6.4.3 Identifiers

The rules for identifiers given in this subclause correspond exactly to those recommended by the Unicode Standard Annex 15 except that underscore is allowed as an initial character (as is traditional in the C programming language), Unicode escape sequences are permitted in identifiers, and the "`@`" character is allowed as a prefix to enable keywords to be used as identifiers.

ANTLR

```
identifier
    : Simple_Identifier
    | contextual_keyword
    ;

Simple_Identifier
    : Available_Identifier
    | Escaped_Identifier
    ;

fragment Available_Identifier
    // excluding keywords or contextual keywords, see note below
    : Basic_Identifier
    ;

fragment Escaped_Identifier
    // Includes keywords and contextual keywords prefixed by '@'.
    // See note below.
    : '@' Basic_Identifier
    ;

fragment Basic_Identifier
    : Identifier_Start_Character Identifier_Part_Character*
    ;

fragment Identifier_Start_Character
    : Letter_Character
    | Underscore_Character
    ;

fragment Underscore_Character
    : '_'                // underscore
    | '\\u005' [fF]      // Unicode_Escape_Sequence for underscore
    | '\\U0000005' [fF] // Unicode_Escape_Sequence for underscore
    ;
```

```
fragment Identifier_Part_Character
    : Letter_Character
    | Decimal_Digit_Character
    | Connecting_Character
    | Combining_Character
    | Formatting_Character
    ;

fragment Letter_Character
    // Category Letter, all subcategories; category Number, subcategory
letter.
    : [\p{L}\p{Nl}]
    // Only escapes for categories L & Nl allowed. See note below.
    | Unicode_Escape_Sequence
    ;

fragment Combining_Character
    // Category Mark, subcategories non-spacing and spacing combining.
    : [\p{Mn}\p{Mc}]
    // Only escapes for categories Mn & Mc allowed. See note below.
    | Unicode_Escape_Sequence
    ;

fragment Decimal_Digit_Character
    // Category Number, subcategory decimal digit.
    : [\p{Nd}]
    // Only escapes for category Nd allowed. See note below.
    | Unicode_Escape_Sequence
    ;

fragment Connecting_Character
    // Category Punctuation, subcategory connector.
    : [\p{Pc}]
    // Only escapes for category Pc allowed. See note below.
    | Unicode_Escape_Sequence
    ;

fragment Formatting_Character
    // Category Other, subcategory format.
    : [\p{Cf}]
    // Only escapes for category Cf allowed, see note below.
    | Unicode_Escape_Sequence
    ;
```

Note:

- For information on the Unicode character classes mentioned above, see *The Unicode Standard*.
- The fragment *Available_Identifier* requires the exclusion of keywords and contextual keywords. If the grammar in this specification is processed with ANTLR then this exclusion is handled automatically by the semantics of ANTLR:

    - Keywords and contextual keywords occur in the grammar as literal strings.
    - ANTLR creates implicit lexical token rules are created from these literal strings.
    - ANTLR considers these implicit rules before the explicit lexical rules in the grammar.
    - Therefore fragment *Available_Identifier* will not match keywords or contextual keywords as the lexical rules for those precede it.
- Fragment *Escaped_Identifier* includes escaped keywords and contextual keywords as they are part of the longer token starting with an `@` and lexical processing always forms the longest possible lexical element (§6.3.1).
- How an implementation enforces the restrictions on the allowable *Unicode_Escape_Sequence* values is an implementation issue.

end note

Example: Examples of valid identifiers are `identifier1`, `_identifier2`, and `@if`. *end example*

An identifier in a conforming program shall be in the canonical format defined by Unicode Normalization Form C, as defined by Unicode Standard Annex 15. The behavior when encountering an identifier not in Normalization Form C is implementation-defined; however, a diagnostic is not required.

The prefix "`@`" enables the use of keywords as identifiers, which is useful when interfacing with other programming languages. The character `@` is not actually part of the identifier, so the identifier might be seen in other languages as a normal identifier, without the prefix. An identifier with an `@` prefix is called a ***verbatim identifier***.

Note: Use of the `@` prefix for identifiers that are not keywords is permitted, but strongly discouraged as a matter of style. *end note*

Example: The example:

```C#
class @class
{
    public static void @static(bool @bool)
    {
        if (@bool)
        {
            System.Console.WriteLine("true");
        }
```

```
        else
        {
            System.Console.WriteLine("false");
        }
    }
}

class Class1
{
    static void M()
    {
        cl\u0061ss.st\u0061tic(true);
    }
}
```

defines a class named "`class`" with a static method named "`static`" that takes a parameter named "`bool`". Note that since Unicode escapes are not permitted in keywords, the token "`cl\u0061ss`" is an identifier, and is the same identifier as "`@class`".

end example

Two identifiers are considered the same if they are identical after the following transformations are applied, in order:

- The prefix "`@`", if used, is removed.
- Each *Unicode_Escape_Sequence* is transformed into its corresponding Unicode character.
- Any *Formatting_Character*s are removed.

The semantics of an identifier named `_` depends on the context in which it appears:

- It can denote a named program element, such as a variable, class, or method, or
- It can denote a discard (§9.2.9.1).

Identifiers containing two consecutive underscore characters (`U+005F`) are reserved for use by the implementation; however, no diagnostic is required if such an identifier is defined.

> *Note*: For example, an implementation might provide extended keywords that begin with two underscores. *end note*

6.4.4 Keywords

A ***keyword*** is an identifier-like sequence of characters that is reserved, and cannot be used as an identifier except when prefaced by the `@` character.

ANTLR

```
keyword
    : 'abstract' | 'as'       | 'base'       | 'bool'      | 'break'
    | 'byte'     | 'case'     | 'catch'      | 'char'      | 'checked'
    | 'class'    | 'const'    | 'continue'   | 'decimal'   | DEFAULT
    | 'delegate' | 'do'       | 'double'     | 'else'      | 'enum'
    | 'event'    | 'explicit' | 'extern'     | FALSE       | 'finally'
    | 'fixed'    | 'float'    | 'for'        | 'foreach'   | 'goto'
    | 'if'       | 'implicit' | 'in'         | 'int'       | 'interface'
    | 'internal' | 'is'       | 'lock'       | 'long'      | 'namespace'
    | 'new'      | NULL       | 'object'     | 'operator'  | 'out'
    | 'override' | 'params'   | 'private'    | 'protected' | 'public'
    | 'readonly' | 'ref'      | 'return'     | 'sbyte'     | 'sealed'
    | 'short'    | 'sizeof'   | 'stackalloc' | 'static'    | 'string'
    | 'struct'   | 'switch'   | 'this'       | 'throw'     | TRUE
    | 'try'      | 'typeof'   | 'uint'       | 'ulong'     | 'unchecked'
    | 'unsafe'   | 'ushort'   | 'using'      | 'virtual'   | 'void'
    | 'volatile' | 'while'
    ;
```

A ***contextual keyword*** is an identifier-like sequence of characters that has special meaning in certain contexts, but is not reserved, and can be used as an identifier outside of those contexts as well as when prefaced by the `@` character.

ANTLR

```
contextual_keyword
    : 'add'    | 'alias'      | 'ascending' | 'async'     | 'await'
    | 'by'     | 'descending' | 'dynamic'   | 'equals'    | 'from'
    | 'get'    | 'global'     | 'group'     | 'into'      | 'join'
    | 'let'    | 'nameof'     | 'on'        | 'orderby'   | 'partial'
    | 'remove' | 'select'     | 'set'       | 'unmanaged' | 'value'
    | 'var'    | 'when'       | 'where'     | 'yield'
    ;
```

> *Note*: The rules *keyword* and *contextual_keyword* are parser rules as they do not introduce new token kinds. All keywords and contextual keywords are defined by implicit lexical rules as they occur as literal strings in the grammar (§6.2.3). *end note*

In most cases, the syntactic location of contextual keywords is such that they can never be confused with ordinary identifier usage. For example, within a property declaration, the `get` and `set` identifiers have special meaning (§15.7.3). An identifier other than `get`

or `set` is never permitted in these locations, so this use does not conflict with a use of these words as identifiers.

In certain cases the grammar is not enough to distinguish contextual keyword usage from identifiers. In all such cases it will be specified how to disambiguate between the two. For example, the contextual keyword `var` in implicitly typed local variable declarations (§13.6.2) might conflict with a declared type called `var`, in which case the declared name takes precedence over the use of the identifier as a contextual keyword.

Another example such disambiguation is the contextual keyword `await` (§12.9.8.1), which is considered a keyword only when inside a method declared `async`, but can be used as an identifier elsewhere.

Just as with keywords, contextual keywords can be used as ordinary identifiers by prefixing them with the `@` character.

> *Note*: When used as contextual keywords, these identifiers cannot contain *Unicode_Escape_Sequence*s. *end note*

6.4.5 Literals

6.4.5.1 General

A ***literal*** (§12.8.2) is a source-code representation of a value.

ANTLR

```
literal
    : boolean_literal
    | Integer_Literal
    | Real_Literal
    | Character_Literal
    | String_Literal
    | null_literal
    ;
```

> *Note*: *literal* is a parser rule as it groups other token kinds and does not introduce a new token kind. *end note*

6.4.5.2 Boolean literals

There are two Boolean literal values: `true` and `false`.

ANTLR

```
boolean_literal
    : TRUE
    | FALSE
    ;
```

> *Note*: *boolean_literal* is a parser rule as it groups other token kinds and does not introduce a new token kind. *end note*

The type of a *boolean_literal* is `bool`.

6.4.5.3 Integer literals

Integer literals are used to write values of types `int`, `uint`, `long`, and `ulong`. Integer literals have three possible forms: decimal, hexadecimal, and binary.

ANTLR

```
Integer_Literal
    : Decimal_Integer_Literal
    | Hexadecimal_Integer_Literal
    | Binary_Integer_Literal
    ;

fragment Decimal_Integer_Literal
    : Decimal_Digit Decorated_Decimal_Digit* Integer_Type_Suffix?
    ;

fragment Decorated_Decimal_Digit
    : '_'* Decimal_Digit
    ;

fragment Decimal_Digit
    : '0'..'9'
    ;

fragment Integer_Type_Suffix
    : 'U' | 'u' | 'L' | 'l' |
      'UL' | 'Ul' | 'uL' | 'ul' | 'LU' | 'Lu' | 'lU' | 'lu'
    ;

fragment Hexadecimal_Integer_Literal
    : ('0x' | '0X') Decorated_Hex_Digit+ Integer_Type_Suffix?
    ;

fragment Decorated_Hex_Digit
    : '_'* Hex_Digit
    ;
```

```
fragment Hex_Digit
    : '0'..'9' | 'A'..'F' | 'a'..'f'
    ;

fragment Binary_Integer_Literal
    : ('0b' | '0B') Decorated_Binary_Digit+ Integer_Type_Suffix?
    ;

fragment Decorated_Binary_Digit
    : '_'* Binary_Digit
    ;

fragment Binary_Digit
    : '0' | '1'
    ;
```

The type of an integer literal is determined as follows:

- If the literal has no suffix, it has the first of these types in which its value can be represented: `int`, `uint`, `long`, `ulong`.
- If the literal is suffixed by `U` or `u`, it has the first of these types in which its value can be represented: `uint`, `ulong`.
- If the literal is suffixed by `L` or `l`, it has the first of these types in which its value can be represented: `long`, `ulong`.
- If the literal is suffixed by `UL`, `Ul`, `uL`, `ul`, `LU`, `Lu`, `lU`, or `lu`, it is of type `ulong`.

If the value represented by an integer literal is outside the range of the `ulong` type, a compile-time error occurs.

> *Note*: As a matter of style, it is suggested that "`L`" be used instead of "`l`" when writing literals of type `long`, since it is easy to confuse the letter "`l`" with the digit "`1`". *end note*

To permit the smallest possible `int` and `long` values to be written as integer literals, the following two rules exist:

- When an *Integer_Literal* representing the value `2147483648` (2^{31}) and no *Integer_Type_Suffix* appears as the token immediately following a unary minus operator token (§12.9.3), the result (of both tokens) is a constant of type int with the value `-2147483648` (-2^{31}). In all other situations, such an *Integer_Literal* is of type `uint`.
- When an *Integer_Literal* representing the value `9223372036854775808` (2^{63}) and no *Integer_Type_Suffix* or the *Integer_Type_Suffix* `L` or `l` appears as the token

immediately following a unary minus operator token (§12.9.3), the result (of both tokens) is a constant of type `long` with the value `-9223372036854775808` (-2^{63}). In all other situations, such an *Integer_Literal* is of type `ulong`.

Example:

```
C#

123                 // decimal, int
10_543_765Lu        // decimal, ulong
1_2__3___4____5     // decimal, int
_123                // not a numeric literal; identifier due to leading
_
123_                // invalid; no trailing _allowed

0xFf                // hex, int
0X1b_a0_44_fEL      // hex, long
0x1ade_3FE1_29AaUL  // hex, ulong
0x_abc              // hex, int
_0x123              // not a numeric literal; identifier due to leading
_
0xabc_              // invalid; no trailing _ allowed

0b101               // binary, int
0B1001_1010u        // binary, uint
0b1111_1111_0000UL  // binary, ulong
0B__111             // binary, int
__0B111             // not a numeric literal; identifier due to leading
_
0B111__             // invalid; no trailing _ allowed
```

end example

6.4.5.4 Real literals

Real literals are used to write values of types `float`, `double`, and `decimal`.

```
ANTLR

Real_Literal
    : Decimal_Digit Decorated_Decimal_Digit* '.'
      Decimal_Digit Decorated_Decimal_Digit* Exponent_Part?
Real_Type_Suffix?
    | '.' Decimal_Digit Decorated_Decimal_Digit* Exponent_Part?
Real_Type_Suffix?
    | Decimal_Digit Decorated_Decimal_Digit* Exponent_Part Real_Type_Suffix?
    | Decimal_Digit Decorated_Decimal_Digit* Real_Type_Suffix
    ;

fragment Exponent_Part
```

```
    : ('e' | 'E') Sign? Decimal_Digit Decorated_Decimal_Digit*
    ;

fragment Sign
    : '+' | '-'
    ;

fragment Real_Type_Suffix
    : 'F' | 'f' | 'D' | 'd' | 'M' | 'm'
    ;
```

If no *Real_Type_Suffix* is specified, the type of the *Real_Literal* is `double`. Otherwise, the *Real_Type_Suffix* determines the type of the real literal, as follows:

- A real literal suffixed by `F` or `f` is of type `float`.

  > *Example*: The literals `1f`, `1.5f`, `1e10f`, and `123.456F` are all of type `float`. *end example*

- A real literal suffixed by `D` or `d` is of type `double`.

  > *Example*: The literals `1d`, `1.5d`, `1e10d`, and `123.456D` are all of type `double`. *end example*

- A real literal suffixed by `M` or `m` is of type `decimal`.

  > *Example*: The literals `1m`, `1.5m`, `1e10m`, and `123.456M` are all of type `decimal`. *end example*
  > This literal is converted to a `decimal` value by taking the exact value, and, if necessary, rounding to the nearest representable value using banker's rounding (§8.3.8). Any scale apparent in the literal is preserved unless the value is rounded. *Note*: Hence, the literal `2.900m` will be parsed to form the `decimal` with sign `0`, coefficient `2900`, and scale `3`. *end note*

If the magnitude of the specified literal is too large to be represented in the indicated type, a compile-time error occurs.

> *Note*: In particular, a *Real_Literal* will never produce a floating-point infinity. A non-zero *Real_Literal* may, however, be rounded to zero. *end note*

The value of a real literal of type `float` or `double` is determined by using the IEC 60559 "round to nearest" mode with ties broken to "even" (a value with the least-significant-bit

zero), and all digits considered significant.

> *Note*: In a real literal, decimal digits are always required after the decimal point. For example, `1.3F` is a real literal but `1.F` is not. *end note*
>
> *Example*:
>
> C#
>
> ```
> 1.234_567      // double
> .3e5f          // float
> 2_345E-2_0     // double
> 15D            // double
> 19.73M         // decimal
> 1.F            // parsed as a member access of F due to non-digit after
> .
> 1_.2F          // invalid; no trailing _ allowed in integer part
> 1._234         // parsed as a member access of _234 due to non-digit
> after .
> 1.234_         // invalid; no trailing _ allowed in fraction
> .3e_5F         // invalid; no leading _ allowed in exponent
> .3e5_F         // invalid; no trailing _ allowed in exponent
> ```
>
> *end example*

6.4.5.5 Character literals

A character literal represents a single character, and consists of a character in quotes, as in `'a'`.

ANTLR

```
Character_Literal
    : '\'' Character '\''
    ;

fragment Character
    : Single_Character
    | Simple_Escape_Sequence
    | Hexadecimal_Escape_Sequence
    | Unicode_Escape_Sequence
    ;

fragment Single_Character
    // anything but ', \, and New_Line_Character
    : ~['\\\u000D\u000A\u0085\u2028\u2029]
    ;

fragment Simple_Escape_Sequence
    : '\\'' | '\\"' | '\\\\' | '\\0' | '\\a' | '\\b' |
```

```
      '\\f' | '\\n' | '\\r' | '\\t' | '\\v'
    ;

fragment Hexadecimal_Escape_Sequence
    : '\\x' Hex_Digit Hex_Digit? Hex_Digit? Hex_Digit?
    ;
```

> *Note*: A character that follows a backslash character (`\`) in a *Character* must be one of the following characters: `'`, `"`, `\`, `0`, `a`, `b`, `f`, `n`, `r`, `t`, `u`, `U`, `x`, `v`. Otherwise, a compile-time error occurs. *end note*

> *Note*: The use of the `\x` *Hexadecimal_Escape_Sequence* production can be error-prone and hard to read due to the variable number of hexadecimal digits following the `\x`. For example, in the code:
>
> C#
>
> ```
> string good = "\x9Good text";
> string bad = "\x9Bad text";
> ```
>
> it might appear at first that the leading character is the same (`U+0009`, a tab character) in both strings. In fact the second string starts with `U+9BAD` as all three letters in the word "Bad" are valid hexadecimal digits. As a matter of style, it is recommended that `\x` is avoided in favour of either specific escape sequences (`\t` in this example) or the fixed-length `\u` escape sequence.
>
> *end note*

A hexadecimal escape sequence represents a single Unicode UTF-16 code unit, with the value formed by the hexadecimal number following "`\x`".

If the value represented by a character literal is greater than `U+FFFF`, a compile-time error occurs.

A Unicode escape sequence (§6.4.2) in a character literal shall be in the range `U+0000` to `U+FFFF`.

A simple escape sequence represents a Unicode character, as described in the table below.

Expand table

Escape sequence	Character name	Unicode code point
`\'`	Single quote	U+0027
`\"`	Double quote	U+0022
`\\`	Backslash	U+005C
`\0`	Null	U+0000
`\a`	Alert	U+0007
`\b`	Backspace	U+0008
`\f`	Form feed	U+000C
`\n`	New line	U+000A
`\r`	Carriage return	U+000D
`\t`	Horizontal tab	U+0009
`\v`	Vertical tab	U+000B

The type of a *Character_Literal* is `char`.

6.4.5.6 String literals

C# supports two forms of string literals: ***regular string literals*** and ***verbatim string literals***. A regular string literal consists of zero or more characters enclosed in double quotes, as in `"hello"`, and can include both simple escape sequences (such as `\t` for the tab character), and hexadecimal and Unicode escape sequences.

A verbatim string literal consists of an `@` character followed by a double-quote character, zero or more characters, and a closing double-quote character.

> *Example*: A simple example is `@"hello"`. *end example*

In a verbatim string literal, the characters between the delimiters are interpreted verbatim, with the only exception being a *Quote_Escape_Sequence*, which represents one double-quote character. In particular, simple escape sequences, and hexadecimal and Unicode escape sequences are not processed in verbatim string literals. A verbatim string literal may span multiple lines.

```
ANTLR
```

```
String_Literal
    : Regular_String_Literal
    | Verbatim_String_Literal
    ;

fragment Regular_String_Literal
    : '"' Regular_String_Literal_Character* '"'
    ;

fragment Regular_String_Literal_Character
    : Single_Regular_String_Literal_Character
    | Simple_Escape_Sequence
    | Hexadecimal_Escape_Sequence
    | Unicode_Escape_Sequence
    ;

fragment Single_Regular_String_Literal_Character
    // anything but ", \, and New_Line_Character
    : ~["\\\u000D\u000A\u0085\u2028\u2029]
    ;

fragment Verbatim_String_Literal
    : '@"' Verbatim_String_Literal_Character* '"'
    ;

fragment Verbatim_String_Literal_Character
    : Single_Verbatim_String_Literal_Character
    | Quote_Escape_Sequence
    ;

fragment Single_Verbatim_String_Literal_Character
    : ~["]     // anything but quotation mark (U+0022)
    ;

fragment Quote_Escape_Sequence
    : '""'
    ;
```

Example: The example

C#

```
string a = "Happy birthday, Joel"; // Happy birthday, Joel
string b = @"Happy birthday, Joel"; // Happy birthday, Joel
string c = "hello \t world"; // hello world
string d = @"hello \t world"; // hello \t world
string e = "Joe said \"Hello\" to me"; // Joe said "Hello" to me
string f = @"Joe said ""Hello"" to me"; // Joe said "Hello" to me
string g = "\\\\server\\share\\file.txt"; // \\server\share\file.txt
string h = @"\\server\share\file.txt"; // \\server\share\file.txt
string i = "one\r\ntwo\r\nthree";
string j = @"one
```

```
two
three";
```

shows a variety of string literals. The last string literal, `j`, is a verbatim string literal that spans multiple lines. The characters between the quotation marks, including white space such as new line characters, are preserved verbatim, and each pair of double-quote characters is replaced by one such character.

end example

Note: Any line breaks within verbatim string literals are part of the resulting string. If the exact characters used to form line breaks are semantically relevant to an application, any tools that translate line breaks in source code to different formats (between "`\n`" and "`\r\n`", for example) will change application behavior. Developers should be careful in such situations. *end note*

Note: Since a hexadecimal escape sequence can have a variable number of hex digits, the string literal `"\x123"` contains a single character with hex value `123`. To create a string containing the character with hex value `12` followed by the character `3`, one could write `"\x00123"` or `"\x12"` + `"3"` instead. *end note*

The type of a *String_Literal* is `string`.

Each string literal does not necessarily result in a new string instance. When two or more string literals that are equivalent according to the string equality operator (§12.12.8), appear in the same assembly, these string literals refer to the same string instance.

Example: For instance, the output produced by

C#

```
class Test
{
    static void Main()
    {
        object a = "hello";
        object b = "hello";
        System.Console.WriteLine(a == b);
    }
}
```

is `True` because the two literals refer to the same string instance.

> *end example*

6.4.5.7 The null literal

ANTLR

```
null_literal
    : NULL
    ;
```

> *Note*: *null_literal* is a parser rule as it does not introduce a new token kind. *end note*

A *null_literal* represents a `null` value. It does not have a type, but can be converted to any reference type or nullable value type through a null literal conversion (§10.2.7).

6.4.6 Operators and punctuators

There are several kinds of operators and punctuators. Operators are used in expressions to describe operations involving one or more operands.

> *Example*: The expression `a + b` uses the `+` operator to add the two operands `a` and `b`. *end example*

Punctuators are for grouping and separating.

ANTLR

```
operator_or_punctuator
    : '{'  | '}'  | '['  | ']'  | '('   | ')'  | '.'  | ','  | ':'  | ';'
    | '+'  | '-'  | ASTERISK    | SLASH | '%'  | '&'  | '|'  | '^'  | '!' |
'~'
    | '='  | '<'  | '>'  | '?'  | '??'  | '::' | '++' | '--' | '&&' | '||'
    | '->' | '==' | '!=' | '<=' | '>='  | '+=' | '-=' | '*=' | '/=' | '%='
    | '&=' | '|=' | '^=' | '<<' | '<<='  | '=>'
    ;

right_shift
    : '>'  '>'
    ;

right_shift_assignment
    : '>' '>='
    ;
```

> *Note*: *right_shift* and *right_shift_assignment* are parser rules as they do not introduce a new token kind but represent a sequence of two tokens. The *operator_or_punctuator* rule exists for descriptive purposes only and is not used elsewhere in the grammar. *end note*

right_shift is made up of the two tokens `>` and `>`. Similarly, *right_shift_assignment* is made up of the two tokens `>` and `>=`. Unlike other productions in the syntactic grammar, no characters of any kind (not even whitespace) are allowed between the two tokens in each of these productions. These productions are treated specially in order to enable the correct handling of *type_parameter_lists* (§15.2.3).

> *Note*: Prior to the addition of generics to C#, `>>` and `>>=` were both single tokens. However, the syntax for generics uses the `<` and `>` characters to delimit type parameters and type arguments. It is often desirable to use nested constructed types, such as `List<Dictionary<string, int>>`. Rather than requiring the programmer to separate the `>` and `>` by a space, the definition of the two *operator_or_punctuator*s was changed. *end note*

6.5 Pre-processing directives

6.5.1 General

The pre-processing directives provide the ability to conditionally skip sections of compilation units, to report error and warning conditions, and to delineate distinct regions of source code.

> *Note*: The term "pre-processing directives" is used only for consistency with the C and C++ programming languages. In C#, there is no separate pre-processing step; pre-processing directives are processed as part of the lexical analysis phase. *end note*

ANTLR

```
PP_Directive
    : PP_Start PP_Kind PP_New_Line
    ;

fragment PP_Kind
    : PP_Declaration
    | PP_Conditional
    | PP_Line
    | PP_Diagnostic
```

```
    | PP_Region
    | PP_Pragma
    ;

// Only recognised at the beginning of a line
fragment PP_Start
    // See note below.
    : { getCharPositionInLine() == 0 }? PP_Whitespace? '#' PP_Whitespace?
    ;

fragment PP_Whitespace
    : ( [\p{Zs}]  // any character with Unicode class Zs
      | '\u0009'  // horizontal tab
      | '\u000B'  // vertical tab
      | '\u000C'  // form feed
      )+
    ;

fragment PP_New_Line
    : PP_Whitespace? Single_Line_Comment? New_Line
    ;
```

> *Note*:
>
> - The pre-processor grammar defines a single lexical token `PP_Directive` used for all pre-processing directives. The semantics of each of the pre-processing directives are defined in this language specification but not how to implement them.
> - The `PP_Start` fragment must only be recognised at the start of a line, the `getCharPositionInLine() == 0` ANTLR lexical predicate above suggests one way in which this may be achieved and is informative *only*, an implementation may use a different strategy.
>
> *end note*

The following pre-processing directives are available:

- `#define` and `#undef`, which are used to define and undefine, respectively, conditional compilation symbols (§6.5.4).
- `#if`, `#elif`, `#else`, and `#endif`, which are used to skip conditionally sections of source code (§6.5.5).
- `#line`, which is used to control line numbers emitted for errors and warnings (§6.5.8).
- `#error`, which is used to issue errors (§6.5.6).
- `#region` and `#endregion`, which are used to explicitly mark sections of source code (§6.5.7).

- `#pragma`, which is used to specify optional contextual information to a compiler (§6.5.9).

A pre-processing directive always occupies a separate line of source code and always begins with a `#` character and a pre-processing directive name. White space may occur before the `#` character and between the `#` character and the directive name.

A source line containing a `#define`, `#undef`, `#if`, `#elif`, `#else`, `#endif`, `#line`, or `#endregion` directive can end with a single-line comment. Delimited comments (the `/* */` style of comments) are not permitted on source lines containing pre-processing directives.

Pre-processing directives are not part of the syntactic grammar of C#. However, pre-processing directives can be used to include or exclude sequences of tokens and can in that way affect the meaning of a C# program.

> *Example*: When compiled, the program
>
> C#
>
> ```
> #define A
> #undef B
> class C
> {
> #if A
>     void F() {}
> #else
>     void G() {}
> #endif
> #if B
>     void H() {}
> #else
>     void I() {}
> #endif
> }
> ```
>
> results in the exact same sequence of tokens as the program
>
> C#
>
> ```
> class C
> {
>     void F() {}
>     void I() {}
> }
> ```

> Thus, whereas lexically, the two programs are quite different, syntactically, they are identical.
>
> *end example*

6.5.2 Conditional compilation symbols

The conditional compilation functionality provided by the `#if`, `#elif`, `#else`, and `#endif` directives is controlled through pre-processing expressions (§6.5.3) and conditional compilation symbols.

ANTLR

```
fragment PP_Conditional_Symbol
    // Must not be equal to tokens TRUE or FALSE. See note below.
    : Basic_Identifier
    ;
```

> *Note* How an implementation enforces the restriction on the allowable *Basic_Identifier* values is an implementation issue. *end note*

Two conditional compilation symbols are considered the same if they are identical after the following transformations are applied, in order:

- Each *Unicode_Escape_Sequence* is transformed into its corresponding Unicode character.
- Any *Formatting_Characters* are removed.

A conditional compilation symbol has two possible states: ***defined*** or ***undefined***. At the beginning of the lexical processing of a compilation unit, a conditional compilation symbol is undefined unless it has been explicitly defined by an external mechanism (such as a command-line compiler option). When a `#define` directive is processed, the conditional compilation symbol named in that directive becomes defined in that compilation unit. The symbol remains defined until a `#undef` directive for that same symbol is processed, or until the end of the compilation unit is reached. An implication of this is that `#define` and `#undef` directives in one compilation unit have no effect on other compilation units in the same program.

When referenced in a pre-processing expression (§6.5.3), a defined conditional compilation symbol has the Boolean value `true`, and an undefined conditional compilation symbol has the Boolean value `false`. There is no requirement that conditional compilation symbols be explicitly declared before they are referenced in

pre-processing expressions. Instead, undeclared symbols are simply undefined and thus have the value `false`.

The namespace for conditional compilation symbols is distinct and separate from all other named entities in a C# program. Conditional compilation symbols can only be referenced in `#define` and `#undef` directives and in pre-processing expressions.

6.5.3 Pre-processing expressions

Pre-processing expressions can occur in `#if` and `#elif` directives. The operators `!`, `==`, `!=`, `&&`, and `||` are permitted in pre-processing expressions, and parentheses may be used for grouping.

ANTLR

```
fragment PP_Expression
    : PP_Whitespace? PP_Or_Expression PP_Whitespace?
    ;

fragment PP_Or_Expression
    : PP_And_Expression (PP_Whitespace? '||' PP_Whitespace?
PP_And_Expression)*
    ;

fragment PP_And_Expression
    : PP_Equality_Expression (PP_Whitespace? '&&' PP_Whitespace?
      PP_Equality_Expression)*
    ;

fragment PP_Equality_Expression
    : PP_Unary_Expression (PP_Whitespace? ('==' | '!=') PP_Whitespace?
      PP_Unary_Expression)*
    ;

fragment PP_Unary_Expression
    : PP_Primary_Expression
    | '!' PP_Whitespace? PP_Unary_Expression
    ;

fragment PP_Primary_Expression
    : TRUE
    | FALSE
    | PP_Conditional_Symbol
    | '(' PP_Whitespace? PP_Expression PP_Whitespace? ')'
    ;
```

When referenced in a pre-processing expression, a defined conditional compilation symbol has the Boolean value `true`, and an undefined conditional compilation symbol

has the Boolean value `false`.

Evaluation of a pre-processing expression always yields a Boolean value. The rules of evaluation for a pre-processing expression are the same as those for a constant expression (§12.23), except that the only user-defined entities that can be referenced are conditional compilation symbols.

6.5.4 Definition directives

The definition directives are used to define or undefine conditional compilation symbols.

```
ANTLR
```

```
fragment PP_Declaration
    : 'define' PP_Whitespace PP_Conditional_Symbol
    | 'undef' PP_Whitespace PP_Conditional_Symbol
    ;
```

The processing of a `#define` directive causes the given conditional compilation symbol to become defined, starting with the source line that follows the directive. Likewise, the processing of a `#undef` directive causes the given conditional compilation symbol to become undefined, starting with the source line that follows the directive.

Any `#define` and `#undef` directives in a compilation unit shall occur before the first *token* (§6.4) in the compilation unit; otherwise a compile-time error occurs. In intuitive terms, `#define` and `#undef` directives shall precede any "real code" in the compilation unit.

> *Example*: The example:
>
> ```
> C#
> ```
>
> ```
> #define Enterprise
> #if Professional || Enterprise
> #define Advanced
> #endif
> namespace Megacorp.Data
> {
> #if Advanced
>     class PivotTable {...}
> #endif
> }
> ```
>
> is valid because the `#define` directives precede the first token (the `namespace` keyword) in the compilation unit.

end example

Example: The following example results in a compile-time error because a #define follows real code:

```
C#

#define A
namespace N
{
#define B
#if B
    class Class1 {}
#endif
}
```

end example

A `#define` may define a conditional compilation symbol that is already defined, without there being any intervening `#undef` for that symbol.

Example: The example below defines a conditional compilation symbol A and then defines it again.

```
C#

#define A
#define A
```

For compilers that allow conditional compilation symbols to be defined as compilation options, an alternative way for such redefinition to occur is to define the symbol as a compiler option as well as in the source.

end example

A `#undef` may "undefine" a conditional compilation symbol that is not defined.

Example: The example below defines a conditional compilation symbol `A` and then undefines it twice; although the second `#undef` has no effect, it is still valid.

```
C#

#define A
#undef A
```

```
#undef A
```

end example

6.5.5 Conditional compilation directives

The conditional compilation directives are used to conditionally include or exclude portions of a compilation unit.

ANTLR

```
fragment PP_Conditional
    : PP_If_Section
    | PP_Elif_Section
    | PP_Else_Section
    | PP_Endif
    ;

fragment PP_If_Section
    : 'if' PP_Whitespace PP_Expression
    ;

fragment PP_Elif_Section
    : 'elif' PP_Whitespace PP_Expression
    ;

fragment PP_Else_Section
    : 'else'
    ;

fragment PP_Endif
    : 'endif'
    ;
```

Conditional compilation directives shall be written in groups consisting of, in order, a `#if` directive, zero or more `#elif` directives, zero or one `#else` directive, and a `#endif` directive. Between the directives are ***conditional sections*** of source code. Each section is controlled by the immediately preceding directive. A conditional section may itself contain nested conditional compilation directives provided these directives form complete groups.

At most one of the contained conditional sections is selected for normal lexical processing:

- The *PP_Expressions* of the `#if` and `#elif` directives are evaluated in order until one yields `true`. If an expression yields `true`, the conditional section following the corresponding directive is selected.

- If all *PP_Expressions* yield `false`, and if a `#else` directive is present, the conditional section following the `#else` directive is selected.
- Otherwise, no conditional section is selected.

The selected conditional section, if any, is processed as a normal *input_section*: the source code contained in the section shall adhere to the lexical grammar; tokens are generated from the source code in the section; and pre-processing directives in the section have the prescribed effects.

Any remaining conditional sections are skipped and no tokens, except those for pre-processing directives, are generated from the source code. Therefore skipped source code, except pre-processing directives, may be lexically incorrect. Skipped pre-processing directives shall be lexically correct but are not otherwise processed. Within a conditional section that is being skipped any nested conditional sections (contained in nested `#if...#endif` constructs) are also skipped.

> *Note*: The above grammar does not capture the allowance that the conditional sections between the pre-processing directives may be malformed lexically. Therefore the grammar is not ANTLR-ready as it only supports lexically correct input. *end note*

> *Example*: The following example illustrates how conditional compilation directives can nest:
>
> C#
>
> ```
> #define Debug // Debugging on
> #undef Trace // Tracing off
> class PurchaseTransaction
> {
>     void Commit()
>     {
> #if Debug
>         CheckConsistency();
>     #if Trace
>         WriteToLog(this.ToString());
>     #endif
> #endif
>         CommitHelper();
>     }
>     ...
> }
> ```
>
> Except for pre-processing directives, skipped source code is not subject to lexical analysis. For example, the following is valid despite the unterminated comment in

the `#else` section:

C#

```
#define Debug // Debugging on
class PurchaseTransaction
{
    void Commit()
    {
#if Debug
        CheckConsistency();
#else
        /* Do something else
#endif
    }
    ...
}
```

Note, however, that pre-processing directives are required to be lexically correct even in skipped sections of source code.

Pre-processing directives are not processed when they appear inside multi-line input elements. For example, the program:

C#

```
class Hello
{
    static void Main()
    {
        System.Console.WriteLine(@"hello,
#if Debug
        world
#else
        Nebraska
#endif
        ");
    }
}
```

results in the output:

Console

```
hello,
#if Debug
        world
#else
```

```
        Nebraska
#endif
```

In peculiar cases, the set of pre-processing directives that is processed might depend on the evaluation of the *pp_expression*. The example:

```
C#

#if X
    /*
#else
    /* */ class Q { }
#endif
```

always produces the same token stream (`class Q { }`), regardless of whether or not `X` is defined. If `X` is defined, the only processed directives are `#if` and `#endif`, due to the multi-line comment. If `X` is undefined, then three directives (`#if`, `#else`, `#endif`) are part of the directive set.

end example

6.5.6 Diagnostic directives

The diagnostic directives are used to generate explicitly error and warning messages that are reported in the same way as other compile-time errors and warnings.

```
ANTLR

fragment PP_Diagnostic
    : 'error' PP_Message?
    | 'warning' PP_Message?
    ;

fragment PP_Message
    : PP_Whitespace Input_Character*
    ;
```

Example: The example

```
C#

#if Debug && Retail
    #error A build can't be both debug and retail
#endif
class Test {...}
```

> produces a compile-time error (“A build can’t be both debug and retail”) if the conditional compilation symbols `Debug` and `Retail` are both defined. Note that a *PP_Message* can contain arbitrary text; specifically, it need not contain well-formed tokens, as shown by the single quote in the word `can't`.
>
> *end example*

6.5.7 Region directives

The region directives are used to mark explicitly regions of source code.

ANTLR

```
fragment PP_Region
    : PP_Start_Region
    | PP_End_Region
    ;

fragment PP_Start_Region
    : 'region' PP_Message?
    ;

fragment PP_End_Region
    : 'endregion' PP_Message?
    ;
```

No semantic meaning is attached to a region; regions are intended for use by the programmer or by automated tools to mark a section of source code. There shall be one `#endregion` directive matching every `#region` directive. The message specified in a `#region` or `#endregion` directive likewise has no semantic meaning; it merely serves to identify the region. Matching `#region` and `#endregion` directives may have different *PP_Message*s.

The lexical processing of a region:

C#

```
#region
...
#endregion
```

corresponds exactly to the lexical processing of a conditional compilation directive of the form:

C#

```
#if true
...
#endif
```

> *Note*: This means that a region can include one or more `#if`/.../`#endif`, or be contained with a conditional section within a `#if`/.../`#endif`; but a region cannot overlap with an just part of an `#if`/.../`#endif`, or start & end in different conditional sections. *end note*

6.5.8 Line directives

Line directives may be used to alter the line numbers and compilation unit names that are reported by the compiler in output such as warnings and errors. These values are also used by caller-info attributes (§22.5.6).

> *Note*: Line directives are most commonly used in meta-programming tools that generate C# source code from some other text input. *end note*

ANTLR

```
fragment PP_Line
    : 'line' PP_Whitespace PP_Line_Indicator
    ;

fragment PP_Line_Indicator
    : Decimal_Digit+ PP_Whitespace PP_Compilation_Unit_Name
    | Decimal_Digit+
    | DEFAULT
    | 'hidden'
    ;

fragment PP_Compilation_Unit_Name
    : '"' PP_Compilation_Unit_Name_Character+ '"'
    ;

fragment PP_Compilation_Unit_Name_Character
    // Any Input_Character except "
    : ~('\u000D' | '\u000A'   | '\u0085' | '\u2028' | '\u2029' | '#')
    ;
```

When no `#line` directives are present, the compiler reports true line numbers and compilation unit names in its output. When processing a `#line` directive that includes a *PP_Line_Indicator* that is not `default`, the compiler treats the line *after* the directive as having the given line number (and compilation unit name, if specified).

The maximum value allowed for `Decimal_Digit+` is implementation-defined.

A `#line default` directive undoes the effect of all preceding `#line` directives. The compiler reports true line information for subsequent lines, precisely as if no `#line` directives had been processed.

A `#line hidden` directive has no effect on the compilation unit and line numbers reported in error messages, or produced by use of `CallerLineNumberAttribute` (§22.5.6.2). It is intended to affect source-level debugging tools so that, when debugging, all lines between a `#line hidden` directive and the subsequent `#line` directive (that is not `#line hidden`) have no line number information, and are skipped entirely when stepping through code.

> *Note*: Although a *PP_Compilation_Unit_Name* might contain text that looks like an escape sequence, such text is not an escape sequence; in this context a '`\`' character simply designates an ordinary backslash character. *end note*

6.5.9 Pragma directives

The `#pragma` preprocessing directive is used to specify contextual information to a compiler.

> *Note*: For example, a compiler might provide `#pragma` directives that
>
> - Enable or disable particular warning messages when compiling subsequent code.
> - Specify which optimizations to apply to subsequent code.
> - Specify information to be used by a debugger.
>
> *end note*

ANTLR

```
fragment PP_Pragma
    : 'pragma' PP_Pragma_Text?
    ;

fragment PP_Pragma_Text
    : PP_Whitespace Input_Character*
    ;
```

The *Input_Character*s in the *PP_Pragma_Text* are interpreted by the compiler in an implementation-defined manner. The information supplied in a `#pragma` directive shall

not change program semantics. A `#pragma` directive shall only change compiler behavior that is outside the scope of this language specification. If the compiler cannot interpret the *Input_Characters*, the compiler can produce a warning; however, it shall not produce a compile-time error.

> *Note*: *PP_Pragma_Text* can contain arbitrary text; specifically, it need not contain well-formed tokens. *end note*

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

7 Basic concepts

Article • 02/07/2024

7.1 Application startup

A program may be compiled either as a ***class library*** to be used as part of other applications, or as an ***application*** that may be started directly. The mechanism for determining this mode of compilation is implementation-specific and external to this specification.

A program compiled as an application shall contain at least one method qualifying as an entry point by satisfying the following requirements:

- It shall have the name `Main`.
- It shall be `static`.
- It shall not be generic.
- It shall be declared in a non-generic type. If the type declaring the method is a nested type, none of its enclosing types may be generic.
- It may have the `async` modifier provided the method's return type is `System.Threading.Tasks.Task` or `System.Threading.Tasks.Task<int>`.
- The return type shall be `void`, `int`, `System.Threading.Tasks.Task`, or `System.Threading.Tasks.Task<int>`.
- It shall not be a partial method (§15.6.9) without an implementation.
- The formal parameter list shall either be empty, or have a single value parameter of type `string[]`.

> *Note*: Methods with the `async` modifier must have exactly one of the two return types specified above in order to qualify as an entry point. An `async void` method, or an `async` method returning a different awaitable type such as `ValueTask` or `ValueTask<int>` does not qualify as an entry point. *end note*

If more than one method qualifying as an entry point is declared within a program, an external mechanism may be used to specify which method is deemed to be the actual entry point for the application. If a qualifying method having a return type of `int` or `void` is found, any qualifying method having a return type of `System.Threading.Tasks.Task` or `System.Threading.Tasks.Task<int>` is not considered an entry point method. It is a compile-time error for a program to be compiled as an application without exactly one entry point. A program compiled as a class library may

contain methods that would qualify as application entry points, but the resulting library has no entry point.

Ordinarily, the declared accessibility (§7.5.2) of a method is determined by the access modifiers (§15.3.6) specified in its declaration, and similarly the declared accessibility of a type is determined by the access modifiers specified in its declaration. In order for a given method of a given type to be callable, both the type and the member shall be accessible. However, the application entry point is a special case. Specifically, the execution environment can access the application's entry point regardless of its declared accessibility and regardless of the declared accessibility of its enclosing type declarations.

When the entry point method has a return type of `System.Threading.Tasks.Task` or `System.Threading.Tasks.Task<int>`, the compiler synthesizes a synchronous entry-point method that calls the corresponding `Main` method. The synthesized method has parameters and return types based on the `Main` method:

- The formal parameter list of the synthesized method is the same as the formal parameter list of the `Main` method
- If the return type of the `Main` method is `System.Threading.Tasks.Task`, the return type of the synthesized method is `void`
- If the return type of the `Main` method is `System.Threading.Tasks.Task<int>`, the return type of the synthesized method is `int`

Execution of the synthesized method proceeds as follows:

- The synthesized method calls the `Main` method, passing its `string[]` parameter value as an argument if the `Main` method has such a parameter.
- If the `Main` method throws an exception, the exception is propagated by the synthesized method.
- Otherwise, the synthesized entry point waits for the returned task to complete, calling `GetAwaiter().GetResult()` on the task, using either the parameterless instance method or the extension method described by §C.3. If the task fails, `GetResult()` will throw an exception, and this exception is propagated by the synthesized method.
- For a `Main` method with a return type of `System.Threading.Tasks.Task<int>`, if the task completes successfully, the `int` value returned by `GetResult()` is returned from the synthesized method.

The ***effective entry point*** of an application is the entry point declared within the program, or the synthesized method if one is required as described above. The return

type of the effective entry point is therefore always `void` or `int`.

When an application is run, a new ***application domain*** is created. Several different instantiations of an application may exist on the same machine at the same time, and each has its own application domain. An application domain enables application isolation by acting as a container for application state. An application domain acts as a container and boundary for the types defined in the application and the class libraries it uses. Types loaded into one application domain are distinct from the same types loaded into another application domain, and instances of objects are not directly shared between application domains. For instance, each application domain has its own copy of static variables for these types, and a static constructor for a type is run at most once per application domain. Implementations are free to provide implementation-specific policy or mechanisms for the creation and destruction of application domains.

Application startup occurs when the execution environment calls the application's effective entry point. If the effective entry point declares a parameter, then during application startup, the implementation shall ensure that the initial value of that parameter is a non-null reference to a string array. This array shall consist of non-null references to strings, called ***application parameters***, which are given implementation-defined values by the host environment prior to application startup. The intent is to supply to the application information determined prior to application startup from elsewhere in the hosted environment.

> *Note*: On systems supporting a command line, application parameters correspond to what are generally known as command-line arguments. *end note*

If the effective entry point's return type is `int`, the return value from the method invocation by the execution environment is used in application termination (§7.2).

Other than the situations listed above, entry point methods behave like those that are not entry points in every respect. In particular, if the entry point is invoked at any other point during the application's lifetime, such as by regular method invocation, there is no special handling of the method: if there is a parameter, it may have an initial value of `null`, or a non-`null` value referring to an array that contains null references. Likewise, the return value of the entry point has no special significance other than in the invocation from the execution environment.

7.2 Application termination

Application termination returns control to the execution environment.

If the return type of the application's effective entry point method is `int` and execution completes without resulting in an exception, the value of the `int` returned serves as the application's ***termination status code***. The purpose of this code is to allow communication of success or failure to the execution environment. If the return type of the effective entry point method is `void` and execution completes without resulting in an exception, the termination status code is `0`.

If the effective entry point method terminates due to an exception (§21.4), the exit code is implementation-specific. Additionally, the implementation may provide alternative APIs for specifying the exit code.

Whether or not finalizers (§15.13) are run as part of application termination is implementation-specific.

> *Note*: The .NET Framework implementation makes every reasonable effort to call finalizers (§15.13) for all of its objects that have not yet been garbage collected, unless such cleanup has been suppressed (by a call to the library method `GC.SuppressFinalize`, for example). *end note*

7.3 Declarations

Declarations in a C# program define the constituent elements of the program. C# programs are organized using namespaces. These are introduced using namespace declarations (§14), which can contain type declarations and nested namespace declarations. Type declarations (§14.7) are used to define classes (§15), structs (§16), interfaces (§18), enums (§19), and delegates (§20). The kinds of members permitted in a type declaration depend on the form of the type declaration. For instance, class declarations can contain declarations for constants (§15.4), fields (§15.5), methods (§15.6), properties (§15.7), events (§15.8), indexers (§15.9), operators (§15.10), instance constructors (§15.11), static constructors (§15.12), finalizers (§15.13), and nested types (§15.3.9).

A declaration defines a name in the ***declaration space*** to which the declaration belongs. It is a compile-time error to have two or more declarations that introduce members with the same name in a declaration space, except in the following cases:

- Two or more namespace declarations with the same name are allowed in the same declaration space. Such namespace declarations are aggregated to form a single logical namespace and share a single declaration space.
- Declarations in separate programs but in the same namespace declaration space are allowed to share the same name.

> *Note*: However, these declarations could introduce ambiguities if included in the same application. *end note*

- Two or more methods with the same name but distinct signatures are allowed in the same declaration space (§7.6).
- Two or more type declarations with the same name but distinct numbers of type parameters are allowed in the same declaration space (§7.8.2).
- Two or more type declarations with the partial modifier in the same declaration space may share the same name, same number of type parameters and same classification (class, struct or interface). In this case, the type declarations contribute to a single type and are themselves aggregated to form a single declaration space (§15.2.7).
- A namespace declaration and a type declaration in the same declaration space can share the same name as long as the type declaration has at least one type parameter (§7.8.2).

There are several different types of declaration spaces, as described in the following.

- Within all compilation units of a program, *namespace_member_declaration*s with no enclosing *namespace_declaration* are members of a single combined declaration space called the ***global declaration space***.
- Within all compilation units of a program, *namespace_member_declaration*s within *namespace_declaration*s that have the same fully qualified namespace name are members of a single combined declaration space.
- Each *compilation_unit* and *namespace_body* has an ***alias declaration space***. Each *extern_alias_directive* and *using_alias_directive* of the *compilation_unit* or *namespace_body* contributes a member to the alias declaration space (§14.5.2).
- Each non-partial class, struct, or interface declaration creates a new declaration space. Each partial class, struct, or interface declaration contributes to a declaration space shared by all matching parts in the same program (§16.2.4). Names are introduced into this declaration space through *class_member_declaration*s, *struct_member_declaration*s, *interface_member_declaration*s, or *type_parameter*s. Except for overloaded instance constructor declarations and static constructor declarations, a class or struct cannot contain a member declaration with the same name as the class or struct. A class, struct, or interface permits the declaration of overloaded methods and indexers. Furthermore, a class or struct permits the declaration of overloaded instance constructors and operators. For example, a class, struct, or interface may contain multiple method declarations with the same name, provided these method declarations differ in their signature (§7.6). Note that base classes do not contribute to the declaration space of a class, and base interfaces do not contribute to the declaration space of an interface. Thus, a

derived class or interface is allowed to declare a member with the same name as an inherited member. Such a member is said to ***hide*** the inherited member.

- Each delegate declaration creates a new declaration space. Names are introduced into this declaration space through formal parameters (*fixed_parameter*s and *parameter_array*s) and *type_parameter*s.
- Each enumeration declaration creates a new declaration space. Names are introduced into this declaration space through *enum_member_declarations*.
- Each method declaration, property declaration, property accessor declaration, indexer declaration, indexer accessor declaration, operator declaration, instance constructor declaration, anonymous function, and local function creates a new declaration space called a ***local variable declaration space***. Names are introduced into this declaration space through formal parameters (*fixed_parameter*s and *parameter_array*s) and *type_parameter*s. The set accessor for a property or an indexer introduces the name `value` as a formal parameter.
- Additional local variable declaration spaces may occur within member declarations, anonymous functions and local functions. Names are introduced into these declaration spaces through *pattern*s, *declaration_expression*s, *declaration_statement*s and *exception_specifier*s. Local variable declaration spaces may be nested, but it is an error for a local variable declaration space and a nested local variable declaration space to contain elements with the same name. Thus, within a nested declaration space it is not possible to declare a local variable, local function or constant with the same name as a parameter, type parameter, local variable, local function or constant in an enclosing declaration space. It is possible for two declaration spaces to contain elements with the same name as long as neither declaration space contains the other. Local declaration spaces are created by the following constructs:
  - Each *variable_initializer* in a field and property declaration introduces its own local variable declaration space, that is not nested within any other local variable declaration space.
  - The body of a function member, anonymous function, or local function, if any, creates a local variable declaration space that is considered to be nested within the function's local variable declaration space.
  - Each *constructor_initializer* creates a local variable declaration space nested within the instance constructor declaration. The local variable declaration space for the constructor body is in turn nested within this local variable declaration space.
  - Each *block*, *switch_block*, *specific_catch_clause*, *iteration_statement* and *using_statement* creates a nested local variable declaration space.
  - Each *embedded_statement* that is not directly part of a *statement_list* creates a nested local variable declaration space.

    - Each *switch_section* creates a nested local variable declaration space. However, variables declared directly within the *statement_list* of the *switch_section* (but not within a nested local variable declaration space inside the *statement_list*) are added directly to the local variable declaration space of the enclosing *switch_block*, instead of that of the *switch_section*.
    - The syntactic translation of a *query_expression* (§12.20.3) may introduce one or more lambda expressions. As anonymous functions, each of these creates a local variable declaration space as described above.
- Each *block* or *switch_block* creates a separate declaration space for labels. Names are introduced into this declaration space through *labeled_statement*s, and the names are referenced through *goto_statement*s. The ***label declaration space*** of a block includes any nested blocks. Thus, within a nested block it is not possible to declare a label with the same name as a label in an enclosing block.

> *Note*: The fact that variables declared directly within a *switch_section* are added to the local variable declaration space of the *switch_block* instead of the *switch_section* can lead to surprising code. In the example below, the local variable `y` is in scope within the switch section for the default case, despite the declaration appearing in the switch section for case 0. The local variable `z` is not in scope within the switch section for the default case, as it is introduced in the local variable declaration space for the switch section in which the declaration occurs.
>
> C#
>
> ```
> int x = 1;
> switch (x)
> {
>     case 0:
>         int y;
>         break;
>     case var z when z < 10:
>         break;
>     default:
>         y = 10;
>         // Valid: y is in scope
>         Console.WriteLine(x + y);
>         // Invalid: z is not scope
>         Console.WriteLine(x + z);
>         break;
> }
> ```
>
> *end note*

The textual order in which names are declared is generally of no significance. In particular, textual order is not significant for the declaration and use of namespaces,

constants, methods, properties, events, indexers, operators, instance constructors, finalizers, static constructors, and types. Declaration order is significant in the following ways:

- Declaration order for field declarations determines the order in which their initializers (if any) are executed (§15.5.6.2, §15.5.6.3).
- Local variables shall be defined before they are used (§7.7).
- Declaration order for enum member declarations (§19.4) is significant when *constant_expression* values are omitted.

> *Example*: The declaration space of a namespace is "open ended", and two namespace declarations with the same fully qualified name contribute to the same declaration space. For example
>
> ```
> C#
> ```
>
> ```
> namespace Megacorp.Data
> {
>     class Customer
>     {
>         ...
>     }
> }
>
> namespace Megacorp.Data
> {
>     class Order
>     {
>         ...
>     }
> }
> ```
>
> The two namespace declarations above contribute to the same declaration space, in this case declaring two classes with the fully qualified names `Megacorp.Data.Customer` and `Megacorp.Data.Order`. Because the two declarations contribute to the same declaration space, it would have caused a compile-time error if each contained a declaration of a class with the same name.
>
> *end example*

> *Note*: As specified above, the declaration space of a block includes any nested blocks. Thus, in the following example, the `F` and `G` methods result in a compile-time error because the name `i` is declared in the outer block and cannot be redeclared in the inner block. However, the `H` and `I` methods are valid since the two `i`'s are declared in separate non-nested blocks.

C#

```
class A
{
    void F()
    {
        int i = 0;
        if (true)
        {
            int i = 1;
        }
    }

    void G()
    {
        if (true)
        {
            int i = 0;
        }
        int i = 1;
    }

    void H()
    {
        if (true)
        {
            int i = 0;
        }
        if (true)
        {
            int i = 1;
        }
    }

    void I()
    {
        for (int i = 0; i < 10; i++)
        {
            H();
        }
        for (int i = 0; i < 10; i++)
        {
            H();
        }
    }
}
```

end note

7.4 Members

7.4.1 General

Namespaces and types have ***members***.

> *Note*: The members of an entity are generally available through the use of a qualified name that starts with a reference to the entity, followed by a "`.`" token, followed by the name of the member. *end note*

Members of a type are either declared in the type declaration or ***inherited*** from the base class of the type. When a type inherits from a base class, all members of the base class, except instance constructors, finalizers, and static constructors become members of the derived type. The declared accessibility of a base class member does not control whether the member is inherited—inheritance extends to any member that isn't an instance constructor, static constructor, or finalizer.

> *Note*: However, an inherited member might not be accessible in a derived type, for example because of its declared accessibility (§7.5.2). *end note*

7.4.2 Namespace members

Namespaces and types that have no enclosing namespace are members of the ***global namespace***. This corresponds directly to the names declared in the global declaration space.

Namespaces and types declared within a namespace are members of that namespace. This corresponds directly to the names declared in the declaration space of the namespace.

Namespaces have no access restrictions. It is not possible to declare private, protected, or internal namespaces, and namespace names are always publicly accessible.

7.4.3 Struct members

The members of a struct are the members declared in the struct and the members inherited from the struct's direct base class `System.ValueType` and the indirect base class `object`.

The members of a simple type correspond directly to the members of the struct type aliased by the simple type (§8.3.5).

7.4.4 Enumeration members

The members of an enumeration are the constants declared in the enumeration and the members inherited from the enumeration's direct base class `System.Enum` and the indirect base classes `System.ValueType` and `object`.

7.4.5 Class members

The members of a class are the members declared in the class and the members inherited from the base class (except for class `object` which has no base class). The members inherited from the base class include the constants, fields, methods, properties, events, indexers, operators, and types of the base class, but not the instance constructors, finalizers, and static constructors of the base class. Base class members are inherited without regard to their accessibility.

A class declaration may contain declarations of constants, fields, methods, properties, events, indexers, operators, instance constructors, finalizers, static constructors, and types.

The members of `object` (§8.2.3) and `string` (§8.2.5) correspond directly to the members of the class types they alias.

7.4.6 Interface members

The members of an interface are the members declared in the interface and in all base interfaces of the interface.

> *Note*: The members in class `object` are not, strictly speaking, members of any interface (§18.4). However, the members in class `object` are available via member lookup in any interface type (§12.5). *end note*

7.4.7 Array members

The members of an array are the members inherited from class `System.Array`.

7.4.8 Delegate members

A delegate inherits members from class `System.Delegate`. Additionally, it contains a method named `Invoke` with the same return type and formal parameter list specified in its declaration (§20.2). An invocation of this method shall behave identically to a delegate invocation (§20.6) on the same delegate instance.

An implementation may provide additional members, either through inheritance or directly in the delegate itself.

7.5 Member access

7.5.1 General

Declarations of members allow control over member access. The accessibility of a member is established by the declared accessibility (§7.5.2) of the member combined with the accessibility of the immediately containing type, if any.

When access to a particular member is allowed, the member is said to be ***accessible***. Conversely, when access to a particular member is disallowed, the member is said to be ***inaccessible***. Access to a member is permitted when the textual location in which the access takes place is included in the accessibility domain (§7.5.3) of the member.

7.5.2 Declared accessibility

The ***declared accessibility*** of a member can be one of the following:

- Public, which is selected by including a `public` modifier in the member declaration. The intuitive meaning of `public` is "access not limited".
- Protected, which is selected by including a `protected` modifier in the member declaration. The intuitive meaning of `protected` is "access limited to the containing class or types derived from the containing class".
- Internal, which is selected by including an `internal` modifier in the member declaration. The intuitive meaning of `internal` is "access limited to this assembly".
- Protected internal, which is selected by including both a `protected` and an `internal` modifier in the member declaration. The intuitive meaning of `protected internal` is "accessible within this assembly as well as types derived from the containing class".
- Private protected, which is selected by including both a `private` and a `protected` modifier in the member declaration. The intuitive meaning of `private protected` is "accessible within this assembly by the containing class and types derived from the containing class."
- Private, which is selected by including a `private` modifier in the member declaration. The intuitive meaning of `private` is "access limited to the containing type".

Depending on the context in which a member declaration takes place, only certain types of declared accessibility are permitted. Furthermore, when a member declaration does not include any access modifiers, the context in which the declaration takes place determines the default declared accessibility.

- Namespaces implicitly have `public` declared accessibility. No access modifiers are allowed on namespace declarations.
- Types declared directly in compilation units or namespaces (as opposed to within other types) can have `public` or `internal` declared accessibility and default to `internal` declared accessibility.
- Class members can have any of the permitted kinds of declared accessibility and default to `private` declared accessibility.

  > *Note*: A type declared as a member of a class can have any of the permitted kinds of declared accessibility, whereas a type declared as a member of a namespace can have only `public` or `internal` declared accessibility. *end note*

- Struct members can have `public`, `internal`, or `private` declared accessibility and default to `private` declared accessibility because structs are implicitly sealed. Struct members introduced in a `struct` (that is, not inherited by that struct) cannot have `protected`, `protected internal`, or `private protected` declared accessibility.

  > *Note*: A type declared as a member of a struct can have `public`, `internal`, or `private` declared accessibility, whereas a type declared as a member of a namespace can have only `public` or `internal` declared accessibility. *end note*

- Interface members implicitly have `public` declared accessibility. No access modifiers are allowed on interface member declarations.
- Enumeration members implicitly have `public` declared accessibility. No access modifiers are allowed on enumeration member declarations.

7.5.3 Accessibility domains

The ***accessibility domain*** of a member consists of the (possibly disjoint) sections of program text in which access to the member is permitted. For purposes of defining the accessibility domain of a member, a member is said to be ***top-level*** if it is not declared within a type, and a member is said to be ***nested*** if it is declared within another type. Furthermore, the ***program text*** of a program is defined as all text contained in all compilation units of the program, and the program text of a type is defined as all text

contained in the *type_declarations* of that type (including, possibly, types that are nested within the type).

The accessibility domain of a predefined type (such as `object`, `int`, or `double`) is unlimited.

The accessibility domain of a top-level unbound type `T` (§8.4.4) that is declared in a program `P` is defined as follows:

- If the declared accessibility of `T` is public, the accessibility domain of `T` is the program text of `P` and any program that references `P`.
- If the declared accessibility of `T` is internal, the accessibility domain of `T` is the program text of `P`.

> *Note*: From these definitions, it follows that the accessibility domain of a top-level unbound type is always at least the program text of the program in which that type is declared. *end note*

The accessibility domain for a constructed type $T<A_1, ..., A_e>$ is the intersection of the accessibility domain of the unbound generic type `T` and the accessibility domains of the type arguments $A_1, ..., A_e$.

The accessibility domain of a nested member `M` declared in a type `T` within a program `P`, is defined as follows (noting that `M` itself might possibly be a type):

- If the declared accessibility of `M` is `public`, the accessibility domain of `M` is the accessibility domain of `T`.
- If the declared accessibility of `M` is `protected internal`, let `D` be the union of the program text of `P` and the program text of any type derived from `T`, which is declared outside `P`. The accessibility domain of `M` is the intersection of the accessibility domain of `T` with `D`.
- If the declared accessibility of `M` is `private protected`, let `D` be the intersection of the program text of `P` and the program text of `T` and any type derived from `T`. The accessibility domain of `M` is the intersection of the accessibility domain of `T` with `D`.
- If the declared accessibility of `M` is `protected`, let `D` be the union of the program text of `T` and the program text of any type derived from `T`. The accessibility domain of `M` is the intersection of the accessibility domain of `T` with `D`.
- If the declared accessibility of `M` is `internal`, the accessibility domain of `M` is the intersection of the accessibility domain of `T` with the program text of `P`.

- If the declared accessibility of `M` is `private`, the accessibility domain of `M` is the program text of `T`.

> *Note*: From these definitions it follows that the accessibility domain of a nested member is always at least the program text of the type in which the member is declared. Furthermore, it follows that the accessibility domain of a member is never more inclusive than the accessibility domain of the type in which the member is declared. *end note*

> *Note*: In intuitive terms, when a type or member `M` is accessed, the following steps are evaluated to ensure that the access is permitted:
>
> - First, if `M` is declared within a type (as opposed to a compilation unit or a namespace), a compile-time error occurs if that type is not accessible.
> - Then, if `M` is `public`, the access is permitted.
> - Otherwise, if `M` is `protected internal`, the access is permitted if it occurs within the program in which `M` is declared, or if it occurs within a class derived from the class in which `M` is declared and takes place through the derived class type (§7.5.4).
> - Otherwise, if `M` is `protected`, the access is permitted if it occurs within the class in which `M` is declared, or if it occurs within a class derived from the class in which `M` is declared and takes place through the derived class type (§7.5.4).
> - Otherwise, if `M` is `internal`, the access is permitted if it occurs within the program in which `M` is declared.
> - Otherwise, if `M` is `private`, the access is permitted if it occurs within the type in which `M` is declared.
> - Otherwise, the type or member is inaccessible, and a compile-time error occurs. *end note*

> *Example*: In the following code
>
> C#
>
> ```
> public class A
> {
>     public static int X;
>     internal static int Y;
>     private static int Z;
> }
>
> internal class B
> {
> ```

```
    public static int X;
    internal static int Y;
    private static int Z;

    public class C
    {
        public static int X;
        internal static int Y;
        private static int Z;
    }

    private class D
    {
        public static int X;
        internal static int Y;
        private static int Z;
    }
}
```

the classes and members have the following accessibility domains:

- The accessibility domain of `A` and `A.X` is unlimited.
- The accessibility domain of `A.Y`, `B`, `B.X`, `B.Y`, `B.C`, `B.C.X`, and `B.C.Y` is the program text of the containing program.
- The accessibility domain of `A.Z` is the program text of `A`.
- The accessibility domain of `B.Z` and `B.D` is the program text of `B`, including the program text of `B.C` and `B.D`.
- The accessibility domain of `B.C.Z` is the program text of `B.C`.
- The accessibility domain of `B.D.X` and `B.D.Y` is the program text of `B`, including the program text of `B.C` and `B.D`.
- The accessibility domain of `B.D.Z` is the program text of `B.D`. As the example illustrates, the accessibility domain of a member is never larger than that of a containing type. For example, even though all `X` members have public declared accessibility, all but `A.X` have accessibility domains that are constrained by a containing type.

end example

As described in §7.4, all members of a base class, except for instance constructors, finalizers, and static constructors, are inherited by derived types. This includes even private members of a base class. However, the accessibility domain of a private member includes only the program text of the type in which the member is declared.

Example: In the following code

```C#
class A
{
    int x;

    static void F(B b)
    {
        b.x = 1;          // Ok
    }
}

class B : A
{
    static void F(B b)
    {
        b.x = 1;          // Error, x not accessible
    }
}
```

the `B` class inherits the private member `x` from the `A` class. Because the member is private, it is only accessible within the *class_body* of `A`. Thus, the access to `b.x` succeeds in the `A.F` method, but fails in the `B.F` method.

end example

7.5.4 Protected access

When a `protected` or `private protected` instance member is accessed outside the program text of the class in which it is declared, and when a `protected internal` instance member is accessed outside the program text of the program in which it is declared, the access shall take place within a class declaration that derives from the class in which it is declared. Furthermore, the access is required to take place *through* an instance of that derived class type or a class type constructed from it. This restriction prevents one derived class from accessing protected members of other derived classes, even when the members are inherited from the same base class.

Let `B` be a base class that declares a protected instance member `M`, and let `D` be a class that derives from `B`. Within the *class_body* of `D`, access to `M` can take one of the following forms:

- An unqualified *type_name* or *primary_expression* of the form `M`.
- A *primary_expression* of the form `E.M`, provided the type of `E` is `T` or a class derived from `T`, where `T` is the class `D`, or a class type constructed from `D`.
- A *primary_expression* of the form `base.M`.

- A *primary_expression* of the form `base[` *argument_list* `]`.

In addition to these forms of access, a derived class can access a protected instance constructor of a base class in a *constructor_initializer* (§15.11.2).

> *Example*: In the following code
>
> C#
>
> ```
> public class A
> {
>     protected int x;
>
>     static void F(A a, B b)
>     {
>         a.x = 1; // Ok
>         b.x = 1; // Ok
>     }
> }
>
> public class B : A
> {
>     static void F(A a, B b)
>     {
>         a.x = 1; // Error, must access through instance of B
>         b.x = 1; // Ok
>     }
> }
> ```
>
> within `A`, it is possible to access `x` through instances of both `A` and `B`, since in either case the access takes place *through* an instance of `A` or a class derived from `A`. However, within `B`, it is not possible to access `x` through an instance of `A`, since `A` does not derive from `B`.
>
> *end example*

> *Example*:
>
> C#
>
> ```
> class C<T>
> {
>     protected T x;
> }
>
> class D<T> : C<T>
> {
>     static void F()
>     {
> ```

```
        D<T> dt = new D<T>();
        D<int> di = new D<int>();
        D<string> ds = new D<string>();
        dt.x = default(T);
        di.x = 123;
        ds.x = "test";
    }
}
```

Here, the three assignments to `x` are permitted because they all take place through instances of class types constructed from the generic type.

end example

Note: The accessibility domain (§7.5.3) of a protected member declared in a generic class includes the program text of all class declarations derived from any type constructed from that generic class. In the example:

C#

```
class C<T>
{
    protected static T x;
}

class D : C<string>
{
    static void Main()
    {
        C<int>.x = 5;
    }
}
```

the reference to `protected` member `C<int>.x` in `D` is valid even though the class `D` derives from `C<string>`. *end note*

7.5.5 Accessibility constraints

Several constructs in the C# language require a type to be at least as accessible as a member or another type. A type `T` is said to be at least as accessible as a member or type `M` if the accessibility domain of `T` is a superset of the accessibility domain of `M`. In other words, `T` is at least as accessible as `M` if `T` is accessible in all contexts in which `M` is accessible.

The following accessibility constraints exist:

- The direct base class of a class type shall be at least as accessible as the class type itself.
- The explicit base interfaces of an interface type shall be at least as accessible as the interface type itself.
- The return type and parameter types of a delegate type shall be at least as accessible as the delegate type itself.
- The type of a constant shall be at least as accessible as the constant itself.
- The type of a field shall be at least as accessible as the field itself.
- The return type and parameter types of a method shall be at least as accessible as the method itself.
- The type of a property shall be at least as accessible as the property itself.
- The type of an event shall be at least as accessible as the event itself.
- The type and parameter types of an indexer shall be at least as accessible as the indexer itself.
- The return type and parameter types of an operator shall be at least as accessible as the operator itself.
- The parameter types of an instance constructor shall be at least as accessible as the instance constructor itself.
- An interface or class type constraint on a type parameter shall be at least as accessible as the member which declares the constraint.

> *Example*: In the following code
>
> ```
> C#
> ```
>
> ```
> class A {...}
> public class B: A {...}
> ```
>
> the `B` class results in a compile-time error because `A` is not at least as accessible as `B`.
>
> *end example*

> *Example*: Likewise, in the following code
>
> ```
> C#
> ```
>
> ```
> class A {...}
>
> public class B
> {
>     A F() {...}
>     internal A G() {...}
> ```

```
    public A H() {...}
}
```

the `H` method in `B` results in a compile-time error because the return type `A` is not at least as accessible as the method.

end example

7.6 Signatures and overloading

Methods, instance constructors, indexers, and operators are characterized by their ***signatures***:

- The signature of a method consists of the name of the method, the number of type parameters, and the type and parameter-passing mode of each of its formal parameters, considered in the order left to right. For these purposes, any type parameter of the method that occurs in the type of a formal parameter is identified not by its name, but by its ordinal position in the type parameter list of the method. The signature of a method specifically does not include the return type, parameter names, type parameter names, type parameter constraints, the `params` or `this` parameter modifiers, nor whether parameters are required or optional.
- The signature of an instance constructor consists of the type and parameter-passing mode of each of its formal parameters, considered in the order left to right. The signature of an instance constructor specifically does not include the `params` modifier that may be specified for the right-most parameter, nor whether parameters are required or optional.
- The signature of an indexer consists of the type of each of its formal parameters, considered in the order left to right. The signature of an indexer specifically does not include the element type, nor does it include the `params` modifier that may be specified for the right-most parameter, nor whether parameters are required or optional.
- The signature of an operator consists of the name of the operator and the type of each of its formal parameters, considered in the order left to right. The signature of an operator specifically does not include the result type.
- The signature of a conversion operator consists of the source type and the target type. The implicit or explicit classification of a conversion operator is not part of the signature.
- Two signatures of the same member kind (method, instance constructor, indexer or operator) are considered to be the *same signatures* if they have the same name, number of type parameters, number of parameters, and parameter-passing

modes, and an identity conversion exists between the types of their corresponding parameters (§10.2.2).

Signatures are the enabling mechanism for ***overloading*** of members in classes, structs, and interfaces:

- Overloading of methods permits a class, struct, or interface to declare multiple methods with the same name, provided their signatures are unique within that class, struct, or interface.
- Overloading of instance constructors permits a class or struct to declare multiple instance constructors, provided their signatures are unique within that class or struct.
- Overloading of indexers permits a class, struct, or interface to declare multiple indexers, provided their signatures are unique within that class, struct, or interface.
- Overloading of operators permits a class or struct to declare multiple operators with the same name, provided their signatures are unique within that class or struct.

Although `in`, `out`, and `ref` parameter modifiers are considered part of a signature, members declared in a single type cannot differ in signature solely by `in`, `out`, and `ref`. A compile-time error occurs if two members are declared in the same type with signatures that would be the same if all parameters in both methods with `out` or `in` modifiers were changed to `ref` modifiers. For other purposes of signature matching (e.g., hiding or overriding), `in`, `out`, and `ref` are considered part of the signature and do not match each other.

> *Note*: This restriction is to allow C# programs to be easily translated to run on the Common Language Infrastructure (CLI), which does not provide a way to define methods that differ solely in `in`, `out`, and `ref`. *end note*

The types `object` and `dynamic` are not distinguished when comparing signatures. Therefore members declared in a single type whose signatures differ only by replacing `object` with `dynamic` are not allowed.

> *Example*: The following example shows a set of overloaded method declarations along with their signatures.

```
C#

interface ITest
{
    void F();                   // F()
    void F(int x);              // F(int)
```

```
    void F(ref int x);          // F(ref int)
    void F(out int x);          // F(out int) error
    void F(object o);           // F(object)
    void F(dynamic d);          // error.
    void F(int x, int y);       // F(int, int)
    int F(string s);            // F(string)
    int F(int x);               // F(int) error
    void F(string[] a);         // F(string[])
    void F(params string[] a);  // F(string[]) error
    void F<S>(S s);             // F<0>(0)
    void F<T>(T t);             // F<0>(0) error
    void F<S,T>(S s);           // F<0,1>(0)
    void F<T,S>(S s);           // F<0,1>(1) ok
}
```

Note that any `in`, `out`, and `ref` parameter modifiers (§15.6.2) are part of a signature. Thus, `F(int)`, `F(in int)`, `F(out int)` , and `F(ref int)` are all unique signatures. However, `F(in int)`, `F(out int)` , and `F(ref int)` cannot be declared within the same interface because their signatures differ solely by `in`, `out`, and `ref`. Also, note that the return type and the `params` modifier are not part of a signature, so it is not possible to overload solely based on return type or on the inclusion or exclusion of the `params` modifier. As such, the declarations of the methods `F(int)` and `F(params string[])` identified above, result in a compile-time error. *end example*

7.7 Scopes

7.7.1 General

The ***scope*** of a name is the region of program text within which it is possible to refer to the entity declared by the name without qualification of the name. Scopes can be *nested*, and an inner scope may redeclare the meaning of a name from an outer scope. (This does not, however, remove the restriction imposed by §7.3 that within a nested block it is not possible to declare a local variable or local constant with the same name as a local variable or local constant in an enclosing block.) The name from the outer scope is then said to be ***hidden*** in the region of program text covered by the inner scope, and access to the outer name is only possible by qualifying the name.

- The scope of a namespace member declared by a *namespace_member_declaration* (§14.6) with no enclosing *namespace_declaration* is the entire program text.
- The scope of a namespace member declared by a *namespace_member_declaration* within a *namespace_declaration* whose fully qualified name is `N`, is the

namespace_body of every *namespace_declaration* whose fully qualified name is `N` or starts with `N`, followed by a period.

- The scope of a name defined by an *extern_alias_directive* (§14.4) extends over the *using_directive*s, *global_attributes* and *namespace_member_declaration*s of its immediately containing *compilation_unit* or *namespace_body*. An *extern_alias_directive* does not contribute any new members to the underlying declaration space. In other words, an *extern_alias_directive* is not transitive, but, rather, affects only the *compilation_unit* or *namespace_body* in which it occurs.

- The scope of a name defined or imported by a *using_directive* (§14.5) extends over the *global_attributes* and *namespace_member_declaration*s of the *compilation_unit* or *namespace_body* in which the *using_directive* occurs. A *using_directive* may make zero or more namespace or type names available within a particular *compilation_unit* or *namespace_body*, but does not contribute any new members to the underlying declaration space. In other words, a *using_directive* is not transitive but rather affects only the *compilation_unit* or *namespace_body* in which it occurs.

- The scope of a type parameter declared by a *type_parameter_list* on a *class_declaration* (§15.2) is the *class_base*, *type_parameter_constraints_clauses*, and *class_body* of that *class_declaration*.

  > *Note*: Unlike members of a class, this scope does not extend to derived classes. *end note*

- The scope of a type parameter declared by a *type_parameter_list* on a *struct_declaration* (§16.2) is the *struct_interfaces*, *type_parameter_constraints_clause*s, and *struct_body* of that *struct_declaration*.

- The scope of a type parameter declared by a *type_parameter_list* on an *interface_declaration* (§18.2) is the *interface_base*, *type_parameter_constraints_clause*s, and *interface_body* of that *interface_declaration*.

- The scope of a type parameter declared by a *type_parameter_list* on a *delegate_declaration* (§20.2) is the *return_type*, *formal_parameter_list*, and *type_parameter_constraints_clause*s of that *delegate_declaration*.

- The scope of a type parameter declared by a *type_parameter_list* on a *method_declaration* (§15.6.1) is the *method_declaration*.

- The scope of a member declared by a *class_member_declaration* (§15.3.1) is the *class_body* in which the declaration occurs. In addition, the scope of a class

member extends to the *class_body* of those derived classes that are included in the accessibility domain (§7.5.3) of the member.

- The scope of a member declared by a *struct_member_declaration* (§16.3) is the *struct_body* in which the declaration occurs.
- The scope of a member declared by an *enum_member_declaration* (§19.4) is the *enum_body* in which the declaration occurs.
- The scope of a parameter declared in a *method_declaration* (§15.6) is the *method_body* or *ref_method_body* of that *method_declaration*.
- The scope of a parameter declared in an *indexer_declaration* (§15.9) is the *indexer_body* of that *indexer_declaration*.
- The scope of a parameter declared in an *operator_declaration* (§15.10) is the *operator_body* of that *operator_declaration*.
- The scope of a parameter declared in a *constructor_declaration* (§15.11) is the *constructor_initializer* and *block* of that *constructor_declaration*.
- The scope of a parameter declared in a *lambda_expression* (§12.19) is the *lambda_expression_body* of that *lambda_expression*.
- The scope of a parameter declared in an *anonymous_method_expression* (§12.19) is the *block* of that *anonymous_method_expression*.
- The scope of a label declared in a *labeled_statement* (§13.5) is the *block* in which the declaration occurs.
- The scope of a local variable declared in a *local_variable_declaration* (§13.6.2) is the *block* in which the declaration occurs.
- The scope of a local variable declared in a *switch_block* of a `switch` statement (§13.8.3) is the *switch_block*.
- The scope of a local variable declared in a *for_initializer* of a `for` statement (§13.9.4) is the *for_initializer*, *for_condition*, *for_iterator*, and *embedded_statement* of the `for` statement.
- The scope of a local constant declared in a *local_constant_declaration* (§13.6.3) is the *block* in which the declaration occurs. It is a compile-time error to refer to a local constant in a textual position that precedes its *constant_declarator*.
- The scope of a variable declared as part of a *foreach_statement*, *using_statement*, *lock_statement* or *query_expression* is determined by the expansion of the given

construct.

Within the scope of a namespace, class, struct, or enumeration member it is possible to refer to the member in a textual position that precedes the declaration of the member.

> *Example*:
>
> C#
>
> ```
> class A
> {
>     void F()
>     {
>         i = 1;
>     }
>
>     int i = 0;
> }
> ```
>
> Here, it is valid for `F` to refer to `i` before it is declared.
>
> *end example*

Within the scope of a local variable, it is a compile-time error to refer to the local variable in a textual position that precedes its declarator.

> *Example*:
>
> C#
>
> ```
> class A
> {
>     int i = 0;
>
>     void F()
>     {
>         i = 1;                // Error, use precedes declaration
>         int i;
>         i = 2;
>     }
>
>     void G()
>     {
>         int j = (j = 1);      // Valid
>     }
>
>     void H()
>     {
>         int a = 1, b = ++a; // Valid
> ```

```
    }
}
```

In the `F` method above, the first assignment to `i` specifically does not refer to the field declared in the outer scope. Rather, it refers to the local variable and it results in a compile-time error because it textually precedes the declaration of the variable. In the `G` method, the use of `j` in the initializer for the declaration of `j` is valid because the use does not precede the declarator. In the `H` method, a subsequent declarator correctly refers to a local variable declared in an earlier declarator within the same *local_variable_declaration*.

end example

Note: The scoping rules for local variables and local constants are designed to guarantee that the meaning of a name used in an expression context is always the same within a block. If the scope of a local variable were to extend only from its declaration to the end of the block, then in the example above, the first assignment would assign to the instance variable and the second assignment would assign to the local variable, possibly leading to compile-time errors if the statements of the block were later to be rearranged.)

The meaning of a name within a block may differ based on the context in which the name is used. In the example

C#

```
class A {}

class Test
{
    static void Main()
    {
        string A = "hello, world";
        string s = A;                       // expression context
        Type t = typeof(A);                 // type context
        Console.WriteLine(s);               // writes "hello, world"
        Console.WriteLine(t);               // writes "A"
    }
}
```

the name `A` is used in an expression context to refer to the local variable `A` and in a type context to refer to the class `A`.

end note

7.7.2 Name hiding

7.7.2.1 General

The scope of an entity typically encompasses more program text than the declaration space of the entity. In particular, the scope of an entity may include declarations that introduce new declaration spaces containing entities of the same name. Such declarations cause the original entity to become *hidden*. Conversely, an entity is said to be ***visible*** when it is not hidden.

Name hiding occurs when scopes overlap through nesting and when scopes overlap through inheritance. The characteristics of the two types of hiding are described in the following subclauses.

7.7.2.2 Hiding through nesting

Name hiding through nesting can occur as a result of nesting namespaces or types within namespaces, as a result of nesting types within classes or structs, and as a result of parameter, local variable, and local constant declarations.

> *Example*: In the following code
>
> C#
>
> ```
> class A
> {
>     int i = 0;
>     void F()
>     {
>         int i = 1;
>     }
>
>     void G()
>     {
>         i = 1;
>     }
> }
> ```
>
> within the `F` method, the instance variable `i` is hidden by the local variable `i`, but within the `G` method, `i` still refers to the instance variable.
>
> *end example*

When a name in an inner scope hides a name in an outer scope, it hides all overloaded occurrences of that name.

> *Example*: In the following code
>
> ```
> C#
> ```
>
> ```
> class Outer
> {
>     static void F(int i) {}
>     static void F(string s) {}
>
>     class Inner
>     {
>         static void F(long l) {}
>
>         void G()
>         {
>             F(1); // Invokes Outer.Inner.F
>             F("Hello"); // Error
>         }
>     }
> }
> ```
>
> the call `F(1)` invokes the `F` declared in `Inner` because all outer occurrences of `F` are hidden by the inner declaration. For the same reason, the call `F("Hello")` results in a compile-time error.
>
> *end example*

7.7.2.3 Hiding through inheritance

Name hiding through inheritance occurs when classes or structs redeclare names that were inherited from base classes. This type of name hiding takes one of the following forms:

- A constant, field, property, event, or type introduced in a class or struct hides all base class members with the same name.
- A method introduced in a class or struct hides all non-method base class members with the same name, and all base class methods with the same signature (§7.6).
- An indexer introduced in a class or struct hides all base class indexers with the same signature (§7.6) .

The rules governing operator declarations (§15.10) make it impossible for a derived class to declare an operator with the same signature as an operator in a base class. Thus, operators never hide one another.

Contrary to hiding a name from an outer scope, hiding a visible name from an inherited scope causes a warning to be reported.

> *Example*: In the following code
>
> ```C#
> class Base
> {
>     public void F() {}
> }
>
> class Derived : Base
> {
>     public void F() {} // Warning, hiding an inherited name
> }
> ```
>
> the declaration of `F` in `Derived` causes a warning to be reported. Hiding an inherited name is specifically not an error, since that would preclude separate evolution of base classes. For example, the above situation might have come about because a later version of `Base` introduced an `F` method that wasn't present in an earlier version of the class.
>
> *end example*

The warning caused by hiding an inherited name can be eliminated through use of the `new` modifier:

> *Example*:
>
> ```C#
> class Base
> {
>     public void F() {}
> }
>
> class Derived : Base
> {
>     public new void F() {}
> }
> ```
>
> The `new` modifier indicates that the `F` in `Derived` is "new", and that it is indeed intended to hide the inherited member.
>
> *end example*

A declaration of a new member hides an inherited member only within the scope of the new member.

> *Example*:
>
> C#
>
> ```
> class Base
> {
>     public static void F() {}
> }
>
> class Derived : Base
> {
>     private new static void F() {} // Hides Base.F in Derived only
> }
>
> class MoreDerived : Derived
> {
>     static void G()
>     {
>         F();                       // Invokes Base.F
>     }
> }
> ```
>
> In the example above, the declaration of `F` in `Derived` hides the `F` that was inherited from `Base`, but since the new `F` in `Derived` has private access, its scope does not extend to `MoreDerived`. Thus, the call `F()` in `MoreDerived.G` is valid and will invoke `Base.F`.
>
> *end example*

7.8 Namespace and type names

7.8.1 General

Several contexts in a C# program require a *namespace_name* or a *type_name* to be specified.

ANTLR

```
namespace_name
    : namespace_or_type_name
    ;

type_name
    : namespace_or_type_name
```

```
    ;

namespace_or_type_name
    : identifier type_argument_list?
    | namespace_or_type_name '.' identifier type_argument_list?
    | qualified_alias_member
    ;
```

A *namespace_name* is a *namespace_or_type_name* that refers to a namespace.

Following resolution as described below, the *namespace_or_type_name* of a *namespace_name* shall refer to a namespace, or otherwise a compile-time error occurs. No type arguments (§8.4.2) can be present in a *namespace_name* (only types can have type arguments).

A *type_name* is a *namespace_or_type_name* that refers to a type. Following resolution as described below, the *namespace_or_type_name* of a *type_name* shall refer to a type, or otherwise a compile-time error occurs.

If the *namespace_or_type_name* is a *qualified_alias_member* its meaning is as described in §14.8.1. Otherwise, a *namespace_or_type_name* has one of four forms:

- `I`
- `I<A₁, ..., Aₓ>`
- `N.I`
- `N.I<A₁, ..., Aₓ>`

where `I` is a single identifier, `N` is a *namespace_or_type_name* and `<A₁, ..., Aₓ>` is an optional *type_argument_list*. When no *type_argument_list* is specified, consider `x` to be zero.

The meaning of a *namespace_or_type_name* is determined as follows:

- If the *namespace_or_type_name* is a *qualified_alias_member*, the meaning is as specified in §14.8.1.
- Otherwise, if the *namespace_or_type_name* is of the form `I` or of the form `I<A₁, ..., Aₓ>`:
  - If `x` is zero and the *namespace_or_type_name* appears within a generic method declaration (§15.6) but outside the *attributes* of its *method-header,* and if that declaration includes a type parameter (§15.2.3) with name `I`, then the *namespace_or_type_name* refers to that type parameter.
  - Otherwise, if the *namespace_or_type_name* appears within a type declaration, then for each instance type `T` (§15.3.2), starting with the instance type of that

type declaration and continuing with the instance type of each enclosing class or struct declaration (if any):

- If `x` is zero and the declaration of `T` includes a type parameter with name `I`, then the *namespace_or_type_name* refers to that type parameter.
- Otherwise, if the *namespace_or_type_name* appears within the body of the type declaration, and `T` or any of its base types contain a nested accessible type having name `I` and `x` type parameters, then the *namespace_or_type_name* refers to that type constructed with the given type arguments. If there is more than one such type, the type declared within the more derived type is selected.

> *Note*: Non-type members (constants, fields, methods, properties, indexers, operators, instance constructors, finalizers, and static constructors) and type members with a different number of type parameters are ignored when determining the meaning of the *namespace_or_type_name*. *end note*

- Otherwise, for each namespace `N`, starting with the namespace in which the *namespace_or_type_name* occurs, continuing with each enclosing namespace (if any), and ending with the global namespace, the following steps are evaluated until an entity is located:
  - If `x` is zero and `I` is the name of a namespace in `N`, then:
    - If the location where the *namespace_or_type_name* occurs is enclosed by a namespace declaration for `N` and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I` with a namespace or type, then the *namespace_or_type_name* is ambiguous and a compile-time error occurs.
    - Otherwise, the *namespace_or_type_name* refers to the namespace named `I` in `N`.
  - Otherwise, if `N` contains an accessible type having name `I` and `x` type parameters, then:
    - If `x` is zero and the location where the *namespace_or_type_name* occurs is enclosed by a namespace declaration for `N` and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I` with a namespace or type, then the *namespace_or_type_name* is ambiguous and a compile-time error occurs.
    - Otherwise, the *namespace_or_type_name* refers to the type constructed with the given type arguments.
  - Otherwise, if the location where the *namespace_or_type_name* occurs is enclosed by a namespace declaration for `N`:
    - If `x` is zero and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I`

with an imported namespace or type, then the *namespace_or_type_name* refers to that namespace or type.
    - Otherwise, if the namespaces imported by the *using_namespace_directive*s of the namespace declaration contain exactly one type having name `I` and `x` type parameters, then the *namespace_or_type_name* refers to that type constructed with the given type arguments.
    - Otherwise, if the namespaces imported by the *using_namespace_directive*s of the namespace declaration contain more than one type having name `I` and `x` type parameters, then the *namespace_or_type_name* is ambiguous and an error occurs.
  - Otherwise, the *namespace_or_type_name* is undefined and a compile-time error occurs.
- Otherwise, the *namespace_or_type_name* is of the form `N.I` or of the form `N.I<A₁, ..., Aₓ>`. `N` is first resolved as a *namespace_or_type_name*. If the resolution of `N` is not successful, a compile-time error occurs. Otherwise, `N.I` or `N.I<A₁, ..., Aₓ>` is resolved as follows:
  - If `x` is zero and `N` refers to a namespace and `N` contains a nested namespace with name `I`, then the *namespace_or_type_name* refers to that nested namespace.
  - Otherwise, if `N` refers to a namespace and `N` contains an accessible type having name `I` and `x` type parameters, then the *namespace_or_type_name* refers to that type constructed with the given type arguments.
  - Otherwise, if `N` refers to a (possibly constructed) class or struct type and `N` or any of its base classes contain a nested accessible type having name `I` and `x` type parameters, then the *namespace_or_type_name* refers to that type constructed with the given type arguments. If there is more than one such type, the type declared within the more derived type is selected.

    > *Note*: If the meaning of `N.I` is being determined as part of resolving the base class specification of `N` then the direct base class of `N` is considered to be `object` (§15.2.4.2). *end note*

  - Otherwise, `N.I` is an invalid *namespace_or_type_name*, and a compile-time error occurs.

A *namespace_or_type_name* is permitted to reference a static class (§15.2.2.4) only if

- The *namespace_or_type_name* is the `T` in a *namespace_or_type_name* of the form `T.I`, or

- The *namespace_or_type_name* is the `T` in a *typeof_expression* (§12.8.17) of the form `typeof(T)`

7.8.2 Unqualified names

Every namespace declaration and type declaration has an ***unqualified name*** determined as follows:

- For a namespace declaration, the unqualified name is the *qualified_identifier* specified in the declaration.
- For a type declaration with no *type_parameter_list*, the unqualified name is the *identifier* specified in the declaration.
- For a type declaration with K type parameters, the unqualified name is the *identifier* specified in the declaration, followed by the *generic_dimension_specifier* (§12.8.17) for K type parameters.

7.8.3 Fully qualified names

Every namespace and type declaration has a ***fully qualified name,*** which uniquely identifies the namespace or type declaration amongst all others within the program. The fully qualified name of a namespace or type declaration with unqualified name `N` is determined as follows:

- If `N` is a member of the global namespace, its fully qualified name is `N`.
- Otherwise, its fully qualified name is `S.N`, where `S` is the fully qualified name of the namespace or type declaration in which `N` is declared.

In other words, the fully qualified name of `N` is the complete hierarchical path of identifiers and *generic_dimension_specifier*s that lead to `N`, starting from the global namespace. Because every member of a namespace or type shall have a unique name, it follows that the fully qualified name of a namespace or type declaration is always unique. It is a compile-time error for the same fully qualified name to refer to two distinct entities. In particular:

- It is an error for both a namespace declaration and a type declaration to have the same fully qualified name.
- It is an error for two different kinds of type declarations to have the same fully qualified name (for example, if both a struct and class declaration have the same fully qualified name).
- It is an error for a type declaration without the partial modifier to have the same fully qualified name as another type declaration (§15.2.7).

Example: The example below shows several namespace and type declarations along with their associated fully qualified names.

```
C#
```

```
class A {}                  // A
namespace X                 // X
{
    class B                 // X.B
    {
        class C {}          // X.B.C
    }
    namespace Y             // X.Y
    {
        class D {}          // X.Y.D
    }
}
namespace X.Y               // X.Y
{
    class E {}              // X.Y.E
    class G<T>              // X.Y.G<>
    {
        class H {}          // X.Y.G<>.H
    }
    class G<S,T>            // X.Y.G<,>
    {
        class H<U> {}       // X.Y.G<,>.H<>
    }
}
```

end example

7.9 Automatic memory management

C# employs automatic memory management, which frees developers from manually allocating and freeing the memory occupied by objects. Automatic memory management policies are implemented by a garbage collector. The memory management life cycle of an object is as follows:

1. When the object is created, memory is allocated for it, the constructor is run, and the object is considered ***live***.
2. If neither the object nor any of its instance fields can be accessed by any possible continuation of execution, other than the running of finalizers, the object is considered ***no longer in use*** and it becomes eligible for finalization.

> *Note*: The C# compiler and the garbage collector might choose to analyze code to determine which references to an object might be used in the future.

> For instance, if a local variable that is in scope is the only existing reference to an object, but that local variable is never referred to in any possible continuation of execution from the current execution point in the procedure, the garbage collector might (but is not required to) treat the object as no longer in use. *end note*

3. Once the object is eligible for finalization, at some unspecified later time the finalizer (§15.13) (if any) for the object is run. Under normal circumstances the finalizer for the object is run once only, though implementation-specific APIs may allow this behavior to be overridden.
4. Once the finalizer for an object is run, if neither the object nor any of its instance fields can be accessed by any possible continuation of execution, including the running of finalizers, the object is considered inaccessible and the object becomes eligible for collection.

   > *Note*: An object which could previously not be accessed may become accessible again due to its finalizer. An example of this is provided below. *end note*

5. Finally, at some time after the object becomes eligible for collection, the garbage collector frees the memory associated with that object.

The garbage collector maintains information about object usage, and uses this information to make memory management decisions, such as where in memory to locate a newly created object, when to relocate an object, and when an object is no longer in use or inaccessible.

Like other languages that assume the existence of a garbage collector, C# is designed so that the garbage collector might implement a wide range of memory management policies. C# specifies neither a time constraint within that span, nor an order in which finalizers are run. Whether or not finalizers are run as part of application termination is implementation-specific (§7.2).

The behavior of the garbage collector can be controlled, to some degree, via static methods on the class `System.GC`. This class can be used to request a collection to occur, finalizers to be run (or not run), and so forth.

> *Example*: Since the garbage collector is allowed wide latitude in deciding when to collect objects and run finalizers, a conforming implementation might produce output that differs from that shown by the following code. The program
>
> ```
> C#
> ```

```
class A
{
    ~A()
    {
        Console.WriteLine("Finalize instance of A");
    }
}

class B
{
    object Ref;
    public B(object o)
    {
        Ref = o;
    }

    ~B()
    {
        Console.WriteLine("Finalize instance of B");
    }
}

class Test
{
    static void Main()
    {
        B b = new B(new A());
        b = null;
        GC.Collect();
        GC.WaitForPendingFinalizers();
    }
}
```

creates an instance of class `A` and an instance of class `B`. These objects become eligible for garbage collection when the variable `b` is assigned the value `null`, since after this time it is impossible for any user-written code to access them. The output could be either

Console

```
Finalize instance of A
Finalize instance of B
```

or

Console

```
Finalize instance of B
```

```
Finalize instance of A
```

because the language imposes no constraints on the order in which objects are garbage collected.

In subtle cases, the distinction between "eligible for finalization" and "eligible for collection" can be important. For example,

C#

```
class A
{
    ~A()
    {
        Console.WriteLine("Finalize instance of A");
    }

    public void F()
    {
        Console.WriteLine("A.F");
        Test.RefA = this;
    }
}

class B
{
    public A Ref;

    ~B()
    {
        Console.WriteLine("Finalize instance of B");
        Ref.F();
    }
}

class Test
{
    public static A RefA;
    public static B RefB;

    static void Main()
    {
        RefB = new B();
        RefA = new A();
        RefB.Ref = RefA;
        RefB = null;
        RefA = null;
        // A and B now eligible for finalization
        GC.Collect();
        GC.WaitForPendingFinalizers();
        // B now eligible for collection, but A is not
        if (RefA != null)
```

```
        {
            Console.WriteLine("RefA is not null");
        }
    }
}
```

In the above program, if the garbage collector chooses to run the finalizer of `A` before the finalizer of `B`, then the output of this program might be:

Console

```
Finalize instance of A
Finalize instance of B
A.F
RefA is not null
```

Note that although the instance of `A` was not in use and `A`'s finalizer was run, it is still possible for methods of `A` (in this case, `F`) to be called from another finalizer. Also, note that running of a finalizer might cause an object to become usable from the mainline program again. In this case, the running of `B`'s finalizer caused an instance of `A` that was previously not in use, to become accessible from the live reference `Test.RefA`. After the call to `WaitForPendingFinalizers`, the instance of `B` is eligible for collection, but the instance of `A` is not, because of the reference `Test.RefA`.

end example

7.10 Execution order

Execution of a C# program proceeds such that the side effects of each executing thread are preserved at critical execution points. A ***side effect*** is defined as a read or write of a volatile field, a write to a non-volatile variable, a write to an external resource, and the throwing of an exception. The critical execution points at which the order of these side effects shall be preserved are references to volatile fields (§15.5.4), `lock` statements (§13.13), and thread creation and termination. The execution environment is free to change the order of execution of a C# program, subject to the following constraints:

- Data dependence is preserved within a thread of execution. That is, the value of each variable is computed as if all statements in the thread were executed in original program order.
- Initialization ordering rules are preserved (§15.5.5, §15.5.6).

- The ordering of side effects is preserved with respect to volatile reads and writes (§15.5.4). Additionally, the execution environment need not evaluate part of an expression if it can deduce that that expression's value is not used and that no needed side effects are produced (including any caused by calling a method or accessing a volatile field). When program execution is interrupted by an asynchronous event (such as an exception thrown by another thread), it is not guaranteed that the observable side effects are visible in the original program order.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

8 Types

Article • 02/07/2024

8.1 General

The types of the C# language are divided into two main categories: ***reference types*** and ***value types***. Both value types and reference types may be ***generic types***, which take one or more ***type parameters***. Type parameters can designate both value types and reference types.

ANTLR

```
type
    : reference_type
    | value_type
    | type_parameter
    | pointer_type     // unsafe code support
    ;
```

pointer_type (§23.3) is available only in unsafe code (§23).

Value types differ from reference types in that variables of the value types directly contain their data, whereas variables of the reference types store ***references*** to their data, the latter being known as ***objects***. With reference types, it is possible for two variables to reference the same object, and thus possible for operations on one variable to affect the object referenced by the other variable. With value types, the variables each have their own copy of the data, and it is not possible for operations on one to affect the other.

> *Note*: When a variable is a ref or out parameter, it does not have its own storage but references the storage of another variable. In this case, the ref or out variable is effectively an alias for another variable and not a distinct variable. *end note*

C#'s type system is unified such that *a value of any type can be treated as an object*. Every type in C# directly or indirectly derives from the `object` class type, and `object` is the ultimate base class of all types. Values of reference types are treated as objects simply by viewing the values as type `object`. Values of value types are treated as objects by performing boxing and unboxing operations (§8.3.13).

For convenience, throughout this specification, some library type names are written without using their full name qualification. Refer to §C.5 for more information.

8.2 Reference types

8.2.1 General

A reference type is a class type, an interface type, an array type, a delegate type, or the `dynamic` type.

ANTLR

```
reference_type
    : class_type
    | interface_type
    | array_type
    | delegate_type
    | 'dynamic'
    ;

class_type
    : type_name
    | 'object'
    | 'string'
    ;

interface_type
    : type_name
    ;

array_type
    : non_array_type rank_specifier+
    ;

non_array_type
    : value_type
    | class_type
    | interface_type
    | delegate_type
    | 'dynamic'
    | type_parameter
    | pointer_type      // unsafe code support
    ;

rank_specifier
    : '[' ','* ']'
    ;

delegate_type
    : type_name
    ;
```

pointer_type is available only in unsafe code (§23.3).

A reference type value is a reference to an ***instance*** of the type, the latter known as an object. The special value `null` is compatible with all reference types and indicates the absence of an instance.

8.2.2 Class types

A class type defines a data structure that contains ***data members*** (constants and fields), ***function members*** (methods, properties, events, indexers, operators, instance constructors, finalizers, and static constructors), and nested types. Class types support inheritance, a mechanism whereby derived classes can extend and specialize base classes. Instances of class types are created using *object_creation_expression*s (§12.8.16.2).

Class types are described in §15.

Certain predefined class types have special meaning in the C# language, as described in the table below.

Expand table

Class type	Description
`System.Object`	The ultimate base class of all other types. See §8.2.3.
`System.String`	The string type of the C# language. See §8.2.5.
`System.ValueType`	The base class of all value types. See §8.3.2.
`System.Enum`	The base class of all `enum` types. See §19.5.
`System.Array`	The base class of all array types. See §17.2.2.
`System.Delegate`	The base class of all `delegate` types. See §20.1.
`System.Exception`	The base class of all exception types. See §21.3.

8.2.3 The object type

The `object` class type is the ultimate base class of all other types. Every type in C# directly or indirectly derives from the `object` class type.

The keyword `object` is simply an alias for the predefined class `System.Object`.

8.2.4 The dynamic type

The `dynamic` type, like `object`, can reference any object. When operations are applied to expressions of type `dynamic`, their resolution is deferred until the program is run. Thus, if the operation cannot legitimately be applied to the referenced object, no error is given during compilation. Instead, an exception will be thrown when resolution of the operation fails at run-time.

The `dynamic` type is further described in §8.7, and dynamic binding in §12.3.1.

8.2.5 The string type

The `string` type is a sealed class type that inherits directly from `object`. Instances of the `string` class represent Unicode character strings.

Values of the `string` type can be written as string literals (§6.4.5.6).

The keyword `string` is simply an alias for the predefined class `System.String`.

8.2.6 Interface types

An interface defines a contract. A class or struct that implements an interface shall adhere to its contract. An interface may inherit from multiple base interfaces, and a class or struct may implement multiple interfaces.

Interface types are described in §18.

8.2.7 Array types

An array is a data structure that contains zero or more variables, which are accessed through computed indices. The variables contained in an array, also called the elements of the array, are all of the same type, and this type is called the element type of the array.

Array types are described in §17.

8.2.8 Delegate types

A delegate is a data structure that refers to one or more methods. For instance methods, it also refers to their corresponding object instances.

> *Note*: The closest equivalent of a delegate in C or C++ is a function pointer, but whereas a function pointer can only reference static functions, a delegate can reference both static and instance methods. In the latter case, the delegate stores

> not only a reference to the method's entry point, but also a reference to the object instance on which to invoke the method. *end note*

Delegate types are described in §20.

8.3 Value types

8.3.1 General

A value type is either a struct type or an enumeration type. C# provides a set of predefined struct types called the ***simple types***. The simple types are identified through keywords.

ANTLR

```
value_type
    : non_nullable_value_type
    | nullable_value_type
    ;

non_nullable_value_type
    : struct_type
    | enum_type
    ;

struct_type
    : type_name
    | simple_type
    | tuple_type
    ;

simple_type
    : numeric_type
    | 'bool'
    ;

numeric_type
    : integral_type
    | floating_point_type
    | 'decimal'
    ;

integral_type
    : 'sbyte'
    | 'byte'
    | 'short'
    | 'ushort'
    | 'int'
    | 'uint'
```

```
    | 'long'
    | 'ulong'
    | 'char'
    ;

floating_point_type
    : 'float'
    | 'double'
    ;

tuple_type
    : '(' tuple_type_element (',' tuple_type_element)+ ')'
    ;

tuple_type_element
    : type identifier?
    ;

enum_type
    : type_name
    ;

nullable_value_type
    : non_nullable_value_type '?'
    ;
```

Unlike a variable of a reference type, a variable of a value type can contain the value `null` only if the value type is a nullable value type (§8.3.12). For every non-nullable value type there is a corresponding nullable value type denoting the same set of values plus the value `null`.

Assignment to a variable of a value type creates a *copy* of the value being assigned. This differs from assignment to a variable of a reference type, which copies the reference but not the object identified by the reference.

8.3.2 The System.ValueType type

All value types implicitly inherit from the `class` `System.ValueType`, which, in turn, inherits from class `object`. It is not possible for any type to derive from a value type, and value types are thus implicitly sealed (§15.2.2.3).

Note that `System.ValueType` is not itself a *value_type*. Rather, it is a *class_type* from which all *value_types* are automatically derived.

8.3.3 Default constructors

All value types implicitly declare a public parameterless instance constructor called the ***default constructor***. The default constructor returns a zero-initialized instance known as the ***default value*** for the value type:

- For all *simple_types*, the default value is the value produced by a bit pattern of all zeros:
  - For `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, and `ulong`, the default value is `0`.
  - For `char`, the default value is `'\x0000'`.
  - For `float`, the default value is `0.0f`.
  - For `double`, the default value is `0.0d`.
  - For `decimal`, the default value is `0m` (that is, value zero with scale 0).
  - For `bool`, the default value is `false`.
  - For an *enum_type* `E`, the default value is `0`, converted to the type `E`.
- For a *struct_type*, the default value is the value produced by setting all value type fields to their default value and all reference type fields to `null`.
- For a *nullable_value_type* the default value is an instance for which the `HasValue` property is false. The default value is also known as the ***null value*** of the nullable value type. Attempting to read the `Value` property of such a value causes an exception of type `System.InvalidOperationException` to be thrown (§8.3.12).

Like any other instance constructor, the default constructor of a value type is invoked using the `new` operator.

> *Note*: For efficiency reasons, this requirement is not intended to actually have the implementation generate a constructor call. For value types, the default value expression (§12.8.20) produces the same result as using the default constructor. *end note*

> *Example*: In the code below, variables `i`, `j` and `k` are all initialized to zero.

C#

```
class A
{
    void F()
    {
        int i = 0;
        int j = new int();
        int k = default(int);
    }
}
```

> *end example*

Because every value type implicitly has a public parameterless instance constructor, it is not possible for a struct type to contain an explicit declaration of a parameterless constructor. A struct type is however permitted to declare parameterized instance constructors (§16.4.9).

8.3.4 Struct types

A struct type is a value type that can declare constants, fields, methods, properties, events, indexers, operators, instance constructors, static constructors, and nested types. The declaration of struct types is described in §16.

8.3.5 Simple types

C# provides a set of predefined `struct` types called the simple types. The simple types are identified through keywords, but these keywords are simply aliases for predefined `struct` types in the `System` namespace, as described in the table below.

Expand table

Keyword	Aliased type
`sbyte`	`System.SByte`
`byte`	`System.Byte`
`short`	`System.Int16`
`ushort`	`System.UInt16`
`int`	`System.Int32`
`uint`	`System.UInt32`
`long`	`System.Int64`
`ulong`	`System.UInt64`
`char`	`System.Char`
`float`	`System.Single`
`double`	`System.Double`
`bool`	`System.Boolean`

Keyword	Aliased type
`decimal`	`System.Decimal`

Because a simple type aliases a struct type, every simple type has members.

> *Example*: `int` has the members declared in `System.Int32` and the members inherited from `System.Object`, and the following statements are permitted:
>
> ```
> C#
> ```
>
> ```
> int i = int.MaxValue;      // System.Int32.MaxValue constant
> string s = i.ToString();   // System.Int32.ToString() instance method
> string t = 123.ToString(); // System.Int32.ToString() instance method
> ```
>
> *end example*

> *Note*: The simple types differ from other struct types in that they permit certain additional operations:
>
> - Most simple types permit values to be created by writing *literals* (§6.4.5), although C# makes no provision for literals of struct types in general. *Example*: `123` is a literal of type `int` and `'a'` is a literal of type `char`. *end example*
> - When the operands of an expression are all simple type constants, it is possible for the compiler to evaluate the expression at compile-time. Such an expression is known as a *constant_expression* (§12.23). Expressions involving operators defined by other struct types are not considered to be constant expressions
> - Through `const` declarations, it is possible to declare constants of the simple types (§15.4). It is not possible to have constants of other struct types, but a similar effect is provided by static readonly fields.
> - Conversions involving simple types can participate in evaluation of conversion operators defined by other struct types, but a user-defined conversion operator can never participate in evaluation of another user-defined conversion operator (§10.5.3).
>
> *end note*.

8.3.6 Integral types

C# supports nine integral types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, and `char`. The integral types have the following sizes and ranges of values:

- The `sbyte` type represents signed 8-bit integers with values from `-128` to `127`, inclusive.
- The `byte` type represents unsigned 8-bit integers with values from `0` to `255`, inclusive.
- The `short` type represents signed 16-bit integers with values from `-32768` to `32767`, inclusive.
- The `ushort` type represents unsigned 16-bit integers with values from `0` to `65535`, inclusive.
- The `int` type represents signed 32-bit integers with values from `-2147483648` to `2147483647`, inclusive.
- The `uint` type represents unsigned 32-bit integers with values from `0` to `4294967295`, inclusive.
- The `long` type represents signed 64-bit integers with values from `-9223372036854775808` to `9223372036854775807`, inclusive.
- The `ulong` type represents unsigned 64-bit integers with values from `0` to `18446744073709551615`, inclusive.
- The `char` type represents unsigned 16-bit integers with values from `0` to `65535`, inclusive. The set of possible values for the `char` type corresponds to the Unicode character set.

> *Note*: Although `char` has the same representation as `ushort`, not all operations permitted on one type are permitted on the other. *end note*

All signed integral types are represented using two's complement format.

The *integral_type* unary and binary operators always operate with signed 32-bit precision, unsigned 32-bit precision, signed 64-bit precision, or unsigned 64-bit precision, as detailed in §12.4.7.

The `char` type is classified as an integral type, but it differs from the other integral types in two ways:

- There are no predefined implicit conversions from other types to the `char` type. In particular, even though the `byte` and `ushort` types have ranges of values that are fully representable using the `char` type, implicit conversions from sbyte, byte, or `ushort` to `char` do not exist.
- Constants of the `char` type shall be written as *character_literal*s or as *integer_literal*s in combination with a cast to type char.

> *Example*: `(char)10` is the same as `'\x000A'`. *end example*

The `checked` and `unchecked` operators and statements are used to control overflow checking for integral-type arithmetic operations and conversions (§12.8.19). In a `checked` context, an overflow produces a compile-time error or causes a `System.OverflowException` to be thrown. In an `unchecked` context, overflows are ignored and any high-order bits that do not fit in the destination type are discarded.

8.3.7 Floating-point types

C# supports two floating-point types: `float` and `double`. The `float` and `double` types are represented using the 32-bit single-precision and 64-bit double-precision IEC 60559 formats, which provide the following sets of values:

- Positive zero and negative zero. In most situations, positive zero and negative zero behave identically as the simple value zero, but certain operations distinguish between the two (§12.10.3).
- Positive infinity and negative infinity. Infinities are produced by such operations as dividing a non-zero number by zero.

  > *Example*: `1.0 / 0.0` yields positive infinity, and `-1.0 / 0.0` yields negative infinity. *end example*

- The ***Not-a-Number*** value, often abbreviated NaN. NaNs are produced by invalid floating-point operations, such as dividing zero by zero.
- The finite set of non-zero values of the form $s \times m \times 2^e$, where s is 1 or −1, and m and e are determined by the particular floating-point type: For `float`, $0 < m < 2^{24}$ and $-149 \le e \le 104$, and for `double`, $0 < m < 2^{53}$ and $-1075 \le e \le 970$. Denormalized floating-point numbers are considered valid non-zero values. C# neither requires nor forbids that a conforming implementation support denormalized floating-point numbers.

The `float` type can represent values ranging from approximately 1.5×10^{-45} to 3.4×10^{38} with a precision of 7 digits.

The `double` type can represent values ranging from approximately 5.0×10^{-324} to 1.7×10^{308} with a precision of 15-16 digits.

If either operand of a binary operator is a floating-point type then standard numeric promotions are applied, as detailed in §12.4.7, and the operation is performed with `float` or `double` precision.

The floating-point operators, including the assignment operators, never produce exceptions. Instead, in exceptional situations, floating-point operations produce zero, infinity, or NaN, as described below:

- The result of a floating-point operation is rounded to the nearest representable value in the destination format.
- If the magnitude of the result of a floating-point operation is too small for the destination format, the result of the operation becomes positive zero or negative zero.
- If the magnitude of the result of a floating-point operation is too large for the destination format, the result of the operation becomes positive infinity or negative infinity.
- If a floating-point operation is invalid, the result of the operation becomes NaN.
- If one or both operands of a floating-point operation is NaN, the result of the operation becomes NaN.

Floating-point operations may be performed with higher precision than the result type of the operation. To force a value of a floating-point type to the exact precision of its type, an explicit cast (§12.9.7) can be used.

> *Example*: Some hardware architectures support an "extended" or "long double" floating-point type with greater range and precision than the `double` type, and implicitly perform all floating-point operations using this higher precision type. Only at excessive cost in performance can such hardware architectures be made to perform floating-point operations with *less* precision, and rather than require an implementation to forfeit both performance and precision, C# allows a higher precision type to be used for all floating-point operations. Other than delivering more precise results, this rarely has any measurable effects. However, in expressions of the form `x * y / z`, where the multiplication produces a result that is outside the `double` range, but the subsequent division brings the temporary result back into the `double` range, the fact that the expression is evaluated in a higher range format can cause a finite result to be produced instead of an infinity. *end example*

8.3.8 The Decimal type

The `decimal` type is a 128-bit data type suitable for financial and monetary calculations. The `decimal` type can represent values including those in the range at least -7.9×10^{-28} to 7.9×10^{28}, with at least 28-digit precision.

The finite set of values of type `decimal` are of the form $(-1)^v \times c \times 10^{-e}$, where the sign v is 0 or 1, the coefficient c is given by $0 \le c < Cmax$, and the scale e is such that $Emin \le e$

≤ *Emax*, where *Cmax* is at least 1×10^{28}, *Emin* ≤ 0, and *Emax* ≥ 28. The `decimal` type does not necessarily support signed zeros, infinities, or NaN's.

A `decimal` is represented as an integer scaled by a power of ten. For `decimal`s with an absolute value less than `1.0m`, the value is exact to at least the 28th decimal place. For `decimal`s with an absolute value greater than or equal to `1.0m`, the value is exact to at least 28 digits. Contrary to the `float` and `double` data types, decimal fractional numbers such as `0.1` can be represented exactly in the decimal representation. In the `float` and `double` representations, such numbers often have non-terminating binary expansions, making those representations more prone to round-off errors.

If either operand of a binary operator is of `decimal` type then standard numeric promotions are applied, as detailed in §12.4.7, and the operation is performed with `double` precision.

The result of an operation on values of type `decimal` is that which would result from calculating an exact result (preserving scale, as defined for each operator) and then rounding to fit the representation. Results are rounded to the nearest representable value, and, when a result is equally close to two representable values, to the value that has an even number in the least significant digit position (this is known as "banker's rounding"). That is, results are exact to at least the 28th decimal place. Note that rounding may produce a zero value from a non-zero value.

If a `decimal` arithmetic operation produces a result whose magnitude is too large for the `decimal` format, a `System.OverflowException` is thrown.

The `decimal` type has greater precision but may have a smaller range than the floating-point types. Thus, conversions from the floating-point types to `decimal` might produce overflow exceptions, and conversions from `decimal` to the floating-point types might cause loss of precision or overflow exceptions. For these reasons, no implicit conversions exist between the floating-point types and `decimal`, and without explicit casts, a compile-time error occurs when floating-point and `decimal` operands are directly mixed in the same expression.

8.3.9 The Bool type

The `bool` type represents Boolean logical quantities. The possible values of type `bool` are `true` and `false`. The representation of `false` is described in §8.3.3. Although the representation of `true` is unspecified, it shall be different from that of `false`.

No standard conversions exist between `bool` and other value types. In particular, the `bool` type is distinct and separate from the integral types, a `bool` value cannot be used in place of an integral value, and vice versa.

> *Note*: In the C and C++ languages, a zero integral or floating-point value, or a null pointer can be converted to the Boolean value `false`, and a non-zero integral or floating-point value, or a non-null pointer can be converted to the Boolean value `true`. In C#, such conversions are accomplished by explicitly comparing an integral or floating-point value to zero, or by explicitly comparing an object reference to `null`. *end note*

8.3.10 Enumeration types

An enumeration type is a distinct type with named constants. Every enumeration type has an underlying type, which shall be `byte`, `sbyte`, `short`, `ushort`, `int`, `uint`, `long` or `ulong`. The set of values of the enumeration type is the same as the set of values of the underlying type. Values of the enumeration type are not restricted to the values of the named constants. Enumeration types are defined through enumeration declarations (§19.2).

8.3.11 Tuple types

A tuple type represents an ordered, fixed-length sequence of values with optional names and individual types. The number of elements in a tuple type is referred to as its ***arity***. A tuple type is written `(T1 I1, ..., Tn In)` with n ≥ 2, where the identifiers `I1...In` are optional ***tuple element names***.

This syntax is shorthand for a type constructed with the types `T1...Tn` from `System.ValueTuple<...>`, which shall be a set of generic struct types capable of directly expressing tuple types of any arity between two and seven inclusive. There does not need to exist a `System.ValueTuple<...>` declaration that directly matches the arity of any tuple type with a corresponding number of type parameters. Instead, tuples with an arity greater than seven are represented with a generic struct type `System.ValueTuple<T1, ..., T7, TRest>` that in addition to tuple elements has a `Rest` field containing a nested value of the remaining elements, using another `System.ValueTuple<...>` type. Such nesting may be observable in various ways, e.g. via the presence of a `Rest` field. Where only a single additional field is required, the generic struct type `System.ValueTuple<T1>` is used; this type is not considered a tuple type in itself. Where more than seven

additional fields are required, `System.ValueTuple<T1, ..., T7, TRest>` is used recursively.

Element names within a tuple type shall be distinct. A tuple element name of the form `ItemX`, where `X` is any sequence of non-`0`-initiated decimal digits that could represent the position of a tuple element, is only permitted at the position denoted by `X`.

The optional element names are not represented in the `ValueTuple<...>` types, and are not stored in the runtime representation of a tuple value. Identity conversions (§10.2.2) exist between tuples with identity-convertible sequences of element types.

The `new` operator §12.8.16.2 cannot be applied with the tuple type syntax `new (T1, ..., Tn)`. Tuple values can be created from tuple expressions (§12.8.6), or by applying the `new` operator directly to a type constructed from `ValueTuple<...>`.

Tuple elements are public fields with the names `Item1`, `Item2`, etc., and can be accessed via a member access on a tuple value (§12.8.7. Additionally, if the tuple type has a name for a given element, that name can be used to access the element in question.

> *Note*: Even when large tuples are represented with nested `System.ValueTuple<...>` values, each tuple element can still be accessed directly with the `Item...` name corresponding to its position. *end note*

> *Example*: Given the following examples:

```C#
(int, string) pair1 = (1, "One");
(int, string word) pair2 = (2, "Two");
(int number, string word) pair3 = (3, "Three");
(int Item1, string Item2) pair4 = (4, "Four");
// Error: "Item" names do not match their position
(int Item2, string Item123) pair5 = (5, "Five");
(int, string) pair6 = new ValueTuple<int, string>(6, "Six");
ValueTuple<int, string> pair7 = (7, "Seven");
Console.WriteLine($"{pair2.Item1}, {pair2.Item2}, {pair2.word}");
```

> The tuple types for `pair1`, `pair2`, and `pair3` are all valid, with names for no, some or all of the tuple type elements.
>
> The tuple type for `pair4` is valid because the names `Item1` and `Item2` match their positions, whereas the tuple type for `pair5` is disallowed, because the names `Item2` and `Item123` do not.

> The declarations for `pair6` and `pair7` demonstrate that tuple types are interchangeable with constructed types of the form `ValueTuple<...>`, and that the `new` operator is allowed with the latter syntax.
>
> The last line shows that tuple elements can be accessed by the `Item` name corresponding to their position, as well as by the corresponding tuple element name, if present in the type. *end example*

8.3.12 Nullable value types

A nullable value type can represent all values of its underlying type plus an additional null value. A nullable value type is written `T?`, where `T` is the underlying type. This syntax is shorthand for `System.Nullable<T>`, and the two forms can be used interchangeably.

Conversely, a ***non-nullable value type*** is any value type other than `System.Nullable<T>` and its shorthand `T?` (for any `T`), plus any type parameter that is constrained to be a non-nullable value type (that is, any type parameter with a value type constraint (§15.2.5)). The `System.Nullable<T>` type specifies the value type constraint for `T`, which means that the underlying type of a nullable value type can be any non-nullable value type. The underlying type of a nullable value type cannot be a nullable value type or a reference type. For example, `int??` and `string?` are invalid types.

An instance of a nullable value type `T?` has two public read-only properties:

- A `HasValue` property of type `bool`
- A `Value` property of type `T`

An instance for which `HasValue` is `true` is said to be non-null. A non-null instance contains a known value and `Value` returns that value.

An instance for which `HasValue` is `false` is said to be null. A null instance has an undefined value. Attempting to read the `Value` of a null instance causes a `System.InvalidOperationException` to be thrown. The process of accessing the Value property of a nullable instance is referred to as ***unwrapping***.

In addition to the default constructor, every nullable value type `T?` has a public constructor with a single parameter of type `T`. Given a value `x` of type `T`, a constructor invocation of the form

```
C#
```

```
new T?(x)
```

creates a non-null instance of `T?` for which the `Value` property is `x`. The process of creating a non-null instance of a nullable value type for a given value is referred to as ***wrapping***.

Implicit conversions are available from the `null` literal to `T?` (§10.2.7) and from `T` to `T?` (§10.2.6).

The nullable type `T?` implements no interfaces (§18). In particular, this means it does not implement any interface that the underlying type `T` does.

8.3.13 Boxing and unboxing

The concept of boxing and unboxing provide a bridge between *value_type*s and *reference_type*s by permitting any value of a *value_type* to be converted to and from type `object`. Boxing and unboxing enables a unified view of the type system wherein a value of any type can ultimately be treated as an `object`.

Boxing is described in more detail in §10.2.9 and unboxing is described in §10.3.7.

8.4 Constructed types

8.4.1 General

A generic type declaration, by itself, denotes an ***unbound generic type*** that is used as a "blueprint" to form many different types, by way of applying ***type arguments***. The type arguments are written within angle brackets (`<` and `>`) immediately following the name of the generic type. A type that includes at least one type argument is called a ***constructed type***. A constructed type can be used in most places in the language in which a type name can appear. An unbound generic type can only be used within a *typeof_expression* (§12.8.17).

Constructed types can also be used in expressions as simple names (§12.8.4) or when accessing a member (§12.8.7).

When a *namespace_or_type_name* is evaluated, only generic types with the correct number of type parameters are considered. Thus, it is possible to use the same identifier to identify different types, as long as the types have different numbers of type

parameters. This is useful when mixing generic and non-generic classes in the same program.

> *Example*:
>
> C#
>
> ```
> namespace Widgets
> {
>     class Queue {...}
>     class Queue<TElement> {...}
> }
>
> namespace MyApplication
> {
>     using Widgets;
>
>     class X
>     {
>         Queue q1;      // Non-generic Widgets.Queue
>         Queue<int> q2; // Generic Widgets.Queue
>     }
> }
> ```
>
> *end example*

The detailed rules for name lookup in the *namespace_or_type_name* productions is described in §7.8. The resolution of ambiguities in these productions is described in §6.2.5. A *type_name* might identify a constructed type even though it doesn't specify type parameters directly. This can occur where a type is nested within a generic `class` declaration, and the instance type of the containing declaration is implicitly used for name lookup (§15.3.9.7).

> *Example*:
>
> C#
>
> ```
> class Outer<T>
> {
>     public class Inner {...}
>
>     public Inner i; // Type of i is Outer<T>.Inner
> }
> ```
>
> *end example*

A non-enum constructed type shall not be used as an *unmanaged_type* (§8.8).

8.4.2 Type arguments

Each argument in a type argument list is simply a *type*.

ANTLR

```
type_argument_list
    : '<' type_arguments '>'
    ;

type_arguments
    : type_argument (',' type_argument)*
    ;

type_argument
    : type
    ;
```

Each type argument shall satisfy any constraints on the corresponding type parameter (§15.2.5).

8.4.3 Open and closed types

All types can be classified as either ***open types*** or ***closed types***. An open type is a type that involves type parameters. More specifically:

- A type parameter defines an open type.
- An array type is an open type if and only if its element type is an open type.
- A constructed type is an open type if and only if one or more of its type arguments is an open type. A constructed nested type is an open type if and only if one or more of its type arguments or the type arguments of its containing type(s) is an open type.

A closed type is a type that is not an open type.

At run-time, all of the code within a generic type declaration is executed in the context of a closed constructed type that was created by applying type arguments to the generic declaration. Each type parameter within the generic type is bound to a particular run-time type. The run-time processing of all statements and expressions always occurs with closed types, and open types occur only during compile-time processing.

Two closed constructed types are identity convertible (§10.2.2) if they are constructed from the same unbound generic type, and an identity conversion exists between each of

their corresponding type arguments. The corresponding type arguments may themselves be closed constructed types or tuples that are identity convertible. Closed constructed types that are identity convertible share a single set of static variables. Otherwise, each closed constructed type has its own set of static variables. Since an open type does not exist at run-time, there are no static variables associated with an open type.

8.4.4 Bound and unbound types

The term ***unbound type*** refers to a non-generic type or an unbound generic type. The term ***bound type*** refers to a non-generic type or a constructed type.

An unbound type refers to the entity declared by a type declaration. An unbound generic type is not itself a type, and cannot be used as the type of a variable, argument or return value, or as a base type. The only construct in which an unbound generic type can be referenced is the `typeof` expression (§12.8.17).

8.4.5 Satisfying constraints

Whenever a constructed type or generic method is referenced, the supplied type arguments are checked against the type parameter constraints declared on the generic type or method (§15.2.5). For each `where` clause, the type argument `A` that corresponds to the named type parameter is checked against each constraint as follows:

- If the constraint is a `class` type, an interface type, or a type parameter, let `C` represent that constraint with the supplied type arguments substituted for any type parameters that appear in the constraint. To satisfy the constraint, it shall be the case that type `A` is convertible to type `C` by one of the following:
  - An identity conversion (§10.2.2)
  - An implicit reference conversion (§10.2.8)
  - A boxing conversion (§10.2.9), provided that type `A` is a non-nullable value type.
  - An implicit reference, boxing or type parameter conversion from a type parameter `A` to `C`.
- If the constraint is the reference type constraint (`class`), the type `A` shall satisfy one of the following:
  - `A` is an interface type, class type, delegate type, array type or the dynamic type.

> *Note*: `System.ValueType` and `System.Enum` are reference types that satisfy this constraint. *end note*

    - `A` is a type parameter that is known to be a reference type (§8.2).
- If the constraint is the value type constraint (`struct`), the type `A` shall satisfy one of the following:
    - `A` is a `struct` type or `enum` type, but not a nullable value type.

    > *Note*: `System.ValueType` and `System.Enum` are reference types that do not satisfy this constraint. *end note*

    - `A` is a type parameter having the value type constraint (§15.2.5).
- If the constraint is the constructor constraint `new()`, the type `A` shall not be `abstract` and shall have a public parameterless constructor. This is satisfied if one of the following is true:
    - `A` is a value type, since all value types have a public default constructor (§8.3.3).
    - `A` is a type parameter having the constructor constraint (§15.2.5).
    - `A` is a type parameter having the value type constraint (§15.2.5).
    - `A` is a `class` that is not abstract and contains an explicitly declared public constructor with no parameters.
    - `A` is not `abstract` and has a default constructor (§15.11.5).

A compile-time error occurs if one or more of a type parameter's constraints are not satisfied by the given type arguments.

Since type parameters are not inherited, constraints are never inherited either.

> *Example*: In the following, `D` needs to specify the constraint on its type parameter `T` so that `T` satisfies the constraint imposed by the base `class` `B<T>`. In contrast, `class` `E` need not specify a constraint, because `List<T>` implements `IEnumerable` for any `T`.
>
> ```
> C#
> ```
>
> ```
> class B<T> where T: IEnumerable {...}
> class D<T> : B<T> where T: IEnumerable {...}
> class E<T> : B<List<T>> {...}
> ```
>
> *end example*

8.5 Type parameters

A type parameter is an identifier designating a value type or reference type that the parameter is bound to at run-time.

ANTLR

```
type_parameter
    : identifier
    ;
```

Since a type parameter can be instantiated with many different type arguments, type parameters have slightly different operations and restrictions than other types.

> *Note*: These include:
>
> - A type parameter cannot be used directly to declare a base class (§15.2.4.2) or interface (§18.2.4).
> - The rules for member lookup on type parameters depend on the constraints, if any, applied to the type parameter. They are detailed in §12.5.
> - The available conversions for a type parameter depend on the constraints, if any, applied to the type parameter. They are detailed in §10.2.12 and §10.3.9.
> - The literal `null` cannot be converted to a type given by a type parameter, except if the type parameter is known to be a reference type (§10.2.12). However, a default expression (§12.8.20) can be used instead. In addition, a value with a type given by a type parameter *can* be compared with null using `==` and `!=` (§12.12.7) unless the type parameter has the value type constraint.
> - A `new` expression (§12.8.16.2) can only be used with a type parameter if the type parameter is constrained by a *constructor_constraint* or the value type constraint (§15.2.5).
> - A type parameter cannot be used anywhere within an attribute.
> - A type parameter cannot be used in a member access (§12.8.7) or type name (§7.8) to identify a static member or a nested type.
> - A type parameter cannot be used as an *unmanaged_type* (§8.8).
>
> *end note*

As a type, type parameters are purely a compile-time construct. At run-time, each type parameter is bound to a run-time type that was specified by supplying a type argument to the generic type declaration. Thus, the type of a variable declared with a type parameter will, at run-time, be a closed constructed type §8.4.3. The run-time execution of all statements and expressions involving type parameters uses the type that was supplied as the type argument for that parameter.

8.6 Expression tree types

Expression trees permit lambda expressions to be represented as data structures instead of executable code. Expression trees are values of ***expression tree types*** of the form `System.Linq.Expressions.Expression<TDelegate>`, where `TDelegate` is any delegate type. For the remainder of this specification these types will be referred to using the shorthand `Expression<TDelegate>`.

If a conversion exists from a lambda expression to a delegate type `D`, a conversion also exists to the expression tree type `Expression<TDelegate>`. Whereas the conversion of a lambda expression to a delegate type generates a delegate that references executable code for the lambda expression, conversion to an expression tree type creates an expression tree representation of the lambda expression. More details of this conversion are provided in §10.7.3.

> *Example*: The following program represents a lambda expression both as executable code and as an expression tree. Because a conversion exists to `Func<int,int>`, a conversion also exists to `Expression<Func<int,int>>`:
>
> ```
> C#
> ```
>
> ```
> Func<int,int> del = x => x + 1;              // Code
> Expression<Func<int,int>> exp = x => x + 1; // Data
> ```
>
> Following these assignments, the delegate `del` references a method that returns `x + 1`, and the expression tree exp references a data structure that describes the expression `x => x + 1`.
>
> *end example*

`Expression<TDelegate>` provides an instance method `Compile` which produces a delegate of type `TDelegate`:

```
C#
```

```
Func<int,int> del2 = exp.Compile();
```

Invoking this delegate causes the code represented by the expression tree to be executed. Thus, given the definitions above, `del` and `del2` are equivalent, and the following two statements will have the same effect:

```
C#
```

```
int i1 = del(1);
int i2 = del2(1);
```

After executing this code, `i1` and `i2` will both have the value `2`.

The API surface provided by `Expression<TDelegate>` is implementation-specific beyond the requirement for a `Compile` method described above.

> *Note*: While the details of the API provided for expression trees are implementation-specific, it is expected that an implementation will:
>
> - Enable code to inspect and respond to the structure of an expression tree created as the result of a conversion from a lambda expression
> - Enable expression trees to be created programatically within user code
>
> *end note*

8.7 The dynamic type

The type `dynamic` uses dynamic binding, as described in detail in §12.3.2, as opposed to static binding which is used by all other types.

The type `dynamic` is considered identical to `object` except in the following respects:

- Operations on expressions of type `dynamic` can be dynamically bound (§12.3.3).
- Type inference (§12.6.3) will prefer `dynamic` over `object` if both are candidates.
- `dynamic` cannot be used as
  - the type in an *object_creation_expression* (§12.8.16.2)
  - a *class_base* (§15.2.4)
  - a *predefined_type* in a *member_access* (§12.8.7.1)
  - the operand of the `typeof` operator
  - an attribute argument
  - a constraint
  - an extension method type
  - any part of a type argument within *struct_interfaces* (§16.2.5) or *interface_type_list* (§15.2.4.1).

Because of this equivalence, the following holds:

- There is an implicit identity conversion
  - between `object` and `dynamic`

    - between constructed types that are the same when replacing `dynamic` with `object`
    - between tuple types that are the same when replacing `dynamic` with `object`
- Implicit and explicit conversions to and from `object` also apply to and from `dynamic`.
- Signatures that are the same when replacing `dynamic` with `object` are considered the same signature.
- The type `dynamic` is indistinguishable from the type `object` at run-time.
- An expression of the type `dynamic` is referred to as a ***dynamic expression***.

8.8 Unmanaged types

ANTLR

```
unmanaged_type
    : value_type
    | pointer_type     // unsafe code support
    ;
```

An *unmanaged_type* is any type that isn't a *reference_type*, a *type_parameter*, or a constructed type, and contains no instance fields whose type is not an *unmanaged_type*. In other words, an *unmanaged_type* is one of the following:

- `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, `double`, `decimal`, or `bool`.
- Any *enum_type*.
- Any user-defined *struct_type* that is not a constructed type and contains instance fields of *unmanaged_type*s only.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

9 Variables

Article • 02/07/2024

9.1 General

Variables represent storage locations. Every variable has a type that determines what values can be stored in the variable. C# is a type-safe language, and the C# compiler guarantees that values stored in variables are always of the appropriate type. The value of a variable can be changed through assignment or through use of the `++` and `--` operators.

A variable shall be *definitely assigned* (§9.4) before its value can be obtained.

As described in the following subclauses, variables are either ***initially assigned*** or ***initially unassigned***. An initially assigned variable has a well-defined initial value and is always considered definitely assigned. An initially unassigned variable has no initial value. For an initially unassigned variable to be considered definitely assigned at a certain location, an assignment to the variable shall occur in every possible execution path leading to that location.

9.2 Variable categories

9.2.1 General

C# defines eight categories of variables: static variables, instance variables, array elements, value parameters, input parameters, reference parameters, output parameters, and local variables. The subclauses that follow describe each of these categories.

> *Example*: In the following code
>
> C#
>
> ```
> class A
> {
>     public static int x;
>     int y;
>
>     void F(int[] v, int a, ref int b, out int c, in int d)
>     {
>         int i = 1;
>         c = a + b++ + d;
> ```

```
    }
}
```

`x` is a static variable, `y` is an instance variable, `v[0]` is an array element, `a` is a value parameter, `b` is a reference parameter, `c` is an output parameter, `d` is an input parameter, and `i` is a local variable. *end example*

9.2.2 Static variables

A field declared with the `static` modifier is a static variable. A static variable comes into existence before execution of the `static` constructor (§15.12) for its containing type, and ceases to exist when the associated application domain ceases to exist.

The initial value of a static variable is the default value (§9.3) of the variable's type.

For the purposes of definite-assignment checking, a static variable is considered initially assigned.

9.2.3 Instance variables

9.2.3.1 General

A field declared without the `static` modifier is an instance variable.

9.2.3.2 Instance variables in classes

An instance variable of a class comes into existence when a new instance of that class is created, and ceases to exist when there are no references to that instance and the instance's finalizer (if any) has executed.

The initial value of an instance variable of a class is the default value (§9.3) of the variable's type.

For the purpose of definite-assignment checking, an instance variable of a class is considered initially assigned.

9.2.3.3 Instance variables in structs

An instance variable of a struct has exactly the same lifetime as the struct variable to which it belongs. In other words, when a variable of a struct type comes into existence or ceases to exist, so too do the instance variables of the struct.

The initial assignment state of an instance variable of a struct is the same as that of the containing `struct` variable. In other words, when a struct variable is considered initially assigned, so too are its instance variables, and when a struct variable is considered initially unassigned, its instance variables are likewise unassigned.

9.2.4 Array elements

The elements of an array come into existence when an array instance is created, and cease to exist when there are no references to that array instance.

The initial value of each of the elements of an array is the default value (§9.3) of the type of the array elements.

For the purpose of definite-assignment checking, an array element is considered initially assigned.

9.2.5 Value parameters

A parameter declared without an `in`, `out`, or `ref` modifier is a ***value parameter***.

A value parameter comes into existence upon invocation of the function member (method, instance constructor, accessor, or operator) or anonymous function to which the parameter belongs, and is initialized with the value of the argument given in the invocation. A value parameter normally ceases to exist when execution of the function body completes. However, if the value parameter is captured by an anonymous function (§12.19.6.2), its lifetime extends at least until the delegate or expression tree created from that anonymous function is eligible for garbage collection.

For the purpose of definite-assignment checking, a value parameter is considered initially assigned.

9.2.6 Reference parameters

A parameter declared with a `ref` modifier is a ***reference parameter***.

A reference parameter is a reference variable (§9.7) which comes into existence upon invocation of the function member, delegate, anonymous function, or local function and its referent is initialized to the variable given as the argument in that invocation. A reference parameter ceases to exist when execution of the function body completes. Unlike value parameters a reference parameter shall not be captured (§9.7.2.9).

The following definite-assignment rules apply to reference parameters.

> *Note*: The rules for output parameters are different, and are described in (§9.2.7). *end note*

- A variable shall be definitely assigned (§9.4) before it can be passed as a reference parameter in a function member or delegate invocation.
- Within a function member or anonymous function, a reference parameter is considered initially assigned.

For a `struct` type, within an instance method or instance accessor (§12.2.1) or instance constructor with a constructor initializer, the `this` keyword behaves exactly as a reference parameter of the struct type (§12.8.13).

9.2.7 Output parameters

A parameter declared with an `out` modifier is an ***output parameter***.

An output parameter is a reference variable (§9.7) which comes into existence upon invocation of the function member, delegate, anonymous function, or local function and its referent is initialized to the variable given as the argument in that invocation. An output parameter ceases to exist when execution of the function body completes. Unlike value parameters an output parameter shall not be captured (§9.7.2.9).

The following definite-assignment rules apply to output parameters.

> *Note*: The rules for reference parameters are different, and are described in (§9.2.6). *end note*

- A variable need not be definitely assigned before it can be passed as an output parameter in a function member or delegate invocation.
- Following the normal completion of a function member or delegate invocation, each variable that was passed as an output parameter is considered assigned in that execution path.
- Within a function member or anonymous function, an output parameter is considered initially unassigned.
- Every output parameter of a function member, anonymous function, or local function shall be definitely assigned (§9.4) before the function member, anonymous function, or local function returns normally.

9.2.8 Input parameters

A parameter declared with an `in` modifier is an ***input parameter***.

An input parameter is a reference variable (§9.7) which comes into existence upon invocation of the function member, delegate, anonymous function, or local function and its referent is initialized to the *variable_reference* given as the argument in that invocation. An input parameter ceases to exist when execution of the function body completes. Unlike value parameters an input parameter shall not be captured (§9.7.2.9).

The following definite assignment rules apply to input parameters.

- A variable shall be definitely assigned (§9.4) before it can be passed as an input parameter in a function member or delegate invocation.
- Within a function member, anonymous function, or local function an input parameter is considered initially assigned.

9.2.9 Local variables

A ***local variable*** is declared by a *local_variable_declaration*, *declaration_expression*, *foreach_statement*, or *specific_catch_clause* of a *try_statement*. A local variable can also be declared by certain kinds of *patterns* (§11). For a *foreach_statement*, the local variable is an iteration variable (§13.9.5). For a *specific_catch_clause*, the local variable is an exception variable (§13.11). A local variable declared by a *foreach_statement* or *specific_catch_clause* is considered initially assigned.

A *local_variable_declaration* can occur in a *block*, a *for_statement*, a *switch_block*, or a *using_statement*. A *declaration_expression* can occur as an `out` *argument_value*, and as a *tuple_element* that is the target of a deconstructing assignment (§12.21.2).

The lifetime of a local variable is the portion of program execution during which storage is guaranteed to be reserved for it. This lifetime extends from entry into the scope with which it is associated, at least until execution of that scope ends in some way. (Entering an enclosed *block*, calling a method, or yielding a value from an iterator block suspends, but does not end, execution of the current scope.) If the local variable is captured by an anonymous function (§12.19.6.2), its lifetime extends at least until the delegate or expression tree created from the anonymous function, along with any other objects that come to reference the captured variable, are eligible for garbage collection. If the parent scope is entered recursively or iteratively, a new instance of the local variable is created each time, and its initializer, if any, is evaluated each time.

> *Note*: A local variable is instantiated each time its scope is entered. This behavior is visible to user code containing anonymous methods. *end note*

> *Note*: The lifetime of an *iteration variable* (§13.9.5) declared by a *foreach_statement* is a single iteration of that statement. Each iteration creates a new variable. *end note*

> *Note*: The actual lifetime of a local variable is implementation-dependent. For example, a compiler might statically determine that a local variable in a block is only used for a small portion of that block. Using this analysis, the compiler could generate code that results in the variable's storage having a shorter lifetime than its containing block.
>
> The storage referred to by a local reference variable is reclaimed independently of the lifetime of that local reference variable (§7.9).
>
> *end note*

A local variable introduced by a *local_variable_declaration* or *declaration_expression* is not automatically initialized and thus has no default value. Such a local variable is considered initially unassigned.

> *Note*: A *local_variable_declaration* that includes an initializer is still initially unassigned. Execution of the declaration behaves exactly like an assignment to the variable (§9.4.4.5). Using a variable before its initializer has been executed; e.g., within the initializer expression itself or by using a *goto_statement* which bypasses the initializer; is a compile-time error:
>
> ```
> C#
> ```
>
> ```
> goto L;
>
> int x = 1; // never executed
>
> L: x += 1; // error: x not definitely assigned
> ```
>
> Within the scope of a local variable, it is a compile-time error to refer to that local variable in a textual position that precedes its declarator.
>
> *end note*

9.2.9.1 Discards

A ***discard*** is a local variable that has no name. A discard is introduced by a declaration expression (§12.17) with the identifier `_`; and is either implicitly typed (`_` or `var _`) or explicitly typed (`T _`).

> *Note*: `_` is a valid identifier in many forms of declarations. *end note*

Because a discard has no name, the only reference to the variable it represents is the expression that introduces it.

> *Note*: A discard can however be passed as an out argument, allowing the out parameter to denote its associated storage location. *end note*

A discard is not initially assigned, so it is always an error to access its value.

> *Example*:
>
> ```
> C#
> ```
>
> ```
> _ = "Hello".Length;
> (int, int, int) M(out int i1, out int i2, out int i3) { ... }
> (int _, var _, _) = M(out int _, out var _, out _);
> ```
>
> The example assumes that there is no declaration of the name `_` in scope.
>
> The assignment to `_` shows a simple pattern for ignoring the result of an expression. The call of `M` shows the different forms of discards available in tuples and as out parameters.
>
> *end example*

9.3 Default values

The following categories of variables are automatically initialized to their default values:

- Static variables.
- Instance variables of class instances.
- Array elements.

The default value of a variable depends on the type of the variable and is determined as follows:

- For a variable of a *value_type*, the default value is the same as the value computed by the *value_type*'s default constructor (§8.3.3).
- For a variable of a *reference_type*, the default value is `null`.

> *Note*: Initialization to default values is typically done by having the memory manager or garbage collector initialize memory to all-bits-zero before it is allocated for use. For this reason, it is convenient to use all-bits-zero to represent the null reference. *end note*

9.4 Definite assignment

9.4.1 General

At a given location in the executable code of a function member or an anonymous function, a variable is said to be ***definitely assigned*** if the compiler can prove, by a particular static flow analysis (§9.4.4), that the variable has been automatically initialized or has been the target of at least one assignment.

> *Note*: Informally stated, the rules of definite assignment are:
>
> - An initially assigned variable (§9.4.2) is always considered definitely assigned.
> - An initially unassigned variable (§9.4.3) is considered definitely assigned at a given location if all possible execution paths leading to that location contain at least one of the following:
>   - A simple assignment (§12.21.2) in which the variable is the left operand.
>   - An invocation expression (§12.8.9) or object creation expression (§12.8.16.2) that passes the variable as an output parameter.
>   - For a local variable, a local variable declaration for the variable (§13.6.2) that includes a variable initializer.
>
> The formal specification underlying the above informal rules is described in §9.4.2, §9.4.3, and §9.4.4.
>
> *end note*

The definite-assignment states of instance variables of a *struct_type* variable are tracked individually as well as collectively. In additional to the rules described in §9.4.2, §9.4.3, and §9.4.4, the following rules apply to *struct_type* variables and their instance variables:

- An instance variable is considered definitely assigned if its containing *struct_type* variable is considered definitely assigned.
- A *struct_type* variable is considered definitely assigned if each of its instance variables is considered definitely assigned.

Definite assignment is a requirement in the following contexts:

- A variable shall be definitely assigned at each location where its value is obtained.

  > *Note*: This ensures that undefined values never occur. *end note*

The occurrence of a variable in an expression is considered to obtain the value of the variable, except when
    - the variable is the left operand of a simple assignment,
    - the variable is passed as an output parameter, or
    - the variable is a *struct_type* variable and occurs as the left operand of a member access.

- A variable shall be definitely assigned at each location where it is passed as a reference parameter.

  > *Note*: This ensures that the function member being invoked can consider the reference parameter initially assigned. *end note*

- A variable shall be definitely assigned at each location where it is passed as an input parameter.

  > *Note*: This ensures that the function member being invoked can consider the input parameter initially assigned. *end note*

- All output parameters of a function member shall be definitely assigned at each location where the function member returns (through a return statement or through execution reaching the end of the function member body).

  > *Note*: This ensures that function members do not return undefined values in output parameters, thus enabling the compiler to consider a function member invocation that takes a variable as an output parameter equivalent to an assignment to the variable. *end note*

- The `this` variable of a *struct_type* instance constructor shall be definitely assigned at each location where that instance constructor returns.

9.4.2 Initially assigned variables

The following categories of variables are classified as initially assigned:

- Static variables.
- Instance variables of class instances.
- Instance variables of initially assigned struct variables.
- Array elements.
- Value parameters.
- Reference parameters.

- Input parameters.
- Variables declared in a `catch` clause or a `foreach` statement.

9.4.3 Initially unassigned variables

The following categories of variables are classified as initially unassigned:

- Instance variables of initially unassigned struct variables.
- Output parameters, including the `this` variable of struct instance constructors without a constructor initializer.
- Local variables, except those declared in a `catch` clause or a `foreach` statement.

9.4.4 Precise rules for determining definite assignment

9.4.4.1 General

In order to determine that each used variable is definitely assigned, the compiler shall use a process that is equivalent to the one described in this subclause.

The compiler processes the body of each function member that has one or more initially unassigned variables. For each initially unassigned variable *v*, the compiler determines a ***definite-assignment state*** for *v* at each of the following points in the function member:

- At the beginning of each statement
- At the end point (§13.2) of each statement
- On each arc which transfers control to another statement or to the end point of a statement
- At the beginning of each expression
- At the end of each expression

The definite-assignment state of *v* can be either:

- Definitely assigned. This indicates that on all possible control flows to this point, *v* has been assigned a value.
- Not definitely assigned. For the state of a variable at the end of an expression of type `bool`, the state of a variable that isn't definitely assigned might (but doesn't necessarily) fall into one of the following sub-states:
  - Definitely assigned after true expression. This state indicates that *v* is definitely assigned if the Boolean expression evaluated as true, but is not necessarily assigned if the Boolean expression evaluated as false.
  - Definitely assigned after false expression. This state indicates that *v* is definitely assigned if the Boolean expression evaluated as false, but is not necessarily

assigned if the Boolean expression evaluated as true.

The following rules govern how the state of a variable *v* is determined at each location.

9.4.4.2 General rules for statements

- *v* is not definitely assigned at the beginning of a function member body.
- The definite-assignment state of *v* at the beginning of any other statement is determined by checking the definite-assignment state of *v* on all control flow transfers that target the beginning of that statement. If (and only if) *v* is definitely assigned on all such control flow transfers, then *v* is definitely assigned at the beginning of the statement. The set of possible control flow transfers is determined in the same way as for checking statement reachability (§13.2).
- The definite-assignment state of *v* at the end point of a `block`, `checked`, `unchecked`, `if`, `while`, `do`, `for`, `foreach`, `lock`, `using`, or `switch` statement is determined by checking the definite-assignment state of *v* on all control flow transfers that target the end point of that statement. If *v* is definitely assigned on all such control flow transfers, then *v* is definitely assigned at the end point of the statement. Otherwise, *v* is not definitely assigned at the end point of the statement. The set of possible control flow transfers is determined in the same way as for checking statement reachability (§13.2).

> *Note*: Because there are no control paths to an unreachable statement, *v* is definitely assigned at the beginning of any unreachable statement. *end note*

9.4.4.3 Block statements, checked, and unchecked statements

The definite-assignment state of *v* on the control transfer to the first statement of the statement list in the block (or to the end point of the block, if the statement list is empty) is the same as the definite-assignment statement of *v* before the block, `checked`, or `unchecked` statement.

9.4.4.4 Expression statements

For an expression statement *stmt* that consists of the expression *expr*:

- *v* has the same definite-assignment state at the beginning of *expr* as at the beginning of *stmt*.
- If *v* if definitely assigned at the end of *expr*, it is definitely assigned at the end point of *stmt*; otherwise, it is not definitely assigned at the end point of *stmt*.

9.4.4.5 Declaration statements

- If *stmt* is a declaration statement without initializers, then *v* has the same definite-assignment state at the end point of *stmt* as at the beginning of *stmt*.
- If *stmt* is a declaration statement with initializers, then the definite-assignment state for *v* is determined as if *stmt* were a statement list, with one assignment statement for each declaration with an initializer (in the order of declaration).

9.4.4.6 If statements

For a statement *stmt* of the form:

```
C#
if ( «expr» ) «then_stmt» else «else_stmt»
```

- *v* has the same definite-assignment state at the beginning of *expr* as at the beginning of *stmt*.
- If *v* is definitely assigned at the end of *expr*, then it is definitely assigned on the control flow transfer to *then_stmt* and to either *else_stmt* or to the end-point of *stmt* if there is no else clause.
- If *v* has the state "definitely assigned after true expression" at the end of *expr*, then it is definitely assigned on the control flow transfer to *then_stmt*, and not definitely assigned on the control flow transfer to either *else_stmt* or to the end-point of *stmt* if there is no else clause.
- If *v* has the state "definitely assigned after false expression" at the end of *expr*, then it is definitely assigned on the control flow transfer to *else_stmt*, and not definitely assigned on the control flow transfer to *then_stmt*. It is definitely assigned at the end-point of *stmt* if and only if it is definitely assigned at the end-point of *then_stmt*.
- Otherwise, *v* is considered not definitely assigned on the control flow transfer to either the *then_stmt* or *else_stmt*, or to the end-point of *stmt* if there is no else clause.

9.4.4.7 Switch statements

For a `switch` statement *stmt* with a controlling expression *expr*:

The definite-assignment state of *v* at the beginning of *expr* is the same as the state of *v* at the beginning of *stmt*.

The definite-assignment state of *v* at the beginning of a case's guard clause is

- If *v* is a pattern variable declared in the *switch_label*: "definitely assigned".
- If the switch label containing that guard clause (§13.8.3) is not reachable: "definitely assigned".
- Otherwise, the state of *v* is the same as the state of *v* after *expr*.

> *Example*: The second rule eliminates the need for the compiler to issue an error if an unassigned variable is accessed in unreachable code. The state of *b* is "definitely assigned" in the unreachable switch label `case 2 when b`.
>
> ```
> C#
> ```
>
> ```
> bool b;
> switch (1)
> {
>     case 2 when b: // b is definitely assigned here.
>     break;
> }
> ```
>
> *end example*

The definite-assignment state of *v* on the control flow transfer to a reachable switch block statement list is

- If the control transfer was due to a 'goto case' or 'goto default' statement, then the state of *v* is the same as the state at the beginning of that 'goto' statement.
- If the control transfer was due to the `default` label of the switch, then the state of *v* is the same as the state of *v* after *expr*.
- If the control transfer was due to an unreachable switch label, then the state of *v* is "definitely assigned".
- If the control transfer was due to a reachable switch label with a guard clause, then the state of *v* is the same as the state of *v* after the guard clause.
- If the control transfer was due to a reachable switch label without a guard clause, then the state of *v* is
  - If *v* is a pattern variable declared in the *switch_label*: "definitely assigned".
  - Otherwise, the state of *v* is the same as the stat of *v* after *expr*.

A consequence of these rules is that a pattern variable declared in a *switch_label* will be "not definitely assigned" in the statements of its switch section if it is not the only reachable switch label in its section.

> *Example*:
>
> ```
> C#
> ```

```
public static double ComputeArea(object shape)
{
    switch (shape)
    {
        case Square s when s.Side == 0:
        case Circle c when c.Radius == 0:
        case Triangle t when t.Base == 0 || t.Height == 0:
        case Rectangle r when r.Length == 0 || r.Height == 0:
            // none of s, c, t, or r is definitely assigned
            return 0;
        case Square s:
            // s is definitely assigned
            return s.Side * s.Side;
        case Circle c:
            // c is definitely assigned
            return c.Radius * c.Radius * Math.PI;
           …
    }
}
```

end example

9.4.4.8 While statements

For a statement *stmt* of the form:

```
C#

while ( «expr» ) «while_body»
```

- *v* has the same definite-assignment state at the beginning of *expr* as at the beginning of *stmt*.
- If *v* is definitely assigned at the end of *expr*, then it is definitely assigned on the control flow transfer to *while_body* and to the end point of *stmt*.
- If *v* has the state "definitely assigned after true expression" at the end of *expr*, then it is definitely assigned on the control flow transfer to *while_body*, but not definitely assigned at the end-point of *stmt*.
- If *v* has the state "definitely assigned after false expression" at the end of *expr*, then it is definitely assigned on the control flow transfer to the end point of *stmt*, but not definitely assigned on the control flow transfer to *while_body*.

9.4.4.9 Do statements

For a statement *stmt* of the form:

```
C#

do «do_body» while ( «expr» ) ;
```

- *v* has the same definite-assignment state on the control flow transfer from the beginning of *stmt* to *do_body* as at the beginning of *stmt*.
- *v* has the same definite-assignment state at the beginning of *expr* as at the end point of *do_body*.
- If *v* is definitely assigned at the end of *expr*, then it is definitely assigned on the control flow transfer to the end point of *stmt*.
- If *v* has the state "definitely assigned after false expression" at the end of *expr*, then it is definitely assigned on the control flow transfer to the end point of *stmt*, but not definitely assigned on the control flow transfer to *do_body*.

9.4.4.10 For statements

For a statement of the form:

```
C#

for ( «for_initializer» ; «for_condition» ; «for_iterator» )
    «embedded_statement»
```

definite-assignment checking is done as if the statement were written:

```
C#

{
    «for_initializer» ;
    while ( «for_condition» )
    {
        «embedded_statement» ;
        LLoop: «for_iterator» ;
    }
}
```

with `continue` statements that target the `for` statement being translated to `goto` statements targeting the label `LLoop`. If the *for_condition* is omitted from the `for` statement, then evaluation of definite-assignment proceeds as if *for_condition* were replaced with true in the above expansion.

9.4.4.11 Break, continue, and goto statements

The definite-assignment state of *v* on the control flow transfer caused by a `break`, `continue`, or `goto` statement is the same as the definite-assignment state of *v* at the beginning of the statement.

9.4.4.12 Throw statements

For a statement *stmt* of the form:

```
C#

throw «expr» ;
```

the definite-assignment state of *v* at the beginning of *expr* is the same as the definite-assignment state of *v* at the beginning of *stmt*.

9.4.4.13 Return statements

For a statement *stmt* of the form:

```
C#

return «expr» ;
```

- The definite-assignment state of *v* at the beginning of *expr* is the same as the definite-assignment state of *v* at the beginning of *stmt*.
- If *v* is an output parameter, then it shall be definitely assigned either:
  - after *expr*
  - or at the end of the `finally` block of a `try`-`finally` or `try`-`catch`-`finally` that encloses the `return` statement.

For a statement *stmt* of the form:

```
C#

return ;
```

- If *v* is an output parameter, then it shall be definitely assigned either:
  - before *stmt*
  - or at the end of the `finally` block of a `try`-`finally` or `try`-`catch`-`finally` that encloses the `return` statement.

9.4.4.14 Try-catch statements

For a statement *stmt* of the form:

```
C#

try «try_block»
catch ( ... ) «catch_block_1»
...
catch ( ... ) «catch_block_n»
```

- The definite-assignment state of *v* at the beginning of *try_block* is the same as the definite-assignment state of *v* at the beginning of *stmt*.
- The definite-assignment state of *v* at the beginning of *catch_block_i* (for any *i*) is the same as the definite-assignment state of *v* at the beginning of *stmt*.
- The definite-assignment state of *v* at the end-point of *stmt* is definitely assigned if (and only if) *v* is definitely assigned at the end-point of *try_block* and every *catch_block_i* (for every *i* from 1 to *n*).

9.4.4.15 Try-finally statements

For a statement *stmt* of the form:

```
C#

try «try_block» finally «finally_block»
```

- The definite-assignment state of *v* at the beginning of *try_block* is the same as the definite-assignment state of *v* at the beginning of *stmt*.
- The definite-assignment state of *v* at the beginning of *finally_block* is the same as the definite-assignment state of *v* at the beginning of *stmt*.
- The definite-assignment state of *v* at the end-point of *stmt* is definitely assigned if (and only if) at least one of the following is true:
  - *v* is definitely assigned at the end-point of *try_block*
  - *v* is definitely assigned at the end-point of *finally_block*

If a control flow transfer (such as a `goto` statement) is made that begins within *try_block*, and ends outside of *try_block*, then *v* is also considered definitely assigned on that control flow transfer if *v* is definitely assigned at the end-point of *finally_block*. (This is not an only if—if *v* is definitely assigned for another reason on this control flow transfer, then it is still considered definitely assigned.)

9.4.4.16 Try-catch-finally statements

For a statement of the form:

```
C#

try «try_block»
catch ( ... ) «catch_block_1»
...
catch ( ... ) «catch_block_n»
finally «finally_block»
```

definite-assignment analysis is done as if the statement were a `try`-`finally` statement enclosing a `try`-`catch` statement:

```
C#

try
{
    try «try_block»
    catch ( ... ) «catch_block_1»
    ...
    catch ( ... ) «catch_block_n»
}
finally «finally_block»
```

Example: The following example demonstrates how the different blocks of a `try` statement (§13.11) affect definite assignment.

```
C#

class A
{
    static void F()
    {
        int i, j;
        try
        {
            goto LABEL;
            // neither i nor j definitely assigned
            i = 1;
            // i definitely assigned
        }
        catch
        {
            // neither i nor j definitely assigned
            i = 3;
            // i definitely assigned
        }
```

```
        finally
        {
            // neither i nor j definitely assigned
            j = 5;
            // j definitely assigned
        }
        // i and j definitely assigned
        LABEL: ;
        // j definitely assigned
    }
}
```

end example

9.4.4.17 Foreach statements

For a statement *stmt* of the form:

```
C#

foreach ( «type» «identifier» in «expr» ) «embedded_statement»
```

- The definite-assignment state of *v* at the beginning of *expr* is the same as the state of *v* at the beginning of *stmt*.
- The definite-assignment state of *v* on the control flow transfer to *embedded_statement* or to the end point of *stmt* is the same as the state of *v* at the end of *expr*.

9.4.4.18 Using statements

For a statement *stmt* of the form:

```
C#

using ( «resource_acquisition» ) «embedded_statement»
```

- The definite-assignment state of *v* at the beginning of *resource_acquisition* is the same as the state of *v* at the beginning of *stmt*.
- The definite-assignment state of *v* on the control flow transfer to *embedded_statement* is the same as the state of *v* at the end of *resource_acquisition*.

9.4.4.19 Lock statements

For a statement *stmt* of the form:

```
C#

lock ( «expr» ) «embedded_statement»
```

- The definite-assignment state of *v* at the beginning of *expr* is the same as the state of *v* at the beginning of *stmt*.
- The definite-assignment state of *v* on the control flow transfer to *embedded_statement* is the same as the state of *v* at the end of *expr*.

9.4.4.20 Yield statements

For a statement *stmt* of the form:

```
C#

yield return «expr» ;
```

- The definite-assignment state of *v* at the beginning of *expr* is the same as the state of *v* at the beginning of *stmt*.
- The definite-assignment state of *v* at the end of *stmt* is the same as the state of *v* at the end of *expr*.

A `yield break` statement has no effect on the definite-assignment state.

9.4.4.21 General rules for constant expressions

The following applies to any constant expression, and takes priority over any rules from the following sections that might apply:

For a constant expression with value `true`:

- If *v* is definitely assigned before the expression, then *v* is definitely assigned after the expression.
- Otherwise *v* is "definitely assigned after false expression" after the expression.

> *Example*:
>
> ```
> C#
>
> int x;
> if (true) {}
> else
> ```

```
{
    Console.WriteLine(x);
}
```

end example

For a constant expression with value `false`:

- If *v* is definitely assigned before the expression, then *v* is definitely assigned after the expression.
- Otherwise *v* is "definitely assigned after true expression" after the expression.

Example:

```
C#

int x;
if (false)
{
    Console.WriteLine(x);
}
```

end example

For all other constant expressions, the definite-assignment state of *v* after the expression is the same as the definite-assignment state of *v* before the expression.

9.4.4.22 General rules for simple expressions

The following rule applies to these kinds of expressions: literals (§12.8.2), simple names (§12.8.4), member access expressions (§12.8.7), non-indexed base access expressions (§12.8.14), `typeof` expressions (§12.8.17), default value expressions (§12.8.20), `nameof` expressions (§12.8.22), and declaration expressions (§12.17).

- The definite-assignment state of *v* at the end of such an expression is the same as the definite-assignment state of *v* at the beginning of the expression.

9.4.4.23 General rules for expressions with embedded expressions

The following rules apply to these kinds of expressions: parenthesized expressions (§12.8.5), tuple expressions (§12.8.6), element access expressions (§12.8.11), base access expressions with indexing (§12.8.14), increment and decrement expressions (§12.8.15, §12.9.6), cast expressions (§12.9.7), unary `+`, `-`, `~`, `*` expressions, binary `+`, `-`, `*`, `/`, `%`,

`<<`, `>>`, `<`, `<=`, `>`, `>=`, `==`, `!=`, `is`, `as`, `&`, `|`, `^` expressions (§12.10, §12.11, §12.12, §12.13), compound assignment expressions (§12.21.4), `checked` and `unchecked` expressions (§12.8.19), array and delegate creation expressions (§12.8.16) , and `await` expressions (§12.9.8).

Each of these expressions has one or more subexpressions that are unconditionally evaluated in a fixed order.

> *Example*: The binary `%` operator evaluates the left hand side of the operator, then the right hand side. An indexing operation evaluates the indexed expression, and then evaluates each of the index expressions, in order from left to right. *end example*

For an expression *expr*, which has subexpressions $expr_1$, $expr_2$, …, $expr_x$, evaluated in that order:

- The definite-assignment state of *v* at the beginning of $expr_1$ is the same as the definite-assignment state at the beginning of *expr*.
- The definite-assignment state of *v* at the beginning of $expr_i$ (*i* greater than one) is the same as the definite-assignment state at the end of $expr_{i-1}$.
- The definite-assignment state of *v* at the end of *expr* is the same as the definite-assignment state at the end of $expr_x$.

9.4.4.24 Invocation expressions and object creation expressions

If the method to be invoked is a partial method that has no implementing partial method declaration, or is a conditional method for which the call is omitted (§22.5.3.2), then the definite-assignment state of *v* after the invocation is the same as the definite-assignment state of *v* before the invocation. Otherwise the following rules apply:

For an invocation expression *expr* of the form:

```
C#
«primary_expression» ( «arg₁», «arg₂», … , «argₓ» )
```

or an object-creation expression *expr* of the form:

```
C#
new «type» ( «arg₁», «arg₂», … , «argₓ» )
```

- For an invocation expression, the definite assignment state of *v* before *primary_expression* is the same as the state of *v* before *expr*.
- For an invocation expression, the definite assignment state of *v* before arg_1 is the same as the state of *v* after *primary_expression*.
- For an object creation expression, the definite assignment state of *v* before arg_1 is the same as the state of *v* before *expr*.
- For each argument arg_i, the definite assignment state of *v* after arg_i is determined by the normal expression rules, ignoring any `in`, `out`, or `ref` modifiers.
- For each argument arg_i for any *i* greater than one, the definite assignment state of *v* before arg_i is the same as the state of *v* after arg_{i-1}.
- If the variable *v* is passed as an `out` argument (i.e., an argument of the form "out *v*") in any of the arguments, then the state of *v* after *expr* is definitely assigned. Otherwise, the state of *v* after *expr* is the same as the state of *v* after arg_x.
- For array initializers (§12.8.16.5), object initializers (§12.8.16.3), collection initializers (§12.8.16.4) and anonymous object initializers (§12.8.16.7), the definite-assignment state is determined by the expansion that these constructs are defined in terms of.

9.4.4.25 Simple assignment expressions

Let the set of *assignment targets* in an expression *e* be defined as follows:

- If *e* is a tuple expression, then the assignment targets in *e* are the union of the assignment targets of the elements of *e*.
- Otherwise, the assignment targets in *e* are *e*.

For an expression *expr* of the form:

```
C#
«expr_lhs» = «expr_rhs»
```

- The definite-assignment state of *v* before *expr_lhs* is the same as the definite-assignment state of *v* before *expr*.
- The definite-assignment state of *v* before *expr_rhs* is the same as the definite-assignment state of *v* after *expr_lhs*.
- If *v* is an assignment target of *expr_lhs*, then the definite-assignment state of *v* after *expr* is definitely assigned. Otherwise, if the assignment occurs within the instance constructor of a struct type, and *v* is the hidden backing field of an automatically implemented property *P* on the instance being constructed, and a property access designating *P* is an assigment target of *expr_lhs*, then the definite-assignment state of *v* after *expr* is definitely assigned. Otherwise, the definite-

assignment state of *v* after *expr* is the same as the definite-assignment state of *v* after *expr_rhs*.

Example: In the following code

```
C#

class A
{
    static void F(int[] arr)
    {
        int x;
        arr[x = 1] = x; // ok
    }
}
```

the variable `x` is considered definitely assigned after `arr[x = 1]` is evaluated as the left hand side of the second simple assignment.

end example

9.4.4.26 && expressions

For an expression *expr* of the form:

```
C#

«expr_first» && «expr_second»
```

- The definite-assignment state of *v* before *expr_first* is the same as the definite-assignment state of *v* before *expr*.
- The definite-assignment state of *v* before *expr_second* is definitely assigned if and only if the state of *v* after *expr_first* is either definitely assigned or "definitely assigned after true expression". Otherwise, it is not definitely assigned.
- The definite-assignment state of *v* after *expr* is determined by:
  - If the state of *v* after *expr_first* is definitely assigned, then the state of *v* after *expr* is definitely assigned.
  - Otherwise, if the state of *v* after *expr_second* is definitely assigned, and the state of *v* after *expr_first* is "definitely assigned after false expression", then the state of *v* after *expr* is definitely assigned.
  - Otherwise, if the state of *v* after *expr_second* is definitely assigned or "definitely assigned after true expression", then the state of *v* after *expr* is "definitely assigned after true expression".

- Otherwise, if the state of *v* after *expr_first* is "definitely assigned after false expression", and the state of *v* after *expr_second* is "definitely assigned after false expression", then the state of *v* after *expr* is "definitely assigned after false expression".
- Otherwise, the state of *v* after *expr* is not definitely assigned.

> *Example*: In the following code
>
> ```
> C#
> ```
>
> ```
> class A
> {
>     static void F(int x, int y)
>     {
>         int i;
>         if (x >= 0 && (i = y) >= 0)
>         {
>             // i definitely assigned
>         }
>         else
>         {
>             // i not definitely assigned
>         }
>         // i not definitely assigned
>     }
> }
> ```
>
> the variable `i` is considered definitely assigned in one of the embedded statements of an `if` statement but not in the other. In the `if` statement in method `F`, the variable `i` is definitely assigned in the first embedded statement because execution of the expression `(i = y)` always precedes execution of this embedded statement. In contrast, the variable `i` is not definitely assigned in the second embedded statement, since `x >= 0` might have tested false, resulting in the variable `i`'s being unassigned.
>
> *end example*

9.4.4.27 || expressions

For an expression *expr* of the form:

```
C#
```

```
«expr_first» || «expr_second»
```

- The definite-assignment state of *v* before *expr_first* is the same as the definite-assignment state of *v* before *expr*.
- The definite-assignment state of *v* before *expr_second* is definitely assigned if and only if the state of *v* after *expr_first* is either definitely assigned or "definitely assigned after true expression". Otherwise, it is not definitely assigned.
- The definite-assignment statement of *v* after *expr* is determined by:
  - If the state of *v* after *expr_first* is definitely assigned, then the state of *v* after *expr* is definitely assigned.
  - Otherwise, if the state of *v* after *expr_second* is definitely assigned, and the state of *v* after *expr_first* is "definitely assigned after true expression", then the state of *v* after *expr* is definitely assigned.
  - Otherwise, if the state of *v* after *expr_second* is definitely assigned or "definitely assigned after false expression", then the state of *v* after *expr* is "definitely assigned after false expression".
  - Otherwise, if the state of *v* after *expr_first* is "definitely assigned after true expression", and the state of *v* after *expr_ second* is "definitely assigned after true expression", then the state of *v* after *expr* is "definitely assigned after true expression".
  - Otherwise, the state of *v* after *expr* is not definitely assigned.

> *Example*: In the following code
>
> C#
>
> ```
> class A
> {
>     static void G(int x, int y)
>     {
>         int i;
>         if (x >= 0 || (i = y) >= 0)
>         {
>             // i not definitely assigned
>         }
>         else
>         {
>             // i definitely assigned
>         }
>         // i not definitely assigned
>     }
> }
> ```
>
> the variable `i` is considered definitely assigned in one of the embedded statements of an `if` statement but not in the other. In the `if` statement in method `G`, the variable `i` is definitely assigned in the second embedded statement because execution of the expression `(i = y)` always precedes execution of this embedded

> statement. In contrast, the variable `i` is not definitely assigned in the first embedded statement, since `x >= 0` might have tested true, resulting in the variable `i`'s being unassigned.
>
> *end example*

9.4.4.28 ! expressions

For an expression *expr* of the form:

```
C#

! «expr_operand»
```

- The definite-assignment state of *v* before *expr_operand* is the same as the definite-assignment state of *v* before *expr*.
- The definite-assignment state of *v* after *expr* is determined by:
  - If the state of `v` after *expr_operand* is definitely assigned, then the state of `v` after *expr* is definitely assigned.
  - Otherwise, if the state of `v` after *expr_operand* is "definitely assigned after false expression", then the state of `v` after *expr* is "definitely assigned after true expression".
  - Otherwise, if the state of `v` after *expr_operand* is "definitely assigned after true expression", then the state of v after *expr* is "definitely assigned after false expression".
  - Otherwise, the state of `v` after *expr* is not definitely assigned.

9.4.4.29 ?? expressions

For an expression *expr* of the form:

```
C#

«expr_first» ?? «expr_second»
```

- The definite-assignment state of *v* before *expr_first* is the same as the definite-assignment state of *v* before *expr*.
- The definite-assignment state of *v* before *expr_second* is the same as the definite-assignment state of *v* after *expr_first*.
- The definite-assignment statement of *v* after *expr* is determined by:

- If *expr_first* is a constant expression (§12.23) with value `null`, then the state of *v* after *expr* is the same as the state of *v* after *expr_second*.
- Otherwise, the state of *v* after *expr* is the same as the definite-assignment state of *v* after *expr_first*.

9.4.4.30 ?: expressions

For an expression *expr* of the form:

```
C#
«expr_cond» ? «expr_true» : «expr_false»
```

- The definite-assignment state of *v* before *expr_cond* is the same as the state of *v* before *expr*.
- The definite-assignment state of *v* before *expr_true* is definitely assigned if the state of *v* after *expr_cond* is definitely assigned or "definitely assigned after true expression".
- The definite-assignment state of *v* before *expr_false* is definitely assigned if the state of *v* after *expr_cond* is definitely assigned or "definitely assigned after false expression".
- The definite-assignment state of *v* after *expr* is determined by:
  - If *expr_cond* is a constant expression (§12.23) with value `true` then the state of *v* after *expr* is the same as the state of *v* after *expr_true*.
  - Otherwise, if *expr_cond* is a constant expression (§12.23) with value `false` then the state of *v* after *expr* is the same as the state of *v* after *expr_false*.
  - Otherwise, if the state of *v* after *expr_true* is definitely assigned and the state of *v* after *expr_false* is definitely assigned, then the state of *v* after *expr* is definitely assigned.
  - Otherwise, the state of *v* after *expr* is not definitely assigned.

9.4.4.31 Anonymous functions

For a *lambda_expression* or *anonymous_method_expression expr* with a body (either *block* or *expression*) *body*:

- The definite assignment state of a parameter is the same as for a parameter of a named method (§9.2.6, §9.2.7, §9.2.8).
- The definite assignment state of an outer variable *v* before *body* is the same as the state of *v* before *expr*. That is, definite assignment state of outer variables is inherited from the context of the anonymous function.

- The definite assignment state of an outer variable *v* after *expr* is the same as the state of *v* before *expr*.

> *Example*: The example
>
> C#
>
> ```
> class A
> {
>     delegate bool Filter(int i);
>     void F()
>     {
>         int max;
>         // Error, max is not definitely assigned
>         Filter f = (int n) => n < max;
>         max = 5;
>         DoWork(f);
>     }
>     void DoWork(Filter f) { ... }
> }
> ```
>
> generates a compile-time error since max is not definitely assigned where the anonymous function is declared.
>
> *end example*

> *Example*: The example
>
> C#
>
> ```
> class A
> {
>     delegate void D();
>     void F()
>     {
>         int n;
>         D d = () => { n = 1; };
>         d();
>         // Error, n is not definitely assigned
>         Console.WriteLine(n);
>     }
> }
> ```
>
> also generates a compile-time error since the assignment to `n` in the anonymous function has no affect on the definite-assignment state of `n` outside the anonymous function.
>
> *end example*

9.4.4.32 Throw expressions

For an expression *expr* of the form:

`throw` *thrown_expr*

- The definite assignment state of *v* before *thrown_expr* is the same as the state of *v* before *expr*.
- The definite assignment state of *v* after *expr* is "definitely assigned".

9.4.4.33 Rules for variables in local functions

Local functions are analyzed in the context of their parent method. There are two control flow paths that matter for local functions: function calls and delegate conversions.

Definite assignment for the body of each local function is defined separately for each call site. At each invocation, variables captured by the local function are considered definitely assigned if they were definitely assigned at the point of call. A control flow path also exists to the local function body at this point and is considered reachable. After a call to the local function, captured variables that were definitely assigned at every control point leaving the function (`return` statements, `yield` statements, `await` expressions) are considered definitely assigned after the call location.

Delegate conversions have a control flow path to the local function body. Captured variables are definitely assigned for the body if they are definitely assigned before the conversion. Variables assigned by the local function are not considered assigned after the conversion.

> *Note*: the above implies that bodies are re-analyzed for definite assignment at every local function invocation or delegate conversion. Compilers are not required to re-analyze the body of a local function at each invocation or delegate conversion. The implementation must produce results equivalent to that description. *end note*

> *Example*: The following example demonstrates definite assignment for captured variables in local functions. If a local function reads a captured variable before writing it, the captured variable must be definitely assigned before calling the local function. The local function `F1` reads `s` without assigning it. It is an error if `F1` is called before `s` is definitely assigned. `F2` assigns `i` before reading it. It may be called before `i` is definitely assigned. Furthermore, `F3` may be called after `F2` because `s2` is definitely assigned in `F2`.
>
> ```
> ```

C#

```
void M()
{
    string s;
    int i;
    string s2;

    // Error: Use of unassigned local variable s:
    F1();
    // OK, F2 assigns i before reading it.
    F2();

    // OK, i is definitely assigned in the body of F2:
    s = i.ToString();

    // OK. s is now definitely assigned.
    F1();

    // OK, F3 reads s2, which is definitely assigned in F2.
    F3();

    void F1()
    {
        Console.WriteLine(s);
    }

    void F2()
    {
        i = 5;
        // OK. i is definitely assigned.
        Console.WriteLine(i);
        s2 = i.ToString();
    }

    void F3()
    {
        Console.WriteLine(s2);
    }
}
```

end example

9.4.4.34 is-pattern expressions

For an expression *expr* of the form:

expr_operand is *pattern*

- The definite-assignment state of *v* before *expr_operand* is the same as the definite-assignment state of *v* before *expr*.

- If the variable 'v' is declared in *pattern*, then the definite-assignment state of 'v' after *expr* is "definitely assigned when true".
- Otherwise the definite assignment state of 'v' after *expr* is the same as the definite assignment state of 'v' after *expr_operand*.

9.5 Variable references

A *variable_reference* is an *expression* that is classified as a variable. A *variable_reference* denotes a storage location that can be accessed both to fetch the current value and to store a new value.

ANTLR

```
variable_reference
    : expression
    ;
```

> *Note*: In C and C++, a *variable_reference* is known as an *lvalue*. *end note*

9.6 Atomicity of variable references

Reads and writes of the following data types shall be atomic: `bool`, `char`, `byte`, `sbyte`, `short`, `ushort`, `uint`, `int`, `float`, and reference types. In addition, reads and writes of enum types with an underlying type in the previous list shall also be atomic. Reads and writes of other types, including `long`, `ulong`, `double`, and `decimal`, as well as user-defined types, need not be atomic. Aside from the library functions designed for that purpose, there is no guarantee of atomic read-modify-write, such as in the case of increment or decrement.

9.7 Reference variables and returns

9.7.1 General

A ***reference variable*** is a variable that refers to another variable, called the referent (§9.2.6). A reference variable is a local variable declared with the `ref` modifier.

A reference variable stores a *variable_reference* (§9.5) to its referent and not the value of its referent. When a reference variable is used where a value is required its referent's value is returned; similarly when a reference variable is the target of an assignment it is the referent which is assigned to. The variable to which a reference variable refers, i.e.

the stored *variable_reference* for its referent, can be changed using a ref assignment (`= ref`).

> *Example:* The following example demonstrates a local reference variable whose referent is an element of an array:
>
> ```
> C#
> ```
>
> ```
> public class C
> {
>     public void M()
>     {
>         int[] arr = new int[10];
>         // element is a reference variable that refers to arr[5]
>         ref int element = ref arr[5];
>         element += 5; // arr[5] has been incremented by 5
>     }
> }
> ```
>
> *end example*

A ***reference return*** is the *variable_reference* returned from a returns-by-ref method (§15.6.1). This *variable_reference* is the referent of the reference return.

> *Example:* The following example demonstrates a reference return whose referent is an element of an array field:
>
> ```
> C#
> ```
>
> ```
> public class C
> {
>     private int[] arr = new int[10];
>
>     public ref readonly int M()
>     {
>         // element is a reference variable that refers to arr[5]
>         ref int element = ref arr[5];
>         return ref element; // return reference to arr[5];
>     }
> }
> ```
>
> *end example*

9.7.2 Ref safe contexts

9.7.2.1 General

All reference variables obey safety rules that ensure the ref-safe-context of the reference variable is not greater than the ref-safe-context of its referent.

> *Note*: The related notion of a *safe-context* is defined in (§16.4.12), along with associated constraints. *end note*

For any variable, the ***ref-safe-context*** of that variable is the context where a *variable_reference* (§9.5) to that variable is valid. The referent of a reference variable shall have a ref-safe-context that is at least as wide as the ref-safe-context of the reference variable itself.

> *Note*: The compiler determines the ref-safe-context through a static analysis of the program text. The ref-safe-context reflects the lifetime of a variable at runtime. *end note*

There are three ref-safe-contexts:

- ***declaration-block***: The ref-safe-context of a *variable_reference* to a local variable (§9.2.9) is that local variable's scope (§13.6.2), including any nested *embedded-statement*s in that scope.

  A *variable_reference* to a local variable is a valid referent for a reference variable only if the reference variable is declared within the ref-safe-context of that variable.

- ***function-member***: Within a function a *variable_reference* to any of the following has a ref-safe-context of function-member:
  - Value parameters (§9.2.5) on a function member declaration, including the implicit `this` of class member functions; and
  - The implicit reference (`ref`) parameter (§9.2.6) `this` of a struct member function, along with its fields.

  A *variable_reference* with ref-safe-context of function-member is a valid referent only if the reference variable is declared in the same function member.

- ***caller-context***: Within a function a *variable_reference* to any of the following has a ref-safe-context of caller-context:
  - Reference (`ref`) parameters (§9.2.6) other than the implicit `this` of a struct member function;
  - Member fields and elements of such parameters;
  - Member fields of parameters of class type; and
  - Elements of parameters of array type.

A *variable_reference* with ref-safe-context of caller-context can be the referent of a reference return.

These values form a nesting relationship from narrowest (declaration-block) to widest (caller-context). Each nested block represents a different context.

> *Example*: The following code shows examples of the different ref-safe-contexts. The declarations show the ref-safe-context for a referent to be the initializing expression for a `ref` variable. The examples show the ref-safe-context for a reference return:

C#

```
public class C
{
    // ref safe context of arr is "caller-context".
    // ref safe context of arr[i] is "caller-context".
    private int[] arr = { 0, 1, 2, 3, 4, 5, 6, 7, 8, 9 };

    // ref safe context is "caller-context"
    public ref int M1(ref int r1)
    {
        return ref r1; // r1 is safe to ref return
    }

    // ref safe context is "function-member"
    public ref int M2(int v1)
    {
        return ref v1; // error: v1 isn't safe to ref return
    }

    public ref int M3()
    {
        int v2 = 5;

        return ref arr[v2]; // arr[v2] is safe to ref return
    }

    public void M4(int p)
    {
        int v3 = 6;

        // context of r2 is declaration-block,
        // ref safe context of p is function-member
        ref int r2 = ref p;

        // context of r3 is declaration-block,
        // ref safe context of v3 is declaration-block
        ref int r3 = ref v3;

        // context of r4 is declaration-block,
        // ref safe context of arr[v3] is caller-context
        ref int r4 = ref arr[v3];
```

```
    }
}
```

end example.

Example: For `struct` types, the implicit `this` parameter is passed as a `ref` parameter. The ref-safe-context of the fields of a `struct` type as function-member prevents returning those fields by reference return. This rule prevents the following code:

```C#
public struct S
{
     private int n;

     // Disallowed: returning ref of a field.
     public ref int GetN() => ref n;
}

class Test
{
    public ref int M()
    {
        S s = new S();
        ref int numRef = ref s.GetN();
        return ref numRef; // reference to local variable 'numRef'
returned
    }
}
```

end example.

9.7.2.2 Local variable ref safe context

For a local variable `v`:

- If `v` is a reference variable, its ref-safe-context is the same as the ref-safe-context of its initializing expression.
- Otherwise its ref-safe-context is declaration-block.

9.7.2.3 Parameter ref safe context

For a formal parameter `p`:

- If `p` is a `ref`, or `in` parameter, its ref-safe-context is the caller-context. If `p` is an `in` parameter, it can't be returned as a writable `ref` but can be returned as `ref readonly`.
- If `p` is an `out` parameter, its ref-safe-context is the caller-context.
- Otherwise, if `p` is the `this` parameter of a struct type, its ref-safe-context is the function-member.
- Otherwise, the parameter is a value parameter, and its ref-safe-context is the function-member.

9.7.2.4 Field ref safe context

For a variable designating a reference to a field, `e.F`:

- If `e` is of a reference type, its ref-safe-context is the caller-context.
- Otherwise, if `e` is of a value type, its ref-safe-context is the same as the ref-safe-context of `e`.

9.7.2.5 Operators

The conditional operator (§12.18), `c ? ref e1 : ref e2`, and reference assignment operator, `= ref e` (§12.21.1) have reference variables as operands and yield a reference variable. For those operators, the ref-safe-context of the result is the narrowest context among the ref-safe-contexts of all `ref` operands.

9.7.2.6 Function invocation

For a variable `c` resulting from a ref-returning function invocation, its ref-safe-context is the narrowest of the following contexts:

- The caller-context.
- The ref-safe-context of all `ref`, `out`, and `in` argument expressions (excluding the receiver).
- For each `in` parameter, if there is a corresponding expression that is a variable and there exists an identity conversion between the type of the variable and the type of the parameter, the variable's ref-safe-context, otherwise the nearest enclosing context.
- The safe-context (§16.4.12) of all argument expressions (including the receiver).

> *Example*: the last bullet is necessary to handle code such as
>
> ```
> 
> ```

```C#
ref int M2()
{
    int v = 5;
    // Not valid.
    // ref safe context of "v" is block.
    // Therefore, ref safe context of the return value of M() is block.
    return ref M(ref v);
}

ref int M(ref int p)
{
    return ref p;
}
```

end example

A property invocation and an indexer invocation (either `get` or `set`) is treated as a function invocation of the underlying accessor by the above rules. A local function invocation is a function invocation.

9.7.2.7 Values

A value's ref-safe-context is the nearest enclosing context.

> *Note*: This occurs in an invocation such as `M(ref d.Length)` where `d` is of type `dynamic`. It is also consistent with arguments corresponding to `in` parameters. *end note*

9.7.2.8 Constructor invocations

A `new` expression that invokes a constructor obeys the same rules as a method invocation (§9.7.2.6) that is considered to return the type being constructed.

9.7.2.9 Limitations on reference variables

- Neither a reference parameter, nor an output parameter, nor an input parameter, nor a `ref` local, nor a parameter or local of a `ref struct` type shall be captured by lambda expression or local function.
- Neither a reference parameter, nor an output parameter, nor an input parameter, nor a parameter of a `ref struct` type shall be an argument for an iterator method or an `async` method.

- Neither a `ref` local, nor a local of a `ref struct` type shall be in context at the point of a `yield return` statement or an `await` expression.
- For a ref reassignment `e1 = ref e2`, the ref-safe-context of `e2` shall be at least as wide a context as the *ref-safe-context* of `e1`.
- For a ref return statement `return ref e1`, the ref-safe-context of `e1` shall be the caller-context.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

10 Conversions

Article • 02/07/2024

10.1 General

A ***conversion*** causes an expression to be converted to, or treated as being of, a particular type; in the former case a conversion may involve a change in representation. Conversions can be ***implicit*** or ***explicit***, and this determines whether an explicit cast is required.

> *Example*: For instance, the conversion from type `int` to type `long` is implicit, so expressions of type `int` can implicitly be treated as type `long`. The opposite conversion, from type `long` to type `int`, is explicit and so an explicit cast is required.
>
> ```
> C#
> ```
>
> ```
> int a = 123;
> long b = a;      // implicit conversion from int to long
> int c = (int) b; // explicit conversion from long to int
> ```
>
> *end example*

Some conversions are defined by the language. Programs may also define their own conversions (§10.5).

Some conversions in the language are defined from expressions to types, others from types to types. A conversion from a type applies to all expressions that have that type.

> *Example*:
>
> ```
> C#
> ```
>
> ```
> enum Color { Red, Blue, Green }
>
> // The expression 0 converts implicitly to enum types
> Color c0 = 0;
>
> // Other int expressions need explicit conversion
> Color c1 = (Color)1;
>
> // Conversion from null expression (no type) to string
> string x = null;
> ```

```
// Conversion from lambda expression to delegate type
Func<int, int> square = x => x * x;
```

end example

10.2 Implicit conversions

10.2.1 General

The following conversions are classified as implicit conversions:

- Identity conversions
- Implicit numeric conversions
- Implicit enumeration conversions
- Implicit interpolated string conversions
- Implicit reference conversions
- Boxing conversions
- Implicit dynamic conversions
- Implicit type parameter conversions
- Implicit constant expression conversions
- User-defined implicit conversions
- Anonymous function conversions
- Method group conversions
- Null literal conversions
- Implicit nullable conversions
- Implicit tuple conversions
- Lifted user-defined implicit conversions
- Default literal conversions
- Implicit throw conversion

Implicit conversions can occur in a variety of situations, including function member invocations (§12.6.6), cast expressions (§12.9.7), and assignments (§12.21).

The pre-defined implicit conversions always succeed and never cause exceptions to be thrown.

> *Note*: Properly designed user-defined implicit conversions should exhibit these characteristics as well. *end note*

For the purposes of conversion, the types `object` and `dynamic` are identity convertible (§10.2.2).

However, dynamic conversions (§10.2.10 and §10.3.8) apply only to expressions of type `dynamic` (§8.2.4).

10.2.2 Identity conversion

An identity conversion converts from any type to the same type or a type that is equivalent at runtime. One reason this conversion exists is so that a type `T` or an expression of type `T` can be said to be convertible to `T` itself. The following identity conversions exist:

- Between `T` and `T`, for any type `T`.
- Between `object` and `dynamic`.
- Between all tuple types with the same arity, and the corresponding constructed `ValueTuple<...>` type, when an identity conversion exists between each pair of corresponding element types.
- Between types constructed from the same generic type where there exists an identity conversion between each corresponding type argument.

> *Example*: The following illustrates the recursive nature of the third rule:
>
> ```C#
> (int a , string b) t1 = (1, "two");
> (int c, string d) t2 = (3, "four");
>
> // Identity conversions exist between
> // the types of t1, t2, and t3.
> var t3 = (5, "six");
> t3 = t2;
> t2 = t1;
>
> var t4 = (t1, 7);
> var t5 = (t2, 8);
>
> // Identity conversions exist between
> // the types of t4, t5, and t6.
> var t6 =((8, "eight"), 9);
> t6 = t5;
> t5 = t4;
> ```
>
> The types of tuples `t1`, `t2` and `t3` all have two elements: an `int` followed by a `string`. Tuple element types may themselves by tuples, as in `t4`, `t5`, and `t6`. An identity conversion exists between each pair of corresponding element types, including nested tuples, therefore an identity conversion exists between the types of tuples `t4`, `t5`, and `t6`.

> *end example*

All identity conversions are symmetric. If an identity conversion exists from `T₁` to `T₂`, then an identity conversion exists from `T₂` to `T₁`. Two types are *identity convertible* when an identity conversion exists between two types.

In most cases, an identity conversion has no effect at runtime. However, since floating point operations may be performed at higher precision than prescribed by their type (§8.3.7), assignment of their results may result in a loss of precision, and explicit casts are guaranteed to reduce precision to what is prescribed by the type (§12.9.7).

10.2.3 Implicit numeric conversions

The implicit numeric conversions are:

- From `sbyte` to `short`, `int`, `long`, `float`, `double`, or `decimal`.
- From `byte` to `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `float`, `double`, or `decimal`.
- From `short` to `int`, `long`, `float`, `double`, or `decimal`.
- From `ushort` to `int`, `uint`, `long`, `ulong`, `float`, `double`, or `decimal`.
- From `int` to `long`, `float`, `double`, or `decimal`.
- From `uint` to `long`, `ulong`, `float`, `double`, or `decimal`.
- From `long` to `float`, `double`, or `decimal`.
- From `ulong` to `float`, `double`, or `decimal`.
- From `char` to `ushort`, `int`, `uint`, `long`, `ulong`, `float`, `double`, or `decimal`.
- From `float` to `double`.

Conversions from `int`, `uint`, `long` or `ulong` to `float` and from `long` or `ulong` to `double` may cause a loss of precision, but will never cause a loss of magnitude. The other implicit numeric conversions never lose any information.

There are no predefined implicit conversions to the `char` type, so values of the other integral types do not automatically convert to the `char` type.

10.2.4 Implicit enumeration conversions

An implicit enumeration conversion permits a *constant_expression* (§12.23) with any integer type and the value zero to be converted to any *enum_type* and to any *nullable_value_type* whose underlying type is an *enum_type*. In the latter case the conversion is evaluated by converting to the underlying *enum_type* and wrapping the result (§8.3.12).

10.2.5 Implicit interpolated string conversions

An implicit interpolated string conversion permits an *interpolated_string_expression* (§12.8.3) to be converted to `System.IFormattable` or `System.FormattableString` (which implements `System.IFormattable`). When this conversion is applied, a string value is not composed from the interpolated string. Instead an instance of `System.FormattableString` is created, as further described in §12.8.3.

10.2.6 Implicit nullable conversions

The implicit nullable conversions are those nullable conversions (§10.6.1) derived from implicit predefined conversions.

10.2.7 Null literal conversions

An implicit conversion exists from the `null` literal to any reference type or nullable value type. This conversion produces a null reference if the target type is a reference type, or the null value (§8.3.12) of the given nullable value type.

10.2.8 Implicit reference conversions

The implicit reference conversions are:

- From any *reference_type* to `object` and `dynamic`.
- From any *class_type* `S` to any *class_type* `T`, provided `S` is derived from `T`.
- From any *class_type* `S` to any *interface_type* `T`, provided `S` implements `T`.
- From any *interface_type* `S` to any *interface_type* `T`, provided `S` is derived from `T`.
- From an *array_type* `S` with an element type S_i to an *array_type* `T` with an element type T_i, provided all of the following are true:
  - `S` and `T` differ only in element type. In other words, `S` and `T` have the same number of dimensions.
  - An implicit reference conversion exists from S_i to T_i.
- From a single-dimensional array type `S[]` to `System.Collections.Generic.IList<T>`, `System.Collections.Generic.IReadOnlyList<T>`, and their base interfaces, provided that there is an implicit identity or reference conversion from `S` to `T`.
- From any *array_type* to `System.Array` and the interfaces it implements.
- From any *delegate_type* to `System.Delegate` and the interfaces it implements.
- From the null literal (§6.4.5.7) to any reference-type.

- From any *reference_type* to a *reference_type* `T` if it has an implicit identity or reference conversion to a *reference_type* `T₀` and `T₀` has an identity conversion to `T`.
- From any *reference_type* to an interface or delegate type `T` if it has an implicit identity or reference conversion to an interface or delegate type `T₀` and `T₀` is variance-convertible (§18.2.3.3) to `T`.
- Implicit conversions involving type parameters that are known to be reference types. See §10.2.12 for more details on implicit conversions involving type parameters.

The implicit reference conversions are those conversions between *reference_type*s that can be proven to always succeed, and therefore require no checks at run-time.

Reference conversions, implicit or explicit, never change the referential identity of the object being converted.

> *Note*: In other words, while a reference conversion can change the type of the reference, it never changes the type or value of the object being referred to. *end note*

10.2.9 Boxing conversions

A boxing conversion permits a *value_type* to be implicitly converted to a *reference_type*. The following boxing conversions exist:

- From any *value_type* to the type `object`.
- From any *value_type* to the type `System.ValueType`.
- From any *enum_type* to the type `System.Enum`.
- From any *non_nullable_value_type* to any *interface_type* implemented by the *non_nullable_value_type*.
- From any *non_nullable_value_type* to any *interface_type* `I` such that there is a boxing conversion from the *non_nullable_value_type* to another *interface_type* `I₀`, and `I₀` has an identity conversion to `I`.
- From any *non_nullable_value_type* to any *interface_type* `I` such that there is a boxing conversion from the *non_nullable_value_type* to another *interface_type* `I₀`, and `I₀` is variance-convertible (§18.2.3.3) to `I`.
- From any *nullable_value_type* to any *reference_type* where there is a boxing conversion from the underlying type of the *nullable_value_type* to the *reference_type*.
- From a type parameter that is not known to be a reference type to any type such that the conversion is permitted by §10.2.12.

Boxing a value of a *non-nullable-value-type* consists of allocating an object instance and copying the value into that instance.

Boxing a value of a *nullable_value_type* produces a null reference if it is the null value (`HasValue` is false), or the result of unwrapping and boxing the underlying value otherwise.

> *Note*: The process of boxing may be imagined in terms of the existence of a boxing class for every value type. For example, consider a `struct S` implementing an interface `I`, with a boxing class called `S_Boxing`.
>
> C#
>
> ```
> interface I
> {
>     void M();
> }
>
> struct S : I
> {
>     public void M() { ... }
> }
>
> sealed class S_Boxing : I
> {
>     S value;
>
>     public S_Boxing(S value)
>     {
>         this.value = value;
>     }
>
>     public void M()
>     {
>         value.M();
>     }
> }
> ```
>
> Boxing a value `v` of type `S` now consists of executing the expression `new S_Boxing(v)` and returning the resulting instance as a value of the target type of the conversion. Thus, the statements
>
> C#
>
> ```
> S s = new S();
> object box = s;
> ```
>
> can be thought of as similar to:

C#

```
S s = new S();
object box = new S_Boxing(s);
```

The imagined boxing type described above does not actually exist. Instead, a boxed value of type `S` has the runtime type `S`, and a runtime type check using the `is` operator with a value type as the right operand tests whether the left operand is a boxed version of the right operand. For example,

C#

```
int i = 123;
object box = i;
if (box is int)
{
    Console.Write("Box contains an int");
}
```

will output the following:

Console

```
Box contains an int
```

A boxing conversion implies making a copy of the value being boxed. This is different from a conversion of a *reference_type* to type `object`, in which the value continues to reference the same instance and simply is regarded as the less derived type `object`. For example, the following

C#

```
struct Point
{
    public int x, y;

    public Point(int x, int y)
    {
        this.x = x;
        this.y = y;
    }
}

class A
{
    void M()
    {
```

```
        Point p = new Point(10, 10);
        object box = p;
        p.x = 20;
        Console.Write(((Point)box).x);
    }
}
```

will output the value 10 on the console because the implicit boxing operation that occurs in the assignment of `p` to `box` causes the value of `p` to be copied. Had `Point` been declared a `class` instead, the value 20 would be output because `p` and `box` would reference the same instance.

The analogy of a boxing class should not be used as more than a helpful tool for picturing how boxing works conceptually. There are numerous subtle differences between the behavior described by this specification and the behavior that would result from boxing being implemented in precisely this manner.

end note

10.2.10 Implicit dynamic conversions

An implicit dynamic conversion exists from an expression of type dynamic to any type `T`. The conversion is dynamically bound §12.3.3, which means that an implicit conversion will be sought at run-time from the run-time type of the expression to `T`. If no conversion is found, a run-time exception is thrown.

This implicit conversion seemingly violates the advice in the beginning of §10.2 that an implicit conversion should never cause an exception. However, it is not the conversion itself, but the *finding* of the conversion that causes the exception. The risk of run-time exceptions is inherent in the use of dynamic binding. If dynamic binding of the conversion is not desired, the expression can be first converted to `object`, and then to the desired type.

Example: The following illustrates implicit dynamic conversions:

C#

```
object o = "object";
dynamic d = "dynamic";
string s1 = o;          // Fails at compile-time – no conversion exists
string s2 = d;          // Compiles and succeeds at run-time
int i = d;              // Compiles but fails at run-time – no conversion
exists
```

> The assignments to `s2` and `i` both employ implicit dynamic conversions, where the binding of the operations is suspended until run-time. At run-time, implicit conversions are sought from the run-time type of `d` (`string`) to the target type. A conversion is found to `string` but not to `int`.
>
> *end example*

10.2.11 Implicit constant expression conversions

An implicit constant expression conversion permits the following conversions:

- A *constant_expression* (§12.23) of type `int` can be converted to type `sbyte`, `byte`, `short`, `ushort`, `uint`, or `ulong`, provided the value of the *constant_expression* is within the range of the destination type.
- A *constant_expression* of type `long` can be converted to type `ulong`, provided the value of the *constant_expression* is not negative.

10.2.12 Implicit conversions involving type parameters

For a *type_parameter* `T` that is known to be a reference type (§15.2.5), the following implicit reference conversions (§10.2.8) exist:

- From `T` to its effective base class `C`, from `T` to any base class of `C`, and from `T` to any interface implemented by `C`.
- From `T` to an *interface_type* `I` in `T`'s effective interface set and from `T` to any base interface of `I`.
- From `T` to a type parameter `U` provided that `T` depends on `U` (§15.2.5).

  > *Note*: Since `T` is known to be a reference type, within the scope of `T`, the run-time type of `U` will always be a reference type, even if `U` is not known to be a reference type at compile-time. *end note*

- From the null literal (§6.4.5.7) to T.

For a *type_parameter* `T` that is *not* known to be a reference type §15.2.5, the following conversions involving `T` are considered to be boxing conversions (§10.2.9) at compile-time. At run-time, if `T` is a value type, the conversion is executed as a boxing conversion. At run-time, if `T` is a reference type, the conversion is executed as an implicit reference conversion or identity conversion.

- From `T` to its effective base class `C`, from `T` to any base class of `C`, and from `T` to any interface implemented by `C`.

  > *Note*: `C` will be one of the types `System.Object`, `System.ValueType`, or `System.Enum` (otherwise `T` would be known to be a reference type). *end note*

- From `T` to an *interface_type* `I` in `T`'s effective interface set and from `T` to any base interface of `I`.

For a *type_parameter* `T` that is *not* known to be a reference type, there is an implicit conversion from `T` to a type parameter `U` provided `T` depends on `U`. At run-time, if `T` is a value type and `U` is a reference type, the conversion is executed as a boxing conversion. At run-time, if both `T` and `U` are value types, then `T` and `U` are necessarily the same type and no conversion is performed. At run-time, if `T` is a reference type, then `U` is necessarily also a reference type and the conversion is executed as an implicit reference conversion or identity conversion (§15.2.5).

The following further implicit conversions exist for a given type parameter `T`:

- From `T` to a reference type `S` if it has an implicit conversion to a reference type S_o and S_o has an identity conversion to `S`. At run-time, the conversion is executed the same way as the conversion to S_o.
- From `T` to an interface type `I` if it has an implicit conversion to an interface type I_o, and I_o is variance-convertible to `I` (§18.2.3.3). At run-time, if `T` is a value type, the conversion is executed as a boxing conversion. Otherwise, the conversion is executed as an implicit reference conversion or identity conversion.

In all cases, the rules ensure that a conversion is executed as a boxing conversion if and only if at run-time the conversion is from a value type to a reference type.

10.2.13 Implicit tuple conversions

An implicit conversion exists from a tuple expression `E` to a tuple type `T` if `E` has the same arity as `T` and an implicit conversion exists from each element in `E` to the corresponding element type in `T`. The conversion is performed by creating an instance of `T`'s corresponding `System.ValueTuple<...>` type, and initializing each of its fields in order from left to right by evaluating the corresponding tuple element expression of `E`, converting it to the corresponding element type of `T` using the implicit conversion found, and initializing the field with the result.

If an element name in the tuple expression does not match a corresponding element name in the tuple type, a warning shall be issued.

> *Example*:
>
> ```
> C#
> ```
>
> ```
> (int, string) t1 = (1, "One");
> (byte, string) t2 = (2, null);
> (int, string) t3 = (null, null);        // Error: No conversion
> (int i, string s) t4 = (i: 4, "Four");
> (int i, string) t5 = (x: 5, s: "Five"); // Warning: Names are ignored
> ```
>
> The declarations of `t1`, `t2`, `t4` and `t5` are all valid, since implicit conversions exist from the element expressions to the corresponding element types. The declaration of `t3` is invalid, because there is no conversion from `null` to `int`. The declaration of `t5` causes a warning because the element names in the tuple expression differs from those in the tuple type.
>
> *end example*

10.2.14 User-defined implicit conversions

A user-defined implicit conversion consists of an optional standard implicit conversion, followed by execution of a user-defined implicit conversion operator, followed by another optional standard implicit conversion. The exact rules for evaluating user-defined implicit conversions are described in §10.5.4.

10.2.15 Anonymous function conversions and method group conversions

Anonymous functions and method groups do not have types in and of themselves, but they may be implicitly converted to delegate types. Additionally, some lambda expressions may be implicitly converted to expression tree types. Anonymous function conversions are described in more detail in §10.7 and method group conversions in §10.8.

10.2.16 Default literal conversions

An implicit conversion exists from a *default_literal* (§12.8.20) to any type. This conversion produces the default value (§9.3) of the inferred type.

10.2.17 Implicit throw conversions

While throw expressions do not have a type, they may be implicitly converted to any type.

10.3 Explicit conversions

10.3.1 General

The following conversions are classified as explicit conversions:

- All implicit conversions
- Explicit numeric conversions
- Explicit enumeration conversions
- Explicit nullable conversions
- Explicit tuple conversions
- Explicit reference conversions
- Explicit interface conversions
- Unboxing conversions
- Explicit type parameter conversions
- Explicit dynamic conversions
- User-defined explicit conversions

Explicit conversions can occur in cast expressions (§12.9.7).

The set of explicit conversions includes all implicit conversions.

> *Note*: This, for example, allows an explicit cast to be used when an implicit identity conversion exists, in order to force the selection of a particular method overload. *end note*

The explicit conversions that are not implicit conversions are conversions that cannot be proven always to succeed, conversions that are known possibly to lose information, and conversions across domains of types sufficiently different to merit explicit notation.

10.3.2 Explicit numeric conversions

The explicit numeric conversions are the conversions from a *numeric_type* to another *numeric_type* for which an implicit numeric conversion (§10.2.3) does not already exist:

- From `sbyte` to `byte`, `ushort`, `uint`, `ulong`, or `char`.

- From `byte` to `sbyte` or `char`.
- From `short` to `sbyte`, `byte`, `ushort`, `uint`, `ulong`, or `char`.
- From `ushort` to `sbyte`, `byte`, `short`, or `char`.
- From `int` to `sbyte`, `byte`, `short`, `ushort`, `uint`, `ulong`, or `char`.
- From `uint` to `sbyte`, `byte`, `short`, `ushort`, `int`, or `char`.
- From `long` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `ulong`, or `char`.
- From `ulong` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, or `char`.
- From `char` to `sbyte`, `byte`, or `short`.
- From `float` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, or `decimal`.
- From `double` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, or `decimal`.
- From `decimal` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, or `double`.

Because the explicit conversions include all implicit and explicit numeric conversions, it is always possible to convert from any *numeric_type* to any other *numeric_type* using a cast expression (§12.9.7).

The explicit numeric conversions possibly lose information or possibly cause exceptions to be thrown. An explicit numeric conversion is processed as follows:

- For a conversion from an integral type to another integral type, the processing depends on the overflow checking context (§12.8.19) in which the conversion takes place:
  - In a `checked` context, the conversion succeeds if the value of the source operand is within the range of the destination type, but throws a `System.OverflowException` if the value of the source operand is outside the range of the destination type.
  - In an `unchecked` context, the conversion always succeeds, and proceeds as follows.
    - If the source type is larger than the destination type, then the source value is truncated by discarding its “extra” most significant bits. The result is then treated as a value of the destination type.
    - If the source type is the same size as the destination type, then the source value is treated as a value of the destination type
- For a conversion from `decimal` to an integral type, the source value is rounded towards zero to the nearest integral value, and this integral value becomes the result of the conversion. If the resulting integral value is outside the range of the destination type, a `System.OverflowException` is thrown.

- For a conversion from `float` or `double` to an integral type, the processing depends on the overflow-checking context (§12.8.19) in which the conversion takes place:
  - In a checked context, the conversion proceeds as follows:
    - If the value of the operand is NaN or infinite, a `System.OverflowException` is thrown.
    - Otherwise, the source operand is rounded towards zero to the nearest integral value. If this integral value is within the range of the destination type then this value is the result of the conversion.
    - Otherwise, a `System.OverflowException` is thrown.
  - In an unchecked context, the conversion always succeeds, and proceeds as follows.
    - If the value of the operand is NaN or infinite, the result of the conversion is an unspecified value of the destination type.
    - Otherwise, the source operand is rounded towards zero to the nearest integral value. If this integral value is within the range of the destination type then this value is the result of the conversion.
    - Otherwise, the result of the conversion is an unspecified value of the destination type.
- For a conversion from `double` to `float`, the `double` value is rounded to the nearest `float` value. If the `double` value is too small to represent as a `float`, the result becomes zero with the same sign as the value. If the magnitude of the `double` value is too large to represent as a `float`, the result becomes infinity with the same sign as the value. If the `double` value is NaN, the result is also NaN.
- For a conversion from `float` or `double` to `decimal`, the source value is converted to `decimal` representation and rounded to the nearest number if required (§8.3.8).
  - If the source value is too small to represent as a `decimal`, the result becomes zero, preserving the sign of the original value if `decimal` supports signed zero values.
  - If the source value's magnitude is too large to represent as a `decimal`, or that value is infinity, the result is infinity preserving the sign of the original value, if the decimal representation supports infinities; otherwise a System.OverflowException is thrown.
  - If the source value is NaN, the result is NaN if the decimal representation supports NaNs; otherwise a System.OverflowException is thrown.
- For a conversion from `decimal` to `float` or `double`, the `decimal` value is rounded to the nearest `double` or `float` value. If the source value's magnitude is too large to represent in the target type, or that value is infinity, the result is infinity preserving the sign of the original value. If the source value is NaN, the result is

NaN. While this conversion may lose precision, it never causes an exception to be thrown.

> *Note*: The `decimal` type is not required to support infinities or NaN values but may do so; its range may be smaller than the range of `float` and `double`, but is not guaranteed to be. For `decimal` representations without infinities or NaN values, and with a range smaller than `float`, the result of a conversion from `decimal` to either `float` or `double` will never be infinity or NaN. *end note*

10.3.3 Explicit enumeration conversions

The explicit enumeration conversions are:

- From `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, `double`, or `decimal` to any *enum_type*.
- From any *enum_type* to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, `double`, or `decimal`.
- From any *enum_type* to any other *enum_type*.

An explicit enumeration conversion between two types is processed by treating any participating *enum_type* as the underlying type of that *enum_type*, and then performing an implicit or explicit numeric conversion between the resulting types.

> *Example*: Given an *enum_type* `E` with and underlying type of `int`, a conversion from `E` to `byte` is processed as an explicit numeric conversion (§10.3.2) from `int` to `byte`, and a conversion from `byte` to `E` is processed as an implicit numeric conversion (§10.2.3) from `byte` to `int`. *end example*

10.3.4 Explicit nullable conversions

The explicit nullable conversions are those nullable conversions (§10.6.1) derived from explicit and implicit predefined conversions.

10.3.5 Explicit reference conversions

The explicit reference conversions are:

- From object and dynamic to any other *reference_type*.
- From any *class_type* `S` to any *class_type* `T`, provided `S` is a base class of `T`.

- From any *class_type* `S` to any *interface_type* `T`, provided `S` is not sealed and provided `S` does not implement `T`.
- From any *interface_type* `S` to any *class_type* `T`, provided `T` is not sealed or provided `T` implements `S`.
- From any *interface_type* `S` to any *interface_type* `T`, provided `S` is not derived from `T`.
- From an *array_type* `S` with an element type S_i to an *array_type* `T` with an element type T_i, provided all of the following are true:
  - `S` and `T` differ only in element type. In other words, `S` and `T` have the same number of dimensions.
  - An explicit reference conversion exists from S_i to T_i.
- From `System.Array` and the interfaces it implements, to any *array_type*.
- From a single-dimensional *array_type* `S[]` to `System.Collections.Generic.IList<T>`, `System.Collections.Generic.IReadOnlyList<T>`, and its base interfaces, provided that there is an identity conversion or explicit reference conversion from `S` to `T`.
- From `System.Collections.Generic.IList<S>`, `System.Collections.Generic.IReadOnlyList<S>`, and their base interfaces to a single-dimensional array type `T[]`, provided that there is an identity conversion or explicit reference conversion from `S` to T.
- From `System.Delegate` and the interfaces it implements to any *delegate_type*.
- From a reference type `S` to a reference type `T` if it has an explicit reference conversion from `S` to a reference type T_0 and T_0 and there is an identity conversion from T_0 to `T`.
- From a reference type `S` to an interface or delegate type `T` if it there is an explicit reference conversion from `S` to an interface or delegate type T_0 and either T_0 is variance-convertible to `T` or `T` is variance-convertible to T_0 §18.2.3.3.
- From $D<S_1...S_v>$ to $D<T_1...T_v>$ where $D<X_1...X_v>$ is a generic delegate type, $D<S_1...S_v>$ is not compatible with or identical to $D<T_1...T_v>$, and for each type parameter X_i of `D` the following holds:
  - If X_i is invariant, then S_i is identical to T_i.
  - If X_i is covariant, then there is an identity conversion, implicit reference conversion or explicit reference conversion from S_i to T_i.
  - If X_i is contravariant, then S_i and T_i are either identical or both reference types.
- Explicit conversions involving type parameters that are known to be reference types. For more details on explicit conversions involving type parameters, see §10.3.9.

The explicit reference conversions are those conversions between *reference_type*s that require run-time checks to ensure they are correct.

For an explicit reference conversion to succeed at run-time, the value of the source operand shall be `null`, or the type of the object referenced by the source operand shall be a type that can be converted to the destination type by an implicit reference conversion (§10.2.8). If an explicit reference conversion fails, a `System.InvalidCastException` is thrown.

> *Note*: Reference conversions, implicit or explicit, never change the value of the reference itself (§8.2.1), only its type; neither does it change the type or value of the object being referenced. *end note*

10.3.6 Explicit tuple conversions

An explicit conversion exists from a tuple expression `E` to a tuple type `T` if `E` has the same arity as `T` and an implicit or explicit conversion exists from each element in `E` to the corresponding element type in `T`. The conversion is performed by creating an instance of `T`'s corresponding `System.ValueTuple<...>` type, and initializing each of its fields in order from left to right by evaluating the corresponding tuple element expression of `E`, converting it to the corresponding element type of `T` using the explicit conversion found, and initializing the field with the result.

10.3.7 Unboxing conversions

An unboxing conversion permits a *reference_type* to be explicitly converted to a *value_type*. The following unboxing conversions exist:

- From the type `object` to any *value_type*.
- From the type `System.ValueType` to any *value_type*.
- From the type `System.Enum` to any *enum_type*.
- From any *interface_type* to any *non-nullable_value_type* that implements the *interface_type*.
- From any *interface_type* `I` to any *non_nullable_value_type* where there is an unboxing conversion from an *interface_type* I_o to the *non_nullable_value-type* and an identity conversion from `I` to I_o.
- From any *interface_type* `I` to any *non_nullable_value_type* where there is an unboxing conversion from an *interface_type* I_o to the *non_nullable_value_type* and either either I_o is variance_convertible to `I` or `I` is variance-convertible to I_o (§18.2.3.3).

- From any *reference_type* to any *nullable_value_type* where there is an unboxing conversion from *reference_type* to the underlying *non_nullable_value_type* of the *nullable_value_type*.
- From a type parameter which is not known to be a value type to any type such that the conversion is permitted by §10.3.9.

An unboxing operation to a *non_nullable_value_type* consists of first checking that the object instance is a boxed value of the given *non_nullable_value_type*, and then copying the value out of the instance.

Unboxing to a *nullable_value_type* produces the null value of the *nullable_value_type* if the source operand is `null`, or the wrapped result of unboxing the object instance to the underlying type of the *nullable_value_type* otherwise.

> *Note*: Referring to the imaginary boxing class described in §10.2.9, an unboxing conversion of an object box to a *value_type* `S` consists of executing the expression `((S_Boxing)box).value`. Thus, the statements
>
> ```C#
> object box = new S();
> S s = (S)box;
> ```
>
> conceptually correspond to
>
> ```C#
> object box = new S_Boxing(new S());
> S s = ((S_Boxing)box).value;
> ```
>
> *end note*

For an unboxing conversion to a given *non_nullable_value_type* to succeed at run-time, the value of the source operand shall be a reference to a boxed value of that *non_nullable_value_type*. If the source operand is `null` a `System.NullReferenceException` is thrown. If the source operand is a reference to an incompatible object, a `System.InvalidCastException` is thrown.

For an unboxing conversion to a given *nullable_value_type* to succeed at run-time, the value of the source operand shall be either null or a reference to a boxed value of the underlying *non_nullable_value_type* of the *nullable_value_type*. If the source operand is a reference to an incompatible object, a `System.InvalidCastException` is thrown.

10.3.8 Explicit dynamic conversions

An explicit dynamic conversion exists from an expression of type `dynamic` to any type `T`. The conversion is dynamically bound (§12.3.3), which means that an explicit conversion will be sought at run-time from the run-time type of the expression to `T`. If no conversion is found, a run-time exception is thrown.

If dynamic binding of the conversion is not desired, the expression can be first converted to `object`, and then to the desired type.

> *Example*: Assume the following class is defined:
>
> ```
> C#
> ```
>
> ```
> class C
> {
>     int i;
>
>     public C(int i)
>     {
>         this.i = i;
>     }
>
>     public static explicit operator C(string s)
>     {
>         return new C(int.Parse(s));
>     }
> }
> ```
>
> The following illustrates explicit dynamic conversions:
>
> ```
> C#
> ```
>
> ```
> object o = "1";
> dynamic d = "2";
> var c1 = (C)o; // Compiles, but explicit reference conversion fails
> var c2 = (C)d; // Compiles and user defined conversion succeeds
> ```
>
> The best conversion of `o` to `C` is found at compile-time to be an explicit reference conversion. This fails at run-time, because `"1"` is not in fact a `C`. The conversion of `d` to `C` however, as an explicit dynamic conversion, is suspended to run-time, where a user defined conversion from the run-time type of `d` (`string`) to `C` is found, and succeeds.
>
> *end example*

10.3.9 Explicit conversions involving type parameters

For a *type_parameter* `T` that is known to be a reference type (§15.2.5), the following explicit reference conversions (§10.3.5) exist:

- From the effective base class `C` of `T` to `T` and from any base class of `C` to `T`.
- From any *interface_type* to `T`.
- From `T` to any *interface_type* `I` provided there isn't already an implicit reference conversion from `T` to `I`.
- From a *type_parameter* `U` to `T` provided that `T` depends on `U` (§15.2.5).

> *Note*: Since `T` is known to be a reference type, within the scope of `T`, the run-time type of U will always be a reference type, even if `U` is not known to be a reference type at compile-time. *end note*

For a *type_parameter* `T` that is *not* known to be a reference type (§15.2.5), the following conversions involving `T` are considered to be unboxing conversions (§10.3.7) at compile-time. At run-time, if `T` is a value type, the conversion is executed as an unboxing conversion. At run-time, if `T` is a reference type, the conversion is executed as an explicit reference conversion or identity conversion.

- From the effective base class `C` of `T` to `T` and from any base class of `C` to `T`.

> *Note*: C will be one of the types `System.Object`, `System.ValueType`, or `System.Enum` (otherwise `T` would be known to be a reference type). *end note*

- From any *interface_type* to `T`.

For a *type_parameter* `T` that is *not* known to be a reference type (§15.2.5), the following explicit conversions exist:

- From `T` to any *interface_type* `I` provided there is not already an implicit conversion from `T` to `I`. This conversion consists of an implicit boxing conversion (§10.2.9) from `T` to `object` followed by an explicit reference conversion from `object` to `I`. At run-time, if `T` is a value type, the conversion is executed as a boxing conversion followed by an explicit reference conversion. At run-time, if `T` is a reference type, the conversion is executed as an explicit reference conversion.
- From a type parameter `U` to `T` provided that `T` depends on `U` (§15.2.5). At run-time, if `T` is a value type and `U` is a reference type, the conversion is executed as an unboxing conversion. At run-time, if both `T` and `U` are value types, then `T` and

`U` are necessarily the same type and no conversion is performed. At run-time, if `T` is a reference type, then `U` is necessarily also a reference type and the conversion is executed as an explicit reference conversion or identity conversion.

In all cases, the rules ensure that a conversion is executed as an unboxing conversion if and only if at run-time the conversion is from a reference type to a value type.

The above rules do not permit a direct explicit conversion from an unconstrained type parameter to a non-interface type, which might be surprising. The reason for this rule is to prevent confusion and make the semantics of such conversions clear.

> *Example*: Consider the following declaration:
>
> ```C#
> class X<T>
> {
>     public static long F(T t)
>     {
>         return (long)t;         // Error
>     }
> }
> ```
>
> If the direct explicit conversion of `t` to `long` were permitted, one might easily expect that `X<int>.F(7)` would return `7L`. However, it would not, because the standard numeric conversions are only considered when the types are known to be numeric at binding-time. In order to make the semantics clear, the above example must instead be written:
>
> ```C#
> class X<T>
> {
>     public static long F(T t)
>     {
>         return (long)(object)t;         // Ok, but will only work when T
> is long
>     }
> }
> ```
>
> This code will now compile but executing `X<int>.F(7)` would then throw an exception at run-time, since a boxed `int` cannot be converted directly to a `long`.
>
> *end example*

10.3.10 User-defined explicit conversions

A user-defined explicit conversion consists of an optional standard explicit conversion, followed by execution of a user-defined implicit or explicit conversion operator, followed by another optional standard explicit conversion. The exact rules for evaluating user-defined explicit conversions are described in §10.5.5.

10.4 Standard conversions

10.4.1 General

The standard conversions are those pre-defined conversions that can occur as part of a user-defined conversion.

10.4.2 Standard implicit conversions

The following implicit conversions are classified as standard implicit conversions:

- Identity conversions (§10.2.2)
- Implicit numeric conversions (§10.2.3)
- Implicit nullable conversions (§10.2.6)
- Null literal conversions (§10.2.7)
- Implicit reference conversions (§10.2.8)
- Boxing conversions (§10.2.9)
- Implicit constant expression conversions (§10.2.11)
- Implicit conversions involving type parameters (§10.2.12)

The standard implicit conversions specifically exclude user-defined implicit conversions.

10.4.3 Standard explicit conversions

The standard explicit conversions are all standard implicit conversions plus the subset of the explicit conversions for which an opposite standard implicit conversion exists.

> *Note*: In other words, if a standard implicit conversion exists from a type `A` to a type `B`, then a standard explicit conversion exists from type `A` to type `B` and from type `B` to type `A`. *end note*

10.5 User-defined conversions

10.5.1 General

C# allows the pre-defined implicit and explicit conversions to be augmented by user-defined conversions. User-defined conversions are introduced by declaring conversion operators (§15.10.4) in class and struct types.

10.5.2 Permitted user-defined conversions

C# permits only certain user-defined conversions to be declared. In particular, it is not possible to redefine an already existing implicit or explicit conversion.

For a given source type `S` and target type `T`, if `S` or `T` are nullable value types, let S_0 and T_0 refer to their underlying types, otherwise S_0 and T_0 are equal to `S` and `T` respectively. A class or struct is permitted to declare a conversion from a source type `S` to a target type `T` only if all of the following are true:

- S_0 and T_0 are different types.
- Either S_0 or T_0 is the class or struct type in which the operator declaration takes place.
- Neither S_0 nor T_0 is an *interface_type*.
- Excluding user-defined conversions, a conversion does not exist from `S` to `T` or from `T` to `S`.

The restrictions that apply to user-defined conversions are specified in §15.10.4.

10.5.3 Evaluation of user-defined conversions

A user-defined conversion converts a ***source expression***, which may have a **source type**, to another type, called the ***target type***. Evaluation of a user-defined conversion centers on finding the ***most-specific*** user-defined conversion operator for the source expression and target type. This determination is broken into several steps:

- Finding the set of classes and structs from which user-defined conversion operators will be considered. This set consists of the source type and its base classes, if the source type exists, along with the target type and its base classes. For this purpose it is assumed that only classes and structs can declare user-defined operators, and that non-class types have no base classes. Also, if either the source or target type is a nullable-value-type, their underlying type is used instead.
- From that set of types, determining which user-defined and lifted conversion operators are applicable. For a conversion operator to be applicable, it shall be possible to perform a standard conversion (§10.4) from the source expression to

the operand type of the operator, and it shall be possible to perform a standard conversion from the result type of the operator to the target type.
- From the set of applicable user-defined operators, determining which operator is unambiguously the most-specific. In general terms, the most-specific operator is the operator whose operand type is “closest” to the source expression and whose result type is “closest” to the target type. User-defined conversion operators are preferred over lifted conversion operators. The exact rules for establishing the most-specific user-defined conversion operator are defined in the following subclauses.

Once a most-specific user-defined conversion operator has been identified, the actual execution of the user-defined conversion involves up to three steps:

- First, if required, performing a standard conversion from the source expression to the operand type of the user-defined or lifted conversion operator.
- Next, invoking the user-defined or lifted conversion operator to perform the conversion.
- Finally, if required, performing a standard conversion from the result type of the user-defined conversion operator to the target type.

Evaluation of a user-defined conversion never involves more than one user-defined or lifted conversion operator. In other words, a conversion from type `S` to type `T` will never first execute a user-defined conversion from `S` to `X` and then execute a user-defined conversion from `X` to `T`.

- Exact definitions of evaluation of user-defined implicit or explicit conversions are given in the following subclauses. The definitions make use of the following terms:
- If a standard implicit conversion (§10.4.2) exists from a type `A` to a type `B`, and if neither `A` nor `B` are *interface_type* `s`, then `A` is said to be ***encompassed by*** `B`, and `B` is said to ***encompass*** `A`.
- If a standard implicit conversion (§10.4.2) exists from an expression `E` to a type `B`, and if neither `B` nor the type of `E` (if it has one) are *interface_type* `s`, then `E` is said to be *encompassed by* `B`, and `B` is said to *encompass* `E`.
- The ***most-encompassing type*** in a set of types is the one type that encompasses all other types in the set. If no single type encompasses all other types, then the set has no most-encompassing type. In more intuitive terms, the most-encompassing type is the “largest” type in the set—the one type to which each of the other types can be implicitly converted.
- The ***most-encompassed type*** in a set of types is the one type that is encompassed by all other types in the set. If no single type is encompassed by all other types, then the set has no most-encompassed type. In more intuitive terms, the most-

encompassed type is the “smallest” type in the set—the one type that can be implicitly converted to each of the other types.

10.5.4 User-defined implicit conversions

A user-defined implicit conversion from an expression `E` to a type `T` is processed as follows:

- Determine the types `S`, S_0 and T_0.
  - If `E` has a type, let `S` be that type.
  - If `S` or `T` are nullable value types, let S_i and T_i be their underlying types, otherwise let S_i and T_i be `S` and `T`, respectively.
  - If S_i or T_i are type parameters, let S_0 and T_0 be their effective base classes, otherwise let S_0 and T_0 be S_x and T_i, respectively.

- Find the set of types, `D`, from which user-defined conversion operators will be considered. This set consists of S_0 (if S_0 exists and is a class or struct), the base classes of S_0 (if S_0 exists and is a class), and T_0 (if T_0 is a class or struct). A type is added to the set `D` only if an identity conversion to another type already included in the set doesn’t exist.

- Find the set of applicable user-defined and lifted conversion operators, `U`. This set consists of the user-defined and lifted implicit conversion operators declared by the classes or structs in `D` that convert from a type encompassing `E` to a type encompassed by `T`. If `U` is empty, the conversion is undefined and a compile-time error occurs.
  - If `S` exists and any of the operators in `U` convert from `S`, then S_x is `S`.
  - Otherwise, S_x is the most-encompassed type in the combined set of source types of the operators in `U`. If exactly one most-encompassed type cannot be found, then the conversion is ambiguous and a compile-time error occurs.

- Find the most-specific target type, T_x, of the operators in `U`:
  - If any of the operators in `U` convert to `T`, then T_x is `T`.
  - Otherwise, T_x is the most-encompassing type in the combined set of target types of the operators in `U`. If exactly one most-encompassing type cannot be found, then the conversion is ambiguous and a compile-time error occurs.

- Find the most-specific conversion operator:
  - If `U` contains exactly one user-defined conversion operator that converts from S_x to T_x, then this is the most-specific conversion operator.

    - Otherwise, if `U` contains exactly one lifted conversion operator that converts from S_x to T_x, then this is the most-specific conversion operator.
    - Otherwise, the conversion is ambiguous and a compile-time error occurs.

- Finally, apply the conversion:
    - If E does not already have the type S_x, then a standard implicit conversion from `E` to S_x is performed.
    - The most-specific conversion operator is invoked to convert from S_x to T_x.
    - If T_x is not `T`, then a standard implicit conversion from T_x to `T` is performed.

A user-defined implicit conversion from a type `S` to a type `T` exists if a user-defined implicit conversion exists from a variable of type `S` to `T`.

10.5.5 User-defined explicit conversions

A user-defined explicit conversion from an expression `E` to a type `T` is processed as follows:

- Determine the types `S`, S_o and T_o.
    - If `E` has a type, let `S` be that type.
    - If `S` or `T` are nullable value types, let S_i and T_i be their underlying types, otherwise let S_i and T_i be `S` and `T`, respectively.
    - If S_i or T_i are type parameters, let S_o and T_o be their effective base classes, otherwise let S_o and T_o be S_i and T_i, respectively.
- Find the set of types, `D`, from which user-defined conversion operators will be considered. This set consists of S_o (if S_o exists and is a class or struct), the base classes of S_o (if S_o exists and is a class), T_o (if T_o is a class or struct), and the base classes of T_o (if T_o is a class). `A` type is added to the set `D` only if an identity conversion to another type already included in the set doesn't exist.
- Find the set of applicable user-defined and lifted conversion operators, `U`. This set consists of the user-defined and lifted implicit or explicit conversion operators declared by the classes or structs in `D` that convert from a type encompassing `E` or encompassed by `S` (if it exists) to a type encompassing or encompassed by `T`. If `U` is empty, the conversion is undefined and a compile-time error occurs.
- Find the most-specific source type, S_x, of the operators in `U`:
    - If S exists and any of the operators in `U` convert from `S`, then S_x is `S`.
    - Otherwise, if any of the operators in `U` convert from types that encompass `E`, then S_x is the most-encompassed type in the combined set of source types of

those operators. If no most-encompassed type can be found, then the conversion is ambiguous and a compile-time error occurs.
    - Otherwise, S_x is the most-encompassing type in the combined set of source types of the operators in U. If exactly one most-encompassing type cannot be found, then the conversion is ambiguous and a compile-time error occurs.
- Find the most-specific target type, T_x, of the operators in U:
    - If any of the operators in U convert to T, then T_x is T.
    - Otherwise, if any of the operators in U convert to types that are encompassed by T, then T_x is the most-encompassing type in the combined set of target types of those operators. If exactly one most-encompassing type cannot be found, then the conversion is ambiguous and a compile-time error occurs.
    - Otherwise, T_x is the most-encompassed type in the combined set of target types of the operators in U. If no most-encompassed type can be found, then the conversion is ambiguous and a compile-time error occurs.
- Find the most-specific conversion operator:
    - If U contains exactly one user-defined conversion operator that converts from S_x to T_x, then this is the most-specific conversion operator.
    - Otherwise, if U contains exactly one lifted conversion operator that converts from S_x to T_x, then this is the most-specific conversion operator.
    - Otherwise, the conversion is ambiguous and a compile-time error occurs.
- Finally, apply the conversion:
    - If E does not already have the type S_x, then a standard explicit conversion from E to S_x is performed.
    - The most-specific user-defined conversion operator is invoked to convert from S_x to T_x.
    - If T_x is not T, then a standard explicit conversion from T_x to T is performed.

A user-defined explicit conversion from a type S to a type T exists if a user-defined explicit conversion exists from a variable of type S to T.

10.6 Conversions involving nullable types

10.6.1 Nullable Conversions

Nullable conversions permit predefined conversions that operate on non-nullable value types to also be used with nullable forms of those types. For each of the predefined implicit or explicit conversions that convert from a non-nullable value type S to a non-nullable value type T (§10.2.2, §10.2.3, §10.2.4, §10.2.11, §10.3.2 and §10.3.3), the following nullable conversions exist:

- An implicit or explicit conversion from `S?` to `T?`
- An implicit or explicit conversion from `S` to `T?`
- An explicit conversion from `S?` to `T`.

A nullable conversion is itself classified as an implicit or explicit conversion.

Certain nullable conversions are classified as standard conversions and can occur as part of a user-defined conversion. Specifically, all implicit nullable conversions are classified as standard implicit conversions (§10.4.2), and those explicit nullable conversions that satisfy the requirements of §10.4.3 are classified as standard explicit conversions.

Evaluation of a nullable conversion based on an underlying conversion from `S` to `T` proceeds as follows:

- If the nullable conversion is from `S?` to `T?`:
  - If the source value is null (`HasValue` property is `false`), the result is the null value of type `T?`.
  - Otherwise, the conversion is evaluated as an unwrapping from `S?` to `S`, followed by the underlying conversion from `S` to `T`, followed by a wrapping from `T` to `T?`.
- If the nullable conversion is from `S` to `T?`, the conversion is evaluated as the underlying conversion from `S` to `T` followed by a wrapping from `T` to `T?`.
- If the nullable conversion is from `S?` to `T`, the conversion is evaluated as an unwrapping from `S?` to `S` followed by the underlying conversion from `S` to `T`.

10.6.2 Lifted conversions

Given a user-defined conversion operator that converts from a non-nullable value type `S` to a non-nullable value type `T`, a ***lifted conversion operator*** exists that converts from `S?` to `T?`. This lifted conversion operator performs an unwrapping from `S?` to `S` followed by the user-defined conversion from `S` to `T` followed by a wrapping from `T` to `T?`, except that a null valued `S?` converts directly to a null valued `T?`. A lifted conversion operator has the same implicit or explicit classification as its underlying user-defined conversion operator.

10.7 Anonymous function conversions

10.7.1 General

An *anonymous_method_expression* or *lambda_expression* is classified as an anonymous function (§12.19). The expression does not have a type, but can be implicitly converted to a compatible delegate type. Some lambda expressions may also be implicitly converted to a compatible expression tree type.

Specifically, an anonymous function `F` is compatible with a delegate type `D` provided:

- If `F` contains an *anonymous_function_signature*, then `D` and `F` have the same number of parameters.
- If `F` does not contain an *anonymous_function_signature*, then `D` may have zero or more parameters of any type, as long as no parameter of `D` has the out parameter modifier.
- If `F` has an explicitly typed parameter list, each parameter in `D` has the same modifiers as the corresponding parameter in `F` and an identity conversion exists between the corresponding parameter in `F`.
- If `F` has an implicitly typed parameter list, `D` has no ref or out parameters.
- If the body of `F` is an expression, and *either* `D` has a void return type *or* `F` is async and `D` has a `«TaskType»` return type (§15.15.1), then when each parameter of `F` is given the type of the corresponding parameter in `D`, the body of `F` is a valid expression (w.r.t §12) that would be permitted as a *statement_expression* (§13.7).
- If the body of `F` is a block, and *either* `D` has a void return type *or* `F` is async and `D` has a `«TaskType»` return type , then when each parameter of `F` is given the type of the corresponding parameter in `D`, the body of `F` is a valid block (w.r.t §13.3) in which no `return` statement specifies an expression.
- If the body of `F` is an expression, and *either* `F` is non-async and `D` has a non-`void` return type `T`, *or* `F` is async and `D` has a `«TaskType»<T>` return type (§15.15.1), then when each parameter of `F` is given the type of the corresponding parameter in `D`, the body of `F` is a valid expression (w.r.t §12) that is implicitly convertible to `T`.
- If the body of `F` is a block, and *either* `F` is non-async and `D` has a non-void return type `T`, *or* `F` is async and `D` has a `«TaskType»<T>` return type, then when each parameter of `F` is given the type of the corresponding parameter in `D`, the body of `F` is a valid statement block (w.r.t §13.3) with a non-reachable end point in which each return statement specifies an expression that is implicitly convertible to `T`.

> *Example*: The following examples illustrate these rules:

```
C#

delegate void D(int x);
D d1 = delegate { };                            // Ok
```

```
D d2 = delegate() { };                          // Error, signature
mismatch
D d3 = delegate(long x) { };                    // Error, signature
mismatch
D d4 = delegate(int x) { };                     // Ok
D d5 = delegate(int x) { return; };             // Ok
D d6 = delegate(int x) { return x; };           // Error, return type
mismatch

delegate void E(out int x);
E e1 = delegate { };                            // Error, E has an out
parameter
E e2 = delegate(out int x) { x = 1; };          // Ok
E e3 = delegate(ref int x) { x = 1; };          // Error, signature
mismatch

delegate int P(params int[] a);
P p1 = delegate { };                            // Error, end of block
reachable
P p2 = delegate { return; };                    // Error, return type
mismatch
P p3 = delegate { return 1; };                  // Ok
P p4 = delegate { return "Hello"; };            // Error, return type
mismatch
P p5 = delegate(int[] a)                        // Ok
{
    return a[0];
};
P p6 = delegate(params int[] a)                 // Error, params modifier
{
    return a[0];
};
P p7 = delegate(int[] a)                        // Error, return type
mismatch
{
    if (a.Length > 0) return a[0];
    return "Hello";
};

delegate object Q(params int[] a);
Q q1 = delegate(int[] a)                       // Ok
{
    if (a.Length > 0) return a[0];
    return "Hello";
};
```

end example

Example: The examples that follow use a generic delegate type `Func<A,R>` that represents a function that takes an argument of type `A` and returns a value of type `R`:

```
C#
```

```
delegate R Func<A,R>(A arg);
```

In the assignments

```
C#
```

```
Func<int,int> f1 = x => x + 1; // Ok
Func<int,double> f2 = x => x + 1; // Ok
Func<double,int> f3 = x => x + 1; // Error
Func<int, Task<int>> f4 = async x => x + 1; // Ok
```

the parameter and return types of each anonymous function are determined from the type of the variable to which the anonymous function is assigned.

The first assignment successfully converts the anonymous function to the delegate type `Func<int,int>` because, when `x` is given type `int`, `x + 1` is a valid expression that is implicitly convertible to type `int`.

Likewise, the second assignment successfully converts the anonymous function to the delegate type Func<int,double> because the result of `x + 1` (of type `int`) is implicitly convertible to type `double`.

However, the third assignment is a compile-time error because, when `x` is given type `double`, the result of `x + 1` (of type `double`) is not implicitly convertible to type `int`.

The fourth assignment successfully converts the anonymous async function to the delegate type `Func<int, Task<int>>` because the result of `x + 1` (of type `int`) is implicitly convertible to the effective return type `int` of the async lambda, which has a return type `Task<int>`.

end example

A lambda expression `F` is compatible with an expression tree type `Expression<D>` if `F` is compatible with the delegate type `D`. This does not apply to anonymous methods, only lambda expressions.

Anonymous functions may influence overload resolution, and participate in type inference. See §12.6 for further details.

10.7.2 Evaluation of anonymous function conversions to delegate types

Conversion of an anonymous function to a delegate type produces a delegate instance that references the anonymous function and the (possibly empty) set of captured outer variables that are active at the time of the evaluation. When the delegate is invoked, the body of the anonymous function is executed. The code in the body is executed using the set of captured outer variables referenced by the delegate. A *delegate_creation_expression* (§12.8.16.6) can be used as an alternate syntax for converting an anonymous method to a delegate type.

The invocation list of a delegate produced from an anonymous function contains a single entry. The exact target object and target method of the delegate are unspecified. In particular, it is unspecified whether the target object of the delegate is `null`, the `this` value of the enclosing function member, or some other object.

Conversions of semantically identical anonymous functions with the same (possibly empty) set of captured outer variable instances to the same delegate types are permitted (but not required) to return the same delegate instance. The term semantically identical is used here to mean that execution of the anonymous functions will, in all cases, produce the same effects given the same arguments. This rule permits code such as the following to be optimized.

C#

```
delegate double Function(double x);

class Test
{
    static double[] Apply(double[] a, Function f)
    {
        double[] result = new double[a.Length];
        for (int i = 0; i < a.Length; i++)
        {
            result[i] = f(a[i]);
        }
        return result;
    }

    static void F(double[] a, double[] b)
    {
        a = Apply(a, (double x) => Math.Sin(x));
        b = Apply(b, (double y) => Math.Sin(y));
        ...
    }
}
```

Since the two anonymous function delegates have the same (empty) set of captured outer variables, and since the anonymous functions are semantically identical, the compiler is permitted to have the delegates refer to the same target method. Indeed, the compiler is permitted to return the very same delegate instance from both anonymous function expressions.

10.7.3 Evaluation of lambda expression conversions to expression tree types

Conversion of a lambda expression to an expression tree type produces an expression tree (§8.6). More precisely, evaluation of the lambda expression conversion produces an object structure that represents the structure of the lambda expression itself.

Not every lambda expression can be converted to expression tree types. The conversion to a compatible delegate type always *exists*, but it may fail at compile-time for implementation-specific reasons.

> *Note*: Common reasons for a lambda expression to fail to convert to an expression tree type include:
>
> - It has a block body
> - It has the `async` modifier
> - It contains an assignment operator
> - It contains an `out` or `ref` parameter
> - It contains a dynamically bound expression
>
> *end note*

10.8 Method group conversions

An implicit conversion exists from a method group (§12.2) to a compatible delegate type (§20.4). If `D` is a delegate type, and `E` is an expression that is classified as a method group, then `D` is compatible with `E` if and only if `E` contains at least one method that is applicable in its normal form (§12.6.4.2) to any argument list (§12.6.2) having types and modifiers matching the parameter types and modifiers of `D`, as described in the following.

The compile-time application of the conversion from a method group `E` to a delegate type `D` is described in the following.

- A single method `M` is selected corresponding to a method invocation (§12.8.9.2) of the form `E(A)`, with the following modifications:
  - The argument list `A` is a list of expressions, each classified as a variable and with the type and modifier (`in`, `out`, or `ref`) of the corresponding parameter in the *formal_parameter_list* of `D` — excepting parameters of type `dynamic`, where the corresponding expression has the type `object` instead of `dynamic`.
  - The candidate methods considered are only those methods that are applicable in their normal form and do not omit any optional parameters (§12.6.4.2). Thus, candidate methods are ignored if they are applicable only in their expanded form, or if one or more of their optional parameters do not have a corresponding parameter in `D`.
- A conversion is considered to exist if the algorithm of §12.8.9.2 produces a single best method `M` which is compatible (§20.4) with `D`.
- If the selected method `M` is an instance method, the instance expression associated with `E` determines the target object of the delegate.
- If the selected method `M` is an extension method which is denoted by means of a member access on an instance expression, that instance expression determines the target object of the delegate.
- The result of the conversion is a value of type `D`, namely a delegate that refers to the selected method and target object.

> *Example*: The following demonstrates method group conversions:
>
> ```
> C#
> ```
>
> ```
> delegate string D1(object o);
> delegate object D2(string s);
> delegate object D3();
> delegate string D4(object o, params object[] a);
> delegate string D5(int i);
> class Test
> {
>     static string F(object o) {...}
>
>     static void G()
>     {
>         D1 d1 = F;          // Ok
>         D2 d2 = F;          // Ok
>         D3 d3 = F;          // Error – not applicable
>         D4 d4 = F;          // Error – not applicable in normal form
>         D5 d5 = F;          // Error – applicable but not compatible
>     }
> }
> ```
>
> The assignment to `d1` implicitly converts the method group `F` to a value of type `D1`.

> The assignment to `d2` shows how it is possible to create a delegate to a method that has less derived (contravariant) parameter types and a more derived (covariant) return type.
>
> The assignment to `d3` shows how no conversion exists if the method is not applicable.
>
> The assignment to `d4` shows how the method must be applicable in its normal form.
>
> The assignment to `d5` shows how parameter and return types of the delegate and method are allowed to differ only for reference types.
>
> *end example*

As with all other implicit and explicit conversions, the cast operator can be used to explicitly perform a particular conversion.

> *Example*: Thus, the example
>
> ```
> C#
> object obj = new EventHandler(myDialog.OkClick);
> ```
>
> could instead be written
>
> ```
> C#
> object obj = (EventHandler)myDialog.OkClick;
> ```
>
> *end example*

A method group conversion can refer to a generic method, either by explicitly specifying type arguments within `E`, or via type inference (§12.6.3). If type inference is used, the parameter types of the delegate are used as argument types in the inference process. The return type of the delegate is not used for inference. Whether the type arguments are specified or inferred, they are part of the method group conversion process; these are the type arguments used to invoke the target method when the resulting delegate is invoked.

> *Example*:
>
> ```
> C#
> ```

```
delegate int D(string s, int i);
delegate int E();

class X
{
    public static T F<T>(string s, T t) {...}
    public static T G<T>() {...}

    static void Main()
    {
        D d1 = F<int>;        // Ok, type argument given explicitly
        D d2 = F;             // Ok, int inferred as type argument
        E e1 = G<int>;        // Ok, type argument given explicitly
        E e2 = G;             // Error, cannot infer from return type
    }
}
```

end example

Method groups may influence overload resolution, and participate in type inference. See §12.6 for further details.

The run-time evaluation of a method group conversion proceeds as follows:

- If the method selected at compile-time is an instance method, or it is an extension method which is accessed as an instance method, the target object of the delegate is determined from the instance expression associated with `E`:
  - The instance expression is evaluated. If this evaluation causes an exception, no further steps are executed.
  - If the instance expression is of a *reference_type*, the value computed by the instance expression becomes the target object. If the selected method is an instance method and the target object is `null`, a `System.NullReferenceException` is thrown and no further steps are executed.
  - If the instance expression is of a *value_type*, a boxing operation (§10.2.9) is performed to convert the value to an object, and this object becomes the target object.
- Otherwise, the selected method is part of a static method call, and the target object of the delegate is `null`.
- A delegate instance of delegate type `D` is obtained with a reference to the method that was determined at compile-time and a reference to the target object computed above, as follows:
- The conversion is permitted (but not required) to use an existing delegate instance that already contains these references.

- If an existing instance was not reused, a new one is created (§20.5). If there is not enough memory available to allocate the new instance, a `System.OutOfMemoryException` is thrown. Otherwise the instance is initialized with the given references.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

11 Patterns and pattern matching

Article • 02/07/2024

11.1 General

A ***pattern*** is a syntactic form that can be used with the `is` operator (§12.12.12) and in a *switch_statement* (§13.8.3) to express the shape of data against which incoming data is to be compared. A pattern is tested against the *expression* of a switch statement, or against a *relational_expression* that is on the left-hand side of an `is` operator, each of which is referred to as a ***pattern input value***.

11.2 Pattern forms

11.2.1 General

A pattern may have one of the following forms:

ANTLR

```
pattern
    : declaration_pattern
    | constant_pattern
    | var_pattern
    ;
```

A *declaration_pattern* and a *var_pattern* can result in the declaration of a local variable.

Each pattern form defines the set of types for input values that the pattern may be applied to. A pattern `P` is *applicable to* a type `T` if `T` is among the types whose values the pattern may match. It is a compile-time error if a pattern `P` appears in a program to match a pattern input value (§11.1) of type `T` if `P` is not applicable to `T`.

> *Example*: The following example generates a compile-time error because the compile-time type of `v` is `TextReader`. A variable of type `TextReader` can never have a value that is reference-compatible with `string`:
>
> C#
>
> ```
> TextReader v = Console.In; // compile-time type of 'v' is 'TextReader'
> if (v is string) // compile-time error
> {
> ```

```
    // code assuming v is a string
}
```

However, the following doesn't generate a compile-time error because the compile-time type of `v` is `object`. A variable of type `object` could have a value that is reference-compatible with `string`:

C#

```
object v = Console.In;
if (v is string s)
{
    // code assuming v is a string
}
```

end example

Each pattern form defines the set of values for which the pattern *matches* the value at runtime.

11.2.2 Declaration pattern

A *declaration_pattern* is used to test that a value has a given type and, if the test succeeds, provide the value in a variable of that type.

ANTLR

```
declaration_pattern
    : type simple_designation
    ;
simple_designation
    : single_variable_designation
    ;
single_variable_designation
    : identifier
    ;
```

The runtime type of the value is tested against the *type* in the pattern using the same rules specified in the is-type operator (§12.12.12.1). If the test succeeds, the pattern *matches* that value. It is a compile-time error if the *type* is a nullable value type (§8.3.12). This pattern form never matches a `null` value.

> *Note*: The is-type expression `e is T` and the declaration pattern `e is T _` are equivalent when `T` isn't a nullable type. *end note*

Given a pattern input value (§11.1) *e*, if the *simple_designation* is the *identifier* `_`, it denotes a discard (§9.2.9.1) and the value of *e* is not bound to anything. (Although a declared variable with the name `_` may be in scope at that point, that named variable is not seen in this context.) If *simple_designation* is any other identifier, a local variable (§9.2.9) of the given type named by the given identifier is introduced. That local variable is assigned the value of the pattern input value when the pattern *matches* the value.

Certain combinations of static type of the pattern input value and the given type are considered incompatible and result in a compile-time error. A value of static type `E` is said to be ***pattern compatible*** with the type `T` if there exists an identity conversion, an implicit or explicit reference conversion, a boxing conversion, or an unboxing conversion from `E` to `T`, or if either `E` or `T` is an open type (§8.4.3). A declaration pattern naming a type `T` is *applicable to* every type `E` for which `E` is pattern compatible with `T`.

> *Note*: The support for open types can be most useful when checking types that may be either struct or class types, and boxing is to be avoided. *end note*

> *Example*: The declaration pattern is useful for performing run-time type tests of reference types, and replaces the idiom
>
> ```
> C#
> ```
>
> ```
> var v = expr as Type;
> if (v != null) { /* code using v */ }
> ```
>
> with the slightly more concise
>
> ```
> C#
> ```
>
> ```
> if (expr is Type v) { /* code using v */ }
> ```
>
> *end example*

It is an error if *type* is a nullable value type.

> *Example*: The declaration pattern can be used to test values of nullable types: a value of type `Nullable<T>` (or a boxed `T`) matches a type pattern `T2 id` if the value is non-null and `T2` is `T`, or some base type or interface of `T`. For example, in the code fragment
>
> ```
> C#
> ```

```
int? x = 3;
if (x is int v) { /* code using v */ }
```

The condition of the `if` statement is `true` at runtime and the variable `v` holds the value `3` of type `int` inside the block. *end example*

11.2.3 Constant pattern

A *constant_pattern* is used to test the value of a pattern input value (§11.1) against the given constant value.

ANTLR

```
constant_pattern
    : constant_expression
    ;
```

A constant pattern `P` is *applicable to* a type `T` if there is an implicit conversion from the constant expression of `P` to the type `T`.

For a constant pattern `P`, its *converted value* is

- if the pattern input value's type is an integral type or an enum type, the pattern's constant value converted to that type; otherwise
- if the pattern input value's type is the nullable version of an integral type or an enum type, the pattern's constant value converted to its underlying type; otherwise
- the value of the pattern's constant value.

Given a pattern input value *e* and a constant pattern `P` with converted value *v*,

- if *e* has integral type or enum type, or a nullable form of one of those, and *v* has integral type, the pattern `P` *matches* the value *e* if result of the expression `e == v` is `true`; otherwise
- the pattern `P` *matches* the value *e* if `object.Equals(e, v)` returns `true`.

Example: The `switch` statement in the following method uses five constant patterns in its case labels.

C#

```
static decimal GetGroupTicketPrice(int visitorCount)
{
    switch (visitorCount)
```

```
    {
        case 1: return 12.0m;
        case 2: return 20.0m;
        case 3: return 27.0m;
        case 4: return 32.0m;
        case 0: return 0.0m;
        default: throw new ArgumentException(...);
    }
}
```

end example

11.2.4 Var pattern

A *var_pattern matches* every value. That is, a pattern-matching operation with a *var_pattern* always succeeds.

A *var_pattern* is *applicable to* every type.

ANTLR

```
var_pattern
    : 'var' designation
    ;
designation
    : simple_designation
    ;
```

Given a pattern input value (§11.1) *e*, if *designation* is the *identifier* `_`, it denotes a discard (§9.2.9.1), and the value of *e* is not bound to anything. (Although a declared variable with that name may be in scope at that point, that named variable is not seen in this context.) If *designation* is any other identifier, at runtime the value of *e* is bound to a newly introduced local variable (§9.2.9) of that name whose type is the static type of *e*, and the pattern input value is assigned to that local variable.

It is an error if the name `var` would bind to a type where a *var_pattern* is used.

11.3 Pattern subsumption

In a switch statement, it is an error if a case's pattern is *subsumed* by the preceding set of unguarded cases (§13.8.3). Informally, this means that any input value would have been matched by one of the previous cases. The following rules define when a set of patterns subsumes a given pattern:

A pattern `P` *would match* a constant `K` if the specification for that pattern's runtime behavior is that `P` matches `K`.

A set of patterns `Q` *subsumes* a pattern `P` if any of the following conditions hold:

- `P` is a constant pattern and any of the patterns in the set `Q` would match `P`'s *converted value*
- `P` is a var pattern and the set of patterns `Q` is *exhaustive* (§11.4) for the type of the pattern input value (§11.1), and either the pattern input value is not of a nullable type or some pattern in `Q` would match `null`.
- `P` is a declaration pattern with type `T` and the set of patterns `Q` is *exhaustive* for the type `T` (§11.4).

11.4 Pattern exhaustiveness

Informally, a set of patterns is exhaustive for a type if, for every possible value of that type other than null, some pattern in the set is applicable. The following rules define when a set of patterns is *exhaustive* for a type:

A set of patterns `Q` is *exhaustive* for a type `T` if any of the following conditions hold:

1. `T` is an integral or enum type, or a nullable version of one of those, and for every possible value of `T`'s non-nullable underlying type, some pattern in `Q` would match that value; or
2. Some pattern in `Q` is a *var pattern*; or
3. Some pattern in `Q` is a *declaration pattern* for type `D`, and there is an identity conversion, an implicit reference conversion, or a boxing conversion from `T` to `D`.

> *Example*:
>
> C#
>
> ```
> static void M(byte b)
> {
>     switch (b) {
>         case 0: case 1: case 2: ... // handle every specific value of
> byte
>             break;
>         // error: the pattern 'byte other' is subsumed by the
> (exhaustive)
>         // previous cases
>         case byte other:
>             break;
> ```

```
    }
}
```

end example

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

12 Expressions

Article • 02/07/2024

12.1 General

An expression is a sequence of operators and operands. This clause defines the syntax, order of evaluation of operands and operators, and meaning of expressions.

12.2 Expression classifications

12.2.1 General

The result of an expression is classified as one of the following:

- A value. Every value has an associated type.
- A variable. Unless otherwise specified, a variable is explicitly typed and has an associated type, namely the declared type of the variable. An implicitly typed variable has no associated type.
- A null literal. An expression with this classification can be implicitly converted to a reference type or nullable value type.
- An anonymous function. An expression with this classification can be implicitly converted to a compatible delegate type or expression tree type.
- A tuple. Every tuple has a fixed number of elements, each with an expression and an optional tuple element name.
- A property access. Every property access has an associated type, namely the type of the property. Furthermore, a property access may have an associated instance expression. When an accessor of an instance property access is invoked, the result of evaluating the instance expression becomes the instance represented by `this` (§12.8.13).
- An indexer access. Every indexer access has an associated type, namely the element type of the indexer. Furthermore, an indexer access has an associated instance expression and an associated argument list. When an accessor of an indexer access is invoked, the result of evaluating the instance expression becomes the instance represented by `this` (§12.8.13), and the result of evaluating the argument list becomes the parameter list of the invocation.
- Nothing. This occurs when the expression is an invocation of a method with a return type of `void`. An expression classified as nothing is only valid in the context of a *statement_expression* (§13.7) or as the body of a *lambda_expression* (§12.19).

For expressions which occur as subexpressions of larger expressions, with the noted restrictions, the result can also be classified as one of the following:

- A namespace. An expression with this classification can only appear as the left-hand side of a *member_access* (§12.8.7). In any other context, an expression classified as a namespace causes a compile-time error.
- A type. An expression with this classification can only appear as the left-hand side of a *member_access* (§12.8.7). In any other context, an expression classified as a type causes a compile-time error.
- A method group, which is a set of overloaded methods resulting from a member lookup (§12.5). A method group may have an associated instance expression and an associated type argument list. When an instance method is invoked, the result of evaluating the instance expression becomes the instance represented by `this` (§12.8.13). A method group is permitted in an *invocation_expression* (§12.8.9) or a *delegate_creation_expression* (§12.8.16.6), and can be implicitly converted to a compatible delegate type (§10.8). In any other context, an expression classified as a method group causes a compile-time error.
- An event access. Every event access has an associated type, namely the type of the event. Furthermore, an event access may have an associated instance expression. An event access may appear as the left operand of the `+=` and `-=` operators (§12.21.5). In any other context, an expression classified as an event access causes a compile-time error. When an accessor of an instance event access is invoked, the result of evaluating the instance expression becomes the instance represented by `this` (§12.8.13).
- A throw expression, which may be used is several contexts to throw an exception in an expression. A throw expression may be converted by an implicit conversion to any type.

A property access or indexer access is always reclassified as a value by performing an invocation of the get accessor or the set accessor. The particular accessor is determined by the context of the property or indexer access: If the access is the target of an assignment, the set accessor is invoked to assign a new value (§12.21.2). Otherwise, the get accessor is invoked to obtain the current value (§12.2.2).

An ***instance accessor*** is a property access on an instance, an event access on an instance, or an indexer access.

12.2.2 Values of expressions

Most of the constructs that involve an expression ultimately require the expression to denote a ***value***. In such cases, if the actual expression denotes a namespace, a type, a

method group, or nothing, a compile-time error occurs. However, if the expression denotes a property access, an indexer access, or a variable, the value of the property, indexer, or variable is implicitly substituted:

- The value of a variable is simply the value currently stored in the storage location identified by the variable. A variable shall be considered definitely assigned (§9.4) before its value can be obtained, or otherwise a compile-time error occurs.
- The value of a property access expression is obtained by invoking the get accessor of the property. If the property has no get accessor, a compile-time error occurs. Otherwise, a function member invocation (§12.6.6) is performed, and the result of the invocation becomes the value of the property access expression.
- The value of an indexer access expression is obtained by invoking the get accessor of the indexer. If the indexer has no get accessor, a compile-time error occurs. Otherwise, a function member invocation (§12.6.6) is performed with the argument list associated with the indexer access expression, and the result of the invocation becomes the value of the indexer access expression.
- The value of a tuple expression is obtained by applying an implicit tuple conversion (§10.2.13) to the type of the tuple expression. It is an error to obtain the value of a tuple expression that does not have a type.

12.3 Static and Dynamic Binding

12.3.1 General

Binding is the process of determining what an operation refers to, based on the type or value of expressions (arguments, operands, receivers). For instance, the binding of a method call is determined based on the type of the receiver and arguments. The binding of an operator is determined based on the type of its operands.

In C# the binding of an operation is usually determined at compile-time, based on the compile-time type of its subexpressions. Likewise, if an expression contains an error, the error is detected and reported by the compiler. This approach is known as ***static binding***.

However, if an expression is a *dynamic expression* (i.e., has the type `dynamic`) this indicates that any binding that it participates in should be based on its run-time type rather than the type it has at compile-time. The binding of such an operation is therefore deferred until the time where the operation is to be executed during the running of the program. This is referred to as ***dynamic binding***.

When an operation is dynamically bound, little or no checking is performed by the compiler. Instead if the run-time binding fails, errors are reported as exceptions at run-time.

The following operations in C# are subject to binding:

- Member access: `e.M`
- Method invocation: `e.M(e₁,...,eᵥ)`
- Delegate invocation: `e(e₁,...,eᵥ)`
- Element access: `e[e₁,...,eᵥ]`
- Object creation: new `C(e₁,...,eᵥ)`
- Overloaded unary operators: `+`, `-`, `!`, `~`, `++`, `--`, `true`, `false`
- Overloaded binary operators: `+`, `-`, `*`, `/`, `%`, `&`, `&&`, `|`, `||`, `??`, `^`, `<<`, `>>`, `==`, `!=`, `>`, `<`, `>=`, `<=`
- Assignment operators: `=`, `= ref`, `+=`, `-=`, `*=`, `/=`, `%=`, `&=`, `|=`, `^=`, `<<=`, `>>=`
- Implicit and explicit conversions

When no dynamic expressions are involved, C# defaults to static binding, which means that the compile-time types of subexpressions are used in the selection process. However, when one of the subexpressions in the operations listed above is a dynamic expression, the operation is instead dynamically bound.

It is a compile time error if a method invocation is dynamically bound and any of the parameters, including the receiver, has the `in` modifier.

12.3.2 Binding-time

Static binding takes place at compile-time, whereas dynamic binding takes place at run-time. In the following subclauses, the term ***binding-time*** refers to either compile-time or run-time, depending on when the binding takes place.

> *Example*: The following illustrates the notions of static and dynamic binding and of binding-time:
>
> C#
>
> ```
> object o = 5;
> dynamic d = 5;
> Console.WriteLine(5); // static binding to Console.WriteLine(int)
> Console.WriteLine(o); // static binding to Console.WriteLine(object)
> Console.WriteLine(d); // dynamic binding to Console.WriteLine(int)
> ```

> The first two calls are statically bound: the overload of `Console.WriteLine` is picked based on the compile-time type of their argument. Thus, the binding-time is *compile-time*.
>
> The third call is dynamically bound: the overload of `Console.WriteLine` is picked based on the run-time type of its argument. This happens because the argument is a dynamic expression – its compile-time type is dynamic. Thus, the binding-time for the third call is *run-time*.
>
> *end example*

12.3.3 Dynamic binding

This subclause is informative.

Dynamic binding allows C# programs to interact with dynamic objects, i.e., objects that do not follow the normal rules of the C# type system. Dynamic objects may be objects from other programming languages with different types systems, or they may be objects that are programmatically setup to implement their own binding semantics for different operations.

The mechanism by which a dynamic object implements its own semantics is implementation-defined. A given interface – again implementation-defined – is implemented by dynamic objects to signal to the C# run-time that they have special semantics. Thus, whenever operations on a dynamic object are dynamically bound, their own binding semantics, rather than those of C# as specified in this specification, take over.

While the purpose of dynamic binding is to allow interoperation with dynamic objects, C# allows dynamic binding on all objects, whether they are dynamic or not. This allows for a smoother integration of dynamic objects, as the results of operations on them may not themselves be dynamic objects, but are still of a type unknown to the programmer at compile-time. Also, dynamic binding can help eliminate error-prone reflection-based code even when no objects involved are dynamic objects.

12.3.4 Types of subexpressions

When an operation is statically bound, the type of a subexpression (e.g., a receiver, and argument, an index or an operand) is always considered to be the compile-time type of that expression.

When an operation is dynamically bound, the type of a subexpression is determined in different ways depending on the compile-time type of the subexpression:

- A subexpression of compile-time type dynamic is considered to have the type of the actual value that the expression evaluates to at run-time
- A subexpression whose compile-time type is a type parameter is considered to have the type which the type parameter is bound to at run-time
- Otherwise, the subexpression is considered to have its compile-time type.

12.4 Operators

12.4.1 General

Expressions are constructed from operands and operators. The operators of an expression indicate which operations to apply to the operands.

> *Example*: Examples of operators include `+`, `-`, `*`, `/`, and `new`. Examples of operands include literals, fields, local variables, and expressions. *end example*

There are three kinds of operators:

- Unary operators. The unary operators take one operand and use either prefix notation (such as `-x`) or postfix notation (such as `x++`).
- Binary operators. The binary operators take two operands and all use infix notation (such as `x + y`).
- Ternary operator. Only one ternary operator, `?:`, exists; it takes three operands and uses infix notation (`c ? x : y`).

The order of evaluation of operators in an expression is determined by the *precedence* and *associativity* of the operators (§12.4.2).

Operands in an expression are evaluated from left to right.

> *Example*: In `F(i) + G(i++) * H(i)`, method `F` is called using the old value of `i`, then method `G` is called with the old value of `i`, and, finally, method `H` is called with the new value of i. This is separate from and unrelated to operator precedence. *end example*

Certain operators can be ***overloaded***. Operator overloading (§12.4.3) permits user-defined operator implementations to be specified for operations where one or both of the operands are of a user-defined class or struct type.

12.4.2 Operator precedence and associativity

When an expression contains multiple operators, the ***precedence*** of the operators controls the order in which the individual operators are evaluated.

> *Note*: For example, the expression `x + y * z` is evaluated as `x + (y * z)` because the `*` operator has higher precedence than the binary `+` operator. *end note*

The precedence of an operator is established by the definition of its associated grammar production.

> *Note*: For example, an *additive_expression* consists of a sequence of *multiplicative_expression*s separated by `+` or `-` operators, thus giving the `+` and `-` operators lower precedence than the `*`, `/`, and `%` operators. *end note*

> *Note*: The following table summarizes all operators in order of precedence from highest to lowest:

Expand table

Subclause	Category	Operators
§12.8	Primary	`x.y` `x?.y` `f(x)` `a[x]` `a?[x]` `x++` `x--` `new` `typeof` `default` `checked` `unchecked` `delegate` `stackalloc`
§12.9	Unary	`+` `-` `!` `~` `++x` `--x` `(T)x` `await x`
§12.10	Multiplicative	`*` `/` `%`
§12.10	Additive	`+` `-`
§12.11	Shift	`<<` `>>`
§12.12	Relational and type-testing	`<` `>` `<=` `>=` `is` `as`
§12.12	Equality	`==` `!=`
§12.13	Logical AND	`&`
§12.13	Logical XOR	`^`
§12.13	Logical OR	`|`
§12.14	Conditional AND	`&&`

Subclause	Category	Operators
§12.14	Conditional OR	`||`
§12.15 and §12.16	Null coalescing and throw expression	`??` `throw x`
§12.18	Conditional	`?:`
§12.21 and §12.19	Assignment and lambda expression	`=` `= ref` `*=` `/=` `%=` `+=` `-=` `<<=` `>>=` `&=` `^=` `|=` `=>`

end note

When an operand occurs between two operators with the same precedence, the ***associativity*** of the operators controls the order in which the operations are performed:

- Except for the assignment operators and the null coalescing operator, all binary operators are ***left-associative***, meaning that operations are performed from left to right.

  > *Example*: `x + y + z` is evaluated as `(x + y) + z`. *end example*

- The assignment operators, the null coalescing operator and the conditional operator (`?:`) are ***right-associative***, meaning that operations are performed from right to left.

  > *Example*: `x = y = z` is evaluated as `x = (y = z)`. *end example*

Precedence and associativity can be controlled using parentheses.

> *Example*: `x + y * z` first multiplies `y` by `z` and then adds the result to `x`, but `(x + y) * z` first adds `x` and `y` and then multiplies the result by `z`. *end example*

12.4.3 Operator overloading

All unary and binary operators have predefined implementations. In addition, user-defined implementations can be introduced by including operator declarations (§15.10) in classes and structs. User-defined operator implementations always take precedence over predefined operator implementations: Only when no applicable user-defined operator implementations exist will the predefined operator implementations be considered, as described in §12.4.4 and §12.4.5.

The ***overloadable unary operators*** are:

```
C#

+   -   !   ~   ++   --   true   false
```

> *Note*: Although `true` and `false` are not used explicitly in expressions (and therefore are not included in the precedence table in §12.4.2), they are considered operators because they are invoked in several expression contexts: Boolean expressions (§12.24) and expressions involving the conditional (§12.18) and conditional logical operators (§12.14). *end note*

The ***overloadable binary operators*** are:

```
C#

+   -   *   /   %   &   |   ^   <<   >>   ==   !=   >   <   <=   >=
```

Only the operators listed above can be overloaded. In particular, it is not possible to overload member access, method invocation, or the `=`, `&&`, `||`, `??`, `?:`, `=>`, `checked`, `unchecked`, `new`, `typeof`, `default`, `as`, and `is` operators.

When a binary operator is overloaded, the corresponding compound assignment operator, if any, is also implicitly overloaded.

> *Example*: An overload of operator `*` is also an overload of operator `*=`. This is described further in §12.21. *end example*

The assignment operator itself `(=)` cannot be overloaded. An assignment always performs a simple store of a value into a variable (§12.21.2).

Cast operations, such as `(T)x`, are overloaded by providing user-defined conversions (§10.5).

> *Note*: User-defined conversions do not affect the behavior of the `is` or `as` operators. *end note*

Element access, such as `a[x]`, is not considered an overloadable operator. Instead, user-defined indexing is supported through indexers (§15.9).

In expressions, operators are referenced using operator notation, and in declarations, operators are referenced using functional notation. The following table shows the relationship between operator and functional notations for unary and binary operators. In the first entry, «op» denotes any overloadable unary prefix operator. In the second entry, «op» denotes the unary postfix `++` and `--` operators. In the third entry, «op» denotes any overloadable binary operator.

> *Note*: For an example of overloading the `++` and `--` operators see §15.10.2. *end note*

Expand table

Operator notation	Functional notation
`«op» x`	`operator «op»(x)`
`x «op»`	`operator «op»(x)`
`x «op» y`	`operator «op»(x, y)`

User-defined operator declarations always require at least one of the parameters to be of the class or struct type that contains the operator declaration.

> *Note*: Thus, it is not possible for a user-defined operator to have the same signature as a predefined operator. *end note*

User-defined operator declarations cannot modify the syntax, precedence, or associativity of an operator.

> *Example*: The `/` operator is always a binary operator, always has the precedence level specified in §12.4.2, and is always left-associative. *end example*

> *Note*: While it is possible for a user-defined operator to perform any computation it pleases, implementations that produce results other than those that are intuitively expected are strongly discouraged. For example, an implementation of operator `==` should compare the two operands for equality and return an appropriate `bool` result. *end note*

The descriptions of individual operators in §12.9 through §12.21 specify the predefined implementations of the operators and any additional rules that apply to each operator. The descriptions make use of the terms ***unary operator overload resolution***, ***binary operator overload resolution***, ***numeric promotion***, and lifted operator definitions of which are found in the following subclauses.

12.4.4 Unary operator overload resolution

An operation of the form `«op» x` or `x «op»`, where «op» is an overloadable unary operator, and `x` is an expression of type `X`, is processed as follows:

- The set of candidate user-defined operators provided by `X` for the operation `operator «op»(x)` is determined using the rules of §12.4.6.
- If the set of candidate user-defined operators is not empty, then this becomes the set of candidate operators for the operation. Otherwise, the predefined binary `operator «op»` implementations, including their lifted forms, become the set of candidate operators for the operation. The predefined implementations of a given operator are specified in the description of the operator. The predefined operators provided by an enum or delegate type are only included in this set when the binding-time type—or the underlying type if it is a nullable type—of either operand is the enum or delegate type.
- The overload resolution rules of §12.6.4 are applied to the set of candidate operators to select the best operator with respect to the argument list `(x)`, and this operator becomes the result of the overload resolution process. If overload resolution fails to select a single best operator, a binding-time error occurs.

12.4.5 Binary operator overload resolution

An operation of the form `x «op» y`, where «op» is an overloadable binary operator, `x` is an expression of type `X`, and `y` is an expression of type `Y`, is processed as follows:

- The set of candidate user-defined operators provided by `X` and `Y` for the operation `operator «op»(x, y)` is determined. The set consists of the union of the candidate operators provided by `X` and the candidate operators provided by `Y`, each determined using the rules of §12.4.6. For the combined set, candidates are merged as follows:
  - If `X` and `Y` are identity convertible, or if `X` and `Y` are derived from a common base type, then shared candidate operators only occur in the combined set once.
  - If there is an identity conversion between `X` and `Y`, an operator `«op»Y` provided by `Y` has the same return type as an `«op»X` provided by `X` and the operand types of `«op»Y` have an identity conversion to the corresponding operand types of `«op»X` then only `«op»X` occurs in the set.
- If the set of candidate user-defined operators is not empty, then this becomes the set of candidate operators for the operation. Otherwise, the predefined binary `operator «op»` implementations, including their lifted forms, become the set of

candidate operators for the operation. The predefined implementations of a given operator are specified in the description of the operator. For predefined enum and delegate operators, the only operators considered are those provided by an enum or delegate type that is the binding-time type of one of the operands.

- The overload resolution rules of §12.6.4 are applied to the set of candidate operators to select the best operator with respect to the argument list `(x, y)`, and this operator becomes the result of the overload resolution process. If overload resolution fails to select a single best operator, a binding-time error occurs.

12.4.6 Candidate user-defined operators

Given a type `T` and an operation `operator «op»(A)`, where «op» is an overloadable operator and `A` is an argument list, the set of candidate user-defined operators provided by `T` for operator `«op»(A)` is determined as follows:

- Determine the type T_o. If `T` is a nullable value type, T_o is its underlying type; otherwise, T_o is equal to `T`.
- For all `operator «op»` declarations in T_o and all lifted forms of such operators, if at least one operator is applicable (§12.6.4.2) with respect to the argument list `A`, then the set of candidate operators consists of all such applicable operators in T_o.
- Otherwise, if T_o is `object`, the set of candidate operators is empty.
- Otherwise, the set of candidate operators provided by T_o is the set of candidate operators provided by the direct base class of T_o, or the effective base class of T_o if T_o is a type parameter.

12.4.7 Numeric promotions

12.4.7.1 General

This subclause is informative.

§12.4.7 and its subclauses are a summary of the combined effect of:

- the rules for implicit numeric conversions (§10.2.3);
- the rules for better conversion (§12.6.4.7); and
- the available arithmetic (§12.10), relational (§12.12), and integral logical (§12.13.2) operators.

Numeric promotion consists of automatically performing certain implicit conversions of the operands of the predefined unary and binary numeric operators. Numeric

promotion is not a distinct mechanism, but rather an effect of applying overload resolution to the predefined operators. Numeric promotion specifically does not affect evaluation of user-defined operators, although user-defined operators can be implemented to exhibit similar effects.

As an example of numeric promotion, consider the predefined implementations of the binary `*` operator:

```
C#
```

```
int operator *(int x, int y);
uint operator *(uint x, uint y);
long operator *(long x, long y);
ulong operator *(ulong x, ulong y);
float operator *(float x, float y);
double operator *(double x, double y);
decimal operator *(decimal x, decimal y);
```

When overload resolution rules (§12.6.4) are applied to this set of operators, the effect is to select the first of the operators for which implicit conversions exist from the operand types.

> *Example*: For the operation `b * s`, where `b` is a `byte` and `s` is a `short`, overload resolution selects `operator *(int, int)` as the best operator. Thus, the effect is that `b` and `s` are converted to `int`, and the type of the result is `int`. Likewise, for the operation `i * d`, where `i` is an `int` and `d` is a `double`, `overload` resolution selects `operator *(double, double)` as the best operator. *end example*

End of informative text.

12.4.7.2 Unary numeric promotions

This subclause is informative.

Unary numeric promotion occurs for the operands of the predefined `+`, `-`, and `~` unary operators. Unary numeric promotion simply consists of converting operands of type `sbyte`, `byte`, `short`, `ushort`, or `char` to type `int`. Additionally, for the unary – operator, unary numeric promotion converts operands of type `uint` to type `long`.

End of informative text.

12.4.7.3 Binary numeric promotions

This subclause is informative.

Binary numeric promotion occurs for the operands of the predefined `+`, `-`, `*`, `/`, `%`, `&`, `|`, `^`, `==`, `!=`, `>`, `<`, `>=`, and `<=` binary operators. Binary numeric promotion implicitly converts both operands to a common type which, in case of the non-relational operators, also becomes the result type of the operation. Binary numeric promotion consists of applying the following rules, in the order they appear here:

- If either operand is of type `decimal`, the other operand is converted to type `decimal`, or a binding-time error occurs if the other operand is of type `float` or `double`.
- Otherwise, if either operand is of type `double`, the other operand is converted to type `double`.
- Otherwise, if either operand is of type `float`, the other operand is converted to type `float`.
- Otherwise, if either operand is of type `ulong`, the other operand is converted to type `ulong`, or a binding-time error occurs if the other operand is of `type sbyte`, `short`, `int`, or `long`.
- Otherwise, if either operand is of type `long`, the other operand is converted to type `long`.
- Otherwise, if either operand is of type `uint` and the other operand is of type `sbyte`, `short`, or `int`, both operands are converted to type `long`.
- Otherwise, if either operand is of type `uint`, the other operand is converted to type `uint`.
- Otherwise, both operands are converted to type `int`.

> *Note*: The first rule disallows any operations that mix the `decimal` type with the `double` and `float` types. The rule follows from the fact that there are no implicit conversions between the `decimal` type and the `double` and `float` types. *end note*

> *Note*: Also note that it is not possible for an operand to be of type `ulong` when the other operand is of a signed integral type. The reason is that no integral type exists that can represent the full range of `ulong` as well as the signed integral types. *end note*

In both of the above cases, a cast expression can be used to explicitly convert one operand to a type that is compatible with the other operand.

> *Example*: In the following code

```C#
decimal AddPercent(decimal x, double percent) =>
    x * (1.0 + percent / 100.0);
```

a binding-time error occurs because a `decimal` cannot be multiplied by a `double`. The error is resolved by explicitly converting the second operand to `decimal`, as follows:

```C#
decimal AddPercent(decimal x, double percent) =>
    x * (decimal)(1.0 + percent / 100.0);
```

end example

End of informative text.

12.4.8 Lifted operators

Lifted operators permit predefined and user-defined operators that operate on non-nullable value types to also be used with nullable forms of those types. Lifted operators are constructed from predefined and user-defined operators that meet certain requirements, as described in the following:

- For the unary operators `+`, `++`, `-`, `--`, `!`, and `~`, a lifted form of an operator exists if the operand and result types are both non-nullable value types. The lifted form is constructed by adding a single `?` modifier to the operand and result types. The lifted operator produces a `null` value if the operand is `null`. Otherwise, the lifted operator unwraps the operand, applies the underlying operator, and wraps the result.
- For the binary operators `+`, `-`, `*`, `/`, `%`, `&`, `|`, `^`, `<<`, and `>>`, a lifted form of an operator exists if the operand and result types are all non-nullable value types. The lifted form is constructed by adding a single `?` modifier to each operand and result type. The lifted operator produces a `null` value if one or both operands are `null` (an exception being the `&` and `|` operators of the `bool?` type, as described in §12.13.5). Otherwise, the lifted operator unwraps the operands, applies the underlying operator, and wraps the result.
- For the equality operators `==` and `!=`, a lifted form of an operator exists if the operand types are both non-nullable value types and if the result type is `bool`. The lifted form is constructed by adding a single `?` modifier to each operand type. The

lifted operator considers two `null` values equal, and a `null` value unequal to any non-`null` value. If both operands are non-`null`, the lifted operator unwraps the operands and applies the underlying operator to produce the `bool` result.

- For the relational operators `<`, `>`, `<=`, and `>=`, a lifted form of an operator exists if the operand types are both non-nullable value types and if the result type is `bool`. The lifted form is constructed by adding a single `?` modifier to each operand type. The lifted operator produces the value `false` if one or both operands are `null`. Otherwise, the lifted operator unwraps the operands and applies the underlying operator to produce the `bool` result.

12.5 Member lookup

12.5.1 General

A member lookup is the process whereby the meaning of a name in the context of a type is determined. A member lookup can occur as part of evaluating a *simple_name* (§12.8.4) or a *member_access* (§12.8.7) in an expression. If the *simple_name* or *member_access* occurs as the *primary_expression* of an *invocation_expression* (§12.8.9.2), the member is said to be *invoked*.

If a member is a method or event, or if it is a constant, field or property of either a delegate type (§20) or the type `dynamic` (§8.2.4), then the member is said to be *invocable*.

Member lookup considers not only the name of a member but also the number of type parameters the member has and whether the member is accessible. For the purposes of member lookup, generic methods and nested generic types have the number of type parameters indicated in their respective declarations and all other members have zero type parameters.

A member lookup of a name `N` with `K` type arguments in a type `T` is processed as follows:

- First, a set of accessible members named `N` is determined:
  - If `T` is a type parameter, then the set is the union of the sets of accessible members named `N` in each of the types specified as a primary constraint or secondary constraint (§15.2.5) for `T`, along with the set of accessible members named `N` in `object`.
  - Otherwise, the set consists of all accessible (§7.5) members named `N` in `T`, including inherited members and the accessible members named `N` in `object`. If

`T` is a constructed type, the set of members is obtained by substituting type arguments as described in §15.3.3. Members that include an `override` modifier are excluded from the set.

- Next, if `K` is zero, all nested types whose declarations include type parameters are removed. If `K` is not zero, all members with a different number of type parameters are removed. When `K` is zero, methods having type parameters are not removed, since the type inference process (§12.6.3) might be able to infer the type arguments.
- Next, if the member is invoked, all non-invocable members are removed from the set.
- Next, members that are hidden by other members are removed from the set. For every member `S.M` in the set, where `S` is the type in which the member `M` is declared, the following rules are applied:
  - If `M` is a constant, field, property, event, or enumeration member, then all members declared in a base type of `S` are removed from the set.
  - If `M` is a type declaration, then all non-types declared in a base type of `S` are removed from the set, and all type declarations with the same number of type parameters as `M` declared in a base type of `S` are removed from the set.
  - If `M` is a method, then all non-method members declared in a base type of `S` are removed from the set.
- Next, interface members that are hidden by class members are removed from the set. This step only has an effect if `T` is a type parameter and `T` has both an effective base class other than `object` and a non-empty effective interface set (§15.2.5). For every member `S.M` in the set, where `S` is the type in which the member `M` is declared, the following rules are applied if `S` is a class declaration other than `object`:
  - If `M` is a constant, field, property, event, enumeration member, or type declaration, then all members declared in an interface declaration are removed from the set.
  - If `M` is a method, then all non-method members declared in an interface declaration are removed from the set, and all methods with the same signature as `M` declared in an interface declaration are removed from the set.
- Finally, having removed hidden members, the result of the lookup is determined:
  - If the set consists of a single member that is not a method, then this member is the result of the lookup.
  - Otherwise, if the set contains only methods, then this group of methods is the result of the lookup.
  - Otherwise, the lookup is ambiguous, and a binding-time error occurs.

For member lookups in types other than type parameters and interfaces, and member lookups in interfaces that are strictly single-inheritance (each interface in the inheritance chain has exactly zero or one direct base interface), the effect of the lookup rules is simply that derived members hide base members with the same name or signature. Such single-inheritance lookups are never ambiguous. The ambiguities that can possibly arise from member lookups in multiple-inheritance interfaces are described in §18.4.6.

> *Note*: This phase only accounts for one kind of ambiguity. If the member lookup results in a method group, further uses of method group may fail due to ambiguity, for example as described in §12.6.4.1 and §12.6.6.2. *end note*

12.5.2 Base types

For purposes of member lookup, a type `T` is considered to have the following base types:

- If `T` is `object` or `dynamic`, then `T` has no base type.
- If `T` is an *enum_type*, the base types of `T` are the class types `System.Enum`, `System.ValueType`, and `object`.
- If `T` is a *struct_type*, the base types of `T` are the class types `System.ValueType` and `object`.

  > *Note*: A *nullable_value_type* is a *struct_type* (§8.3.1). *end note*

- If `T` is a *class_type*, the base types of `T` are the base classes of `T`, including the class type `object`.
- If `T` is an *interface_type*, the base types of `T` are the base interfaces of `T` and the class type `object`.
- If `T` is an *array_type*, the base types of `T` are the class types `System.Array` and `object`.
- If `T` is a *delegate_type*, the base types of `T` are the class types `System.Delegate` and `object`.

12.6 Function members

12.6.1 General

Function members are members that contain executable statements. Function members are always members of types and cannot be members of namespaces. C# defines the

following categories of function members:

- Methods
- Properties
- Events
- Indexers
- User-defined operators
- Instance constructors
- Static constructors
- Finalizers

Except for finalizers and static constructors (which cannot be invoked explicitly), the statements contained in function members are executed through function member invocations. The actual syntax for writing a function member invocation depends on the particular function member category.

The argument list (§12.6.2) of a function member invocation provides actual values or variable references for the parameters of the function member.

Invocations of generic methods may employ type inference to determine the set of type arguments to pass to the method. This process is described in §12.6.3.

Invocations of methods, indexers, operators, and instance constructors employ overload resolution to determine which of a candidate set of function members to invoke. This process is described in §12.6.4.

Once a particular function member has been identified at binding-time, possibly through overload resolution, the actual run-time process of invoking the function member is described in §12.6.6.

> *Note*: The following table summarizes the processing that takes place in constructs involving the six categories of function members that can be explicitly invoked. In the table, `e`, `x`, `y`, and `value` indicate expressions classified as variables or values, `T` indicates an expression classified as a type, `F` is the simple name of a method, and `P` is the simple name of a property.

Expand table

Construct	Example	Description
Method invocation	`F(x, y)`	Overload resolution is applied to select the best method `F` in the containing class or struct. The method is invoked with the argument list `(x, y)`. If the method is not `static`, the instance expression is `this`.

	`T.F(x, y)`	Overload resolution is applied to select the best method `F` in the class or struct `T`. A binding-time error occurs if the method is not `static`. The method is invoked with the argument list `(x, y)`.
	`e.F(x, y)`	Overload resolution is applied to select the best method `F` in the class, struct, or interface given by the type of `e`. A binding-time error occurs if the method is `static`. The method is invoked with the instance expression `e` and the argument list `(x, y)`.
Property access	`P`	The get accessor of the property `P` in the containing class or struct is invoked. A compile-time error occurs if `P` is write-only. If `P` is not `static`, the instance expression is `this`.
	`P = value`	The set accessor of the property `P` in the containing class or struct is invoked with the argument list `(value)`. A compile-time error occurs if `P` is read-only. If `P` is not `static`, the instance expression is `this`.
	`T.P`	The get accessor of the property `P` in the class or struct `T` is invoked. A compile-time error occurs if `P` is not `static` or if `P` is write-only.
	`T.P = value`	The set accessor of the property `P` in the class or struct `T` is invoked with the argument list `(value)`. A compile-time error occurs if `P` is not `static` or if `P` is read-only.
	`e.P`	The get accessor of the property `P` in the class, struct, or interface given by the type of `E` is invoked with the instance expression `e`. A binding-time error occurs if `P` is `static` or if `P` is write-only.
	`e.P = value`	The set accessor of the property `P` in the class, struct, or interface given by the type of `E` is invoked with the instance expression `e` and the argument list `(value)`. A binding-time error occurs if `P` is `static` or if `P` is read-only.
Event access	`E += value`	The add accessor of the event `E` in the containing class or struct is invoked. If `E` is not `static`, the instance expression is `this`.
	`E -= value`	The remove accessor of the event `E` in the containing class or struct is invoked. If `E` is not `static`, the instance expression is `this`.
	`T.E += value`	The add accessor of the event `E` in the class or struct `T` is invoked. A binding-time error occurs if `E` is not `static`.

	`T.E -= value`	The remove accessor of the event `E` in the class or struct `T` is invoked. A binding-time error occurs if `E` is not `static`.
	`e.E += value`	The add accessor of the event `E` in the class, struct, or interface given by the type of `E` is invoked with the instance expression `e`. A binding-time error occurs if `E` is `static`.
	`e.E -= value`	The remove accessor of the event `E` in the class, struct, or interface given by the type of `E` is invoked with the instance expression `e`. A binding-time error occurs if `E` is `static`.
Indexer access	`e[x, y]`	Overload resolution is applied to select the best indexer in the class, struct, or interface given by the type of `e`. The get accessor of the indexer is invoked with the instance expression `e` and the argument list `(x, y)`. A binding-time error occurs if the indexer is write-only.
	`e[x, y] = value`	Overload resolution is applied to select the best indexer in the class, struct, or interface given by the type of `e`. The set accessor of the indexer is invoked with the instance expression `e` and the argument list `(x, y, value)`. A binding-time error occurs if the indexer is read-only.
Operator invocation	`-x`	Overload resolution is applied to select the best unary operator in the class or struct given by the type of `x`. The selected operator is invoked with the argument list `(x)`.
	`x + y`	Overload resolution is applied to select the best binary operator in the classes or structs given by the types of `x` and `y`. The selected operator is invoked with the argument list `(x, y)`.
Instance constructor invocation	`new T(x, y)`	Overload resolution is applied to select the best instance constructor in the class or struct `T`. The instance constructor is invoked with the argument list `(x, y)`.

end note

12.6.2 Argument lists

12.6.2.1 General

Every function member and delegate invocation includes an argument list, which provides actual values or variable references for the parameters of the function member. The syntax for specifying the argument list of a function member invocation depends on the function member category:

- For instance constructors, methods, indexers and delegates, the arguments are specified as an *argument_list*, as described below. For indexers, when invoking the set accessor, the argument list additionally includes the expression specified as the right operand of the assignment operator.

  > *Note*: This additional argument is not used for overload resolution, just during invocation of the set accessor. *end note*

- For properties, the argument list is empty when invoking the get accessor, and consists of the expression specified as the right operand of the assignment operator when invoking the set accessor.
- For events, the argument list consists of the expression specified as the right operand of the `+=` or `-=` operator.
- For user-defined operators, the argument list consists of the single operand of the unary operator or the two operands of the binary operator.

The arguments of properties (§15.7) and events (§15.8) are always passed as value parameters (§15.6.2.2). The arguments of user-defined operators (§15.10) are always passed as value parameters (§15.6.2.2) or input parameters (§9.2.8). The arguments of indexers (§15.9) are always passed as value parameters (§15.6.2.2), input parameters (§9.2.8), or parameter arrays (§15.6.2.4). Output and reference parameters are not supported for these categories of function members.

The arguments of an instance constructor, method, indexer, or delegate invocation are specified as an *argument_list*:

ANTLR

```
argument_list
    : argument (',' argument)*
    ;

argument
    : argument_name? argument_value
    ;

argument_name
    : identifier ':'
    ;

argument_value
    : expression
    | 'in' variable_reference
    | 'ref' variable_reference
    | 'out' variable_reference
    ;
```

An *argument_list* consists of one or more *argument*s, separated by commas. Each argument consists of an optional *argument_name* followed by an *argument_value*. An *argument* with an *argument_name* is referred to as a ***named argument***, whereas an *argument* without an *argument_name* is a ***positional argument***.

The *argument_value* can take one of the following forms:

- An *expression*, indicating that the argument is passed as a value parameter or is transformed into an input parameter and then passed as that, as determined by (§12.6.4.2 and described in §12.6.2.3.
- The keyword `in` followed by a *variable_reference* (§9.5), indicating that the argument is passed as an input parameter (§15.6.2.3.2). A variable shall be definitely assigned (§9.4) before it can be passed as an input parameter.
- The keyword `ref` followed by a *variable_reference* (§9.5), indicating that the argument is passed as a reference parameter (§15.6.2.3.3). A variable shall be definitely assigned (§9.4) before it can be passed as a reference parameter.
- The keyword `out` followed by a *variable_reference* (§9.5), indicating that the argument is passed as an output parameter (§15.6.2.3.4). A variable is considered definitely assigned (§9.4) following a function member invocation in which the variable is passed as an output parameter.

The form determines the ***parameter-passing mode*** of the argument: *value*, *input*, *reference*, or *output*, respectively. However, as mentioned above, an argument with value passing mode, might be transformed into one with input passing mode.

Passing a volatile field (§15.5.4) as an input, output, or reference parameter causes a warning, since the field may not be treated as volatile by the invoked method.

12.6.2.2 Corresponding parameters

For each argument in an argument list there has to be a corresponding parameter in the function member or delegate being invoked.

The parameter list used in the following is determined as follows:

- For virtual methods and indexers defined in classes, the parameter list is picked from the first declaration or override of the function member found when starting with the static type of the receiver, and searching through its base classes.
- For partial methods, the parameter list of the defining partial method declaration is used.
- For all other function members and delegates there is only a single parameter list, which is the one used.

The position of an argument or parameter is defined as the number of arguments or parameters preceding it in the argument list or parameter list.

The corresponding parameters for function member arguments are established as follows:

- Arguments in the *argument_list* of instance constructors, methods, indexers and delegates:
  - A positional argument where a parameter occurs at the same position in the parameter list corresponds to that parameter, unless the parameter is a parameter array and the function member is invoked in its expanded form.
  - A positional argument of a function member with a parameter array invoked in its expanded form, which occurs at or after the position of the parameter array in the parameter list, corresponds to an element in the parameter array.
  - A named argument corresponds to the parameter of the same name in the parameter list.
  - For indexers, when invoking the set accessor, the expression specified as the right operand of the assignment operator corresponds to the implicit `value` parameter of the set accessor declaration.
- For properties, when invoking the get accessor there are no arguments. When invoking the set accessor, the expression specified as the right operand of the assignment operator corresponds to the implicit value parameter of the set accessor declaration.
- For user-defined unary operators (including conversions), the single operand corresponds to the single parameter of the operator declaration.
- For user-defined binary operators, the left operand corresponds to the first parameter, and the right operand corresponds to the second parameter of the operator declaration.
- An unnamed argument corresponds to no parameter when it is after an out-of-position named argument or a named argument that corresponds to a parameter array.

> *Note*: This prevents `void M(bool a = true, bool b = true, bool c = true);` being invoked by `M(c: false, valueB);`. The first argument is used out-of-position (the argument is used in first position, but the parameter named `c` is in third position), so the following arguments should be named. In other words, non-trailing named arguments are only allowed when the name and the position result in finding the same corresponding parameter. *end note*

12.6.2.3 Run-time evaluation of argument lists

During the run-time processing of a function member invocation (§12.6.6), the expressions or variable references of an argument list are evaluated in order, from left to right, as follows:

- For a value argument, if the parameter's passing mode is value
  - the argument expression is evaluated and an implicit conversion (§10.2) to the corresponding parameter type is performed. The resulting value becomes the initial value of the value parameter in the function member invocation.
  - otherwise, the parameter's passing mode is input. If the argument is a variable reference and there exists an identity conversion (§10.2.2) between the argument's type and the parameter's type, the resulting storage location becomes the storage location represented by the parameter in the function member invocation. Otherwise, a storage location is created with the same type as that of the corresponding parameter. The argument expression is evaluated and an implicit conversion (§10.2) to the corresponding parameter type is performed. The resulting value is stored within that storage location. That storage location is represented by the input parameter in the function member invocation.

    > *Example*: Given the following declarations and method calls:
    >
    > C#
    >
    > ```
    > static void M1(in int p1) { ... }
    > int i = 10;
    > M1(i);         // i is passed as an input argument
    > M1(i + 5);     // transformed to a temporary input argument
    > ```
    >
    > In the `M1(i)` method call, `i` itself is passed as an input argument, because it is classified as a variable and has the same type `int` as the input parameter. In the `M1(i + 5)` method call, an unnamed `int` variable is created, initialized with the argument's value, and then passed as an input argument. See §12.6.4.2 and §12.6.4.4.
    >
    > *end example*

- For an input, output, or reference argument, the variable reference is evaluated and the resulting storage location becomes the storage location represented by the parameter in the function member invocation. For an input or reference argument, the variable shall be definitely assigned at the point of the method call. If the variable reference is given as an output argument, or is an array element of a

reference_type, a run-time check is performed to ensure that the element type of the array is identical to the type of the parameter. If this check fails, a `System.ArrayTypeMismatchException` is thrown.

> *Note*: this run-time check is required due to array covariance (§17.6). *end note*

> *Example*: In the following code

```C#
class Test
{
    static void F(ref object x) {...}

    static void Main()
    {
        object[] a = new object[10];
        object[] b = new string[10];
        F(ref a[0]); // Ok
        F(ref b[1]); // ArrayTypeMismatchException
    }
}
```

> the second invocation of `F` causes a `System.ArrayTypeMismatchException` to be thrown because the actual element type of `b` is `string` and not `object`.
>
> *end example*

Methods, indexers, and instance constructors may declare their right-most parameter to be a parameter array (§15.6.2.4). Such function members are invoked either in their normal form or in their expanded form depending on which is applicable (§12.6.4.2):

- When a function member with a parameter array is invoked in its normal form, the argument given for the parameter array shall be a single expression that is implicitly convertible (§10.2) to the parameter array type. In this case, the parameter array acts precisely like a value parameter.
- When a function member with a parameter array is invoked in its expanded form, the invocation shall specify zero or more positional arguments for the parameter array, where each argument is an expression that is implicitly convertible (§10.2) to the element type of the parameter array. In this case, the invocation creates an instance of the parameter array type with a length corresponding to the number of arguments, initializes the elements of the array instance with the given argument values, and uses the newly created array instance as the actual argument.

The expressions of an argument list are always evaluated in textual order.

> *Example*: Thus, the example
>
> C#
>
> ```
> class Test
> {
>     static void F(int x, int y = -1, int z = -2) =>
>         Console.WriteLine($"x = {x}, y = {y}, z = {z}");
>
>     static void Main()
>     {
>         int i = 0;
>         F(i++, i++, i++);
>         F(z: i++, x: i++);
>     }
> }
> ```
>
> produces the output
>
> Console
>
> ```
> x = 0, y = 1, z = 2
> x = 4, y = -1, z = 3
> ```
>
> *end example*

When a function member with a parameter array is invoked in its expanded form with at least one expanded argument, the invocation is processed as if an array creation expression with an array initializer (§12.8.16.5) was inserted around the expanded arguments. An empty array is passed when there are no arguments for the parameter array; it is unspecified whether the reference passed is to a newly allocated or existing empty array.

> *Example*: Given the declaration
>
> C#
>
> ```
> void F(int x, int y, params object[] args);
> ```
>
> the following invocations of the expanded form of the method
>
> C#

```
F(10, 20, 30, 40);
F(10, 20, 1, "hello", 3.0);
```

correspond exactly to

C#

```
F(10, 20, new object[] { 30, 40 });
F(10, 20, new object[] { 1, "hello", 3.0 });
```

end example

When arguments are omitted from a function member with corresponding optional parameters, the default arguments of the function member declaration are implicitly passed. (This can involve the creation of a storage location, as described above.)

Note: Because these are always constant, their evaluation will not impact the evaluation of the remaining arguments. *end note*

12.6.3 Type inference

12.6.3.1 General

When a generic method is called without specifying type arguments, a ***type inference*** process attempts to infer type arguments for the call. The presence of type inference allows a more convenient syntax to be used for calling a generic method, and allows the programmer to avoid specifying redundant type information.

Example:

C#

```
class Chooser
{
    static Random rand = new Random();

    public static T Choose<T>(T first, T second) =>
        rand.Next(2) == 0 ? first : second;
}

class A
{
    static void M()
    {
```

```
        int i = Chooser.Choose(5, 213); // Calls Choose<int>
        string s = Chooser.Choose("apple", "banana"); // Calls
Choose<string>
    }
}
```

Through type inference, the type arguments `int` and `string` are determined from the arguments to the method.

end example

Type inference occurs as part of the binding-time processing of a method invocation (§12.8.9.2) and takes place before the overload resolution step of the invocation. When a particular method group is specified in a method invocation, and no type arguments are specified as part of the method invocation, type inference is applied to each generic method in the method group. If type inference succeeds, then the inferred type arguments are used to determine the types of arguments for subsequent overload resolution. If overload resolution chooses a generic method as the one to invoke, then the inferred type arguments are used as the type arguments for the invocation. If type inference for a particular method fails, that method does not participate in overload resolution. The failure of type inference, in and of itself, does not cause a binding-time error. However, it often leads to a binding-time error when overload resolution then fails to find any applicable methods.

If each supplied argument does not correspond to exactly one parameter in the method (§12.6.2.2), or there is a non-optional parameter with no corresponding argument, then inference immediately fails. Otherwise, assume that the generic method has the following signature:

T_e M<X_1...X_v>(T_1 p_1 ... T_x p_x)

With a method call of the form M(E_1 ...E_x) the task of type inference is to find unique type arguments S_1...S_v for each of the type parameters X_1...X_v so that the call M<S_1...S_v>(E_1...E_x) becomes valid.

The process of type inference is described below as an algorithm. A conformant compiler may be implemented using an alternative approach, provided it reaches the same result in all cases.

During the process of inference each type parameter X_i is either *fixed* to a particular type S_i or *unfixed* with an associated set of *bounds*. Each of the bounds is some type `T`. Initially each type variable X_i is unfixed with an empty set of bounds.

Type inference takes place in phases. Each phase will try to infer type arguments for more type variables based on the findings of the previous phase. The first phase makes some initial inferences of bounds, whereas the second phase fixes type variables to specific types and infers further bounds. The second phase may have to be repeated a number of times.

> *Note*: Type inference is also used in other contexts including for conversion of method groups (§12.6.3.14) and finding the best common type of a set of expressions (§12.6.3.15). *end note*

12.6.3.2 The first phase

For each of the method arguments E_i:

- If E_i is an anonymous function, an *explicit parameter type inference* (§12.6.3.8) is made *from* E_i *to* T_i
- Otherwise, if E_i has a type U and the corresponding parameter is a value parameter (§15.6.2.2) then a *lower-bound inference* (§12.6.3.10) is made *from* U *to* T_i.
- Otherwise, if E_i has a type U and the corresponding parameter is a reference parameter (§15.6.2.3.3), or output parameter (§15.6.2.3.4) then an *exact inference* (§12.6.3.9) is made *from* U *to* T_i.
- Otherwise, if E_i has a type U and the corresponding parameter is an input parameter (§15.6.2.3.2) and E_i is an input argument, then an *exact inference* (§12.6.3.9) is made *from* U *to* T_i.
- Otherwise, if E_i has a type U and the corresponding parameter is an input parameter (§15.6.2.3.2) then a *lower bound inference* (§12.6.3.10) is made *from* U *to* T_i.
- Otherwise, no inference is made for this argument.

12.6.3.3 The second phase

The second phase proceeds as follows:

- All *unfixed* type variables X_i which do not *depend on* (§12.6.3.6) any X_e are fixed (§12.6.3.12).
- If no such type variables exist, all *unfixed* type variables X_i are *fixed* for which all of the following hold:
  - There is at least one type variable X_e that *depends on* X_i
  - X_i has a non-empty set of bounds

- If no such type variables exist and there are still *unfixed* type variables, type inference fails.
- Otherwise, if no further *unfixed* type variables exist, type inference succeeds.
- Otherwise, for all arguments E_i with corresponding parameter type T_i where the *output types* (§12.6.3.5) contain *unfixed* type variables X_e but the *input types* (§12.6.3.4) do not, an *output type inference* (§12.6.3.7) is made *from* E_i *to* T_i. Then the second phase is repeated.

12.6.3.4 Input types

If E is a method group or implicitly typed anonymous function and T is a delegate type or expression tree type then all the parameter types of T are *input types of* E *with type* T.

12.6.3.5 Output types

If E is a method group or an anonymous function and T is a delegate type or expression tree type then the return type of T is an *output type of* E *with type* T.

12.6.3.6 Dependence

An *unfixed* type variable X_i *depends directly on* an *unfixed* type variable X_e if for some argument E_v with type T_v X_e occurs in an *input type* of E_v with type T_v and X_i occurs in an *output type* of E_v with type T_v.

X_e *depends on* X_i if X_e *depends directly on* X_i or if X_i *depends directly on* X_v and X_v *depends on* X_e. Thus "*depends on*" is the transitive but not reflexive closure of "*depends directly on*".

12.6.3.7 Output type inferences

An *output type inference* is made *from* an expression E *to* a type T in the following way:

- If E is an anonymous function with inferred return type U (§12.6.3.13) and T is a delegate type or expression tree type with return type T_x, then a *lower-bound inference* (§12.6.3.10) is made *from* U *to* T_x.
- Otherwise, if E is a method group and T is a delegate type or expression tree type with parameter types $T_1 \ldots T_v$ and return type T_x, and overload resolution of E

with the types $T_1 \ldots T_v$ yields a single method with return type U, then a *lower-bound inference* is made *from* U *to* T_x.
- Otherwise, if E is an expression with type U, then a *lower-bound inference* is made *from* U *to* T.
- Otherwise, no inferences are made.

12.6.3.8 Explicit parameter type inferences

An *explicit parameter type inference* is made *from* an expression E *to* a type T in the following way:

- If E is an explicitly typed anonymous function with parameter types $U_1 \ldots U_v$ and T is a delegate type or expression tree type with parameter types $V_1 \ldots V_v$ then for each U_i an *exact inference* (§12.6.3.9) is made *from* U_i *to* the corresponding V_i.

12.6.3.9 Exact inferences

An *exact inference from* a type U *to* a type V is made as follows:

- If V is one of the *unfixed* X_i then U is added to the set of exact bounds for X_i.
- Otherwise, sets $V_1 \ldots V_e$ and $U_1 \ldots U_e$ are determined by checking if any of the following cases apply:
  - V is an array type $V_1[\ldots]$ and U is an array type $U_1[\ldots]$ of the same rank
  - V is the type $V_1?$ and U is the type U_1
  - V is a constructed type $C<V_1 \ldots V_e>$ and U is a constructed type $C<U_1 \ldots U_e>$
    If any of these cases apply then an *exact inference* is made from each U_i to the corresponding V_i.
- Otherwise, no inferences are made.

12.6.3.10 Lower-bound inferences

A *lower-bound inference from* a type U *to* a type V is made as follows:

- If V is one of the *unfixed* X_i then U is added to the set of lower bounds for X_i.
- Otherwise, if V is the type $V_1?$ and U is the type $U_1?$ then a lower bound inference is made from U_1 to V_1.
- Otherwise, sets $U_1 \ldots U_e$ and $V_1 \ldots V_e$ are determined by checking if any of the following cases apply:
  - V is an array type $V_1[\ldots]$ and U is an array type $U_1[\ldots]$ of the same rank

    - `V` is one of `IEnumerable<V₁>`, `ICollection<V₁>`, `IReadOnlyList<V₁>>`, `IReadOnlyCollection<V₁>` or `IList<V₁>` and `U` is a single-dimensional array type `U₁[]`
    - `V` is a constructed `class`, `struct`, `interface` or `delegate` type `C<V₁...Vₑ>` and there is a unique type `C<U₁...Uₑ>` such that `U` (or, if `U` is a type `parameter`, its effective base class or any member of its effective interface set) is identical to, `inherits` from (directly or indirectly), or implements (directly or indirectly) `C<U₁...Uₑ>`.
    - (The "uniqueness" restriction means that in the case interface `C<T>{} class U: C<X>, C<Y>{}`, then no inference is made when inferring from `U` to `C<T>` because `U₁` could be `X` or `Y`.)
      If any of these cases apply then an inference is made from each `Uᵢ` to the corresponding `Vᵢ` as follows:
    - If `Uᵢ` is not known to be a reference type then an *exact inference* is made
    - Otherwise, if `U` is an array type then a *lower-bound inference* is made
    - Otherwise, if `V` is `C<V₁...Vₑ>` then inference depends on the `i-th` type parameter of `C`:
        - If it is covariant then a *lower-bound inference* is made.
        - If it is contravariant then an *upper-bound inference* is made.
        - If it is invariant then an *exact inference* is made.
- Otherwise, no inferences are made.

12.6.3.11 Upper-bound inferences

An *upper-bound inference from* a type `U` *to* a type `V` is made as follows:

- If `V` is one of the *unfixed* `Xᵢ` then `U` is added to the set of upper bounds for `Xᵢ`.
- Otherwise, sets `V₁...Vₑ` and `U₁...Uₑ` are determined by checking if any of the following cases apply:
    - `U` is an array type `U₁[...]` and `V` is an array type `V₁[...]` of the same rank
    - `U` is one of `IEnumerable<Uₑ>`, `ICollection<Uₑ>`, `IReadOnlyList<Uₑ>`, `IReadOnlyCollection<Uₑ>` or `IList<Uₑ>` and `V` is a single-dimensional array type `Vₑ[]`
    - `U` is the type `U1?` and `V` is the type `V1?`
    - `U` is constructed class, struct, interface or delegate type `C<U₁...Uₑ>` and `V` is a `class, struct, interface` or `delegate` type which is `identical` to, `inherits` from (directly or indirectly), or implements (directly or indirectly) a unique type `C<V₁...Vₑ>`

    - (The "uniqueness" restriction means that given an interface `C<T>{} class V<Z>: C<X<Z>>, C<Y<Z>>{}`, then no inference is made when inferring from `C<U₁>` to `V<Q>`. Inferences are not made from `U₁` to either `X<Q>` or `Y<Q>`.)
      If any of these cases apply then an inference is made from each `Uᵢ` to the corresponding `Vᵢ` as follows:
    - If `Uᵢ` is not known to be a reference type then an *exact inference* is made
    - Otherwise, if `V` is an array type then an *upper-bound inference* is made
    - Otherwise, if `U` is `C<U₁...Uₑ>` then inference depends on the `i-th` type parameter of `C`:
      - If it is covariant then an *upper-bound inference* is made.
      - If it is contravariant then a *lower-bound inference* is made.
      - If it is invariant then an *exact inference* is made.
- Otherwise, no inferences are made.

12.6.3.12 Fixing

An *unfixed* type variable `Xᵢ` with a set of bounds is *fixed* as follows:

- The set of *candidate types* `Uₑ` starts out as the set of all types in the set of bounds for `Xᵢ`.
- Each bound for `Xᵢ` is examined in turn: For each exact bound U of `Xᵢ` all types `Uₑ` that are not identical to `U` are removed from the candidate set. For each lower bound `U` of `Xᵢ` all types `Uₑ` to which there is *not* an implicit conversion from `U` are removed from the candidate set. For each upper-bound U of `Xᵢ` all types `Uₑ` from which there is *not* an implicit conversion to `U` are removed from the candidate set.
- If among the remaining candidate types `Uₑ` there is a unique type `V` to which there is an implicit conversion from all the other candidate types, then `Xᵢ` is fixed to `V`.
- Otherwise, type inference fails.

12.6.3.13 Inferred return type

The inferred return type of an anonymous function `F` is used during type inference and overload resolution. The inferred return type can only be determined for an anonymous function where all parameter types are known, either because they are explicitly given, provided through an anonymous function conversion or inferred during type inference on an enclosing generic method invocation.

The ***inferred effective return type*** is determined as follows:

- If the body of `F` is an *expression* that has a type, then the inferred effective return type of `F` is the type of that expression.
- If the body of `F` is a *block* and the set of expressions in the block's `return` statements has a best common type `T` (§12.6.3.15), then the inferred effective return type of `F` is `T`.
- Otherwise, an effective return type cannot be inferred for `F`.

The ***inferred return type*** is determined as follows:

- If `F` is async and the body of `F` is either an expression classified as nothing (§12.2), or a block where no `return` statements have expressions, the inferred return type is `«TaskType»` (§15.15.1).
- If `F` is async and has an inferred effective return type `T`, the inferred return type is `«TaskType»<T>»` (§15.15.1).
- If `F` is non-async and has an inferred effective return type `T`, the inferred return type is `T`.
- Otherwise, a return type cannot be inferred for `F`.

> *Example*: As an example of type inference involving anonymous functions, consider the `Select` extension method declared in the `System.Linq.Enumerable` class:
>
> C#
>
> ```
> namespace System.Linq
> {
>     public static class Enumerable
>     {
>         public static IEnumerable<TResult> Select<TSource,TResult>(
>             this IEnumerable<TSource> source,
>             Func<TSource,TResult> selector)
>         {
>             foreach (TSource element in source)
>             {
>                 yield return selector(element);
>             }
>         }
>    }
> }
> ```
>
> Assuming the `System.Linq` namespace was imported with a `using namespace` directive, and given a class `Customer` with a `Name` property of type `string`, the `Select` method can be used to select the names of a list of customers:
>
> C#

```
List<Customer> customers = GetCustomerList();
IEnumerable<string> names = customers.Select(c => c.Name);
```

The extension method invocation (§12.8.9.3) of `Select` is processed by rewriting the invocation to a static method invocation:

C#

```
IEnumerable<string> names = Enumerable.Select(customers, c => c.Name);
```

Since type arguments were not explicitly specified, type inference is used to infer the type arguments. First, the customers argument is related to the source parameter, inferring `TSource` to be `Customer`. Then, using the anonymous function type inference process described above, `c` is given type `Customer`, and the expression `c.Name` is related to the return type of the selector parameter, inferring `TResult` to be `string`. Thus, the invocation is equivalent to

C#

```
Sequence.Select<Customer,string>(customers, (Customer c) => c.Name)
```

and the result is of type `IEnumerable<string>`.

The following example demonstrates how anonymous function type inference allows type information to "flow" between arguments in a generic method invocation. Given the following method and invocation:

C#

```
class A
{
    static Z F<X,Y,Z>(X value, Func<X,Y> f1, Func<Y,Z> f2)
    {
        return f2(f1(value));
    }

    static void M()
    {
        double hours = F("1:15:30", s => TimeSpan.Parse(s), t =>
t.TotalHours);
    }
}
```

> type inference for the invocation proceeds as follows: First, the argument "1:15:30" is related to the value parameter, inferring `X` to be string. Then, the parameter of the first anonymous function, `s`, is given the inferred type `string`, and the expression `TimeSpan.Parse(s)` is related to the return type of `f1`, inferring `Y` to be `System.TimeSpan`. Finally, the parameter of the second anonymous function, `t`, is given the inferred type `System.TimeSpan`, and the expression `t.TotalHours` is related to the return type of `f2`, inferring `Z` to be `double`. Thus, the result of the invocation is of type `double`.
>
> *end example*

12.6.3.14 Type inference for conversion of method groups

Similar to calls of generic methods, type inference shall also be applied when a method group `M` containing a generic method is converted to a given delegate type `D` (§10.8). Given a method

$T_e\ M<X_1...X_v>(T_1\ x_1\ ...\ T_e\ x_e)$

and the method group `M` being assigned to the delegate type `D` the task of type inference is to find type arguments $S_1...S_v$ so that the expression:

$M<S_1...S_v>$

becomes compatible (§20.2) with `D`.

Unlike the type inference algorithm for generic method calls, in this case, there are only argument *types*, no argument *expressions*. In particular, there are no anonymous functions and hence no need for multiple phases of inference.

Instead, all X_i are considered *unfixed*, and a *lower-bound inference* is made *from* each argument type U_e of `D` *to* the corresponding parameter type T_e of `M`. If for any of the X_i no bounds were found, type inference fails. Otherwise, all X_i are *fixed* to corresponding S_i, which are the result of type inference.

12.6.3.15 Finding the best common type of a set of expressions

In some cases, a common type needs to be inferred for a set of expressions. In particular, the element types of implicitly typed arrays and the return types of anonymous functions with *block* bodies are found in this way.

The best common type for a set of expressions $E_1...E_v$ is determined as follows:

- A new *unfixed* type variable `X` is introduced.
- For each expression `Ei` an *output type inference* (§12.6.3.7) is performed from it to `X`.
- `X` is *fixed* (§12.6.3.12), if possible, and the resulting type is the best common type.
- Otherwise inference fails.

> *Note*: Intuitively this inference is equivalent to calling a method `void M<X>(X x₁ ... X xᵥ)` with the `Eᵢ` as arguments and inferring `X`. *end note*

12.6.4 Overload resolution

12.6.4.1 General

Overload resolution is a binding-time mechanism for selecting the best function member to invoke given an argument list and a set of candidate function members. Overload resolution selects the function member to invoke in the following distinct contexts within C#:

- Invocation of a method named in an *invocation_expression* (§12.8.9).
- Invocation of an instance constructor named in an *object_creation_expression* (§12.8.16.2).
- Invocation of an indexer accessor through an *element_access* (§12.8.11).
- Invocation of a predefined or user-defined operator referenced in an expression (§12.4.4 and §12.4.5).

Each of these contexts defines the set of candidate function members and the list of arguments in its own unique way. For instance, the set of candidates for a method invocation does not include methods marked override (§12.5), and methods in a base class are not candidates if any method in a derived class is applicable (§12.8.9.2).

Once the candidate function members and the argument list have been identified, the selection of the best function member is the same in all cases:

- First, the set of candidate function members is reduced to those function members that are applicable with respect to the given argument list (§12.6.4.2). If this reduced set is empty, a compile-time error occurs.
- Then, the best function member from the set of applicable candidate function members is located. If the set contains only one function member, then that function member is the best function member. Otherwise, the best function member is the one function member that is better than all other function members with respect to the given argument list, provided that each function member is

compared to all other function members using the rules in §12.6.4.3. If there is not exactly one function member that is better than all other function members, then the function member invocation is ambiguous and a binding-time error occurs.

The following subclauses define the exact meanings of the terms *applicable function member* and *better function member*.

12.6.4.2 Applicable function member

A function member is said to be an ***applicable function member*** with respect to an argument list `A` when all of the following are true:

- Each argument in `A` corresponds to a parameter in the function member declaration as described in §12.6.2.2, at most one argument corresponds to each parameter, and any parameter to which no argument corresponds is an optional parameter.
- For each argument in `A`, the parameter-passing mode of the argument is identical to the parameter-passing mode of the corresponding parameter, and
  - for a value parameter or a parameter array, an implicit conversion (§10.2) exists from the argument expression to the type of the corresponding parameter, or
  - for a `ref` or `out` parameter, there is an identity conversion between the type of the argument expression (if any) and the type of the corresponding parameter, or
  - for an `in` parameter when the corresponding argument has the `in` modifier, there is an identity conversion between the type of the argument expression (if any) and the type of the corresponding parameter, or
  - for an `in` parameter when the corresponding argument omits the `in` modifier, an implicit conversion (§10.2) exists from the argument expression to the type of the corresponding parameter.

For a function member that includes a parameter array, if the function member is applicable by the above rules, it is said to be applicable in its ***normal form***. If a function member that includes a parameter array is not applicable in its normal form, the function member might instead be applicable in its ***expanded form***:

- The expanded form is constructed by replacing the parameter array in the function member declaration with zero or more value parameters of the element type of the parameter array such that the number of arguments in the argument list `A` matches the total number of parameters. If `A` has fewer arguments than the number of fixed parameters in the function member declaration, the expanded form of the function member cannot be constructed and is thus not applicable.

- Otherwise, the expanded form is applicable if for each argument in `A`, one of the following is true:
  - the parameter-passing mode of the argument is identical to the parameter-passing mode of the corresponding parameter, and
    - for a fixed value parameter or a value parameter created by the expansion, an implicit conversion (§10.2) exists from the argument expression to the type of the corresponding parameter, or
    - for an `in`, `out`, or `ref` parameter, the type of the argument expression is identical to the type of the corresponding parameter.
  - the parameter-passing mode of the argument is value, and the parameter-passing mode of the corresponding parameter is input, and an implicit conversion (§10.2) exists from the argument expression to the type of the corresponding parameter

When the implicit conversion from the argument type to the parameter type of an `in` parameter is a dynamic implicit conversion (§10.2.10), the results are undefined.

> *Example*: Given the following declarations and method calls:
>
> C#
>
> ```
> public static void M1(int p1) { ... }
> public static void M1(in int p1) { ... }
> public static void M2(in int p1) { ... }
> public static void Test()
> {
>     int i = 10; uint ui = 34U;
>
>     M1(in i);   // M1(in int) is applicable
>     M1(in ui);  // no exact type match, so M1(in int) is not applicable
>     M1(i);      // M1(int) and M1(in int) are applicable
>     M1(i + 5);  // M1(int) and M1(in int) are applicable
>     M1(100u);   // no implicit conversion exists, so M1(int) is not
> applicable
>
>     M2(in i);   // M2(in int) is applicable
>     M2(i);      // M2(in int) is applicable
>     M2(i + 5);  // M2(in int) is applicable
> }
> ```
>
> *end example*

- A static method is only applicable if the method group results from a *simple_name* or a *member_access* through a type.

- An instance method is only applicable if the method group results from a *simple_name*, a *member_access* through a variable or value, or a *base_access*.
  - If the method group results from a *simple_name*, an instance method is only applicable if `this` access is permitted §12.8.13.
- When the method group results from a *member_access* which could be via either an instance or a type as described in §12.8.7.2, both instance and static methods are applicable.
- A generic method whose type arguments (explicitly specified or inferred) do not all satisfy their constraints is not applicable.
- In the context of a method group conversion, there shall exist an identity conversion (§10.2.2) or an implicit reference conversion (§10.2.8) from the method return type to the delegate's return type. Otherwise, the candidate method is not applicable.

12.6.4.3 Better function member

For the purposes of determining the better function member, a stripped-down argument list `A` is constructed containing just the argument expressions themselves in the order they appear in the original argument list, and leaving out any `out` or `ref` arguments.

Parameter lists for each of the candidate function members are constructed in the following way:

- The expanded form is used if the function member was applicable only in the expanded form.
- Optional parameters with no corresponding arguments are removed from the parameter list
- `ref` and `out` parameters are removed from the parameter list
- The parameters are reordered so that they occur at the same position as the corresponding argument in the argument list.

Given an argument list `A` with a set of argument expressions $\{E_1, E_2, ..., E_v\}$ and two applicable function members M_v and M_x with parameter types $\{P_1, P_2, ..., P_v\}$ and $\{Q_1, Q_2, ..., Q_v\}$, M_v is defined to be a ***better function member*** than M_x if

- for each argument, the implicit conversion from E_v to Q_v is not better than the implicit conversion from E_v to P_v, and
- for at least one argument, the conversion from E_v to P_v is better than the conversion from E_v to Q_v.

In case the parameter type sequences $\{P_1, P_2, ..., P_v\}$ and $\{Q_1, Q_2, ..., Q_v\}$ are equivalent (i.e., each P_i has an identity conversion to the corresponding Q_i), the following tie-breaking rules are applied, in order, to determine the better function member.

- If M_i is a non-generic method and M_e is a generic method, then M_i is better than M_e.
- Otherwise, if M_i is applicable in its normal form and M_e has a params array and is applicable only in its expanded form, then M_i is better than M_e.
- Otherwise, if both methods have params arrays and are applicable only in their expanded forms, and if the params array of M_i has fewer elements than the params array of M_e, then M_i is better than M_e.
- Otherwise, if M_v has more specific parameter types than M_x, then M_v is better than M_x. Let $\{R1, R2, ..., Rn\}$ and $\{S1, S2, ..., Sn\}$ represent the uninstantiated and unexpanded parameter types of M_v and M_x. M_v's parameter types are more specific than M_xs if, for each parameter, Rx is not less specific than Sx, and, for at least one parameter, Rx is more specific than Sx:
  - A type parameter is less specific than a non-type parameter.
  - Recursively, a constructed type is more specific than another constructed type (with the same number of type arguments) if at least one type argument is more specific and no type argument is less specific than the corresponding type argument in the other.
  - An array type is more specific than another array type (with the same number of dimensions) if the element type of the first is more specific than the element type of the second.
- Otherwise if one member is a non-lifted operator and the other is a lifted operator, the non-lifted one is better.
- If neither function member was found to be better, and all parameters of M_v have a corresponding argument whereas default arguments need to be substituted for at least one optional parameter in M_x, then M_v is better than M_x.
- If for at least one parameter M_v uses the ***better parameter-passing choice*** (§12.6.4.4) than the corresponding parameter in M_x and none of the parameters in M_x use the better parameter-passing choice than M_v, M_v is better than M_x.
- Otherwise, no function member is better.

12.6.4.4 Better parameter-passing mode

It is permitted to have corresponding parameters in two overloaded methods differ only by parameter-passing mode provided one of the two parameters has value-passing

mode, as follows:

```
C#
public static void M1(int p1) { ... }
public static void M1(in int p1) { ... }
```

Given `int i = 10;`, according to §12.6.4.2, the calls `M1(i)` and `M1(i + 5)` result in both overloads being applicable. In such cases, the method with the parameter-passing mode of value is the ***better parameter-passing mode choice***.

> *Note*: No such choice need exist for arguments of input, output, or reference passing modes, as those arguments only match the exact same parameter passing modes. *end note*

12.6.4.5 Better conversion from expression

Given an implicit conversion C_1 that converts from an expression E to a type T_1, and an implicit conversion C_2 that converts from an expression E to a type T_2, C_1 is a ***better conversion*** than C_2 if one of the following holds:

- E exactly matches T_1 and E does not exactly match T_2 (§12.6.4.6)
- E exactly matches both or neither of T_1 and T_2, and T_1 is a better conversion target than T_2 (§12.6.4.7)
- E is a method group (§12.2), T_1 is compatible (§20.4) with the single best method from the method group for conversion C_1, and T_2 is not compatible with the single best method from the method group for conversion C_2

12.6.4.6 Exactly matching expression

Given an expression E and a type T, E ***exactly matches*** T if one of the following holds:

- E has a type S, and an identity conversion exists from S to T
- E is an anonymous function, T is either a delegate type D or an expression tree type `Expression<D>` and one of the following holds:
  - An inferred return type X exists for E in the context of the parameter list of D (§12.6.3.12), and an identity conversion exists from X to the return type of D
  - E is an `async` lambda with no return value, and D has a return type which is a non-generic `«TaskType»`

- Either `E` is non-async and `D` has a return type `Y` or `E` is async and `D` has a return type `«TaskType»<Y>` (§15.15.1), and one of the following holds:
  - The body of `E` is an expression that exactly matches `Y`
  - The body of `E` is a block where every return statement returns an expression that exactly matches `Y`

12.6.4.7 Better conversion target

Given two types T_1 and T_2, T_1 is a ***better conversion target*** than T_2 if one of the following holds:

- An implicit conversion from T_1 to T_2 exists and no implicit conversion from T_2 to T_1 exists
- T_1 is `«TaskType»<`S_1`>` (§15.15.1), T_2 is `«TaskType»<`S_2`>`, and S_1 is a better conversion target than S_2
- T_1 is `«TaskType»<`S_1`>` (§15.15.1), T_2 is `«TaskType»<`S_2`>`, and T_1 is more specialized than T_2
- T_1 is S_1 or S_1`?` where S_1 is a signed integral type, and T_2 is S_2 or S_2`?` where S_2 is an unsigned integral type. Specifically:
  - S_1 is `sbyte` and S_2 is `byte`, `ushort`, `uint`, or `ulong`
  - S_1 is `short` and S_2 is `ushort`, `uint`, or `ulong`
  - S_1 is `int` and S_2 is `uint`, or `ulong`
  - S_1 is `long` and S_2 is `ulong`

12.6.4.8 Overloading in generic classes

> *Note*: While signatures as declared shall be unique (§8.6), it is possible that substitution of type arguments results in identical signatures. In such a situation, overload resolution will pick the most specific (§12.6.4.3) of the original signatures (before substitution of type arguments), if it exists, and otherwise report an error. *end note*

> *Example*: The following examples show overloads that are valid and invalid according to this rule:

```
C#

public interface I1<T> { ... }
public interface I2<T> { ... }

public abstract class G1<U>
```

```
{
    public abstract int F1(U u);           // Overload resolution for
G<int>.F1
    public abstract int F1(int i);         // will pick non-generic

    public abstract void F2(I1<U> a);      // Valid overload
    public abstract void F2(I2<U> a);
}

abstract class G2<U,V>
{
    public abstract void F3(U u, V v);     // Valid, but overload
resolution for
    public abstract void F3(V v, U u);     // G2<int,int>.F3 will fail

    public abstract void F4(U u, I1<V> v); // Valid, but overload
resolution for
    public abstract void F4(I1<V> v, U u); // G2<I1<int>,int>.F4 will
fail

    public abstract void F5(U u1, I1<V> v2);   // Valid overload
    public abstract void F5(V v1, U u2);

    public abstract void F6(ref U u);      // Valid overload
    public abstract void F6(out V v);
}
```

end example

12.6.5 Compile-time checking of dynamic member invocation

Even though overload resolution of a dynamically bound operation takes place at run-time, it is sometimes possible at compile-time to know the list of function members from which an overload will be chosen:

- For a delegate invocation (§12.8.9.4), the list is a single function member with the same parameter list as the *delegate_type* of the invocation
- For a method invocation (§12.8.9.2) on a type, or on a value whose static type is not dynamic, the set of accessible methods in the method group is known at compile-time.
- For an object creation expression (§12.8.16.2) the set of accessible constructors in the type is known at compile-time.
- For an indexer access (§12.8.11.3) the set of accessible indexers in the receiver is known at compile-time.

In these cases a limited compile-time check is performed on each member in the known set of function members, to see if it can be known for certain never to be invoked at run-time. For each function member `F` a modified parameter and argument list are constructed:

- First, if `F` is a generic method and type arguments were provided, then those are substituted for the type parameters in the parameter list. However, if type arguments were not provided, no such substitution happens.
- Then, any parameter whose type is open (i.e., contains a type parameter; see §8.4.3) is elided, along with its corresponding parameter(s).

For `F` to pass the check, all of the following shall hold:

- The modified parameter list for `F` is applicable to the modified argument list in terms of §12.6.4.2.
- All constructed types in the modified parameter list satisfy their constraints (§8.4.5).
- If the type parameters of `F` were substituted in the step above, their constraints are satisfied.
- If `F` is a static method, the method group shall not have resulted from a *member_access* whose receiver is known at compile-time to be a variable or value.
- If `F` is an instance method, the method group shall not have resulted from a *member_access* whose receiver is known at compile-time to be a type.

If no candidate passes this test, a compile-time error occurs.

12.6.6 Function member invocation

12.6.6.1 General

This subclause describes the process that takes place at run-time to invoke a particular function member. It is assumed that a binding-time process has already determined the particular member to invoke, possibly by applying overload resolution to a set of candidate function members.

For purposes of describing the invocation process, function members are divided into two categories:

- Static function members. These are static methods, static property accessors, and user-defined operators. Static function members are always non-virtual.
- Instance function members. These are instance methods, instance constructors, instance property accessors, and indexer accessors. Instance function members are either non-virtual or virtual, and are always invoked on a particular instance. The

instance is computed by an instance expression, and it becomes accessible within the function member as `this` (§12.8.13). For an instance constructor, the instance expression is taken to be the newly allocated object.

The run-time processing of a function member invocation consists of the following steps, where `M` is the function member and, if `M` is an instance member, `E` is the instance expression:

- If `M` is a static function member:
  - The argument list is evaluated as described in §12.6.2.
  - `M` is invoked.

- Otherwise, if the type of `E` is a value-type `V`, and `M` is declared or overridden in `V`:
  - `E` is evaluated. If this evaluation causes an exception, then no further steps are executed. For an instance constructor, this evaluation consists of allocating storage (typically from an execution stack) for the new object. In this case `E` is classified as a variable.
  - If `E` is not classified as a variable, or if `V` is not a readonly struct type (§16.2.2), and `E` is one of:
    - an input parameter (§15.6.2.3.2), or
    - a `readonly` field (§15.5.3), or
    - a `readonly` reference variable or return (§9.7),

  then a temporary local variable of `E`'s type is created and the value of `E` is assigned to that variable. `E` is then reclassified as a reference to that temporary local variable. The temporary variable is accessible as `this` within `M`, but not in any other way. Thus, only when `E` can be written is it possible for the caller to observe the changes that `M` makes to `this`.
  - The argument list is evaluated as described in §12.6.2.
  - `M` is invoked. The variable referenced by `E` becomes the variable referenced by `this`.

- Otherwise:
  - `E` is evaluated. If this evaluation causes an exception, then no further steps are executed.
  - The argument list is evaluated as described in §12.6.2.
  - If the type of `E` is a *value_type*, a boxing conversion (§10.2.9) is performed to convert `E` to a *class_type*, and `E` is considered to be of that *class_type* in the following steps. If the *value_type* is an *enum_type*, the *class_type* is `System.Enum;` otherwise, it is `System.ValueType`.

- The value of `E` is checked to be valid. If the value of `E` is null, a `System.NullReferenceException` is thrown and no further steps are executed.
- The function member implementation to invoke is determined:
  - If the binding-time type of `E` is an interface, the function member to invoke is the implementation of `M` provided by the run-time type of the instance referenced by `E`. This function member is determined by applying the interface mapping rules (§18.6.5) to determine the implementation of `M` provided by the run-time type of the instance referenced by `E`.
  - Otherwise, if `M` is a virtual function member, the function member to invoke is the implementation of `M` provided by the run-time type of the instance referenced by `E`. This function member is determined by applying the rules for determining the most derived implementation (§15.6.4) of `M` with respect to the run-time type of the instance referenced by `E`.
  - Otherwise, `M` is a non-virtual function member, and the function member to invoke is `M` itself.
- The function member implementation determined in the step above is invoked. The object referenced by `E` becomes the object referenced by this.

The result of the invocation of an instance constructor (§12.8.16.2) is the value created. The result of the invocation of any other function member is the value, if any, returned (§13.10.5) from its body.

12.6.6.2 Invocations on boxed instances

A function member implemented in a *value_type* can be invoked through a boxed instance of that *value_type* in the following situations:

- When the function member is an override of a method inherited from type *class_type* and is invoked through an instance expression of that *class_type*.

  > *Note*: The *class_type* will always be one of `System.Object`, `System.ValueType` or `System.Enum`. *end note*

- When the function member is an implementation of an interface function member and is invoked through an instance expression of an *interface_type*.
- When the function member is invoked through a delegate.

In these situations, the boxed instance is considered to contain a variable of the *value_type*, and this variable becomes the variable referenced by this within the function member invocation.

> *Note*: In particular, this means that when a function member is invoked on a boxed instance, it is possible for the function member to modify the value contained in the boxed instance. *end note*

12.7 Deconstruction

Deconstruction is a process whereby an expression gets turned into a tuple of individual expressions. Deconstruction is used when the target of a simple assignment is a tuple expression, in order to obtain values to assign to each of that tuple's elements.

An expression `E` is ***deconstructed*** to a tuple expression with `n` elements in the following way:

- If `E` is a tuple expression with `n` elements, the result of deconstruction is the expression `E` itself.
- Otherwise, if `E` has a tuple type `(T1, ..., Tn)` with `n` elements, then `E` is evaluated into a temporary variable `__v`, and the result of deconstruction is the expression `(__v.Item1, ..., __v.Itemn)`.
- Otherwise, if the expression `E.Deconstruct(out var __v1, ..., out var __vn)` resolves at compile-time to a unique instance or extension method, that expression is evaluated, and the result of deconstruction is the expression `(__v1, ..., __vn)`. Such a method is referred to as a ***deconstructor***.
- Otherwise, `E` cannot be deconstructed.

Here, `__v` and `__v1, ..., __vn` refer to otherwise invisible and inaccessible temporary variables.

> *Note*: An expression of type `dynamic` cannot be deconstructed. *end note*

12.8 Primary expressions

12.8.1 General

Primary expressions include the simplest forms of expressions.

ANTLR

```
primary_expression
    : primary_no_array_creation_expression
    | array_creation_expression
    ;
```

```
primary_no_array_creation_expression
    : literal
    | interpolated_string_expression
    | simple_name
    | parenthesized_expression
    | tuple_expression
    | member_access
    | null_conditional_member_access
    | invocation_expression
    | element_access
    | null_conditional_element_access
    | this_access
    | base_access
    | post_increment_expression
    | post_decrement_expression
    | object_creation_expression
    | delegate_creation_expression
    | anonymous_object_creation_expression
    | typeof_expression
    | sizeof_expression
    | checked_expression
    | unchecked_expression
    | default_value_expression
    | nameof_expression
    | anonymous_method_expression
    | pointer_member_access     // unsafe code support
    | pointer_element_access    // unsafe code support
    | stackalloc_expression
    ;
```

> *Note*: These grammar rules are not ANTLR-ready as they are part of a set of mutually left-recursive rules (`primary_expression`, `primary_no_array_creation_expression`, `member_access`, `invocation_expression`, `element_access`, `post_increment_expression`, `post_decrement_expression`, `pointer_member_access` and `pointer_element_access`) which ANTLR does not handle. Standard techniques can be used to transform the grammar to remove the mutual left-recursion. This has not been done as not all parsing strategies require it (e.g. an LALR parser would not) and doing so would obfuscate the structure and description. *end note*

pointer_member_access (§23.6.3) and *pointer_element_access* (§23.6.4) are only available in unsafe code (§23).

Primary expressions are divided between *array_creation_expression*s and *primary_no_array_creation_expression*s. Treating *array_creation_expression* in this way, rather than listing it along with the other simple expression forms, enables the grammar to disallow potentially confusing code such as

C#

```
object o = new int[3][1];
```

which would otherwise be interpreted as

C#

```
object o = (new int[3])[1];
```

12.8.2 Literals

A *primary_expression* that consists of a *literal* (§6.4.5) is classified as a value.

12.8.3 Interpolated string expressions

An *interpolated_string_expression* consists of `$`, `$@`, or `@$`, immediately followed by text within `"` characters. Within the quoted text there are zero or more ***interpolations*** delimited by `{` and `}` characters, each of which encloses an *expression* and optional formatting specifications.

Interpolated string expressions have two forms; regular (*interpolated_regular_string_expression*) and verbatim (*interpolated_verbatim_string_expression*); which are lexically similar to, but differ semantically from, the two forms of string literals (§6.4.5.6).

ANTLR

```
interpolated_string_expression
    : interpolated_regular_string_expression
    | interpolated_verbatim_string_expression
    ;

// interpolated regular string expressions

interpolated_regular_string_expression
    : Interpolated_Regular_String_Start Interpolated_Regular_String_Mid?
      ('{' regular_interpolation '}' Interpolated_Regular_String_Mid?)*
      Interpolated_Regular_String_End
    ;

regular_interpolation
    : expression (',' interpolation_minimum_width)?
      Regular_Interpolation_Format?
    ;
```

```
interpolation_minimum_width
    : constant_expression
    ;

Interpolated_Regular_String_Start
    : '$"'
    ;

// the following three lexical rules are context sensitive, see details
below

Interpolated_Regular_String_Mid
    : Interpolated_Regular_String_Element+
    ;

Regular_Interpolation_Format
    : ':' Interpolated_Regular_String_Element+
    ;

Interpolated_Regular_String_End
    : '"'
    ;

fragment Interpolated_Regular_String_Element
    : Interpolated_Regular_String_Character
    | Simple_Escape_Sequence
    | Hexadecimal_Escape_Sequence
    | Unicode_Escape_Sequence
    | Open_Brace_Escape_Sequence
    | Close_Brace_Escape_Sequence
    ;

fragment Interpolated_Regular_String_Character
    // Any character except " (U+0022), \\ (U+005C),
    // { (U+007B), } (U+007D), and New_Line_Character.
    : ~["\\{}\u000D\u000A\u0085\u2028\u2029]
    ;

// interpolated verbatim string expressions

interpolated_verbatim_string_expression
    : Interpolated_Verbatim_String_Start Interpolated_Verbatim_String_Mid?
      ('{' verbatim_interpolation '}' Interpolated_Verbatim_String_Mid?)*
      Interpolated_Verbatim_String_End
    ;

verbatim_interpolation
    : expression (',' interpolation_minimum_width)?
      Verbatim_Interpolation_Format?
    ;

Interpolated_Verbatim_String_Start
    : '$@"'
    | '@$"'
    ;
```

```
// the following three lexical rules are context sensitive, see details
below

Interpolated_Verbatim_String_Mid
    : Interpolated_Verbatim_String_Element+
    ;

Verbatim_Interpolation_Format
    : ':' Interpolated_Verbatim_String_Element+
    ;

Interpolated_Verbatim_String_End
    : '"'
    ;

fragment Interpolated_Verbatim_String_Element
    : Interpolated_Verbatim_String_Character
    | Quote_Escape_Sequence
    | Open_Brace_Escape_Sequence
    | Close_Brace_Escape_Sequence
    ;

fragment Interpolated_Verbatim_String_Character
    : ~["{}]    // Any character except " (U+0022), { (U+007B) and }
(U+007D)
    ;

// lexical fragments used by both regular and verbatim interpolated strings

fragment Open_Brace_Escape_Sequence
    : '{{'
    ;

fragment Close_Brace_Escape_Sequence
    : '}}'
    ;
```

Six of the lexical rules defined above are *context sensitive* as follows:

Expand table

Rule	Contextual Requirements
Interpolated_Regular_String_Mid	Only recognised after an *Interpolated_Regular_String_Start*, between any following interpolations, and before the corresponding *Interpolated_Regular_String_End*.
Regular_Interpolation_Format	Only recognised within a *regular_interpolation* and when the starting colon (:) is not nested within any kind of bracket (parentheses/braces/square).

Rule	Contextual Requirements
Interpolated_Regular_String_End	Only recognised after an *Interpolated_Regular_String_Start* and only if any intervening tokens are either *Interpolated_Regular_String_Mid*s or tokens that can be part of *regular_interpolation*s, including tokens for any *interpolated_regular_string_expression*s contained within such interpolations.
Interpolated_Verbatim_String_Mid *Verbatim_Interpolation_Format* *Interpolated_Verbatim_String_End*	Recognition of these three rules follows that of the corresponding rules above with each mentioned *regular* grammar rule replaced by the corresponding *verbatim* one.

> *Note*: The above rules are context sensitive as their definitions overlap with those of other tokens in the language. *end note*

> *Note*: The above grammar is not ANTLR-ready due to the context sensitive lexical rules. As with other lexer generators ANTLR supports context sensitive lexical rules, for example using its *lexical modes*, but this is an implementation detail and therefore not part of this specification. *end note*

An *interpolated_string_expression* is classified as a value. If it is immediately converted to `System.IFormattable` or `System.FormattableString` with an implicit interpolated string conversion (§10.2.5), the interpolated string expression has that type. Otherwise, it has the type `string`.

> *Note*: The differences between the possible types an *interpolated_string_expression* may be determined from the documentation for `System.String` (§C.2) and `System.FormattableString` (§C.3). *end note*

The meaning of an interpolation, both *regular_interpolation* and *verbatim_interpolation*, is to format the value of the *expression* as a `string` either according to the format specified by the *Regular_Interpolation_Format* or *Verbatim_Interpolation_Format*, or according to a default format for the type of *expression*. The formatted string is then modified by the *interpolation_minimum_width*, if any, to produce the final `string` to be interpolated into the *interpolated_string_expression*.

> *Note*: How the default format for a type is determined is detailed in the documentation for `System.String` (§C.2) and `System.FormattableString` (§C.3). Descriptions of standard formats, which are identical for

> *Regular_Interpolation_Format* and *Verbatim_Interpolation_Format*, may be found in the documentation for `System.IFormattable` (§C.4) and in other types in the standard library (§C). *end note*

In an *interpolation_minimum_width* the *constant_expression* shall have an implicit conversion to `int`. Let the *field width* be the absolute value of this *constant_expression* and the *alignment* be the sign (positive or negative) of the value of this *constant_expression*:

- If the value of field width is less than or equal to the length of the formatted string the formatted string is not modified.
- Otherwise the formatted string is padded with white space characters so that its length is equal to field width:
  - If the alignment is positive the formatted string is right-aligned by prepending the padding,
  - Otherwise it is left-aligned by appending the padding.

The overall meaning of an *interpolated_string_expression*, including the above formatting and padding of interpolations, is defined by a conversion of the expression to a method invocation: if the type of the expression is `System.IFormattable` or `System.FormattableString` that method is `System.Runtime.CompilerServices.FormattableStringFactory.Create` (§C.3) which returns a value of type `System.FormattableString`; otherwise the type shall be `string` and the method is `string.Format` (§C.2) which returns a value of type `string`.

In both cases, the argument list of the call consists of a *format string literal* with *format specifications* for each interpolation, and an argument for each expression corresponding to the format specifications.

The format string literal is constructed as follows, where `N` is the number of interpolations in the *interpolated_string_expression*. The format string literal consists of, in order:

- The characters of the *Interpolated_Regular_String_Start* or *Interpolated_Verbatim_String_Start*
- The characters of the *Interpolated_Regular_String_Mid* or *Interpolated_Verbatim_String_Mid*, if any
- Then if `N ≥ 1` for each number `I` from `0` to `N-1`:
  - A placeholder specification:
    - A left brace (`{`) character
    - The decimal representation of `I`

    - Then, if the corresponding *regular_interpolation* or *verbatim_interpolation* has a *interpolation_minimum_width*, a comma (`,`) followed by the decimal representation of the value of the *constant_expression*
    - The characters of the *Regular_Interpolation_Format* or *Verbatim_Interpolation_Format*, if any, of the corresponding *regular_interpolation* or *verbatim_interpolation*
    - A right brace (`}`) character
  - The characters of the *Interpolated_Regular_String_Mid* or *Interpolated_Verbatim_String_Mid* immediately following the corresponding interpolation, if any
- Finally the characters of the *Interpolated_Regular_String_End* or *Interpolated_Verbatim_String_End*.

The subsequent arguments are the *expression*s from the interpolations, if any, in order.

When an *interpolated_string_expression* contains multiple interpolations, the expressions in those interpolations are evaluated in textual order from the left to right.

Example:

This example uses the following format specification features:

- the `X` format specification which formats integers as uppercase hexadecimal,
- the default format for a `string` value is the value itself,
- positive alignment values that right-justify within the specified minimum field width,
- negative alignment values that left-justify within the specified minimum field width,
- defined constants for the *interpolation_minimum_width*, and
- that `{{` and `}}` are formatted as `{` and `}` respectively.

Given:

C#

```
string text = "red";
int number = 14;
const int width = -4;
```

Then:

Expand table

Interpolated String Expression	Equivalent Meaning As `string`	Value
`$"{text}"`	`string.Format("{0}", text)`	`"red"`
`$"{{text}}"`	`string.Format("{{text}})`	`"{text}"`
`$"{ text , 4 }"`	`string.Format("{0,4}", text)`	`" red"`
`$"{ text , width }"`	`string.Format("{0,-4}", text)`	`"red "`
`$"{number:X}"`	`string.Format("{0:X}", number)`	`"E"`
`$"{text + '?'} {number % 3}"`	`string.Format("{0} {1}", text + '?', number % 3)`	`"red? 2"`
`$"{text + $"[{number}]"}"`	`string.Format("{0}", text + string.Format("[{0}]", number))`	`"red[14]"`
`$"{(number==0?"Zero":"Non-zero")}"`	`string.Format("{0}", (number==0?"Zero":"Non-zero"))`	`"Non-zero"`

end example

12.8.4 Simple names

A *simple_name* consists of an identifier, optionally followed by a type argument list:

ANTLR

```
simple_name
    : identifier type_argument_list?
    ;
```

A *simple_name* is either of the form `I` or of the form `I<A₁, ..., Aₑ>`, where `I` is a single identifier and `I<A₁, ..., Aₑ>` is an optional *type_argument_list*. When no *type_argument_list* is specified, consider `e` to be zero. The *simple_name* is evaluated and classified as follows:

- If `e` is zero and the *simple_name* appears within a local variable declaration space (§7.3) that directly contains a local variable, parameter or constant with name `I`, then the *simple_name* refers to that local variable, parameter or constant and is classified as a variable or value.
- If `e` is zero and the *simple_name* appears within a generic method declaration but outside the *attributes* of its *method_declaration*, and if that declaration includes a type parameter with name `I`, then the *simple_name* refers to that type parameter.

- Otherwise, for each instance type `T` (§15.3.2), starting with the instance type of the immediately enclosing type declaration and continuing with the instance type of each enclosing class or struct declaration (if any):
  - If `e` is zero and the declaration of `T` includes a type parameter with name `I`, then the *simple_name* refers to that type parameter.
  - Otherwise, if a member lookup (§12.5) of `I` in `T` with `e` type arguments produces a match:
    - If `T` is the instance type of the immediately enclosing class or struct type and the lookup identifies one or more methods, the result is a method group with an associated instance expression of `this`. If a type argument list was specified, it is used in calling a generic method (§12.8.9.2).
    - Otherwise, if `T` is the instance type of the immediately enclosing class or struct type, if the lookup identifies an instance member, and if the reference occurs within the *block* of an instance constructor, an instance method, or an instance accessor (§12.2.1), the result is the same as a member access (§12.8.7) of the form `this.I`. This can only happen when `e` is zero.
    - Otherwise, the result is the same as a member access (§12.8.7) of the form `T.I` or $T.I<A_1, \ldots, A_e>$.
- Otherwise, for each namespace `N`, starting with the namespace in which the *simple_name* occurs, continuing with each enclosing namespace (if any), and ending with the global namespace, the following steps are evaluated until an entity is located:
  - If `e` is zero and `I` is the name of a namespace in `N`, then:
    - If the location where the *simple_name* occurs is enclosed by a namespace declaration for `N` and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I` with a namespace or type, then the *simple_name* is ambiguous and a compile-time error occurs.
    - Otherwise, the *simple_name* refers to the namespace named `I` in `N`.
  - Otherwise, if `N` contains an accessible type having name `I` and `e` type parameters, then:
    - If `e` is zero and the location where the *simple_name* occurs is enclosed by a namespace declaration for `N` and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I` with a namespace or type, then the *simple_name* is ambiguous and a compile-time error occurs.
    - Otherwise, the *namespace_or_type_name* refers to the type constructed with the given type arguments.

    - Otherwise, if the location where the *simple_name* occurs is enclosed by a namespace declaration for `N`:
        - If `e` is zero and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I` with an imported namespace or type, then the *simple_name* refers to that namespace or type.
        - Otherwise, if the namespaces imported by the *using_namespace_directive*s of the namespace declaration contain exactly one type having name `I` and `e` type parameters, then the *simple_name* refers to that type constructed with the given type arguments.
        - Otherwise, if the namespaces imported by the *using_namespace_directive*s of the namespace declaration contain more than one type having name `I` and `e` type parameters, then the *simple_name* is ambiguous and a compile-time error occurs.

    > *Note*: This entire step is exactly parallel to the corresponding step in the processing of a *namespace_or_type_name* (§7.8). *end note*

- Otherwise, if `e` is zero and `I` is the identifier `_`, the *simple_name* is a *simple discard*, which is a form of declaration expression (§12.17).
- Otherwise, the *simple_name* is undefined and a compile-time error occurs.

12.8.5 Parenthesized expressions

A *parenthesized_expression* consists of an *expression* enclosed in parentheses.

ANTLR

```
parenthesized_expression
    : '(' expression ')'
    ;
```

A *parenthesized_expression* is evaluated by evaluating the *expression* within the parentheses. If the *expression* within the parentheses denotes a namespace or type, a compile-time error occurs. Otherwise, the result of the *parenthesized_expression* is the result of the evaluation of the contained *expression*.

12.8.6 Tuple expressions

A *tuple_expression* represents a tuple, and consists of two or more comma-separated and optionally-named *expression*s enclosed in parentheses. A *deconstruction_expression* is a shorthand syntax for a tuple containing implicitly typed declaration expressions.

ANTLR

```
tuple_expression
    : '(' tuple_element (',' tuple_element)+ ')'
    | deconstruction_expression
    ;

tuple_element
    : (identifier ':')? expression
    ;

deconstruction_expression
    : 'var' deconstruction_tuple
    ;

deconstruction_tuple
    : '(' deconstruction_element (',' deconstruction_element)+ ')'
    ;

deconstruction_element
    : deconstruction_tuple
    | identifier
    ;
```

A *tuple_expression* is classified as a tuple.

A *deconstruction_expression* `var (e1, ..., en)` is shorthand for the *tuple_expression* `(var e1, ..., var en)` and follows the same behavior. This applies recursively to any nested *deconstruction_tuple*s in the *deconstruction_expression*. Each identifier nested within a *deconstruction_expression* thus introduces a declaration expression (§12.17). As a result, a *deconstruction_expression* can only occur on the left side of a simple assignment.

A tuple expression has a type if and only if each of its element expressions `Ei` has a type `Ti`. The type shall be a tuple type of the same arity as the tuple expression, where each element is given by the following:

- If the tuple element in the corresponding position has a name `Ni`, then the tuple type element shall be `Ti Ni`.
- Otherwise, if `Ei` is of the form `Ni` or `E.Ni` or `E?.Ni` then the tuple type element shall be `Ti Ni`, *unless* any of the following holds:
  - Another element of the tuple expression has the name `Ni`, or
  - Another tuple element without a name has a tuple element expression of the form `Ni` or `E.Ni` or `E?.Ni`, or
  - `Ni` is of the form `ItemX`, where `X` is a sequence of non-`0`-initiated decimal digits that could represent the position of a tuple element, and `X` does not

represent the position of the element.

- Otherwise, the tuple type element shall be `Ti`.

A tuple expression is evaluated by evaluating each of its element expressions in order from left to right.

A tuple value can be obtained from a tuple expression by converting it to a tuple type (§10.2.13), by reclassifying it as a value (§12.2.2)) or by making it the target of a deconstructing assignment (§12.21.2).

> *Example*:
>
> ```
> C#
> ```
>
> ```
> (int i, string) t1 = (i: 1, "One");
> (long l, string) t2 = (l: 2, null);
> var t3 = (i: 3, "Three");          // (int i, string)
> var t4 = (i: 4, null);             // Error: no type
> ```
>
> In this example, all four tuple expressions are valid. The first two, `t1` and `t2`, do not use the type of the tuple expression, but instead apply an implicit tuple conversion. In the case of `t2`, the implicit tuple conversion relies on the implicit conversions from `2` to `long` and from `null` to `string`. The third tuple expression has a type `(int i, string)`, and can therefore be reclassified as a value of that type. The declaration of `t4`, on the other hand, is an error: The tuple expression has no type because its second element has no type.
>
> ```
> C#
> ```
>
> ```
> if ((x, y).Equals((1, 2))) { ... };
> ```
>
> This example shows that tuples can sometimes lead to multiple layers of parentheses, especially when the tuple expression is the sole argument to a method invocation.
>
> *end example*

12.8.7 Member access

12.8.7.1 General

A *member_access* consists of a *primary_expression*, a *predefined_type*, or a *qualified_alias_member*, followed by a "`.`" token, followed by an *identifier*, optionally followed by a *type_argument_list*.

ANTLR

```
member_access
    : primary_expression '.' identifier type_argument_list?
    | predefined_type '.' identifier type_argument_list?
    | qualified_alias_member '.' identifier type_argument_list?
    ;

predefined_type
    : 'bool' | 'byte' | 'char' | 'decimal' | 'double' | 'float' | 'int'
    | 'long' | 'object' | 'sbyte' | 'short' | 'string' | 'uint' | 'ulong'
    | 'ushort'
    ;
```

The *qualified_alias_member* production is defined in §14.8.

A *member_access* is either of the form `E.I` or of the form `E.I<A₁, ..., Aₑ>`, where `E` is a *primary_expression*, *predefined_type* or *qualified_alias_member*, `I` is a single identifier, and `<A₁, ..., Aₑ>` is an optional *type_argument_list*. When no *type_argument_list* is specified, consider `e` to be zero.

A *member_access* with a *primary_expression* of type `dynamic` is dynamically bound (§12.3.3). In this case, the compiler classifies the member access as a property access of type `dynamic`. The rules below to determine the meaning of the *member_access* are then applied at run-time, using the run-time type instead of the compile-time type of the *primary_expression*. If this run-time classification leads to a method group, then the member access shall be the *primary_expression* of an *invocation_expression*.

The *member_access* is evaluated and classified as follows:

- If `e` is zero and `E` is a namespace and `E` contains a nested namespace with name `I`, then the result is that namespace.
- Otherwise, if `E` is a namespace and `E` contains an accessible type having name `I` and `K` type parameters, then the result is that type constructed with the given type arguments.
- If `E` is classified as a type, if `E` is not a type parameter, and if a member lookup (§12.5) of `I` in `E` with `K` type parameters produces a match, then `E.I` is evaluated and classified as follows:

> *Note*: When the result of such a member lookup is a method group and `K` is zero, the method group can contain methods having type parameters. This allows such methods to be considered for type argument inferencing. *end note*

  - If `I` identifies a type, then the result is that type constructed with any given type arguments.
  - If `I` identifies one or more methods, then the result is a method group with no associated instance expression.
  - If `I` identifies a static property, then the result is a property access with no associated instance expression.
  - If `I` identifies a static field:
    - If the field is readonly and the reference occurs outside the static constructor of the class or struct in which the field is declared, then the result is a value, namely the value of the static field `I` in `E`.
    - Otherwise, the result is a variable, namely the static field `I` in `E`.
  - If `I` identifies a static event:
    - If the reference occurs within the class or struct in which the event is declared, and the event was declared without *event_accessor_declarations* (§15.8.1), then `E.I` is processed exactly as if `I` were a static field.
    - Otherwise, the result is an event access with no associated instance expression.
  - If `I` identifies a constant, then the result is a value, namely the value of that constant.
  - If `I` identifies an enumeration member, then the result is a value, namely the value of that enumeration member.
  - Otherwise, `E.I` is an invalid member reference, and a compile-time error occurs.
- If `E` is a property access, indexer access, variable, or value, the type of which is `T`, and a member lookup (§12.5) of `I` in `T` with `K` type arguments produces a match, then `E.I` is evaluated and classified as follows:
  - First, if `E` is a property or indexer access, then the value of the property or indexer access is obtained (§12.2.2) and E is reclassified as a value.
  - If `I` identifies one or more methods, then the result is a method group with an associated instance expression of `E`.
  - If `I` identifies an instance property, then the result is a property access with an associated instance expression of `E` and an associated type that is the type of the property. If `T` is a class type, the associated type is picked from the first declaration or override of the property found when starting with `T`, and searching through its base classes.

- If `T` is a *class_type* and `I` identifies an instance field of that *class_type*:
  - If the value of `E` is `null`, then a `System.NullReferenceException` is thrown.
  - Otherwise, if the field is readonly and the reference occurs outside an instance constructor of the class in which the field is declared, then the result is a value, namely the value of the field `I` in the object referenced by `E`.
  - Otherwise, the result is a variable, namely the field `I` in the object referenced by `E`.
- If `T` is a *struct_type* and `I` identifies an instance field of that *struct_type*:
  - If `E` is a value, or if the field is readonly and the reference occurs outside an instance constructor of the struct in which the field is declared, then the result is a value, namely the value of the field `I` in the struct instance given by `E`.
  - Otherwise, the result is a variable, namely the field `I` in the struct instance given by `E`.
- If `I` identifies an instance event:
  - If the reference occurs within the class or struct in which the event is declared, and the event was declared without *event_accessor_declarations* (§15.8.1), and the reference does not occur as the left-hand side of `a +=` or `-=` operator, then `E.I` is processed exactly as if `I` was an instance field.
  - Otherwise, the result is an event access with an associated instance expression of `E`.

- Otherwise, an attempt is made to process `E.I` as an extension method invocation (§12.8.9.3). If this fails, `E.I` is an invalid member reference, and a binding-time error occurs.

12.8.7.2 Identical simple names and type names

In a member access of the form `E.I`, if `E` is a single identifier, and if the meaning of `E` as a *simple_name* (§12.8.4) is a constant, field, property, local variable, or parameter with the same type as the meaning of `E` as a *type_name* (§7.8.1), then both possible meanings of `E` are permitted. The member lookup of `E.I` is never ambiguous, since `I` shall necessarily be a member of the type `E` in both cases. In other words, the rule simply permits access to the static members and nested types of `E` where a compile-time error would otherwise have occurred.

> *Example*:
>
> C#

```
struct Color
{
    public static readonly Color White = new Color(...);
    public static readonly Color Black = new Color(...);
    public Color Complement() => new Color(...);
}

class A
{
    public «Color» Color;              // Field Color of type Color

    void F()
    {
        Color = «Color».Black;         // Refers to Color.Black static
member
        Color = Color.Complement();  // Invokes Complement() on Color
field
    }

    static void G()
    {
        «Color» c = «Color».White;       // Refers to Color.White static
member
    }
}
```

For expository purposes only, within the `A` class, those occurrences of the `Color` identifier that reference the `Color` type are delimited by `«...»`, and those that reference the `Color` field are not.

end example

12.8.8 Null Conditional Member Access

A *null_conditional_member_access* is a conditional version of *member_access* (§12.8.7) and it is a binding time error if the result type is `void`. For a null conditional expression where the result type may be `void` see (§12.8.10).

A *null_conditional_member_access* consists of a *primary_expression* followed by the two tokens "`?`" and "`.`", followed by an *identifier* with an optional *type_argument_list*, followed by zero or more *dependent_access*es.

ANTLR

```
null_conditional_member_access
    : primary_expression '?' '.' identifier type_argument_list?
      dependent_access*
    ;
```

```
dependent_access
    : '.' identifier type_argument_list?    // member access
    | '[' argument_list ']'                 // element access
    | '(' argument_list? ')'                // invocation
    ;

null_conditional_projection_initializer
    : primary_expression '?' '.' identifier type_argument_list?
    ;
```

A *null_conditional_member_access* expression `E` is of the form `P?.A`. The meaning of `E` is determined as follows:

- If the type of `P` is a nullable value type:

  Let `T` be the type of `P.Value.A`.

  - If `T` is a type parameter that is not known to be either a reference type or a non-nullable value type, a compile-time error occurs.
  - If `T` is a non-nullable value type, then the type of `E` is `T?`, and the meaning of `E` is the same as the meaning of:

    ```
    C#
    ((object)P == null) ? (T?)null : P.Value.A
    ```

    Except that `P` is evaluated only once.
  - Otherwise the type of `E` is `T`, and the meaning of `E` is the same as the meaning of:

    ```
    C#
    ((object)P == null) ? (T)null : P.Value.A
    ```

    Except that `P` is evaluated only once.

- Otherwise:

  Let `T` be the type of the expression `P.A`.

  - If `T` is a type parameter that is not known to be either a reference type or a non-nullable value type, a compile-time error occurs.

- If `T` is a non-nullable value type, then the type of `E` is `T?`, and the meaning of `E` is the same as the meaning of:

  ```
  C#
  ((object)P == null) ? (T?)null : P.A
  ```

  Except that `P` is evaluated only once.

- Otherwise the type of `E` is `T`, and the meaning of `E` is the same as the meaning of:

  ```
  C#
  ((object)P == null) ? (T)null : P.A
  ```

  Except that `P` is evaluated only once.

> *Note*: In an expression of the form:
>
> ```
> C#
> P?.A₀?.A₁
> ```
>
> then if `P` evaluates to `null` neither `A₀` or `A₁` are evaluated. The same is true if an expression is a sequence of *null_conditional_member_access* or *null_conditional_element_access* §12.8.12 operations.
>
> *end note*

A *null_conditional_projection_initializer* is a restriction of *null_conditional_member_access* and has the same semantics. It only occurs as a projection initializer in an anonymous object creation expression (§12.8.16.7).

12.8.9 Invocation expressions

12.8.9.1 General

An *invocation_expression* is used to invoke a method.

```
ANTLR
```

```
invocation_expression
    : primary_expression '(' argument_list? ')'
    ;
```

An *invocation_expression* is dynamically bound (§12.3.3) if at least one of the following holds:

- The *primary_expression* has compile-time type `dynamic`.
- At least one argument of the optional *argument_list* has compile-time type `dynamic`.

In this case, the compiler classifies the *invocation_expression* as a value of type `dynamic`. The rules below to determine the meaning of the *invocation_expression* are then applied at run-time, using the run-time type instead of the compile-time type of those of the *primary_expression* and arguments that have the compile-time type `dynamic`. If the *primary_expression* does not have compile-time type `dynamic`, then the method invocation undergoes a limited compile-time check as described in §12.6.5.

The *primary_expression* of an *invocation_expression* shall be a method group or a value of a *delegate_type*. If the *primary_expression* is a method group, the *invocation_expression* is a method invocation (§12.8.9.2). If the *primary_expression* is a value of a *delegate_type*, the *invocation_expression* is a delegate invocation (§12.8.9.4). If the *primary_expression* is neither a method group nor a value of a *delegate_type*, a binding-time error occurs.

The optional *argument_list* (§12.6.2) provides values or variable references for the parameters of the method.

The result of evaluating an *invocation_expression* is classified as follows:

- If the *invocation_expression* invokes a returns-no-value method (§15.6.1) or a returns-no-value delegate, the result is nothing. An expression that is classified as nothing is permitted only in the context of a *statement_expression* (§13.7) or as the body of a *lambda_expression* (§12.19). Otherwise, a binding-time error occurs.
- Otherwise, if the *invocation_expression* invokes a returns-by-ref method (§15.6.1) or a returns-by-ref delegate, the result is a variable with an associated type of the return type of the method or delegate. If the invocation is of an instance method, and the receiver is of a class type `T`, the associated type is picked from the first declaration or override of the method found when starting with `T` and searching through its base classes.
- Otherwise, the *invocation_expression* invokes a returns-by-value method (§15.6.1) or returns-by-value delegate, and the result is a value, with an associated type of

the return type of the method or delegate. If the invocation is of an instance method, and the receiver is of a class type `T`, the associated type is picked from the first declaration or override of the method found when starting with `T` and searching through its base classes.

12.8.9.2 Method invocations

For a method invocation, the *primary_expression* of the *invocation_expression* shall be a method group. The method group identifies the one method to invoke or the set of overloaded methods from which to choose a specific method to invoke. In the latter case, determination of the specific method to invoke is based on the context provided by the types of the arguments in the *argument_list*.

The binding-time processing of a method invocation of the form `M(A)`, where `M` is a method group (possibly including a *type_argument_list*), and `A` is an optional *argument_list*, consists of the following steps:

- The set of candidate methods for the method invocation is constructed. For each method `F` associated with the method group `M`:
  - If `F` is non-generic, `F` is a candidate when:
    - `M` has no type argument list, and
    - `F` is applicable with respect to `A` (§12.6.4.2).
  - If `F` is generic and `M` has no type argument list, `F` is a candidate when:
    - Type inference (§12.6.3) succeeds, inferring a list of type arguments for the call, and
    - Once the inferred type arguments are substituted for the corresponding method type parameters, all constructed types in the parameter list of `F` satisfy their constraints (§8.4.5), and the parameter list of `F` is applicable with respect to `A` (§12.6.4.2)
  - If `F` is generic and `M` includes a type argument list, `F` is a candidate when:
    - `F` has the same number of method type parameters as were supplied in the type argument list, and
    - Once the type arguments are substituted for the corresponding method type parameters, all constructed types in the parameter list of `F` satisfy their constraints (§8.4.5), and the parameter list of `F` is applicable with respect to `A` (§12.6.4.2).
- The set of candidate methods is reduced to contain only methods from the most derived types: For each method `C.F` in the set, where `C` is the type in which the method `F` is declared, all methods declared in a base type of `C` are removed from

the set. Furthermore, if `C` is a class type other than `object`, all methods declared in an interface type are removed from the set.

> *Note*: This latter rule only has an effect when the method group was the result of a member lookup on a type parameter having an effective base class other than `object` and a non-empty effective interface set. *end note*

- If the resulting set of candidate methods is empty, then further processing along the following steps are abandoned, and instead an attempt is made to process the invocation as an extension method invocation (§12.8.9.3). If this fails, then no applicable methods exist, and a binding-time error occurs.
- The best method of the set of candidate methods is identified using the overload resolution rules of §12.6.4. If a single best method cannot be identified, the method invocation is ambiguous, and a binding-time error occurs. When performing overload resolution, the parameters of a generic method are considered after substituting the type arguments (supplied or inferred) for the corresponding method type parameters.

Once a method has been selected and validated at binding-time by the above steps, the actual run-time invocation is processed according to the rules of function member invocation described in §12.6.6.

> *Note*: The intuitive effect of the resolution rules described above is as follows: To locate the particular method invoked by a method invocation, start with the type indicated by the method invocation and proceed up the inheritance chain until at least one applicable, accessible, non-override method declaration is found. Then perform type inference and overload resolution on the set of applicable, accessible, non-override methods declared in that type and invoke the method thus selected. If no method was found, try instead to process the invocation as an extension-method invocation. *end note*

12.8.9.3 Extension method invocations

In a method invocation (§12.6.6.2) of one of the forms

```
C#
```

```
«expr» . «identifier» ( )
«expr» . «identifier» ( «args» )
«expr» . «identifier» < «typeargs» > ( )
«expr» . «identifier» < «typeargs» > ( «args» )
```

if the normal processing of the invocation finds no applicable methods, an attempt is made to process the construct as an extension method invocation. If «expr» or any of the «args» has compile-time type `dynamic`, extension methods will not apply.

The objective is to find the best *type_name* `C`, so that the corresponding static method invocation can take place:

```
C#

C . «identifier» ( «expr» )
C . «identifier» ( «expr» , «args» )
C . «identifier» < «typeargs» > ( «expr» )
C . «identifier» < «typeargs» > ( «expr» , «args» )
```

An extension method $C_i.M_e$ is ***eligible*** if:

- C_i is a non-generic, non-nested class
- The name of M_e is *identifier*
- M_e is accessible and applicable when applied to the arguments as a static method as shown above
- An implicit identity, reference or boxing conversion exists from *expr* to the type of the first parameter of M_e.

The search for `C` proceeds as follows:

- Starting with the closest enclosing namespace declaration, continuing with each enclosing namespace declaration, and ending with the containing compilation unit, successive attempts are made to find a candidate set of extension methods:
  - If the given namespace or compilation unit directly contains non-generic type declarations C_i with eligible extension methods M_e, then the set of those extension methods is the candidate set.
  - If namespaces imported by using namespace directives in the given namespace or compilation unit directly contain non-generic type declarations C_i with eligible extension methods M_e, then the set of those extension methods is the candidate set.
- If no candidate set is found in any enclosing namespace declaration or compilation unit, a compile-time error occurs.
- Otherwise, overload resolution is applied to the candidate set as described in §12.6.4. If no single best method is found, a compile-time error occurs.
- `C` is the type within which the best method is declared as an extension method.

Using `C` as a target, the method call is then processed as a static method invocation (§12.6.6).

> *Note*: Unlike an instance method invocation, no exception is thrown when *expr* evaluates to a null reference. Instead, this `null` value is passed to the extension method as it would be via a regular static method invocation. It is up to the extension method implementation to decide how to respond to such a call. *end note*

The preceding rules mean that instance methods take precedence over extension methods, that extension methods available in inner namespace declarations take precedence over extension methods available in outer namespace declarations, and that extension methods declared directly in a namespace take precedence over extension methods imported into that same namespace with a using namespace directive.

> *Example*:

```
C#
```

```
public static class E
{
    public static void F(this object obj, int i) { }
    public static void F(this object obj, string s) { }
}

class A { }

class B
{
    public void F(int i) { }
}

class C
{
    public void F(object obj) { }
}

class X
{
    static void Test(A a, B b, C c)
    {
        a.F(1);             // E.F(object, int)
        a.F("hello");       // E.F(object, string)
        b.F(1);             // B.F(int)
        b.F("hello");       // E.F(object, string)
        c.F(1);             // C.F(object)
        c.F("hello");       // C.F(object)
    }
}
```

> In the example, `B`'s method takes precedence over the first extension method, and `C`'s method takes precedence over both extension methods.

C#

```
public static class C
{
    public static void F(this int i) => Console.WriteLine($"C.F({i})");
    public static void G(this int i) => Console.WriteLine($"C.G({i})");
    public static void H(this int i) => Console.WriteLine($"C.H({i})");
}

namespace N1
{
    public static class D
    {
        public static void F(this int i) =>
Console.WriteLine($"D.F({i})");
        public static void G(this int i) =>
Console.WriteLine($"D.G({i})");
    }
}

namespace N2
{
    using N1;

    public static class E
    {
        public static void F(this int i) =>
Console.WriteLine($"E.F({i})");
    }

    class Test
    {
        static void Main(string[] args)
        {
            1.F();
            2.G();
            3.H();
        }
    }
}
```

The output of this example is:

Console

```
E.F(1)
D.G(2)
C.H(3)
```

`D.G` takes precendece over `C.G`, and `E.F` takes precedence over both `D.F` and `C.F`.

> *end example*

12.8.9.4 Delegate invocations

For a delegate invocation, the *primary_expression* of the *invocation_expression* shall be a value of a *delegate_type*. Furthermore, considering the *delegate_type* to be a function member with the same parameter list as the *delegate_type*, the *delegate_type* shall be applicable (§12.6.4.2) with respect to the *argument_list* of the *invocation_expression*.

The run-time processing of a delegate invocation of the form `D(A)`, where `D` is a *primary_expression* of a *delegate_type* and `A` is an optional *argument_list*, consists of the following steps:

- `D` is evaluated. If this evaluation causes an exception, no further steps are executed.
- The argument list `A` is evaluated. If this evaluation causes an exception, no further steps are executed.
- The value of `D` is checked to be valid. If the value of `D` is `null`, a `System.NullReferenceException` is thrown and no further steps are executed.
- Otherwise, `D` is a reference to a delegate instance. Function member invocations (§12.6.6) are performed on each of the callable entities in the invocation list of the delegate. For callable entities consisting of an instance and instance method, the instance for the invocation is the instance contained in the callable entity.

See §20.6 for details of multiple invocation lists without parameters.

12.8.10 Null Conditional Invocation Expression

A *null_conditional_invocation_expression* is syntactically either a *null_conditional_member_access* (§12.8.8) or *null_conditional_element_access* (§12.8.12) where the final *dependent_access* is an invocation expression (§12.8.9).

A *null_conditional_invocation_expression* occurs within the context of a *statement_expression* (§13.7), *anonymous_function_body* (§12.19.1), or *method_body* (§15.6.1).

Unlike the syntactically equivalent *null_conditional_member_access* or *null_conditional_element_access*, a *null_conditional_invocation_expression* may be classified as nothing.

```
ANTLR
```

```
null_conditional_invocation_expression
    : null_conditional_member_access '(' argument_list? ')'
    | null_conditional_element_access '(' argument_list? ')'
    ;
```

A *null_conditional_invocation_expression* expression `E` is of the form `P?A`; where `A` is the remainder of the syntactically equivalent *null_conditional_member_access* or *null_conditional_element_access*, `A` will therefore start with `.` or `[`. Let `PA` signify the concatention of `P` and `A`.

When `E` occurs as a *statement_expression* the meaning of `E` is the same as the meaning of the *statement*:

C#

```
if ((object)P != null) PA
```

except that `P` is evaluated only once.

When `E` occurs as a *anonymous_function_body* or *method_body* the meaning of `E` depends on its classification:

- If `E` is classified as nothing then its meaning is the same as the meaning of the *block*:

  C#

  ```
  { if ((object)P != null) PA; }
  ```

  except that `P` is evaluated only once.

- Otherwise the meaning of `E` is the same as the meaning of the *block*:

  C#

  ```
  { return E; }
  ```

  and in turn the meaning of this *block* depends on whether `E` is syntactically equivalent to a *null_conditional_member_access* (§12.8.8) or *null_conditional_element_access* (§12.8.12).

12.8.11 Element access

12.8.11.1 General

An *element_access* consists of a *primary_no_array_creation_expression*, followed by a "`[`" token, followed by an *argument_list*, followed by a "`]`" token. The *argument_list* consists of one or more *argument*s, separated by commas.

ANTLR

```
element_access
    : primary_no_array_creation_expression '[' argument_list ']'
    ;
```

The *argument_list* of an *element_access* is not allowed to contain `out` or `ref` arguments.

An *element_access* is dynamically bound (§12.3.3) if at least one of the following holds:

- The *primary_no_array_creation_expression* has compile-time type `dynamic`.
- At least one expression of the *argument_list* has compile-time type `dynamic` and the *primary_no_array_creation_expression* does not have an array type.

In this case, the compiler classifies the *element_access* as a value of type `dynamic`. The rules below to determine the meaning of the *element_access* are then applied at run-time, using the run-time type instead of the compile-time type of those of the *primary_no_array_creation_expression* and *argument_list* expressions which have the compile-time type `dynamic`. If the *primary_no_array_creation_expression* does not have compile-time type `dynamic`, then the element access undergoes a limited compile-time check as described in §12.6.5.

If the *primary_no_array_creation_expression* of an *element_access* is a value of an *array_type*, the *element_access* is an array access (§12.8.11.2). Otherwise, the *primary_no_array_creation_expression* shall be a variable or value of a class, struct, or interface type that has one or more indexer members, in which case the *element_access* is an indexer access (§12.8.11.3).

12.8.11.2 Array access

For an array access, the *primary_no_array_creation_expression* of the *element_access* shall be a value of an *array_type*. Furthermore, the *argument_list* of an array access is not allowed to contain named arguments. The number of expressions in the *argument_list* shall be the same as the rank of the *array_type*, and each expression shall be of type `int`, `uint`, `long`, or `ulong`, or shall be implicitly convertible to one or more of these types.

The result of evaluating an array access is a variable of the element type of the array, namely the array element selected by the value(s) of the expression(s) in the *argument_list*.

The run-time processing of an array access of the form `P[A]`, where `P` is a *primary_no_array_creation_expression* of an *array_type* and `A` is an *argument_list*, consists of the following steps:

- `P` is evaluated. If this evaluation causes an exception, no further steps are executed.
- The index expressions of the *argument_list* are evaluated in order, from left to right. Following evaluation of each index expression, an implicit conversion (§10.2) to one of the following types is performed: `int`, `uint`, `long`, `ulong`. The first type in this list for which an implicit conversion exists is chosen. For instance, if the index expression is of type `short` then an implicit conversion to `int` is performed, since implicit conversions from `short` to `int` and from `short` to `long` are possible. If evaluation of an index expression or the subsequent implicit conversion causes an exception, then no further index expressions are evaluated and no further steps are executed.
- The value of `P` is checked to be valid. If the value of `P` is `null`, a `System.NullReferenceException` is thrown and no further steps are executed.
- The value of each expression in the *argument_list* is checked against the actual bounds of each dimension of the array instance referenced by `P`. If one or more values are out of range, a `System.IndexOutOfRangeException` is thrown and no further steps are executed.
- The location of the array element given by the index expression(s) is computed, and this location becomes the result of the array access.

12.8.11.3 Indexer access

For an indexer access, the *primary_no_array_creation_expression* of the *element_access* shall be a variable or value of a class, struct, or interface type, and this type shall implement one or more indexers that are applicable with respect to the *argument_list* of the *element_access*.

The binding-time processing of an indexer access of the form `P[A]`, where `P` is a *primary_no_array_creation_expression* of a class, struct, or interface type `T`, and `A` is an *argument_list*, consists of the following steps:

- The set of indexers provided by `T` is constructed. The set consists of all indexers declared in `T` or a base type of `T` that are not override declarations and are

accessible in the current context (§7.5).

- The set is reduced to those indexers that are applicable and not hidden by other indexers. The following rules are applied to each indexer `S.I` in the set, where `S` is the type in which the indexer `I` is declared:
  - If `I` is not applicable with respect to `A` (§12.6.4.2), then `I` is removed from the set.
  - If `I` is applicable with respect to `A` (§12.6.4.2), then all indexers declared in a base type of `S` are removed from the set.
  - If `I` is applicable with respect to `A` (§12.6.4.2) and `S` is a class type other than `object`, all indexers declared in an interface are removed from the set.
- If the resulting set of candidate indexers is empty, then no applicable indexers exist, and a binding-time error occurs.
- The best indexer of the set of candidate indexers is identified using the overload resolution rules of §12.6.4. If a single best indexer cannot be identified, the indexer access is ambiguous, and a binding-time error occurs.
- The index expressions of the *argument_list* are evaluated in order, from left to right. The result of processing the indexer access is an expression classified as an indexer access. The indexer access expression references the indexer determined in the step above, and has an associated instance expression of `P` and an associated argument list of `A`, and an associated type that is the type of the indexer. If `T` is a class type, the associated type is picked from the first declaration or override of the indexer found when starting with `T` and searching through its base classes.

Depending on the context in which it is used, an indexer access causes invocation of either the get accessor or the set accessor of the indexer. If the indexer access is the target of an assignment, the set accessor is invoked to assign a new value (§12.21.2). In all other cases, the get accessor is invoked to obtain the current value (§12.2.2).

12.8.12 Null Conditional Element Access

A *null_conditional_element_access* consists of a *primary_no_array_creation_expression* followed by the two tokens "`?`" and "`[`", followed by an *argument_list*, followed by a "`]`" token, followed by zero or more *dependent_access*es.

ANTLR

```
null_conditional_element_access
    : primary_no_array_creation_expression '?' '[' argument_list ']'
      dependent_access*
    ;
```

A *null_conditional_element_access* is a conditional version of *element_access* (§12.8.11) and it is a binding time error if the result type is `void`. For a null conditional expression where the result type may be `void` see (§12.8.10).

A *null_conditional_element_access* expression `E` is of the form `P?[A]B`; where `B` are the *dependent_access*es, if any. The meaning of `E` is determined as follows:

- If the type of `P` is a nullable value type:

  Let `T` be the type of the expression `P.Value[A]B`.

  - If `T` is a type parameter that is not known to be either a reference type or a non-nullable value type, a compile-time error occurs.
  - If `T` is a non-nullable value type, then the type of `E` is `T?`, and the meaning of `E` is the same as the meaning of:

    ```
    C#
    ((object)P == null) ? (T?)null : P.Value[A]B
    ```

    Except that `P` is evaluated only once.
  - Otherwise the type of `E` is `T`, and the meaning of `E` is the same as the meaning of:

    ```
    C#
    ((object)P == null) ? null : P.Value[A]B
    ```

    Except that `P` is evaluated only once.

- Otherwise:

  Let `T` be the type of the expression `P[A]B`.

  - If `T` is a type parameter that is not known to be either a reference type or a non-nullable value type, a compile-time error occurs.
  - If `T` is a non-nullable value type, then the type of `E` is `T?`, and the meaning of `E` is the same as the meaning of:

    ```
    C#
    ```

```
((object)P == null) ? (T?)null : P[A]B
```

Except that `P` is evaluated only once.

- Otherwise the type of `E` is `T`, and the meaning of `E` is the same as the meaning of:

```
C#
((object)P == null) ? null : P[A]B
```

Except that `P` is evaluated only once.

> *Note*: In an expression of the form:
>
> ```
> C#
> P?[A₀]?[A₁]
> ```
>
> if `P` evaluates to `null` neither `A₀` or `A₁` are evaluated. The same is true if an expression is a sequence of *null_conditional_element_access* or *null_conditional_member_access* §12.8.8 operations.
>
> *end note*

12.8.13 This access

A *this_access* consists of the keyword `this`.

```
ANTLR
this_access
    : 'this'
    ;
```

A *this_access* is permitted only in the *block* of an instance constructor, an instance method, an instance accessor (§12.2.1), or a finalizer. It has one of the following meanings:

- When `this` is used in a *primary_expression* within an instance constructor of a class, it is classified as a value. The type of the value is the instance type (§15.3.2) of

the class within which the usage occurs, and the value is a reference to the object being constructed.

- When `this` is used in a *primary_expression* within an instance method or instance accessor of a class, it is classified as a value. The type of the value is the instance type (§15.3.2) of the class within which the usage occurs, and the value is a reference to the object for which the method or accessor was invoked.
- When `this` is used in a *primary_expression* within an instance constructor of a struct, it is classified as a variable. The type of the variable is the instance type (§15.3.2) of the struct within which the usage occurs, and the variable represents the struct being constructed.
  - If the constructor declaration has no constructor initializer, the `this` variable behaves exactly the same as an `out` parameter of the struct type. In particular, this means that the variable shall be definitely assigned in every execution path of the instance constructor.
  - Otherwise, the `this` variable behaves exactly the same as a `ref` parameter of the struct type. In particular, this means that the variable is considered initially assigned.
- When `this` is used in a *primary_expression* within an instance method or instance accessor of a struct, it is classified as a variable. The type of the variable is the instance type (§15.3.2) of the struct within which the usage occurs.
  - If the method or accessor is not an iterator (§15.14) or async function (§15.15), the `this` variable represents the struct for which the method or accessor was invoked.
    - If the struct is a `readonly struct`, the `this` variable behaves exactly the same as an `in` parameter of the struct type
    - Otherwise the `this` variable behaves exactly the same as a `ref` parameter of the struct type
  - If the method or accessor is an iterator or async function, the `this` variable represents a *copy* of the struct for which the method or accessor was invoked, and behaves exactly the same as a *value* parameter of the struct type.

Use of `this` in a *primary_expression* in a context other than the ones listed above is a compile-time error. In particular, it is not possible to refer to `this` in a static method, a static property accessor, or in a *variable_initializer* of a field declaration.

12.8.14 Base access

A *base_access* consists of the keyword base followed by either a "`.`" token and an identifier and optional *type_argument_list* or an *argument_list* enclosed in square brackets:

ANTLR

```
base_access
    : 'base' '.' identifier type_argument_list?
    | 'base' '[' argument_list ']'
    ;
```

A *base_access* is used to access base class members that are hidden by similarly named members in the current class or struct. A *base_access* is permitted only in the body of an instance constructor, an instance method, an instance accessor (§12.2.1), or a finalizer. When `base.I` occurs in a class or struct, I shall denote a member of the base class of that class or struct. Likewise, when `base[E]` occurs in a class, an applicable indexer shall exist in the base class.

At binding-time, *base_access* expressions of the form `base.I` and `base[E]` are evaluated exactly as if they were written `((B)this).I` and `((B)this)[E]`, where `B` is the base class of the class or struct in which the construct occurs. Thus, `base.I` and `base[E]` correspond to `this.I` and `this[E]`, except `this` is viewed as an instance of the base class.

When a *base_access* references a virtual function member (a method, property, or indexer), the determination of which function member to invoke at run-time (§12.6.6) is changed. The function member that is invoked is determined by finding the most derived implementation (§15.6.4) of the function member with respect to `B` (instead of with respect to the run-time type of `this`, as would be usual in a non-base access). Thus, within an override of a virtual function member, a *base_access* can be used to invoke the inherited implementation of the function member. If the function member referenced by a *base_access* is abstract, a binding-time error occurs.

> *Note*: Unlike `this`, `base` is not an expression in itself. It is a keyword only used in the context of a *base_access* or a *constructor_initializer* (§15.11.2). *end note*

12.8.15 Postfix increment and decrement operators

ANTLR

```
post_increment_expression
    : primary_expression '++'
    ;

post_decrement_expression
    : primary_expression '--'
    ;
```

The operand of a postfix increment or decrement operation shall be an expression classified as a variable, a property access, or an indexer access. The result of the operation is a value of the same type as the operand.

If the *primary_expression* has the compile-time type `dynamic` then the operator is dynamically bound (§12.3.3), the *post_increment_expression* or *post_decrement_expression* has the compile-time type `dynamic` and the following rules are applied at run-time using the run-time type of the *primary_expression*.

If the operand of a postfix increment or decrement operation is a property or indexer access, the property or indexer shall have both a get and a set accessor. If this is not the case, a binding-time error occurs.

Unary operator overload resolution (§12.4.4) is applied to select a specific operator implementation. Predefined `++` and `--` operators exist for the following types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, `double`, `decimal`, and any enum type. The predefined `++` operators return the value produced by adding `1` to the operand, and the predefined `--` operators return the value produced by subtracting `1` from the operand. In a checked context, if the result of this addition or subtraction is outside the range of the result type and the result type is an integral type or enum type, a `System.OverflowException` is thrown.

There shall be an implicit conversion from the return type of the selected unary operator to the type of the *primary_expression*, otherwise a compile-time error occurs.

The run-time processing of a postfix increment or decrement operation of the form `x++` or `x--` consists of the following steps:

- If `x` is classified as a variable:
  - `x` is evaluated to produce the variable.
  - The value of `x` is saved.
  - The saved value of `x` is converted to the operand type of the selected operator and the operator is invoked with this value as its argument.
  - The value returned by the operator is converted to the type of `x` and stored in the location given by the earlier evaluation of `x`.
  - The saved value of `x` becomes the result of the operation.
- If `x` is classified as a property or indexer access:
  - The instance expression (if `x` is not `static`) and the argument list (if `x` is an indexer access) associated with `x` are evaluated, and the results are used in the subsequent get and set accessor invocations.
  - The get accessor of `x` is invoked and the returned value is saved.

- The saved value of `x` is converted to the operand type of the selected operator and the operator is invoked with this value as its argument.
- The value returned by the operator is converted to the type of `x` and the set accessor of `x` is invoked with this value as its value argument.
- The saved value of `x` becomes the result of the operation.

The `++` and `--` operators also support prefix notation (§12.9.6). Typically, the result of `x++` or `x--` is the value of `x` *before* the operation, whereas the result of `++x` or `--x` is the value of `x` *after* the operation. In either case, `x` itself has the same value after the operation.

An operator `++` or operator `--` implementation can be invoked using either postfix or prefix notation. It is not possible to have separate operator implementations for the two notations.

12.8.16 The new operator

12.8.16.1 General

The `new` operator is used to create new instances of types.

There are three forms of new expressions:

- Object creation expressions and anonymous object creation expressions are used to create new instances of class types and value types.
- Array creation expressions are used to create new instances of array types.
- Delegate creation expressions are used to obtain instances of delegate types.

The `new` operator implies creation of an instance of a type, but does not necessarily imply allocation of memory. In particular, instances of value types require no additional memory beyond the variables in which they reside, and no allocations occur when `new` is used to create instances of value types.

> *Note*: Delegate creation expressions do not always create new instances. When the expression is processed in the same way as a method group conversion (§10.8) or an anonymous function conversion (§10.7) this may result in an existing delegate instance being reused. *end note*

12.8.16.2 Object creation expressions

An *object_creation_expression* is used to create a new instance of a *class_type* or a *value_type*.

ANTLR

```
object_creation_expression
    : 'new' type '(' argument_list? ')' object_or_collection_initializer?
    | 'new' type object_or_collection_initializer
    ;

object_or_collection_initializer
    : object_initializer
    | collection_initializer
    ;
```

The *type* of an *object_creation_expression* shall be a *class_type*, a *value_type*, or a *type_parameter*. The *type* cannot be a *tuple_type* or an abstract or static *class_type*.

The optional *argument_list* (§12.6.2) is permitted only if the *type* is a *class_type* or a *struct_type*.

An object creation expression can omit the constructor argument list and enclosing parentheses provided it includes an object initializer or collection initializer. Omitting the constructor argument list and enclosing parentheses is equivalent to specifying an empty argument list.

Processing of an object creation expression that includes an object initializer or collection initializer consists of first processing the instance constructor and then processing the member or element initializations specified by the object initializer (§12.8.16.3) or collection initializer (§12.8.16.4).

If any of the arguments in the optional *argument_list* has the compile-time type `dynamic` then the *object_creation_expression* is dynamically bound (§12.3.3) and the following rules are applied at run-time using the run-time type of those arguments of the *argument_list* that have the compile-time type `dynamic`. However, the object creation undergoes a limited compile-time check as described in §12.6.5.

The binding-time processing of an *object_creation_expression* of the form new `T(A)`, where `T` is a *class_type*, or a *value_type*, and `A` is an optional *argument_list*, consists of the following steps:

- If `T` is a *value_type* and `A` is not present:
  - The *object_creation_expression* is a default constructor invocation. The result of the *object_creation_expression* is a value of type `T`, namely the default value for `T` as defined in §8.3.3.

- Otherwise, if `T` is a *type_parameter* and `A` is not present:
  - If no value type constraint or constructor constraint (§15.2.5) has been specified for `T`, a binding-time error occurs.
  - The result of the *object_creation_expression* is a value of the run-time type that the type parameter has been bound to, namely the result of invoking the default constructor of that type. The run-time type may be a reference type or a value type.
- Otherwise, if `T` is a *class_type* or a *struct_type*:
  - If `T` is an abstract or static *class_type*, a compile-time error occurs.
  - The instance constructor to invoke is determined using the overload resolution rules of §12.6.4. The set of candidate instance constructors consists of all accessible instance constructors declared in `T`, which are applicable with respect to A (§12.6.4.2). If the set of candidate instance constructors is empty, or if a single best instance constructor cannot be identified, a binding-time error occurs.
  - The result of the *object_creation_expression* is a value of type `T`, namely the value produced by invoking the instance constructor determined in the step above.
  - Otherwise, the *object_creation_expression* is invalid, and a binding-time error occurs.

Even if the *object_creation_expression* is dynamically bound, the compile-time type is still `T`.

The run-time processing of an *object_creation_expression* of the form new `T(A)`, where `T` is *class_type* or a *struct_type* and `A` is an optional *argument_list*, consists of the following steps:

- If `T` is a *class_type*:
  - A new instance of class `T` is allocated. If there is not enough memory available to allocate the new instance, a `System.OutOfMemoryException` is thrown and no further steps are executed.
  - All fields of the new instance are initialized to their default values (§9.3).
  - The instance constructor is invoked according to the rules of function member invocation (§12.6.6). A reference to the newly allocated instance is automatically passed to the instance constructor and the instance can be accessed from within that constructor as this.
- If `T` is a *struct_type*:
  - An instance of type `T` is created by allocating a temporary local variable. Since an instance constructor of a *struct_type* is required to definitely assign a value to

each field of the instance being created, no initialization of the temporary variable is necessary.

- The instance constructor is invoked according to the rules of function member invocation (§12.6.6). A reference to the newly allocated instance is automatically passed to the instance constructor and the instance can be accessed from within that constructor as this.

12.8.16.3 Object initializers

An ***object initializer*** specifies values for zero or more fields, properties, or indexed elements of an object.

ANTLR

```
object_initializer
    : '{' member_initializer_list? '}'
    | '{' member_initializer_list ',' '}'
    ;

member_initializer_list
    : member_initializer (',' member_initializer)*
    ;

member_initializer
    : initializer_target '=' initializer_value
    ;

initializer_target
    : identifier
    | '[' argument_list ']'
    ;

initializer_value
    : expression
    | object_or_collection_initializer
    ;
```

An object initializer consists of a sequence of member initializers, enclosed by `{` and `}` tokens and separated by commas. Each *member_initializer* shall designate a target for the initialization. An *identifier* shall name an accessible field or property of the object being initialized, whereas an *argument_list* enclosed in square brackets shall specify arguments for an accessible indexer on the object being initialized. It is an error for an object initializer to include more than one member initializer for the same field or property.

> *Note*: While an object initializer is not permitted to set the same field or property more than once, there are no such restrictions for indexers. An object initializer may contain multiple initializer targets referring to indexers, and may even use the same indexer arguments multiple times. *end note*

Each *initializer_target* is followed by an equals sign and either an expression, an object initializer or a collection initializer. It is not possible for expressions within the object initializer to refer to the newly created object it is initializing.

A member initializer that specifies an expression after the equals sign is processed in the same way as an assignment (§12.21.2) to the target.

A member initializer that specifies an object initializer after the equals sign is a ***nested object initializer***, i.e., an initialization of an embedded object. Instead of assigning a new value to the field or property, the assignments in the nested object initializer are treated as assignments to members of the field or property. Nested object initializers cannot be applied to properties with a value type, or to read-only fields with a value type.

A member initializer that specifies a collection initializer after the equals sign is an initialization of an embedded collection. Instead of assigning a new collection to the target field, property, or indexer, the elements given in the initializer are added to the collection referenced by the target. The target shall be of a collection type that satisfies the requirements specified in §12.8.16.4.

When an initializer target refers to an indexer, the arguments to the indexer shall always be evaluated exactly once. Thus, even if the arguments end up never getting used (e.g., because of an empty nested initializer), they are evaluated for their side effects.

> *Example*: The following class represents a point with two coordinates:

C#

```
public class Point
{
    public int X { get; set; }
    public int Y { get; set; }
}
```

> An instance of `Point` can be created and initialized as follows:

C#

```
Point a = new Point { X = 0, Y = 1 };
```

This has the same effect as

C#

```
Point __a = new Point();
__a.X = 0;
__a.Y = 1;
Point a = __a;
```

where `__a` is an otherwise invisible and inaccessible temporary variable.

The following class shows a rectangle created from two points, and the creation and initialization of a `Rectangle` instance:

C#

```
public class Rectangle
{
    public Point P1 { get; set; }
    public Point P2 { get; set; }
}
```

An instance of `Rectangle` can be created and initialized as follows:

C#

```
Rectangle r = new Rectangle
{
    P1 = new Point { X = 0, Y = 1 },
    P2 = new Point { X = 2, Y = 3 }
};
```

This has the same effect as

C#

```
Rectangle __r = new Rectangle();
Point __p1 = new Point();
__p1.X = 0;
__p1.Y = 1;
__r.P1 = __p1;
Point __p2 = new Point();
__p2.X = 2;
__p2.Y = 3;
__r.P2 = __p2;
Rectangle r = __r;
```

where `__r`, `__p1` and `__p2` are temporary variables that are otherwise invisible and inaccessible.

If `Rectangle`'s constructor allocates the two embedded `Point` instances, they can be used to initialize the embedded `Point` instances instead of assigning new instances:

C#

```
public class Rectangle
{
    public Point P1 { get; } = new Point();
    public Point P2 { get; } = new Point();
}
```

the following construct can be used to initialize the embedded `Point` instances instead of assigning new instances:

C#

```
Rectangle r = new Rectangle
{
    P1 = { X = 0, Y = 1 },
    P2 = { X = 2, Y = 3 }
};
```

This has the same effect as

C#

```
Rectangle __r = new Rectangle();
__r.P1.X = 0;
__r.P1.Y = 1;
__r.P2.X = 2;
__r.P2.Y = 3;
Rectangle r = __r;
```

end example

12.8.16.4 Collection initializers

A collection initializer specifies the elements of a collection.

ANTLR

```
collection_initializer
    : '{' element_initializer_list '}'
```

```
    | '{' element_initializer_list ',' '}'
    ;

element_initializer_list
    : element_initializer (',' element_initializer)*
    ;

element_initializer
    : non_assignment_expression
    | '{' expression_list '}'
    ;

expression_list
    : expression
    | expression_list ',' expression
    ;
```

A collection initializer consists of a sequence of element initializers, enclosed by `{` and `}` tokens and separated by commas. Each element initializer specifies an element to be added to the collection object being initialized, and consists of a list of expressions enclosed by `{` and `}` tokens and separated by commas. A single-expression element initializer can be written without braces, but cannot then be an assignment expression, to avoid ambiguity with member initializers. The *non_assignment_expression* production is defined in §12.22.

> *Example*: The following is an example of an object creation expression that includes a collection initializer:
>
> C#
>
> ```
> List<int> digits = new List<int> { 0, 1, 2, 3, 4, 5, 6, 7, 8, 9 };
> ```
>
> *end example*

The collection object to which a collection initializer is applied shall be of a type that implements `System.Collections.IEnumerable` or a compile-time error occurs. For each specified element in order from left to right, normal member lookup is applied to find a member named `Add`. If the result of the member lookup is not a method group, a compile-time error occurs. Otherwise, overload resolution is applied with the expression list of the element initializer as the argument list, and the collection initializer invokes the resulting method. Thus, the collection object shall contain an applicable instance or extension method with the name `Add` for each element initializer.

Example:The following shows a class that represents a contact with a name and a list of phone numbers, and the creation and initialization of a `List<Contact>`:

C#

```
public class Contact
{
    public string Name { get; set; }
    public List<string> PhoneNumbers { get; } = new List<string>();
}

class A
{
    static void M()
    {
        var contacts = new List<Contact>
        {
            new Contact
            {
                Name = "Chris Smith",
                PhoneNumbers = { "206-555-0101", "425-882-8080" }
            },
            new Contact
            {
                Name = "Bob Harris",
                PhoneNumbers = { "650-555-0199" }
            }
        };
    }
}
```

which has the same effect as

C#

```
var __clist = new List<Contact>();
Contact __c1 = new Contact();
__c1.Name = "Chris Smith";
__c1.PhoneNumbers.Add("206-555-0101");
__c1.PhoneNumbers.Add("425-882-8080");
__clist.Add(__c1);
Contact __c2 = new Contact();
__c2.Name = "Bob Harris";
__c2.PhoneNumbers.Add("650-555-0199");
__clist.Add(__c2);
var contacts = __clist;
```

where `__clist`, `__c1` and `__c2` are temporary variables that are otherwise invisible and inaccessible.

> *end example*

12.8.16.5 Array creation expressions

An *array_creation_expression* is used to create a new instance of an *array_type*.

ANTLR

```
array_creation_expression
    : 'new' non_array_type '[' expression_list ']' rank_specifier*
      array_initializer?
    | 'new' array_type array_initializer
    | 'new' rank_specifier array_initializer
    ;
```

An array creation expression of the first form allocates an array instance of the type that results from deleting each of the individual expressions from the expression list.

> *Example*: The array creation expression `new int[10,20]` produces an array instance of type `int[,]`, and the array creation expression new `int[10][,]` produces an array instance of type `int[][,]`. *end example*

Each expression in the expression list shall be of type `int`, `uint`, `long`, or `ulong`, or implicitly convertible to one or more of these types. The value of each expression determines the length of the corresponding dimension in the newly allocated array instance. Since the length of an array dimension shall be nonnegative, it is a compile-time error to have a constant expression with a negative value, in the expression list.

Except in an unsafe context (§23.2), the layout of arrays is unspecified.

If an array creation expression of the first form includes an array initializer, each expression in the expression list shall be a constant and the rank and dimension lengths specified by the expression list shall match those of the array initializer.

In an array creation expression of the second or third form, the rank of the specified array type or rank specifier shall match that of the array initializer. The individual dimension lengths are inferred from the number of elements in each of the corresponding nesting levels of the array initializer. Thus, the initializer expression in the following declaration

C#

```
var a = new int[,] {{0, 1}, {2, 3}, {4, 5}};
```

exactly corresponds to

```
C#
var a = new int[3, 2] {{0, 1}, {2, 3}, {4, 5}};
```

An array creation expression of the third form is referred to as an ***implicitly typed array-creation expression***. It is similar to the second form, except that the element type of the array is not explicitly given, but determined as the best common type (§12.6.3.15) of the set of expressions in the array initializer. For a multidimensional array, i.e., one where the *rank_specifier* contains at least one comma, this set comprises all *expression*s found in nested *array_initializer*s.

Array initializers are described further in §17.7.

The result of evaluating an array creation expression is classified as a value, namely a reference to the newly allocated array instance. The run-time processing of an array creation expression consists of the following steps:

- The dimension length expressions of the *expression_list* are evaluated in order, from left to right. Following evaluation of each expression, an implicit conversion (§10.2) to one of the following types is performed: `int`, `uint`, `long`, `ulong`. The first type in this list for which an implicit conversion exists is chosen. If evaluation of an expression or the subsequent implicit conversion causes an exception, then no further expressions are evaluated and no further steps are executed.
- The computed values for the dimension lengths are validated, as follows: If one or more of the values are less than zero, a `System.OverflowException` is thrown and no further steps are executed.
- An array instance with the given dimension lengths is allocated. If there is not enough memory available to allocate the new instance, a `System.OutOfMemoryException` is thrown and no further steps are executed.
- All elements of the new array instance are initialized to their default values (§9.3).
- If the array creation expression contains an array initializer, then each expression in the array initializer is evaluated and assigned to its corresponding array element. The evaluations and assignments are performed in the order the expressions are written in the array initializer—in other words, elements are initialized in increasing index order, with the rightmost dimension increasing first. If evaluation of a given expression or the subsequent assignment to the corresponding array element causes an exception, then no further elements are initialized (and the remaining elements will thus have their default values).

An array creation expression permits instantiation of an array with elements of an array type, but the elements of such an array shall be manually initialized.

> *Example*: The statement
>
> ```C#
> int[][] a = new int[100][];
> ```
>
> creates a single-dimensional array with 100 elements of type `int[]`. The initial value of each element is `null`. It is not possible for the same array creation expression to also instantiate the sub-arrays, and the statement
>
> ```C#
> int[][] a = new int[100][5]; // Error
> ```
>
> results in a compile-time error. Instantiation of the sub-arrays can instead be performed manually, as in
>
> ```C#
> int[][] a = new int[100][];
> for (int i = 0; i < 100; i++)
> {
>     a[i] = new int[5];
> }
> ```
>
> *end example*

> *Note*: When an array of arrays has a "rectangular" shape, that is when the sub-arrays are all of the same length, it is more efficient to use a multi-dimensional array. In the example above, instantiation of the array of arrays creates 101 objects—one outer array and 100 sub-arrays. In contrast,
>
> ```C#
> int[,] a = new int[100, 5];
> ```
>
> creates only a single object, a two-dimensional array, and accomplishes the allocation in a single statement.
>
> *end note*

Example: The following are examples of implicitly typed array creation expressions:

C#

```
var a = new[] { 1, 10, 100, 1000 };                    // int[]
var b = new[] { 1, 1.5, 2, 2.5 };                      // double[]
var c = new[,] { { "hello", null }, { "world", "!" } }; // string[,]
var d = new[] { 1, "one", 2, "two" };                  // Error
```

The last expression causes a compile-time error because neither `int` nor `string` is implicitly convertible to the other, and so there is no best common type. An explicitly typed array creation expression must be used in this case, for example specifying the type to be `object[]`. Alternatively, one of the elements can be cast to a common base type, which would then become the inferred element type.

end example

Implicitly typed array creation expressions can be combined with anonymous object initializers (§12.8.16.7) to create anonymously typed data structures.

Example:

C#

```
var contacts = new[]
{
    new
    {
        Name = "Chris Smith",
        PhoneNumbers = new[] { "206-555-0101", "425-882-8080" }
    },
    new
    {
        Name = "Bob Harris",
       PhoneNumbers = new[] { "650-555-0199" }
    }
};
```

end example

12.8.16.6 Delegate creation expressions

A *delegate_creation_expression* is used to obtain an instance of a *delegate_type*.

ANTLR

```
delegate_creation_expression
    : 'new' delegate_type '(' expression ')'
    ;
```

The argument of a delegate creation expression shall be a method group, an anonymous function, or a value of either the compile-time type `dynamic` or a *delegate_type*. If the argument is a method group, it identifies the method and, for an instance method, the object for which to create a delegate. If the argument is an anonymous function it directly defines the parameters and method body of the delegate target. If the argument is a value it identifies a delegate instance of which to create a copy.

If the *expression* has the compile-time type `dynamic`, the *delegate_creation_expression* is dynamically bound (§12.8.16.6), and the rules below are applied at run-time using the run-time type of the *expression*. Otherwise, the rules are applied at compile-time.

The binding-time processing of a *delegate_creation_expression* of the form new `D(E)`, where `D` is a *delegate_type* and `E` is an *expression*, consists of the following steps:

- If `E` is a method group, the delegate creation expression is processed in the same way as a method group conversion (§10.8) from `E` to `D`.
- If `E` is an anonymous function, the delegate creation expression is processed in the same way as an anonymous function conversion (§10.7) from `E` to `D`.
- If `E` is a value, `E` shall be compatible (§20.2) with `D`, and the result is a reference to a newly created delegate with a single-entry invocation list that invokes `E`.

The run-time processing of a *delegate_creation_expression* of the form new `D(E)`, where `D` is a *delegate_type* and `E` is an *expression*, consists of the following steps:

- If `E` is a method group, the delegate creation expression is evaluated as a method group conversion (§10.8) from `E` to `D`.
- If `E` is an anonymous function, the delegate creation is evaluated as an anonymous function conversion from `E` to `D` (§10.7).
- If `E` is a value of a *delegate_type*:
  - `E` is evaluated. If this evaluation causes an exception, no further steps are executed.
  - If the value of `E` is `null`, a `System.NullReferenceException` is thrown and no further steps are executed.

- A new instance of the delegate type `D` is allocated. If there is not enough memory available to allocate the new instance, a `System.OutOfMemoryException` is thrown and no further steps are executed.
- The new delegate instance is initialized with a single-entry invocation list that invokes `E`.

The invocation list of a delegate is determined when the delegate is instantiated and then remains constant for the entire lifetime of the delegate. In other words, it is not possible to change the target callable entities of a delegate once it has been created.

> *Note*: Remember, when two delegates are combined or one is removed from another, a new delegate results; no existing delegate has its content changed. *end note*

It is not possible to create a delegate that refers to a property, indexer, user-defined operator, instance constructor, finalizer, or static constructor.

> *Example*: As described above, when a delegate is created from a method group, the formal parameter list and return type of the delegate determine which of the overloaded methods to select. In the example
>
> C#
>
> ```
> delegate double DoubleFunc(double x);
>
> class A
> {
>     DoubleFunc f = new DoubleFunc(Square);
>
>     static float Square(float x) => x * x;
>     static double Square(double x) => x * x;
> }
> ```
>
> the `A.f` field is initialized with a delegate that refers to the second `Square` method because that method exactly matches the formal parameter list and return type of `DoubleFunc`. Had the second `Square` method not been present, a compile-time error would have occurred.
>
> *end example*

12.8.16.7 Anonymous object creation expressions

An *anonymous_object_creation_expression* is used to create an object of an anonymous type.

ANTLR

```
anonymous_object_creation_expression
    : 'new' anonymous_object_initializer
    ;

anonymous_object_initializer
    : '{' member_declarator_list? '}'
    | '{' member_declarator_list ',' '}'
    ;

member_declarator_list
    : member_declarator (',' member_declarator)*
    ;

member_declarator
    : simple_name
    | member_access
    | null_conditional_projection_initializer
    | base_access
    | identifier '=' expression
    ;
```

An anonymous object initializer declares an anonymous type and returns an instance of that type. An anonymous type is a nameless class type that inherits directly from `object`. The members of an anonymous type are a sequence of read-only properties inferred from the anonymous object initializer used to create an instance of the type. Specifically, an anonymous object initializer of the form

`new {` p_1 `=` e_1 `,` p_2 `=` e_2 `,` ... p_v `=` e_v `}`

declares an anonymous type of the form

C#

```
class __Anonymous1
{
    private readonly «T1» «f1»;
    private readonly «T2» «f2»;
    ...
    private readonly «Tn» «fn»;

    public __Anonymous1(«T1» «a1», «T2» «a2»,..., «Tn» «an»)
    {
        «f1» = «a1»;
        «f2» = «a2»;
        ...
        «fn» = «an»;
    }

    public «T1» «p1» { get { return «f1»; } }
```

```
    public «T2» «p2» { get { return «f2»; } }
    ...
    public «Tn» «pn» { get { return «fn»; } }
    public override bool Equals(object __o) { ... }
    public override int GetHashCode() { ... }
}
```

where each «Tx» is the type of the corresponding expression «ex». The expression used in a *member_declarator* shall have a type. Thus, it is a compile-time error for an expression in a *member_declarator* to be `null` or an anonymous function.

The names of an anonymous type and of the parameter to its `Equals` method are automatically generated by the compiler and cannot be referenced in program text.

Within the same program, two anonymous object initializers that specify a sequence of properties of the same names and compile-time types in the same order will produce instances of the same anonymous type.

> *Example*: In the example
>
> ```
> C#
> ```
>
> ```
> var p1 = new { Name = "Lawnmower", Price = 495.00 };
> var p2 = new { Name = "Shovel", Price = 26.95 };
> p1 = p2;
> ```
>
> the assignment on the last line is permitted because `p1` and `p2` are of the same anonymous type.
>
> *end example*

The `Equals` and `GetHashcode` methods on anonymous types override the methods inherited from `object`, and are defined in terms of the `Equals` and `GetHashcode` of the properties, so that two instances of the same anonymous type are equal if and only if all their properties are equal.

A member declarator can be abbreviated to a simple name (§12.8.4), a member access (§12.8.7), a null conditional projection initializer §12.8.8 or a base access (§12.8.14). This is called a ***projection initializer*** and is shorthand for a declaration of and assignment to a property with the same name. Specifically, member declarators of the forms

`«identifier»`, `«expr» . «identifier»` and `«expr» ? . «identifier»`

are precisely equivalent to the following, respectively:

`«identifer» = «identifier»`, `«identifier» = «expr» . «identifier»` and `«identifier» = «expr» ? . «identifier»`

Thus, in a projection initializer the identifier selects both the value and the field or property to which the value is assigned. Intuitively, a projection initializer projects not just a value, but also the name of the value.

12.8.17 The typeof operator

The `typeof` operator is used to obtain the `System.Type` object for a type.

ANTLR

```
typeof_expression
    : 'typeof' '(' type ')'
    | 'typeof' '(' unbound_type_name ')'
    | 'typeof' '(' 'void' ')'
    ;

unbound_type_name
    : identifier generic_dimension_specifier?
    | identifier '::' identifier generic_dimension_specifier?
    | unbound_type_name '.' identifier generic_dimension_specifier?
    ;

generic_dimension_specifier
    : '<' comma* '>'
    ;

comma
    : ','
    ;
```

The first form of *typeof_expression* consists of a `typeof` keyword followed by a parenthesized type. The result of an expression of this form is the `System.Type` object for the indicated type. There is only one `System.Type` object for any given type. This means that for a type `T`, `typeof(T) == typeof(T)` is always true. The type cannot be `dynamic`.

The second form of *typeof_expression* consists of a `typeof` keyword followed by a parenthesized *unbound_type_name*.

> *Note*: An *unbound_type_name* is very similar to a *type_name* (§7.8) except that an *unbound_type_name* contains *generic_dimension_specifier*s where a *type_name* contains *type_argument_list*s. *end note*

When the operand of a *typeof_expression* is a sequence of tokens that satisfies the grammars of both *unbound_type_name* and *type_name*, namely when it contains neither a *generic_dimension_specifier* nor a *type_argument_list*, the sequence of tokens is considered to be a *type_name*. The meaning of an *unbound_type_name* is determined as follows:

- Convert the sequence of tokens to a *type_name* by replacing each *generic_dimension_specifier* with a *type_argument_list* having the same number of commas and the keyword `object` as each *type_argument*.
- Evaluate the resulting *type_name*, while ignoring all type parameter constraints.
- The *unbound_type_name* resolves to the unbound generic type associated with the resulting constructed type (§8.4).

The result of the *typeof_expression* is the `System.Type` object for the resulting unbound generic type.

The third form of *typeof_expression* consists of a `typeof` keyword followed by a parenthesized `void` keyword. The result of an expression of this form is the `System.Type` object that represents the absence of a type. The type object returned by `typeof(void)` is distinct from the type object returned for any type.

> *Note*: This special `System.Type` object is useful in class libraries that allow reflection onto methods in the language, where those methods wish to have a way to represent the return type of any method, including `void` methods, with an instance of `System.Type`. *end note*

The `typeof` operator can be used on a type parameter. The result is the `System.Type` object for the run-time type that was bound to the type parameter. The `typeof` operator can also be used on a constructed type or an unbound generic type (§8.4.4). The `System.Type` object for an unbound generic type is not the same as the `System.Type` object of the instance type (§15.3.2). The instance type is always a closed constructed type at run-time so its `System.Type` object depends on the run-time type arguments in use. The unbound generic type, on the other hand, has no type arguments, and yields the same `System.Type` object regardless of runtime type arguments.

> *Example*: The example

```C#
class X<T>
{
    public static void PrintTypes()
```

```
    {
        Type[] t =
        {
            typeof(int),
            typeof(System.Int32),
            typeof(string),
            typeof(double[]),
            typeof(void),
            typeof(T),
            typeof(X<T>),
            typeof(X<X<T>>),
            typeof(X<>)
        };
        for (int i = 0; i < t.Length; i++)
        {
            Console.WriteLine(t[i]);
        }
    }
}

class Test
{
    static void Main()
    {
        X<int>.PrintTypes();
    }
}
```

produces the following output:

Console

```
System.Int32
System.Int32
System.String
System.Double[]
System.Void
System.Int32
X`1[System.Int32]
X`1[X`1[System.Int32]]
X`1[T]
```

Note that `int` and `System.Int32` are the same type. The result of `typeof(X<>)` does not depend on the type argument but the result of `typeof(X<T>)` does.

end example

12.8.18 The sizeof operator

The `sizeof` operator returns the number of 8-bit bytes occupied by a variable of a given type. The type specified as an operand to sizeof shall be an *unmanaged_type* (§8.8).

ANTLR

```
sizeof_expression
    : 'sizeof' '(' unmanaged_type ')'
    ;
```

For certain predefined types the `sizeof` operator yields a constant `int` value as shown in the table below:

Expand table

Expression	Result
`sizeof(sbyte)`	1
`sizeof(byte)`	1
`sizeof(short)`	2
`sizeof(ushort)`	2
`sizeof(int)`	4
`sizeof(uint)`	4
`sizeof(long)`	8
`sizeof(ulong)`	8
`sizeof(char)`	2
`sizeof(float)`	4
`sizeof(double)`	8
`sizeof(bool)`	1
`sizeof(decimal)`	16

For an enum type `T`, the result of the expression `sizeof(T)` is a constant value equal to the size of its underlying type, as given above. For all other operand types, the `sizeof` operator is specified in §23.6.9.

12.8.19 The checked and unchecked operators

The `checked` and `unchecked` operators are used to control the overflow-checking context for integral-type arithmetic operations and conversions.

ANTLR

```
checked_expression
    : 'checked' '(' expression ')'
    ;

unchecked_expression
    : 'unchecked' '(' expression ')'
    ;
```

The `checked` operator evaluates the contained expression in a checked context, and the `unchecked` operator evaluates the contained expression in an unchecked context. A *checked_expression* or *unchecked_expression* corresponds exactly to a *parenthesized_expression* (§12.8.5), except that the contained expression is evaluated in the given overflow checking context.

The overflow checking context can also be controlled through the `checked` and `unchecked` statements (§13.12).

The following operations are affected by the overflow checking context established by the checked and unchecked operators and statements:

- The predefined `++` and `--` operators (§12.8.15 and §12.9.6), when the operand is of an integral or enum type.
- The predefined `-` unary operator (§12.9.3), when the operand is of an integral type.
- The predefined `+`, `-`, `*`, and `/` binary operators (§12.10), when both operands are of integral or enum types.
- Explicit numeric conversions (§10.3.2) from one integral or enum type to another integral or enum type, or from `float` or `double` to an integral or enum type.

When one of the above operations produces a result that is too large to represent in the destination type, the context in which the operation is performed controls the resulting behavior:

- In a `checked` context, if the operation is a constant expression (§12.23), a compile-time error occurs. Otherwise, when the operation is performed at run-time, a `System.OverflowException` is thrown.
- In an `unchecked` context, the result is truncated by discarding any high-order bits that do not fit in the destination type.

For non-constant expressions (§12.23) (expressions that are evaluated at run-time) that are not enclosed by any `checked` or `unchecked` operators or statements, the default overflow checking context is unchecked, unless external factors (such as compiler switches and execution environment configuration) call for checked evaluation.

For constant expressions (§12.23) (expressions that can be fully evaluated at compile-time), the default overflow checking context is always checked. Unless a constant expression is explicitly placed in an `unchecked` context, overflows that occur during the compile-time evaluation of the expression always cause compile-time errors.

The body of an anonymous function is not affected by `checked` or `unchecked` contexts in which the anonymous function occurs.

> *Example*: In the following code
>
> C#
>
> ```
> class Test
> {
>     static readonly int x = 1000000;
>     static readonly int y = 1000000;
>
>     static int F() => checked(x * y);    // Throws OverflowException
>     static int G() => unchecked(x * y);  // Returns -727379968
>     static int H() => x * y;             // Depends on default
> }
> ```
>
> no compile-time errors are reported since neither of the expressions can be evaluated at compile-time. At run-time, the `F` method throws a `System.OverflowException`, and the `G` method returns –727379968 (the lower 32 bits of the out-of-range result). The behavior of the `H` method depends on the default overflow-checking context for the compilation, but it is either the same as `F` or the same as `G`.
>
> *end example*

> *Example*: In the following code
>
> C#
>
> ```
> class Test
> {
>     const int x = 1000000;
>     const int y = 1000000;
> ```

```
    static int F() => checked(x * y);    // Compile-time error, overflow
    static int G() => unchecked(x * y);  // Returns -727379968
    static int H() => x * y;             // Compile-time error, overflow
}
```

the overflows that occur when evaluating the constant expressions in `F` and `H` cause compile-time errors to be reported because the expressions are evaluated in a `checked` context. An overflow also occurs when evaluating the constant expression in `G`, but since the evaluation takes place in an `unchecked` context, the overflow is not reported.

end example

The `checked` and `unchecked` operators only affect the overflow checking context for those operations that are textually contained within the "`(`" and "`)`" tokens. The operators have no effect on function members that are invoked as a result of evaluating the contained expression.

Example: In the following code

C#

```
class Test
{
    static int Multiply(int x, int y) => x * y;

    static int F() => checked(Multiply(1000000, 1000000));
}
```

the use of `checked` in F does not affect the evaluation of `x * y` in `Multiply`, so `x * y` is evaluated in the default overflow checking context.

end example

The `unchecked` operator is convenient when writing constants of the signed integral types in hexadecimal notation.

Example:

C#

```
class Test
{
    public const int AllBits = unchecked((int)0xFFFFFFFF);
```

```
    public const int HighBit = unchecked((int)0x80000000);
}
```

Both of the hexadecimal constants above are of type `uint`. Because the constants are outside the `int` range, without the `unchecked` operator, the casts to `int` would produce compile-time errors.

end example

Note: The `checked` and `unchecked` operators and statements allow programmers to control certain aspects of some numeric calculations. However, the behavior of some numeric operators depends on their operands' data types. For example, multiplying two decimals always results in an exception on overflow even within an explicitly unchecked construct. Similarly, multiplying two floats never results in an exception on overflow even within an explicitly checked construct. In addition, other operators are never affected by the mode of checking, whether default or explicit. *end note*

12.8.20 Default value expressions

A default value expression is used to obtain the default value (§9.3) of a type.

ANTLR

```
default_value_expression
    : explictly_typed_default
    | default_literal
    ;

explictly_typed_default
    : 'default' '(' type ')'
    ;

default_literal
    : 'default'
    ;
```

A *default_literal* represents a default value (§9.3). It does not have a type, but can be converted to any type through a default literal conversion (§10.2.16).

The result of a *default_value_expression* is the default (§9.3) of the explicit type in an *explictly_typed_default*, or the target type of the conversion for a *default_value_expression*.

A *default_value_expression* is a constant expression (§12.23) if the type is one of:

- a reference type
- a type parameter that is known to be a reference type (§8.2);
- one of the following value types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, `double`, `decimal`, `bool`, ; or
- any enumeration type.

12.8.21 Stack allocation

A stack allocation expression allocates a block of memory from the execution stack. The ***execution stack*** is an area of memory where local variables are stored. The execution stack is not part of the managed heap. The memory used for local variable storage is automatically recovered when the current function returns.

The safe context rules for a stack allocation expression are described in §16.4.12.7.

ANTLR

```
stackalloc_expression
    : 'stackalloc' unmanaged_type '[' expression ']'
    | 'stackalloc' unmanaged_type? '[' constant_expression? ']'
      stackalloc_initializer
    ;

stackalloc_initializer
     : '{' stackalloc_initializer_element_list '}'
     ;

stackalloc_initializer_element_list
     : stackalloc_element_initializer (',' stackalloc_element_initializer)*
','?
     ;

stackalloc_element_initializer
    : expression
    ;
```

A *stackalloc_expression* is only permitted in two contexts:

1. The initializing *expression*, `E`, of a *local_variable_declaration* (§13.6.2); and
2. The right operand *expression*, `E`, of a simple assignment (§12.21.2) which itself occurs as a *expression_statement* (§13.7)

In both contexts the *stackalloc_expression* is only permitted to occur as:

- The whole of `E`; or

- The second and/or third operands of a *conditional_expression* (§12.18) which is itself the whole of `E`.

The *unmanaged_type* (§8.8) indicates the type of the items that will be stored in the newly allocated location, and the *expression* indicates the number of these items. Taken together, these specify the required allocation size. The type of *expression* shall be implicitly convertible to the type `int`.

As the size of a stack allocation cannot be negative, it is a compile-time error to specify the number of items as a *constant_expression* that evaluates to a negative value.

At runtime if the number of items to be allocated is a negative value then the behavior is undefined. If it is zero, then no allocation is made, and the value returned is implementation-defined. If there is not enough memory available to allocate the items a `System.StackOverflowException` is thrown.

When a *stackalloc_initializer* is present:

- If *unmanaged_type* is omitted, it is inferred following the rules for best common type (§12.6.3.15) for the set of *stackalloc_element_initializer*s.
- If *constant_expression* is omitted it is inferred to be the number of *stackalloc_element_initializer*s.
- If *constant_expression* is present it shall equal the number of *stackalloc_element_initializer*s.

Each *stackalloc_element_initializer* shall have an implicit conversion to *unmanaged_type* (§10.2). The *stackalloc_element_initializer*s initialize elements in the allocated memory in increasing order, starting with the element at index zero. In the absence of a *stackalloc_initializer*, the content of the newly allocated memory is undefined.

The result of a *stackalloc_expression* is an instance of type `Span<T>`, where `T` is the *unmanaged_type*:

- `Span<T>` (§C.3) is a ref struct type (§16.2.3), which presents a block of memory, here the block allocated by the *stackalloc_expression*, as an indexable collection of typed (`T`) items.
- The result's `Length` property returns the number of items allocated.
- The result's indexer (§15.9) returns a *variable_reference* (§9.5) to an item of the allocated block and is range checked.

> *Note*: When occurring in unsafe code the result of a *stackalloc_expression* may be of a different type, see (§23.9). *end note*

Stack allocation initializers are not permitted in `catch` or `finally` blocks (§13.11).

> *Note*: There is no way to explicitly free memory allocated using `stackalloc`. *end note*

All stack-allocated memory blocks created during the execution of a function member are automatically discarded when that function member returns.

Except for the `stackalloc` operator, C# provides no predefined constructs for managing non-garbage collected memory. Such services are typically provided by supporting class libraries or imported directly from the underlying operating system.

> *Example*:
>
> C#
>
> ```
> // Memory uninitialized
> Span<int> span1 = stackalloc int[3];
> // Memory initialized
> Span<int> span2 = stackalloc int[3] { -10, -15, -30 };
> // Type int is inferred
> Span<int> span3 = stackalloc[] { 11, 12, 13 };
> // Error; result is int*, not allowed in a safe context
> var span4 = stackalloc[] { 11, 12, 13 };
> // Error; no conversion from Span<int> to Span<long>
> Span<long> span5 = stackalloc[] { 11, 12, 13 };
> // Converts 11 and 13, and returns Span<long>
> Span<long> span6 = stackalloc[] { 11, 12L, 13 };
> // Converts all and returns Span<long>
> Span<long> span7 = stackalloc long[] { 11, 12, 13 };
> // Implicit conversion of Span<T>
> ReadOnlySpan<int> span8 = stackalloc int[] { 10, 22, 30 };
> // Implicit conversion of Span<T>
> Widget<double> span9 = stackalloc double[] { 1.2, 5.6 };
>
> public class Widget<T>
> {
>     public static implicit operator Widget<T>(Span<double> sp) { return
> null; }
> }
> ```
>
> In the case of `span8`, `stackalloc` results in a `Span<int>`, which is converted by an implicit operator to `ReadOnlySpan<int>`. Similarly, for `span9`, the resulting `Span<double>` is converted to the user-defined type `Widget<double>` using the conversion, as shown. *end example*

12.8.22 The nameof operator

A *nameof_expression* is used to obtain the name of a program entity as a constant string.

ANTLR

```
nameof_expression
    : 'nameof' '(' named_entity ')'
    ;

named_entity
    : named_entity_target ('.' identifier type_argument_list?)*
    ;

named_entity_target
    : simple_name
    | 'this'
    | 'base'
    | predefined_type
    | qualified_alias_member
    ;
```

Because `nameof` is not a keyword, a *nameof_expression* is always syntactically ambiguous with an invocation of the simple name `nameof`. For compatibility reasons, if a name lookup (§12.8.4) of the name `nameof` succeeds, the expression is treated as an *invocation_expression* — regardless of whether the invocation is valid. Otherwise it is a *nameof_expression*.

Simple name and member access lookups are performed on the *named_entity* at compile time, following the rules described in §12.8.4 and §12.8.7. However, where the lookup described in §12.8.4 and §12.8.7 results in an error because an instance member was found in a static context, a *nameof_expression* produces no such error.

It is a compile-time error for a *named_entity* designating a method group to have a *type_argument_list*. It is a compile time error for a *named_entity_target* to have the type `dynamic`.

A *nameof_expression* is a constant expression of type `string`, and has no effect at runtime. Specifically, its *named_entity* is not evaluated, and is ignored for the purposes of definite assignment analysis (§9.4.4.22). Its value is the last identifier of the *named_entity* before the optional final *type_argument_list*, transformed in the following way:

- The prefix "`@`", if used, is removed.
- Each *unicode_escape_sequence* is transformed into its corresponding Unicode character.
- Any *formatting_characters* are removed.

These are the same transformations applied in §6.4.3 when testing equality between identifiers.

> *Example*: The following illustrates the results of various `nameof` expressions, assuming a generic type `List<T>` declared within the `System.Collections.Generic` namespace:

C#

```
using TestAlias = System.String;

class Program
{
    static void Main()
    {
        var point = (x: 3, y: 4);

        string n1 = nameof(System);                       // "System"
        string n2 = nameof(System.Collections.Generic);   // "Generic"
        string n3 = nameof(point);                        // "point"
        string n4 = nameof(point.x);                      // "x"
        string n5 = nameof(Program);                      // "Program"
        string n6 = nameof(System.Int32);                 // "Int32"
        string n7 = nameof(TestAlias);                    // "TestAlias"
        string n8 = nameof(List<int>);                    // "List"
        string n9 = nameof(Program.InstanceMethod);       //
"InstanceMethod"
        string n10 = nameof(Program.GenericMethod);       //
"GenericMethod"
        string n11 = nameof(Program.NestedClass);         //
"NestedClass"

        // Invalid
        // string x1 = nameof(List<>);            // Empty type argument
list
        // string x2 = nameof(List<T>);           // T is not in scope
        // string x3 = nameof(GenericMethod<>);   // Empty type argument
list
        // string x4 = nameof(GenericMethod<T>);  // T is not in scope
        // string x5 = nameof(int);               // Keywords not
permitted
        // Type arguments not permitted for method group
        // string x6 = nameof(GenericMethod<Program>);
    }

    void InstanceMethod() { }

    void GenericMethod<T>()
    {
        string n1 = nameof(List<T>); // "List"
        string n2 = nameof(T);       // "T"
    }
```

```
    class NestedClass { }
}
```

Potentially surprising parts of this example are the resolution of `nameof(System.Collections.Generic)` to just "Generic" instead of the full namespace, and of `nameof(TestAlias)` to "TestAlias" rather than "String". *end example*

12.8.23 Anonymous method expressions

An *anonymous_method_expression* is one of two ways of defining an anonymous function. These are further described in §12.19.

12.9 Unary operators

12.9.1 General

The `+`, `-`, `!`, `~`, `++`, `--`, cast, and `await` operators are called the unary operators.

ANTLR

```
unary_expression
    : primary_expression
    | '+' unary_expression
    | '-' unary_expression
    | '!' unary_expression
    | '~' unary_expression
    | pre_increment_expression
    | pre_decrement_expression
    | cast_expression
    | await_expression
    | pointer_indirection_expression    // unsafe code support
    | addressof_expression              // unsafe code support
    ;
```

pointer_indirection_expression (§23.6.2) and *addressof_expression* (§23.6.5) are available only in unsafe code (§23).

If the operand of a *unary_expression* has the compile-time type `dynamic`, it is dynamically bound (§12.3.3). In this case, the compile-time type of the *unary_expression* is `dynamic`, and the resolution described below will take place at run-time using the run-time type of the operand.

12.9.2 Unary plus operator

For an operation of the form `+x`, unary operator overload resolution (§12.4.4) is applied to select a specific operator implementation. The operand is converted to the parameter type of the selected operator, and the type of the result is the return type of the operator. The predefined unary plus operators are:

```
C#

int operator +(int x);
uint operator +(uint x);
long operator +(long x);
ulong operator +(ulong x);
float operator +(float x);
double operator +(double x);
decimal operator +(decimal x);
```

For each of these operators, the result is simply the value of the operand.

Lifted (§12.4.8) forms of the unlifted predefined unary plus operators defined above are also predefined.

12.9.3 Unary minus operator

For an operation of the form `-x`, unary operator overload resolution (§12.4.4) is applied to select a specific operator implementation. The operand is converted to the parameter type of the selected operator, and the type of the result is the return type of the operator. The predefined unary minus operators are:

- Integer negation:

  ```
  C#

  int operator -(int x);
  long operator -(long x);
  ```

  The result is computed by subtracting `x` from zero. If the value of `x` is the smallest representable value of the operand type (-2^{31} for `int` or -2^{63} for `long`), then the mathematical negation of `x` is not representable within the operand type. If this occurs within a `checked` context, a `System.OverflowException` is thrown; if it occurs within an `unchecked` context, the result is the value of the operand and the overflow is not reported.

If the operand of the negation operator is of type `uint`, it is converted to type `long`, and the type of the result is `long`. An exception is the rule that permits the `int` value `-2147483648` (-2^{31}) to be written as a decimal integer literal (§6.4.5.3).

If the operand of the negation operator is of type `ulong`, a compile-time error occurs. An exception is the rule that permits the `long` value `-9223372036854775808` (-2^{63}) to be written as a decimal integer literal (§6.4.5.3)

- Floating-point negation:

```
C#

float operator -(float x);
double operator -(double x);
```

The result is the value of `x` with its sign inverted. If `x` is `NaN`, the result is also `NaN`.

- Decimal negation:

```
C#

decimal operator -(decimal x);
```

The result is computed by subtracting `x` from zero. Decimal negation is equivalent to using the unary minus operator of type `System.Decimal`.

Lifted (§12.4.8) forms of the unlifted predefined unary minus operators defined above are also predefined.

12.9.4 Logical negation operator

For an operation of the form `!x`, unary operator overload resolution (§12.4.4) is applied to select a specific operator implementation. The operand is converted to the parameter type of the selected operator, and the type of the result is the return type of the operator. Only one predefined logical negation operator exists:

```
C#

bool operator !(bool x);
```

This operator computes the logical negation of the operand: If the operand is `true`, the result is `false`. If the operand is `false`, the result is `true`.

Lifted (§12.4.8) forms of the unlifted predefined logical negation operator defined above are also predefined.

12.9.5 Bitwise complement operator

For an operation of the form `~x`, unary operator overload resolution (§12.4.4) is applied to select a specific operator implementation. The operand is converted to the parameter type of the selected operator, and the type of the result is the return type of the operator. The predefined bitwise complement operators are:

C#

```
int operator ~(int x);
uint operator ~(uint x);
long operator ~(long x);
ulong operator ~(ulong x);
```

For each of these operators, the result of the operation is the bitwise complement of `x`.

Every enumeration type `E` implicitly provides the following bitwise complement operator:

C#

```
E operator ~(E x);
```

The result of evaluating `~x`, where `X` is an expression of an enumeration type `E` with an underlying type `U`, is exactly the same as evaluating `(E)(~(U)x)`, except that the conversion to `E` is always performed as if in an `unchecked` context (§12.8.19).

Lifted (§12.4.8) forms of the unlifted predefined bitwise complement operators defined above are also predefined.

12.9.6 Prefix increment and decrement operators

ANTLR

```
pre_increment_expression
    : '++' unary_expression
    ;

pre_decrement_expression
```

```
    : '--' unary_expression
    ;
```

The operand of a prefix increment or decrement operation shall be an expression classified as a variable, a property access, or an indexer access. The result of the operation is a value of the same type as the operand.

If the operand of a prefix increment or decrement operation is a property or indexer access, the property or indexer shall have both a get and a set accessor. If this is not the case, a binding-time error occurs.

Unary operator overload resolution (§12.4.4) is applied to select a specific operator implementation. Predefined `++` and `--` operators exist for the following types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, `double`, `decimal`, and any enum type. The predefined `++` operators return the value produced by adding `1` to the operand, and the predefined `--` operators return the value produced by subtracting `1` from the operand. In a `checked` context, if the result of this addition or subtraction is outside the range of the result type and the result type is an integral type or enum type, a `System.OverflowException` is thrown.

There shall be an implicit conversion from the return type of the selected unary operator to the type of the *unary_expression*, otherwise a compile-time error occurs.

The run-time processing of a prefix increment or decrement operation of the form `++x` or `--x` consists of the following steps:

- If `x` is classified as a variable:
  - `x` is evaluated to produce the variable.
  - The value of `x` is converted to the operand type of the selected operator and the operator is invoked with this value as its argument.
  - The value returned by the operator is converted to the type of `x`. The resulting value is stored in the location given by the evaluation of `x`.
  - and becomes the result of the operation.
- If `x` is classified as a property or indexer access:
  - The instance expression (if `x` is not `static`) and the argument list (if `x` is an indexer access) associated with `x` are evaluated, and the results are used in the subsequent get and set accessor invocations.
  - The get accessor of `X` is invoked.
  - The value returned by the get accessor is converted to the operand type of the selected operator and operator is invoked with this value as its argument.

- The value returned by the operator is converted to the type of `x`. The set accessor of `X` is invoked with this value as its value argument.
- This value also becomes the result of the operation.

The `++` and `--` operators also support postfix notation (§12.8.15). Typically, the result of `x++` or `x--` is the value of `X` before the operation, whereas the result of `++x` or `--x` is the value of `X` after the operation. In either case, `x` itself has the same value after the operation.

An operator `++` or operator `--` implementation can be invoked using either postfix or prefix notation. It is not possible to have separate operator implementations for the two notations.

Lifted (§12.4.8) forms of the unlifted predefined prefix increment and decrement operators defined above are also predefined.

12.9.7 Cast expressions

A *cast_expression* is used to convert explicitly an expression to a given type.

ANTLR

```
cast_expression
    : '(' type ')' unary_expression
    ;
```

A *cast_expression* of the form `(T)E`, where `T` is a type and `E` is a *unary_expression*, performs an explicit conversion (§10.3) of the value of `E` to type `T`. If no explicit conversion exists from `E` to `T`, a binding-time error occurs. Otherwise, the result is the value produced by the explicit conversion. The result is always classified as a value, even if `E` denotes a variable.

The grammar for a *cast_expression* leads to certain syntactic ambiguities.

> *Example*: The expression `(x)-y` could either be interpreted as a *cast_expression* (a cast of `-y` to type `x`) or as an *additive_expression* combined with a *parenthesized_expression* (which computes the value `x - y`). *end example*

To resolve *cast_expression* ambiguities, the following rule exists: A sequence of one or more tokens (§6.4) enclosed in parentheses is considered the start of a *cast_expression* only if at least one of the following are true:

- The sequence of tokens is correct grammar for a type, but not for an expression.
- The sequence of tokens is correct grammar for a type, and the token immediately following the closing parentheses is the token "`~`", the token "`!`", the token "`(`", an identifier (§6.4.3), a literal (§6.4.5), or any keyword (§6.4.4) except `as` and `is`.

The term "correct grammar" above means only that the sequence of tokens shall conform to the particular grammatical production. It specifically does not consider the actual meaning of any constituent identifiers.

> *Example*: If `x` and `y` are identifiers, then `x.y` is correct grammar for a type, even if `x.y` doesn't actually denote a type. *end example*

> *Note*: From the disambiguation rule, it follows that, if `x` and `y` are identifiers, `(x)y`, `(x)(y)`, and `(x)(-y)` are *cast_expression*s, but `(x)-y` is not, even if `x` identifies a type. However, if `x` is a keyword that identifies a predefined type (such as `int`), then all four forms are *cast_expression*s (because such a keyword could not possibly be an expression by itself). *end note*

12.9.8 Await expressions

12.9.8.1 General

The `await` operator is used to suspend evaluation of the enclosing async function until the asynchronous operation represented by the operand has completed.

ANTLR

```
await_expression
    : 'await' unary_expression
    ;
```

An *await_expression* is only allowed in the body of an async function (§15.15). Within the nearest enclosing async function, an *await_expression* shall not occur in these places:

- Inside a nested (non-async) anonymous function
- Inside the block of a *lock_statement*
- In an anonymous function conversion to an expression tree type (§10.7.3)
- In an unsafe context

> *Note*: An *await_expression* cannot occur in most places within a *query_expression*, because those are syntactically transformed to use non-async lambda expressions.

> *end note*

Inside an async function, `await` shall not be used as an *available_identifier* although the verbatim identifier `@await` may be used. There is therefore no syntactic ambiguity between *await_expression*s and various expressions involving identifiers. Outside of async functions, `await` acts as a normal identifier.

The operand of an *await_expression* is called the ***task***. It represents an asynchronous operation that may or may not be complete at the time the *await_expression* is evaluated. The purpose of the `await` operator is to suspend execution of the enclosing async function until the awaited task is complete, and then obtain its outcome.

12.9.8.2 Awaitable expressions

The task of an *await_expression* is required to be ***awaitable***. An expression `t` is awaitable if one of the following holds:

- `t` is of compile-time type `dynamic`
- `t` has an accessible instance or extension method called `GetAwaiter` with no parameters and no type parameters, and a return type `A` for which all of the following hold:
  - `A` implements the interface `System.Runtime.CompilerServices.INotifyCompletion` (hereafter known as `INotifyCompletion` for brevity)
  - `A` has an accessible, readable instance property `IsCompleted` of type `bool`
  - `A` has an accessible instance method `GetResult` with no parameters and no type parameters

The purpose of the `GetAwaiter` method is to obtain an ***awaiter*** for the task. The type `A` is called the ***awaiter type*** for the await expression.

The purpose of the `IsCompleted` property is to determine if the task is already complete. If so, there is no need to suspend evaluation.

The purpose of the `INotifyCompletion.OnCompleted` method is to sign up a "continuation" to the task; i.e., a delegate (of type `System.Action`) that will be invoked once the task is complete.

The purpose of the `GetResult` method is to obtain the outcome of the task once it is complete. This outcome may be successful completion, possibly with a result value, or it may be an exception which is thrown by the `GetResult` method.

12.9.8.3 Classification of await expressions

The expression `await t` is classified the same way as the expression `(t).GetAwaiter().GetResult()`. Thus, if the return type of `GetResult` is `void`, the *await_expression* is classified as nothing. If it has a non-`void` return type `T`, the *await_expression* is classified as a value of type `T`.

12.9.8.4 Run-time evaluation of await expressions

At run-time, the expression `await t` is evaluated as follows:

- An awaiter `a` is obtained by evaluating the expression `(t).GetAwaiter()`.
- A `bool` `b` is obtained by evaluating the expression `(a).IsCompleted`.
- If `b` is `false` then evaluation depends on whether `a` implements the interface `System.Runtime.CompilerServices.ICriticalNotifyCompletion` (hereafter known as `ICriticalNotifyCompletion` for brevity). This check is done at binding time; i.e., at run-time if `a` has the compile-time type `dynamic`, and at compile-time otherwise. Let `r` denote the resumption delegate (§15.15):
  - If `a` does not implement `ICriticalNotifyCompletion`, then the expression `((a) as INotifyCompletion).OnCompleted(r)` is evaluated.
  - If `a` does implement `ICriticalNotifyCompletion`, then the expression `((a) as ICriticalNotifyCompletion).UnsafeOnCompleted(r)` is evaluated.
  - Evaluation is then suspended, and control is returned to the current caller of the async function.
- Either immediately after (if `b` was `true`), or upon later invocation of the resumption delegate (if `b` was `false`), the expression `(a).GetResult()` is evaluated. If it returns a value, that value is the result of the *await_expression*. Otherwise, the result is nothing.

An awaiter's implementation of the interface methods `INotifyCompletion.OnCompleted` and `ICriticalNotifyCompletion.UnsafeOnCompleted` should cause the delegate `r` to be invoked at most once. Otherwise, the behavior of the enclosing async function is undefined.

12.10 Arithmetic operators

12.10.1 General

The `*`, `/`, `%`, `+`, and `-` operators are called the arithmetic operators.

ANTLR

```
multiplicative_expression
    : unary_expression
    | multiplicative_expression '*' unary_expression
    | multiplicative_expression '/' unary_expression
    | multiplicative_expression '%' unary_expression
    ;

additive_expression
    : multiplicative_expression
    | additive_expression '+' multiplicative_expression
    | additive_expression '-' multiplicative_expression
    ;
```

If an operand of an arithmetic operator has the compile-time type `dynamic`, then the expression is dynamically bound (§12.3.3). In this case, the compile-time type of the expression is `dynamic`, and the resolution described below will take place at run-time using the run-time type of those operands that have the compile-time type `dynamic`.

12.10.2 Multiplication operator

For an operation of the form `x * y`, binary operator overload resolution (§12.4.5) is applied to select a specific operator implementation. The operands are converted to the parameter types of the selected operator, and the type of the result is the return type of the operator.

The predefined multiplication operators are listed below. The operators all compute the product of `x` and `y`.

- Integer multiplication:

  C#

  ```
  int operator *(int x, int y);
  uint operator *(uint x, uint y);
  long operator *(long x, long y);
  ulong operator *(ulong x, ulong y);
  ```

  In a `checked` context, if the product is outside the range of the result type, a `System.OverflowException` is thrown. In an `unchecked` context, overflows are not reported and any significant high-order bits outside the range of the result type are discarded.

- Floating-point multiplication:

```
C#

float operator *(float x, float y);
double operator *(double x, double y);
```

The product is computed according to the rules of IEC 60559 arithmetic. The following table lists the results of all possible combinations of nonzero finite values, zeros, infinities, and NaNs. In the table, `x` and `y` are positive finite values. `z` is the result of `x * y`, rounded to the nearest representable value. If the magnitude of the result is too large for the destination type, `z` is infinity. Because of rounding, `z` may be zero even though neither `x` nor `y` is zero.

Expand table

	+y	-y	+0	-0	+∞	-∞	NaN
+x	+z	-z	+0	-0	+∞	-∞	NaN
-x	-z	+z	-0	+0	-∞	+∞	NaN
+0	+0	-0	+0	-0	NaN	NaN	NaN
-0	-0	+0	-0	+0	NaN	NaN	NaN
+∞	+∞	-∞	NaN	NaN	+∞	-∞	NaN
-∞	-∞	+∞	NaN	NaN	-∞	+∞	NaN
NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN

(Except were otherwise noted, in the floating-point tables in §12.10.2–§12.10.6 the use of "`+`" means the value is positive; the use of "`-`" means the value is negative; and the lack of a sign means the value may be positive or negative or has no sign (NaN).)

- Decimal multiplication:

```
C#

decimal operator *(decimal x, decimal y);
```

If the magnitude of the resulting value is too large to represent in the decimal format, a `System.OverflowException` is thrown. Because of rounding, the result may be zero even though neither operand is zero. The scale of the result, before any

rounding, is the sum of the scales of the two operands. Decimal multiplication is equivalent to using the multiplication operator of type `System.Decimal`.

Lifted (§12.4.8) forms of the unlifted predefined multiplication operators defined above are also predefined.

12.10.3 Division operator

For an operation of the form `x / y`, binary operator overload resolution (§12.4.5) is applied to select a specific operator implementation. The operands are converted to the parameter types of the selected operator, and the type of the result is the return type of the operator.

The predefined division operators are listed below. The operators all compute the quotient of `x` and `y`.

- Integer division:

  C#

  ```
  int operator /(int x, int y);
  uint operator /(uint x, uint y);
  long operator /(long x, long y);
  ulong operator /(ulong x, ulong y);
  ```

  If the value of the right operand is zero, a `System.DivideByZeroException` is thrown.

  The division rounds the result towards zero. Thus the absolute value of the result is the largest possible integer that is less than or equal to the absolute value of the quotient of the two operands. The result is zero or positive when the two operands have the same sign and zero or negative when the two operands have opposite signs.

  If the left operand is the smallest representable `int` or `long` value and the right operand is `-1`, an overflow occurs. In a `checked` context, this causes a `System.ArithmeticException` (or a subclass thereof) to be thrown. In an `unchecked` context, it is implementation-defined as to whether a `System.ArithmeticException` (or a subclass thereof) is thrown or the overflow goes unreported with the resulting value being that of the left operand.

- Floating-point division:

  C#

```
float operator /(float x, float y);
double operator /(double x, double y);
```

The quotient is computed according to the rules of IEC 60559 arithmetic. The following table lists the results of all possible combinations of nonzero finite values, zeros, infinities, and NaNs. In the table, `x` and `y` are positive finite values. `z` is the result of `x / y`, rounded to the nearest representable value.

Expand table

	+y	-y	+0	-0	+∞	-∞	NaN
+x	+z	-z	+∞	-∞	+0	-0	NaN
-x	-z	+z	-∞	+∞	-0	+0	NaN
+0	+0	-0	NaN	NaN	+0	-0	NaN
-0	-0	+0	NaN	NaN	-0	+0	NaN
+∞	+∞	-∞	+∞	-∞	NaN	NaN	NaN
-∞	-∞	+∞	-∞	+∞	NaN	NaN	NaN
NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN

- Decimal division:

C#

```
decimal operator /(decimal x, decimal y);
```

If the value of the right operand is zero, a `System.DivideByZeroException` is thrown. If the magnitude of the resulting value is too large to represent in the decimal format, a `System.OverflowException` is thrown. Because of rounding, the result may be zero even though the first operand is not zero. The scale of the result, before any rounding, is the closest scale to the preferred scale that will preserve a result equal to the exact result. The preferred scale is the scale of `x` less the scale of `y`.

Decimal division is equivalent to using the division operator of type `System.Decimal`.

Lifted (§12.4.8) forms of the unlifted predefined division operators defined above are also predefined.

12.10.4 Remainder operator

For an operation of the form `x % y`, binary operator overload resolution (§12.4.5) is applied to select a specific operator implementation. The operands are converted to the parameter types of the selected operator, and the type of the result is the return type of the operator.

The predefined remainder operators are listed below. The operators all compute the remainder of the division between `x` and `y`.

- Integer remainder:

  ```
  C#

  int operator %(int x, int y);
  uint operator %(uint x, uint y);
  long operator %(long x, long y);
  ulong operator %(ulong x, ulong y);
  ```

  The result of `x % y` is the value produced by `x - (x / y) * y`. If `y` is zero, a `System.DivideByZeroException` is thrown.

  If the left operand is the smallest `int` or `long` value and the right operand is `-1`, a `System.OverflowException` is thrown if and only if `x / y` would throw an exception.

- Floating-point remainder:

  ```
  C#

  float operator %(float x, float y);
  double operator %(double x, double y);
  ```

  The following table lists the results of all possible combinations of nonzero finite values, zeros, infinities, and NaNs. In the table, `x` and `y` are positive finite values. `z` is the result of `x % y` and is computed as `x - n * y`, where n is the largest possible integer that is less than or equal to `x / y`. This method of computing the remainder is analogous to that used for integer operands, but differs from the IEC 60559 definition (in which `n` is the integer closest to `x / y`).

  Expand table

	+y	-y	+0	-0	+∞	-∞	NaN
+x	+z	+z	NaN	NaN	+x	+x	NaN

	+y	-y	+0	-0	+∞	-∞	NaN
-x	-z	-z	NaN	NaN	-x	-x	NaN
+0	+0	+0	NaN	NaN	+0	+0	NaN
-0	-0	-0	NaN	NaN	-0	-0	NaN
+∞	NaN	NaN	NaN	NaN	NaN	NaN	NaN
-∞	NaN	NaN	NaN	NaN	NaN	NaN	NaN
NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN

- Decimal remainder:

```
C#

decimal operator %(decimal x, decimal y);
```

If the value of the right operand is zero, a `System.DivideByZeroException` is thrown. It is implementation-defined when a `System.ArithmeticException` (or a subclass thereof) is thrown. A conforming implementation shall not throw an exception for `x % y` in any case where `x / y` does not throw an exception. The scale of the result, before any rounding, is the larger of the scales of the two operands, and the sign of the result, if non-zero, is the same as that of `x`.

Decimal remainder is equivalent to using the remainder operator of type `System.Decimal`.

> *Note*: These rules ensure that for all types, the result never has the opposite sign of the left operand. *end note*

Lifted (§12.4.8) forms of the unlifted predefined remainder operators defined above are also predefined.

12.10.5 Addition operator

For an operation of the form `x + y`, binary operator overload resolution (§12.4.5) is applied to select a specific operator implementation. The operands are converted to the parameter types of the selected operator, and the type of the result is the return type of the operator.

The predefined addition operators are listed below. For numeric and enumeration types, the predefined addition operators compute the sum of the two operands. When one or both operands are of type `string`, the predefined addition operators concatenate the string representation of the operands.

- Integer addition:

```
C#

int operator +(int x, int y);
uint operator +(uint x, uint y);
long operator +(long x, long y);
ulong operator +(ulong x, ulong y
```

  In a `checked` context, if the sum is outside the range of the result type, a `System.OverflowException` is thrown. In an `unchecked` context, overflows are not reported and any significant high-order bits outside the range of the result type are discarded.

- Floating-point addition:

```
C#

float operator +(float x, float y);
double operator +(double x, double y);
```

  The sum is computed according to the rules of IEC 60559 arithmetic. The following table lists the results of all possible combinations of nonzero finite values, zeros, infinities, and NaNs. In the table, `x` and `y` are nonzero finite values, and `z` is the result of `x + y`. If `x` and `y` have the same magnitude but opposite signs, `z` is positive zero. If `x + y` is too large to represent in the destination type, `z` is an infinity with the same sign as `x + y`.

Expand table

	y	**+0**	**-0**	**+∞**	**-∞**	**NaN**
x	z	x	x	+∞	-∞	NaN
+0	y	+0	+0	+∞	−∞	NaN
-0	y	+0	-0	+∞	-∞	NaN
+∞	+∞	+∞	+∞	+∞	NaN	NaN

	y	+0	-0	+∞	-∞	NaN
-∞	-∞	-∞	-∞	NaN	-∞	NaN
NaN	NaN	NaN	NaN	NaN	NaN	NaN

- Decimal addition:

```
C#

decimal operator +(decimal x, decimal y);
```

If the magnitude of the resulting value is too large to represent in the decimal format, a `System.OverflowException` is thrown. The scale of the result, before any rounding, is the larger of the scales of the two operands.

Decimal addition is equivalent to using the addition operator of type `System.Decimal`.

- Enumeration addition. Every enumeration type implicitly provides the following predefined operators, where `E` is the enum type, and `U` is the underlying type of `E`:

```
C#

E operator +(E x, U y);
E operator +(U x, E y);
```

At run-time these operators are evaluated exactly as `(E)((U)x + (U)y)`.

- String concatenation:

```
C#

string operator +(string x, string y);
string operator +(string x, object y);
string operator +(object x, string y);
```

These overloads of the binary `+` operator perform string concatenation. If an operand of string concatenation is `null`, an empty string is substituted. Otherwise, any non-`string` operand is converted to its string representation by invoking the virtual `ToString` method inherited from type `object`. If `ToString` returns `null`, an empty string is substituted.

Example:

```
C#
```

```
class Test
{
    static void Main()
    {
        string s = null;
        Console.WriteLine("s = >" + s + "<");  // Displays s = ><

        int i = 1;
        Console.WriteLine("i = " + i);         // Displays i = 1

        float f = 1.2300E+15F;
        Console.WriteLine("f = " + f);         // Displays f =
1.23E+15

        decimal d = 2.900m;
        Console.WriteLine("d = " + d);         // Displays d =
2.900
   }
}
```

The output shown in the comments is the typical result on a US-English system. The precise output might depend on the regional settings of the execution environment. The string-concatenation operator itself behaves the same way in each case, but the `ToString` methods implicitly called during execution might be affected by regional settings.

end example

The result of the string concatenation operator is a `string` that consists of the characters of the left operand followed by the characters of the right operand. The string concatenation operator never returns a `null` value. A `System.OutOfMemoryException` may be thrown if there is not enough memory available to allocate the resulting string.

- Delegate combination. Every delegate type implicitly provides the following predefined operator, where `D` is the delegate type:

```
C#
```

```
D operator +(D x, D y);
```

If the first operand is `null`, the result of the operation is the value of the second operand (even if that is also `null`). Otherwise, if the second operand is `null`, then the result of the operation is the value of the first operand. Otherwise, the result of the operation is a new delegate instance whose invocation list consists of the elements in the invocation list of the first operand, followed by the elements in the invocation list of the second operand. That is, the invocation list of the resulting delegate is the concatenation of the invocation lists of the two operands.

> *Note*: For examples of delegate combination, see §12.10.6 and §20.6. Since `System.Delegate` is not a delegate type, operator + is not defined for it. *end note*

Lifted (§12.4.8) forms of the unlifted predefined addition operators defined above are also predefined.

12.10.6 Subtraction operator

For an operation of the form `x - y`, binary operator overload resolution (§12.4.5) is applied to select a specific operator implementation. The operands are converted to the parameter types of the selected operator, and the type of the result is the return type of the operator.

The predefined subtraction operators are listed below. The operators all subtract `y` from `x`.

- Integer subtraction:

  C#

  ```
  int operator -(int x, int y);
  uint operator -(uint x, uint y);
  long operator -(long x, long y);
  ulong operator -(ulong x, ulong y
  ```

  In a `checked` context, if the difference is outside the range of the result type, a `System.OverflowException` is thrown. In an `unchecked` context, overflows are not reported and any significant high-order bits outside the range of the result type are discarded.

- Floating-point subtraction:

  C#

```
float operator -(float x, float y);
double operator -(double x, double y);
```

The difference is computed according to the rules of IEC 60559 arithmetic. The following table lists the results of all possible combinations of nonzero finite values, zeros, infinities, and NaNs. In the table, `x` and `y` are nonzero finite values, and `z` is the result of `x - y`. If `x` and `y` are equal, `z` is positive zero. If `x - y` is too large to represent in the destination type, `z` is an infinity with the same sign as `x - y`.

Expand table

	y	**+0**	**-0**	**+∞**	**-∞**	**NaN**
x	z	x	x	-∞	+∞	NaN
+0	-y	+0	+0	-∞	+∞	NaN
-0	-y	-0	+0	-∞	+∞	NaN
+∞	+∞	+∞	+∞	NaN	+∞	NaN
-∞	-∞	-∞	-∞	-∞	NaN	NaN
NaN	NaN	NaN	NaN	NaN	NaN	NaN

(In the above table, the `-y` entries denote the *negation* of `y`, not that the value is negative.)

- Decimal subtraction:

```
C#
```

```
decimal operator -(decimal x, decimal y);
```

If the magnitude of the resulting value is too large to represent in the decimal format, a `System.OverflowException` is thrown. The scale of the result, before any rounding, is the larger of the scales of the two operands.

Decimal subtraction is equivalent to using the subtraction operator of type `System.Decimal`.

- Enumeration subtraction. Every enumeration type implicitly provides the following predefined operator, where `E` is the enum type, and `U` is the underlying type of `E`:

```
C#
```

```
U operator -(E x, E y);
```

This operator is evaluated exactly as `(U)((U)x - (U)y)`. In other words, the operator computes the difference between the ordinal values of `x` and `y`, and the type of the result is the underlying type of the enumeration.

```
C#
```

```
E operator -(E x, U y);
```

This operator is evaluated exactly as `(E)((U)x - y)`. In other words, the operator subtracts a value from the underlying type of the enumeration, yielding a value of the enumeration.

- Delegate removal. Every delegate type implicitly provides the following predefined operator, where `D` is the delegate type:

```
C#
```

```
D operator -(D x, D y);
```

The semantics are as follows:

- If the first operand is `null`, the result of the operation is `null`.
- Otherwise, if the second operand is `null`, then the result of the operation is the value of the first operand.
- Otherwise, both operands represent non-empty invocation lists (§20.2).
  - If the lists compare equal, as determined by the delegate equality operator (§12.12.9), the result of the operation is `null`.
  - Otherwise, the result of the operation is a new invocation list consisting of the first operand's list with the second operand's entries removed from it, provided the second operand's list is a sublist of the first's. (To determine sublist equality, corresponding entries are compared as for the delegate equality operator.) If the second operand's list matches multiple sublists of contiguous entries in the first operand's list, the last matching sublist of contiguous entries is removed.
  - Otherwise, the result of the operation is the value of the left operand.

Neither of the operands' lists (if any) is changed in the process.

> *Example*:

C#

```
delegate void D(int x);

class C
{
    public static void M1(int i) { ... }
    public static void M2(int i) { ... }
}

class Test
{
    static void Main()
    {
        D cd1 = new D(C.M1);
        D cd2 = new D(C.M2);
        D list = null;

        list = null - cd1;                              // null
        list = (cd1 + cd2 + cd2 + cd1) - null;          // M1 + M2 +
M2 + M1
        list = (cd1 + cd2 + cd2 + cd1) - cd1;           // M1 + M2 +
M2
        list = (cd1 + cd2 + cd2 + cd1) - (cd1 + cd2);   // M2 + M1
        list = (cd1 + cd2 + cd2 + cd1) - (cd2 + cd2);   // M1 + M1
        list = (cd1 + cd2 + cd2 + cd1) - (cd2 + cd1);   // M1 + M2
        list = (cd1 + cd2 + cd2 + cd1) - (cd1 + cd1);   // M1 + M2 +
M2 + M1
        list = (cd1 + cd2 + cd2 + cd1) - (cd1 + cd2 + cd2 + cd1);
// null
    }
}
```

end example

Lifted (§12.4.8) forms of the unlifted predefined subtraction operators defined above are also predefined.

12.11 Shift operators

The `<<` and `>>` operators are used to perform bit-shifting operations.

ANTLR

```
shift_expression
    : additive_expression
    | shift_expression '<<' additive_expression
    | shift_expression right_shift additive_expression
    ;
```

If an operand of a *shift_expression* has the compile-time type `dynamic`, then the expression is dynamically bound (§12.3.3). In this case, the compile-time type of the expression is `dynamic`, and the resolution described below will take place at run-time using the run-time type of those operands that have the compile-time type `dynamic`.

For an operation of the form `x << count` or `x >> count`, binary operator overload resolution (§12.4.5) is applied to select a specific operator implementation. The operands are converted to the parameter types of the selected operator, and the type of the result is the return type of the operator.

When declaring an overloaded shift operator, the type of the first operand shall always be the class or struct containing the operator declaration, and the type of the second operand shall always be `int`.

The predefined shift operators are listed below.

- Shift left:

  ```
  C#
  ```

  ```
  int operator <<(int x, int count);
  uint operator <<(uint x, int count);
  long operator <<(long x, int count);
  ulong operator <<(ulong x, int count);
  ```

  The `<<` operator shifts `x` left by a number of bits computed as described below.

  The high-order bits outside the range of the result type of `x` are discarded, the remaining bits are shifted left, and the low-order empty bit positions are set to zero.

- Shift right:

  ```
  C#
  ```

  ```
  int operator >>(int x, int count);
  uint operator >>(uint x, int count);
  long operator >>(long x, int count);
  ulong operator >>(ulong x, int count);
  ```

  The `>>` operator shifts `x` right by a number of bits computed as described below.

  When `x` is of type `int` or `long`, the low-order bits of `x` are discarded, the remaining bits are shifted right, and the high-order empty bit positions are set to zero if `x` is non-negative and set to one if `x` is negative.

When `x` is of type `uint` or `ulong`, the low-order bits of `x` are discarded, the remaining bits are shifted right, and the high-order empty bit positions are set to zero.

For the predefined operators, the number of bits to shift is computed as follows:

- When the type of `x` is `int` or `uint`, the shift count is given by the low-order five bits of `count`. In other words, the shift count is computed from `count & 0x1F`.
- When the type of `x` is `long` or `ulong`, the shift count is given by the low-order six bits of `count`. In other words, the shift count is computed from `count & 0x3F`.

If the resulting shift count is zero, the shift operators simply return the value of `x`.

Shift operations never cause overflows and produce the same results in checked and unchecked contexts.

When the left operand of the `>>` operator is of a signed integral type, the operator performs an *arithmetic* shift right wherein the value of the most significant bit (the sign bit) of the operand is propagated to the high-order empty bit positions. When the left operand of the `>>` operator is of an unsigned integral type, the operator performs a *logical* shift right wherein high-order empty bit positions are always set to zero. To perform the opposite operation of that inferred from the operand type, explicit casts can be used.

> *Example*: If `x` is a variable of type `int`, the operation `unchecked ((int)((uint)x >> y))` performs a logical shift right of `x`. *end example*

Lifted (§12.4.8) forms of the unlifted predefined shift operators defined above are also predefined.

12.12 Relational and type-testing operators

12.12.1 General

The `==`, `!=`, `<`, `>`, `<=`, `>=`, `is`, and `as` operators are called the relational and type-testing operators.

ANTLR

```
relational_expression
    : shift_expression
    | relational_expression '<' shift_expression
```

```
    | relational_expression '>' shift_expression
    | relational_expression '<=' shift_expression
    | relational_expression '>=' shift_expression
    | relational_expression 'is' type
    | relational_expression 'is' pattern
    | relational_expression 'as' type
    ;

equality_expression
    : relational_expression
    | equality_expression '==' relational_expression
    | equality_expression '!=' relational_expression
    ;
```

> *Note*: Lookup for the right operand of the `is` operator must first test as a *type*, then as an *expression* which may span multiple tokens. In the case where the operand is an *expreesion*, the pattern expression must have precedence at least as high as *shift_expression*. *end note*

The `is` operator is described in §12.12.12 and the `as` operator is described in §12.12.13.

The `==`, `!=`, `<`, `>`, `<=` and `>=` operators are ***comparison operators***.

If a *default_literal* (§12.8.20) is used as an operand of a `<`, `>`, `<=`, or `>=` operator, a compile-time error occurs. If a *default_literal* is used as both operands of a `==` or `!=` operator, a compile-time error occurs. If a *default_literal* is used as the left operand of the `is` or `as` operator, a compile-time error occurs.

If an operand of a comparison operator has the compile-time type `dynamic`, then the expression is dynamically bound (§12.3.3). In this case the compile-time type of the expression is `dynamic`, and the resolution described below will take place at run-time using the run-time type of those operands that have the compile-time type `dynamic`.

For an operation of the form `x «op» y`, where «op» is a comparison operator, overload resolution (§12.4.5) is applied to select a specific operator implementation. The operands are converted to the parameter types of the selected operator, and the type of the result is the return type of the operator. If both operands of an *equality_expression* are the `null` literal, then overload resolution is not performed and the expression evaluates to a constant value of `true` or `false` according to whether the operator is `==` or `!=`.

The predefined comparison operators are described in the following subclauses. All predefined comparison operators return a result of type bool, as described in the following table.

Expand table

Operation	Result
`x == y`	`true` if `x` is equal to `y`, `false` otherwise
`x != y`	`true` if `x` is not equal to `y`, `false` otherwise
`x < y`	`true` if `x` is less than `y`, `false` otherwise
`x > y`	`true` if `x` is greater than `y`, `false` otherwise
`x <= y`	`true` if `x` is less than or equal to `y`, `false` otherwise
`x >= y`	`true` if `x` is greater than or equal to `y`, `false` otherwise

12.12.2 Integer comparison operators

The predefined integer comparison operators are:

C#

```
bool operator ==(int x, int y);
bool operator ==(uint x, uint y);
bool operator ==(long x, long y);
bool operator ==(ulong x, ulong y);

bool operator !=(int x, int y);
bool operator !=(uint x, uint y);
bool operator !=(long x, long y);
bool operator !=(ulong x, ulong y);

bool operator <(int x, int y);
bool operator <(uint x, uint y);
bool operator <(long x, long y);
bool operator <(ulong x, ulong y);

bool operator >(int x, int y);
bool operator >(uint x, uint y);
bool operator >(long x, long y);
bool operator >(ulong x, ulong y);

bool operator <=(int x, int y);
bool operator <=(uint x, uint y);
bool operator <=(long x, long y);
bool operator <=(ulong x, ulong y);

bool operator >=(int x, int y);
bool operator >=(uint x, uint y);
bool operator >=(long x, long y);
bool operator >=(ulong x, ulong y);
```

Each of these operators compares the numeric values of the two integer operands and returns a `bool` value that indicates whether the particular relation is `true` or `false`.

Lifted (§12.4.8) forms of the unlifted predefined integer comparison operators defined above are also predefined.

12.12.3 Floating-point comparison operators

The predefined floating-point comparison operators are:

```
C#
```

```
bool operator ==(float x, float y);
bool operator ==(double x, double y);

bool operator !=(float x, float y);
bool operator !=(double x, double y);

bool operator <(float x, float y);
bool operator <(double x, double y);

bool operator >(float x, float y);
bool operator >(double x, double y);

bool operator <=(float x, float y);
bool operator <=(double x, double y);

bool operator >=(float x, float y);
bool operator >=(double x, double y);
```

The operators compare the operands according to the rules of the IEC 60559 standard:

If either operand is NaN, the result is `false` for all operators except `!=`, for which the result is `true`. For any two operands, `x != y` always produces the same result as `!(x == y)`. However, when one or both operands are NaN, the `<`, `>`, `<=`, and `>=` operators do *not* produce the same results as the logical negation of the opposite operator.

> *Example*: If either of `x` and `y` is NaN, then `x < y` is `false`, but `!(x >= y)` is `true`. *end example*

When neither operand is NaN, the operators compare the values of the two floating-point operands with respect to the ordering

```
C#
```

```
-∞ < -max < ... < -min < -0.0 == +0.0 < +min < ... < +max < +∞
```

where `min` and `max` are the smallest and largest positive finite values that can be represented in the given floating-point format. Notable effects of this ordering are:

- Negative and positive zeros are considered equal.
- A negative infinity is considered less than all other values, but equal to another negative infinity.
- A positive infinity is considered greater than all other values, but equal to another positive infinity.

Lifted (§12.4.8) forms of the unlifted predefined floating-point comparison operators defined above are also predefined.

12.12.4 Decimal comparison operators

The predefined decimal comparison operators are:

C#

```
bool operator ==(decimal x, decimal y);
bool operator !=(decimal x, decimal y);
bool operator <(decimal x, decimal y);
bool operator >(decimal x, decimal y);
bool operator <=(decimal x, decimal y);
bool operator >=(decimal x, decimal y);
```

Each of these operators compares the numeric values of the two decimal operands and returns a `bool` value that indicates whether the particular relation is `true` or `false`. Each decimal comparison is equivalent to using the corresponding relational or equality operator of type `System.Decimal`.

Lifted (§12.4.8) forms of the unlifted predefined decimal comparison operators defined above are also predefined.

12.12.5 Boolean equality operators

The predefined Boolean equality operators are:

C#

```
bool operator ==(bool x, bool y);
```

```
bool operator !=(bool x, bool y);
```

The result of `==` is `true` if both `x` and `y` are `true` or if both `x` and `y` are `false`. Otherwise, the result is `false`.

The result of `!=` is `false` if both `x` and `y` are `true` or if both `x` and `y` are `false`. Otherwise, the result is `true`. When the operands are of type `bool`, the `!=` operator produces the same result as the `^` operator.

Lifted (§12.4.8) forms of the unlifted predefined Boolean equality operators defined above are also predefined.

12.12.6 Enumeration comparison operators

Every enumeration type implicitly provides the following predefined comparison operators

C#

```
bool operator ==(E x, E y);
bool operator !=(E x, E y);

bool operator <(E x, E y);
bool operator >(E x, E y);
bool operator <=(E x, E y);
bool operator >=(E x, E y);
```

The result of evaluating `x «op» y`, where x and y are expressions of an enumeration type `E` with an underlying type `U`, and «op» is one of the comparison operators, is exactly the same as evaluating `((U)x) «op» ((U)y)`. In other words, the enumeration type comparison operators simply compare the underlying integral values of the two operands.

Lifted (§12.4.8) forms of the unlifted predefined enumeration comparison operators defined above are also predefined.

12.12.7 Reference type equality operators

Every class type `C` implicitly provides the following predefined reference type equality operators:

C#

```
bool operator ==(C x, C y);
bool operator !=(C x, C y);
```

unless predefined equality operators otherwise exist for `C` (for example, when `C` is `string` or `System.Delegate`).

The operators return the result of comparing the two references for equality or non-equality. `operator ==` returns `true` if and only if `x` and `y` refer to the same instance or are both `null`, while `operator !=` returns `true` if and only if `operator ==` with the same operands would return `false`.

In addition to normal applicability rules (§12.6.4.2), the predefined reference type equality operators require one of the following in order to be applicable:

- Both operands are a value of a type known to be a *reference_type* or the literal `null`. Furthermore, an identity or explicit reference conversion (§10.3.5) exists from either operand to the type of the other operand.
- One operand is the literal `null`, and the other operand is a value of type `T` where `T` is a *type_parameter* that is not known to be a value type, and does not have the value type constraint.
  - If at runtime `T` is a non-nullable value type, the result of `==` is `false` and the result of `!=` is `true`.
  - If at runtime `T` is a nullable value type, the result is computed from the `HasValue` property of the operand, as described in (§12.12.10).
  - If at runtime `T` is a reference type, the result is `true` if the operand is `null`, and `false` otherwise.

Unless one of these conditions is true, a binding-time error occurs.

> *Note*: Notable implications of these rules are:
>
> - It is a binding-time error to use the predefined reference type equality operators to compare two references that are known to be different at binding-time. For example, if the binding-time types of the operands are two class types, and if neither derives from the other, then it would be impossible for the two operands to reference the same object. Thus, the operation is considered a binding-time error.
> - The predefined reference type equality operators do not permit value type operands to be compared (except when type parameters are compared to `null`, which is handled specially).

> - Operands of predefined reference type equality operators are never boxed. It would be meaningless to perform such boxing operations, since references to the newly allocated boxed instances would necessarily differ from all other references.
>
> For an operation of the form `x == y` or `x != y`, if any applicable user-defined `operator ==` or `operator !=` exists, the operator overload resolution rules (§12.4.5) will select that operator instead of the predefined reference type equality operator. It is always possible to select the predefined reference type equality operator by explicitly casting one or both of the operands to type `object`.
>
> *end note*

> *Example*: The following example checks whether an argument of an unconstrained type parameter type is `null`.
>
> C#
>
> ```
> class C<T>
> {
>    void F(T x)
>    {
>       if (x == null)
>       {
>           throw new ArgumentNullException();
>       }
>       ...
>    }
> }
> ```
>
> The `x == null` construct is permitted even though `T` could represent a non-nullable value type, and the result is simply defined to be `false` when `T` is a non-nullable value type.
>
> *end example*

For an operation of the form `x == y` or `x != y`, if any applicable `operator ==` or `operator !=` exists, the operator overload resolution (§12.4.5) rules will select that operator instead of the predefined reference type equality operator.

> *Note*: It is always possible to select the predefined reference type equality operator by explicitly casting both of the operands to type `object`. *end note*

Example: The example

C#

```
class Test
{
    static void Main()
    {
        string s = "Test";
        string t = string.Copy(s);
        Console.WriteLine(s == t);
        Console.WriteLine((object)s == t);
        Console.WriteLine(s == (object)t);
        Console.WriteLine((object)s == (object)t);
    }
}
```

produces the output

Console

```
True
False
False
False
```

The `s` and `t` variables refer to two distinct string instances containing the same characters. The first comparison outputs `True` because the predefined string equality operator (§12.12.8) is selected when both operands are of type `string`. The remaining comparisons all output `False` because the overload of `operator ==` in the `string` type is not applicable when either operand has a binding-time type of `object`.

Note that the above technique is not meaningful for value types. The example

C#

```
class Test
{
    static void Main()
    {
        int i = 123;
        int j = 123;
        Console.WriteLine((object)i == (object)j);
    }
}
```

> outputs `False` because the casts create references to two separate instances of boxed `int` values.
>
> *end example*

12.12.8 String equality operators

The predefined string equality operators are:

```
C#
```

```
bool operator ==(string x, string y);
bool operator !=(string x, string y);
```

Two `string` values are considered equal when one of the following is true:

- Both values are `null`.
- Both values are non-`null` references to string instances that have identical lengths and identical characters in each character position.

The string equality operators compare string values rather than string references. When two separate string instances contain the exact same sequence of characters, the values of the strings are equal, but the references are different.

> *Note*: As described in §12.12.7, the reference type equality operators can be used to compare string references instead of string values. *end note*

12.12.9 Delegate equality operators

The predefined delegate equality operators are:

```
C#
```

```
bool operator ==(System.Delegate x, System.Delegate y);
bool operator !=(System.Delegate x, System.Delegate y);
```

Two delegate instances are considered equal as follows:

- If either of the delegate instances is `null`, they are equal if and only if both are `null`.
- If the delegates have different run-time type, they are never equal.
- If both of the delegate instances have an invocation list (§20.2), those instances are equal if and only if their invocation lists are the same length, and each entry in

one's invocation list is equal (as defined below) to the corresponding entry, in order, in the other's invocation list.

The following rules govern the equality of invocation list entries:

- If two invocation list entries both refer to the same static method then the entries are equal.
- If two invocation list entries both refer to the same non-static method on the same target object (as defined by the reference equality operators) then the entries are equal.
- Invocation list entries produced from evaluation of semantically identical anonymous functions (§12.19) with the same (possibly empty) set of captured outer variable instances are permitted (but not required) to be equal.

If operator overload resolution resolves to either delegate equality operator, and the binding-time types of both operands are delegate types as described in §20 rather than `System.Delegate`, and there is no identity conversion between the binding-type operand types, a binding-time error occurs.

> *Note*: This rule prevents comparisons which can never consider non-`null` values as equal due to being references to instances of different types of delegates. *end note*

12.12.10 Equality operators between nullable value types and the null literal

The `==` and `!=` operators permit one operand to be a value of a nullable value type and the other to be the `null` literal, even if no predefined or user-defined operator (in unlifted or lifted form) exists for the operation.

For an operation of one of the forms

```
C#
x == null    null == x    x != null    null != x
```

where `x` is an expression of a nullable value type, if operator overload resolution (§12.4.5) fails to find an applicable operator, the result is instead computed from the `HasValue` property of `x`. Specifically, the first two forms are translated into `!x.HasValue`, and the last two forms are translated into `x.HasValue`.

12.12.11 Tuple equality operators

The tuple equality operators are applied pairwise to the elements of the tuple operands in lexical order.

If each operand `x` and `y` of a `==` or `!=` operator is classified either as a tuple or as a value with a tuple type (§8.3.11), the operator is a *tuple equality operator*.

If an operand `e` is classified as a tuple, the elements `e1...en` shall be the results of evaluating the element expressions of the tuple expression. Otherwise if `e` is a value of a tuple type, the elements shall be `t.Item1...t.Itemn` where `t` is the result of evaluating `e`.

The operands `x` and `y` of a tuple equality operator shall have the same arity, or a compile time error occurs. For each pair of elements `xi` and `yi`, the same equality operator shall apply, and shall yield a result of type `bool`, `dynamic`, a type that has an implicit conversion to `bool`, or a type that defines the `true` and `false` operators.

The tuple equality operator `x == y` is evaluated as follows:

- The left side operand `x` is evaluated.
- The right side operand `y` is evaluated.
- For each pair of elements `xi` and `yi` in lexical order:
  - The operator `xi == yi` is evaluated, and a result of type `bool` is obtained in the following way:
    - If the comparison yielded a `bool` then that is the result.
    - Otherwise if the comparison yielded a `dynamic` then the operator `false` is dynamically invoked on it, and the resulting `bool` value is negated with the `!` operator.
    - Otherwise, if the type of the comparison has an implicit conversion to `bool`, that conversion is applied.
    - Otherwise, if the type of the comparison has an operator `false`, that operator is invoked and the resulting `bool` value is negated with the `!` operator.
  - If the resulting `bool` is `false`, then no further evaluation occurs, and the result of the tuple equality operator is `false`.
- If all element comparisons yielded `true`, the result of the tuple equality operator is `true`.

The tuple equality operator `x != y` is evaluated as follows:

- The left side operand `x` is evaluated.
- The right side operand `y` is evaluated.

- For each pair of elements `xi` and `yi` in lexical order:
  - The operator `xi != yi` is evaluated, and a result of type `bool` is obtained in the following way:
    - If the comparison yielded a `bool` then that is the result.
    - Otherwise if the comparison yielded a `dynamic` then the operator `true` is dynamically invoked on it, and the resulting `bool` value is the result.
    - Otherwise, if the type of the comparison has an implicit conversion to `bool`, that conversion is applied.
    - Otherwise, if the type of the comparison has an operator `true`, that operator is invoked and the resulting `bool` value is the result.
  - If the resulting `bool` is `true`, then no further evaluation occurs, and the result of the tuple equality operator is `true`.
- If all element comparisons yielded `false`, the result of the tuple equality operator is `false`.

12.12.12 The is operator

There are two forms of the `is` operator. One is the *is-type operator*, which has a type on the right-hand-side. The other is the *is-pattern operator*, which has a pattern on the right-hand-side.

12.12.12.1 The is-type operator

The *is-type operator* is used to check if the run-time type of an object is compatible with a given type. The check is performed at runtime. The result of the operation `E is T`, where `E` is an expression and `T` is a type other than `dynamic`, is a Boolean value indicating whether `E` is non-null and can successfully be converted to type `T` by a reference conversion, a boxing conversion, an unboxing conversion, a wrapping conversion, or an unwrapping conversion.

The operation is evaluated as follows:

1. If `E` is an anonymous function or method group, a compile-time error occurs
2. If `E` is the `null` literal, or if the value of `E` is `null`, the result is `false`.
3. Otherwise:
4. Let `R` be the runtime type of `E`.
5. Let `D` be derived from `R` as follows:
6. If `R` is a nullable value type, `D` is the underlying type of `R`.
7. Otherwise, `D` is `R`.

8. The result depends on `D` and `T` as follows:
9. If `T` is a reference type, the result is `true` if:

    - an identity conversion exists between `D` and `T`,
    - `D` is a reference type and an implicit reference conversion from `D` to `T` exists, or
    - Either: `D` is a value type and a boxing conversion from `D` to `T` exists.
      Or: `D` is a value type and `T` is an interface type implemented by `D`.

10. If `T` is a nullable value type, the result is `true` if `D` is the underlying type of `T`.
11. If `T` is a non-nullable value type, the result is `true` if `D` and `T` are the same type.
12. Otherwise, the result is `false`.

User defined conversions are not considered by the `is` operator.

> *Note*: As the `is` operator is evaluated at runtime, all type arguments have been substituted and there are no open types (§8.4.3) to consider. *end note*

> *Note*: The `is` operator can be understood in terms of compile-time types and conversions as follows, where `C` is the compile-time type of `E`:
>
> - If the compile-time type of `e` is the same as `T`, or if an implicit reference conversion (§10.2.8), boxing conversion (§10.2.9), wrapping conversion (§10.6), or an explicit unwrapping conversion (§10.6) exists from the compile-time type of `E` to `T`:
>   - If `C` is of a non-nullable value type, the result of the operation is `true`.
>   - Otherwise, the result of the operation is equivalent to evaluating `E != null`.
> - Otherwise, if an explicit reference conversion (§10.3.5) or unboxing conversion (§10.3.7) exists from `C` to `T`, or if `C` or `T` is an open type (§8.4.3), then runtime checks as above shall be peformed.
> - Otherwise, no reference, boxing, wrapping, or unwrapping conversion of `E` to type `T` is possible, and the result of the operation is `false`. A compiler may implement optimisations based on the compile-time type.
>
> *end note*

12.12.12.2 The is-pattern operator

The *is-pattern operator* is used to check if the value computed by an expression *matches* a given pattern (§11). The check is performed at runtime. The result of the is-pattern operator is true if the value matches the pattern; otherwise it is false.

For an expression of the form `E is P`, where `E` is a relational expression of type `T` and `P` is a pattern, it is a compile-time error if any of the following hold:

- `E` does not designate a value or does not have a type.
- The pattern `P` is not applicable (§11.2) to the type `T`.

12.12.13 The as operator

The `as` operator is used to explicitly convert a value to a given reference type or nullable value type. Unlike a cast expression (§12.9.7), the `as` operator never throws an exception. Instead, if the indicated conversion is not possible, the resulting value is `null`.

In an operation of the form `E as T`, `E` shall be an expression and `T` shall be a reference type, a type parameter known to be a reference type, or a nullable value type. Furthermore, at least one of the following shall be true, or otherwise a compile-time error occurs:

- An identity (§10.2.2), implicit nullable (§10.2.6), implicit reference (§10.2.8), boxing (§10.2.9), explicit nullable (§10.3.4), explicit reference (§10.3.5), or wrapping (§8.3.12) conversion exists from `E` to `T`.
- The type of `E` or `T` is an open type.
- `E` is the `null` literal.

If the compile-time type of `E` is not `dynamic`, the operation `E as T` produces the same result as

```
C#
E is T ? (T)(E) : (T)null
```

except that `E` is only evaluated once. The compiler can be expected to optimize `E as T` to perform at most one runtime type check as opposed to the two runtime type checks implied by the expansion above.

If the compile-time type of `E` is `dynamic`, unlike the cast operator the `as` operator is not dynamically bound (§12.3.3). Therefore the expansion in this case is:

```
C#
E is T ? (T)(object)(E) : (T)null
```

Note that some conversions, such as user defined conversions, are not possible with the `as` operator and should instead be performed using cast expressions.

> *Example*: In the example
>
> ```C#
> class X
> {
>     public string F(object o)
>     {
>         return o as string;  // OK, string is a reference type
>     }
>
>     public T G<T>(object o)
>         where T : Attribute
>     {
>         return o as T;       // Ok, T has a class constraint
>     }
>
>     public U H<U>(object o)
>     {
>         return o as U;       // Error, U is unconstrained
>     }
> }
> ```
>
> the type parameter `T` of `G` is known to be a reference type, because it has the class constraint. The type parameter `U` of `H` is not however; hence the use of the `as` operator in `H` is disallowed.
>
> *end example*

12.13 Logical operators

12.13.1 General

The `&`, `^`, and `|` operators are called the logical operators.

```ANTLR
and_expression
    : equality_expression
    | and_expression '&' equality_expression
    ;

exclusive_or_expression
    : and_expression
```

```
    | exclusive_or_expression '^' and_expression
    ;

inclusive_or_expression
    : exclusive_or_expression
    | inclusive_or_expression '|' exclusive_or_expression
    ;
```

If an operand of a logical operator has the compile-time type `dynamic`, then the expression is dynamically bound (§12.3.3). In this case the compile-time type of the expression is `dynamic`, and the resolution described below will take place at run-time using the run-time type of those operands that have the compile-time type `dynamic`.

For an operation of the form `x «op» y`, where «op» is one of the logical operators, overload resolution (§12.4.5) is applied to select a specific operator implementation. The operands are converted to the parameter types of the selected operator, and the type of the result is the return type of the operator.

The predefined logical operators are described in the following subclauses.

12.13.2 Integer logical operators

The predefined integer logical operators are:

C#

```
int operator &(int x, int y);
uint operator &(uint x, uint y);
long operator &(long x, long y);
ulong operator &(ulong x, ulong y);

int operator |(int x, int y);
uint operator |(uint x, uint y);
long operator |(long x, long y);
ulong operator |(ulong x, ulong y);

int operator ^(int x, int y);
uint operator ^(uint x, uint y);
long operator ^(long x, long y);
ulong operator ^(ulong x, ulong y);
```

The `&` operator computes the bitwise logical AND of the two operands, the `|` operator computes the bitwise logical OR of the two operands, and the `^` operator computes the bitwise logical exclusive OR of the two operands. No overflows are possible from these operations.

Lifted (§12.4.8) forms of the unlifted predefined integer logical operators defined above are also predefined.

12.13.3 Enumeration logical operators

Every enumeration type `E` implicitly provides the following predefined logical operators:

```
C#

E operator &(E x, E y);
E operator |(E x, E y);
E operator ^(E x, E y);
```

The result of evaluating `x «op» y`, where `x` and `y` are expressions of an enumeration type `E` with an underlying type `U`, and «op» is one of the logical operators, is exactly the same as evaluating `(E)((U)x «op» (U)y)`. In other words, the enumeration type logical operators simply perform the logical operation on the underlying type of the two operands.

Lifted (§12.4.8) forms of the unlifted predefined enumeration logical operators defined above are also predefined.

12.13.4 Boolean logical operators

The predefined Boolean logical operators are:

```
C#

bool operator &(bool x, bool y);
bool operator |(bool x, bool y);
bool operator ^(bool x, bool y);
```

The result of `x & y` is `true` if both `x` and `y` are `true`. Otherwise, the result is `false`.

The result of `x | y` is `true` if either `x` or `y` is `true`. Otherwise, the result is `false`.

The result of `x ^ y` is `true` if `x` is `true` and `y` is `false`, or `x` is `false` and `y` is `true`. Otherwise, the result is `false`. When the operands are of type `bool`, the `^` operator computes the same result as the `!=` operator.

12.13.5 Nullable Boolean & and | operators

The nullable Boolean type `bool?` can represent three values, `true`, `false`, and `null`.

As with the other binary operators, lifted forms of the logical operators `&` and `|` (§12.13.4) are also pre-defined:

```
C#

bool? operator &(bool? x, bool? y);
bool? operator |(bool? x, bool? y);
```

The semantics of the lifted `&` and `|` operators are defined by the following table:

Expand table

x	y	x & y	x | y
true	true	true	true
true	false	false	true
true	null	null	true
false	true	false	true
false	false	false	false
false	null	false	null
null	true	null	true
null	false	false	null
null	null	null	null

> *Note*: The `bool?` type is conceptually similar to the three-valued type used for Boolean expressions in SQL. The table above follows the same semantics as SQL, whereas applying the rules of §12.4.8 to the `&` and `|` operators would not. The rules of §12.4.8 already provide SQL-like semantics for the lifted `^` operator. *end note*

12.14 Conditional logical operators

12.14.1 General

The `&&` and `||` operators are called the conditional logical operators. They are also called the "short-circuiting" logical operators.

ANTLR

```
conditional_and_expression
    : inclusive_or_expression
    | conditional_and_expression '&&' inclusive_or_expression
    ;

conditional_or_expression
    : conditional_and_expression
    | conditional_or_expression '||' conditional_and_expression
    ;
```

The `&&` and `||` operators are conditional versions of the `&` and `|` operators:

- The operation `x && y` corresponds to the operation `x & y`, except that `y` is evaluated only if `x` is not `false`.
- The operation `x || y` corresponds to the operation `x | y`, except that `y` is evaluated only if `x` is not `true`.

> *Note*: The reason that short circuiting uses the 'not true' and 'not false' conditions is to enable user-defined conditional operators to define when short circuiting applies. User-defined types could be in a state where `operator true` returns `false` and `operator false` returns `false`. In those cases, neither `&&` nor `||` would short circuit. *end note*

If an operand of a conditional logical operator has the compile-time type `dynamic`, then the expression is dynamically bound (§12.3.3). In this case the compile-time type of the expression is `dynamic`, and the resolution described below will take place at run-time using the run-time type of those operands that have the compile-time type `dynamic`.

An operation of the form `x && y` or `x || y` is processed by applying overload resolution (§12.4.5) as if the operation was written `x & y` or `x | y`. Then,

- If overload resolution fails to find a single best operator, or if overload resolution selects one of the predefined integer logical operators or nullable Boolean logical operators (§12.13.5), a binding-time error occurs.
- Otherwise, if the selected operator is one of the predefined Boolean logical operators (§12.13.4), the operation is processed as described in §12.14.2.
- Otherwise, the selected operator is a user-defined operator, and the operation is processed as described in §12.14.3.

It is not possible to directly overload the conditional logical operators. However, because the conditional logical operators are evaluated in terms of the regular logical

operators, overloads of the regular logical operators are, with certain restrictions, also considered overloads of the conditional logical operators. This is described further in §12.14.3.

12.14.2 Boolean conditional logical operators

When the operands of `&&` or `||` are of type `bool`, or when the operands are of types that do not define an applicable `operator &` or `operator |`, but do define implicit conversions to `bool`, the operation is processed as follows:

- The operation `x && y` is evaluated as `x ? y : false`. In other words, `x` is first evaluated and converted to type `bool`. Then, if `x` is `true`, `y` is evaluated and converted to type `bool`, and this becomes the result of the operation. Otherwise, the result of the operation is `false`.
- The operation `x || y` is evaluated as `x ? true : y`. In other words, `x` is first evaluated and converted to type `bool`. Then, if `x` is `true`, the result of the operation is `true`. Otherwise, `y` is evaluated and converted to type `bool`, and this becomes the result of the operation.

12.14.3 User-defined conditional logical operators

When the operands of `&&` or `||` are of types that declare an applicable user-defined `operator &` or `operator |`, both of the following shall be true, where `T` is the type in which the selected operator is declared:

- The return type and the type of each parameter of the selected operator shall be `T`. In other words, the operator shall compute the logical AND or the logical OR of two operands of type `T`, and shall return a result of type `T`.
- `T` shall contain declarations of `operator true` and `operator false`.

A binding-time error occurs if either of these requirements is not satisfied. Otherwise, the `&&` or `||` operation is evaluated by combining the user-defined `operator true` or `operator false` with the selected user-defined operator:

- The operation `x && y` is evaluated as `T.false(x) ? x : T.&(x, y)`, where `T.false(x)` is an invocation of the `operator false` declared in `T`, and `T.&(x, y)` is an invocation of the selected `operator &`. In other words, `x` is first evaluated and `operator false` is invoked on the result to determine if `x` is definitely false. Then, if `x` is definitely false, the result of the operation is the value previously computed for `x`. Otherwise, `y` is evaluated, and the selected `operator &` is invoked on the

value previously computed for `x` and the value computed for `y` to produce the result of the operation.

- The operation `x || y` is evaluated as `T.true(x) ? x : T.|(x, y)`, where `T.true(x)` is an invocation of the `operator true` declared in `T`, and `T.|(x, y)` is an invocation of the selected `operator |`. In other words, `x` is first evaluated and `operator true` is invoked on the result to determine if `x` is definitely true. Then, if `x` is definitely true, the result of the operation is the value previously computed for `x`. Otherwise, `y` is evaluated, and the selected `operator |` is invoked on the value previously computed for `x` and the value computed for `y` to produce the result of the operation.

In either of these operations, the expression given by `x` is only evaluated once, and the expression given by `y` is either not evaluated or evaluated exactly once.

12.15 The null coalescing operator

The `??` operator is called the null coalescing operator.

ANTLR

```
null_coalescing_expression
    : conditional_or_expression
    | conditional_or_expression '??' null_coalescing_expression
    | throw_expression
    ;
```

In a null coalescing expression of the form `a ?? b`, if `a` is non-`null`, the result is `a`; otherwise, the result is `b`. The operation evaluates `b` only if `a` is `null`.

The null coalescing operator is right-associative, meaning that operations are grouped from right to left.

> *Example*: An expression of the form `a ?? b ?? c` is evaluated as a `?? (b ?? c)`. In general terms, an expression of the form `E1 ?? E2 ?? ... ?? EN` returns the first of the operands that is non-`null`, or `null` if all operands are `null`. *end example*

The type of the expression `a ?? b` depends on which implicit conversions are available on the operands. In order of preference, the type of `a ?? b` is A_o, `A`, or `B`, where `A` is the type of `a` (provided that `a` has a type), `B` is the type of `b`(provided that `b` has a type), and A_o is the underlying type of `A` if `A` is a nullable value type, or `A` otherwise. Specifically, `a ?? b` is processed as follows:

- If `A` exists and is not a nullable value type or a reference type, a compile-time error occurs.
- Otherwise, if `A` exists and `b` is a dynamic expression, the result type is `dynamic`. At run-time, `a` is first evaluated. If `a` is not `null`, `a` is converted to `dynamic`, and this becomes the result. Otherwise, `b` is evaluated, and this becomes the result.
- Otherwise, if `A` exists and is a nullable value type and an implicit conversion exists from `b` to A_o, the result type is A_o. At run-time, `a` is first evaluated. If `a` is not `null`, `a` is unwrapped to type A_o, and this becomes the result. Otherwise, `b` is evaluated and converted to type A_o, and this becomes the result.
- Otherwise, if `A` exists and an implicit conversion exists from `b` to `A`, the result type is `A`. At run-time, a is first evaluated. If a is not null, a becomes the result. Otherwise, `b` is evaluated and converted to type `A`, and this becomes the result.
- Otherwise, if `A` exists and is a nullable value type, `b` has a type `B` and an implicit conversion exists from A_o to `B`, the result type is `B`. At run-time, `a` is first evaluated. If `a` is not `null`, `a` is unwrapped to type A_o and converted to type `B`, and this becomes the result. Otherwise, `b` is evaluated and becomes the result.
- Otherwise, if `b` has a type `B` and an implicit conversion exists from `a` to `B`, the result type is `B`. At run-time, `a` is first evaluated. If `a` is not `null`, `a` is converted to type `B`, and this becomes the result. Otherwise, `b` is evaluated and becomes the result.

Otherwise, `a` and `b` are incompatible, and `a` compile-time error occurs.

12.16 The throw expression operator

ANTLR

```
throw_expression
    : 'throw' null_coalescing_expression
    ;
```

A *throw_expression* throws the value produced by evaluating the *null_coalescing_expression*. The expression shall be implicitly convertible to `System.Exception`, and the result of evaluating the expression is converted to `System.Exception` before being thrown. The behavior at runtime of the evaluation of a *throw expression* is the same as specified for a *throw statement* (§13.10.6).

A *throw_expression* has no type. A *throw_expression* is convertible to every type by an *implicit throw conversion*.

A *throw expression* shall only occur in the following syntactic contexts:

- As the second or third operand of a ternary conditional operator (`?:`).
- As the second operand of a null coalescing operator (`??`).
- As the body of an expression-bodied lambda or member.

12.17 Declaration expressions

A declaration expression declares a local variable.

ANTLR

```
declaration_expression
    : local_variable_type identifier
    ;

local_variable_type
    : type
    | 'var'
    ;
```

The *simple_name* `_` is also considered a declaration expression if simple name lookup did not find an associated declaration (§12.8.4). When used as a declaration expression, `_` is called a *simple discard*. It is semantically equivalent to `var _`, but is permitted in more places.

A declaration expression shall only occur in the following syntactic contexts:

- As an `out` *argument_value* in an *argument_list*.
- As a simple discard `_` comprising the left side of a simple assignment (§12.21.2).
- As a *tuple_element* in one or more recursively nested *tuple_expression*s, the outermost of which comprises the left side of a deconstructing assignment. A *deconstruction_expression* gives rise to declaration expressions in this position, even though the declaration expressions are not syntactically present.

> *Note*: This means that a declaration expression cannot be parenthesized. *end note*

It is an error for an implicitly typed variable declared with a *declaration_expression* to be referenced within the *argument_list* where it is declared.

It is an error for a variable declared with a *declaration_expression* to be referenced within the deconstructing assignment where it occurs.

A declaration expression that is a simple discard or where the *local_variable_type* is the identifier `var` is classified as an *implicitly typed* variable. The expression has no type, and the type of the local variable is inferred based on the syntactic context as follows:

- In an *argument_list* the inferred type of the variable is the declared type of the corresponding parameter.
- As the left side of a simple assignment, the inferred type of the variable is the type of the right side of the assignment.
- In a *tuple_expression* on the left side of a simple assignment, the inferred type of the variable is the type of the corresponding tuple element on the right side (after deconstruction) of the assignment.

Otherwise, the declaration expression is classified as an *explicitly typed* variable, and the type of the expression as well as the declared variable shall be that given by the *local_variable_type*.

A declaration expression with the identifier `_` is a discard (§9.2.9.1), and does not introduce a name for the variable. A declaration expression with an identifier other than `_` introduces that name into the nearest enclosing local variable declaration space (§7.3).

> *Example*:
>
> C#
>
> ```
> string M(out int i, string s, out bool b) { ... }
>
> var s1 = M(out int i1, "One", out var b1);
> Console.WriteLine($"{i1}, {b1}, {s1}");
> // Error: i2 referenced within declaring argument list
> var s2 = M(out var i2, M(out i2, "Two", out bool b2), out b2);
> var s3 = M(out int _, "Three", out var _);
> ```
>
> The declaration of `s1` shows both explicitly and implicitly typed declaration expressions. The inferred type of `b1` is `bool` because that is the type of the corresponding out parameter in `M1`. The subsequent `WriteLine` is able to access `i1` and `b1`, which have been introduced to the enclosing scope.
>
> The declaration of `s2` shows an attempt to use `i2` in the nested call to `M`, which is disallowed, because the reference occurs within the argument list where `i2` was declared. On the other hand the reference to `b2` in the final argument is allowed, because it occurs after the end of the nested argument list where `b2` was declared.

The declaration of `s3` shows the use of both implicitly and explicitly typed declaration expressions that are discards. Because discards do not declare a named variable, the multiple occurrences of the identifier `_` are allowed.

C#

```
(int i1, int _, (var i2, var _), _) = (1, 2, (3, 4), 5);
```

This example shows the use of implicitly and explicitly typed declaration expressions for both variables and discards in a deconstructing assignment. The *simple_name* `_` is equivalent to `var _` when no declaration of `_` is found.

C#

```
void M1(out int i) { ... }

void M2(string _)
{
    M1(out _);      // Error: `_` is a string
    M1(out var _);
}
```

This examples shows the use of `var _` to provide an implicitly typed discard when `_` is not available, because it designates a variable in the enclosing scope.

end example

12.18 Conditional operator

The `?:` operator is called the conditional operator. It is at times also called the ternary operator.

ANTLR

```
conditional_expression
    : null_coalescing_expression
    | null_coalescing_expression '?' expression ':' expression
    | null_coalescing_expression '?' 'ref' variable_reference ':'
      'ref' variable_reference
    ;
```

A throw expression (§12.16) is not allowed in a conditional operator if `ref` is present.

A conditional expression of the form `b ? x : y` first evaluates the condition `b`. Then, if `b` is `true`, `x` is evaluated and becomes the result of the operation. Otherwise, `y` is evaluated and becomes the result of the operation. A conditional expression never evaluates both `x` and `y`.

The conditional operator is right-associative, meaning that operations are grouped from right to left.

> *Example*: An expression of the form `a ? b : c ? d : e` is evaluated as `a ? b : (c ? d : e)`. *end example*

The first operand of the `?:` operator shall be an expression that can be implicitly converted to `bool`, or an expression of a type that implements `operator true`. If neither of these requirements is satisfied, a compile-time error occurs.

If `ref` is present:

- An identity conversion shall exist between the types of the two *variable_reference*s, and type of the result can be either type. If either type is `dynamic`, type inference prefers `dynamic` (§8.7). If either type is a tuple type (§8.3.11), type inference includes the element names when the element names in the same ordinal position match in both tuples.
- The result is a variable reference, which is writeable if both *variable_reference*s are writeable.

> *Note*: When `ref` is present, the *conditional_expression* returns a variable reference, which can be assigned to a reference variable using the `= ref` operator or passed as a reference/input/output parameter. *end note*

If `ref` is not present, the second and third operands, `x` and `y`, of the `?:` operator control the type of the conditional expression:

- If `x` has type `X` and `y` has type `Y` then,
  - If an identity conversion exists between `X` and `Y`, then the result is the best common type of a set of expressions (§12.6.3.15). If either type is `dynamic`, type inference prefers `dynamic` (§8.7). If either type is a tuple type (§8.3.11), type inference includes the element names when the element names in the same ordinal position match in both tuples.
  - Otherwise, if an implicit conversion (§10.2) exists from `X` to `Y`, but not from `Y` to `X`, then `Y` is the type of the conditional expression.

    - Otherwise, if an implicit enumeration conversion (§10.2.4) exists from `X` to `Y`, then `Y` is the type of the conditional expression.
    - Otherwise, if an implicit enumeration conversion (§10.2.4) exists from `Y` to `X`, then `X` is the type of the conditional expression.
    - Otherwise, if an implicit conversion (§10.2) exists from `Y` to `X`, but not from `X` to `Y`, then `X` is the type of the conditional expression.
    - Otherwise, no expression type can be determined, and a compile-time error occurs.
- If only one of `x` and `y` has a type, and both `x` and `y` are implicitly convertible to that type, then that is the type of the conditional expression.
- Otherwise, no expression type can be determined, and a compile-time error occurs.

The run-time processing of a ref conditional expression of the form `b ? ref x : ref y` consists of the following steps:

- First, `b` is evaluated, and the `bool` value of `b` is determined:
    - If an implicit conversion from the type of `b` to `bool` exists, then this implicit conversion is performed to produce a `bool` value.
    - Otherwise, the `operator true` defined by the type of `b` is invoked to produce a `bool` value.
- If the `bool` value produced by the step above is `true`, then `x` is evaluated and the resulting variable reference becomes the result of the conditional expression.
- Otherwise, `y` is evaluated and the resulting variable reference becomes the result of the conditional expression.

The run-time processing of a conditional expression of the form `b ? x : y` consists of the following steps:

- First, `b` is evaluated, and the `bool` value of `b` is determined:
    - If an implicit conversion from the type of `b` to `bool` exists, then this implicit conversion is performed to produce a `bool` value.
    - Otherwise, the `operator true` defined by the type of `b` is invoked to produce a `bool` value.
- If the `bool` value produced by the step above is `true`, then `x` is evaluated and converted to the type of the conditional expression, and this becomes the result of the conditional expression.
- Otherwise, `y` is evaluated and converted to the type of the conditional expression, and this becomes the result of the conditional expression.

12.19 Anonymous function expressions

12.19.1 General

An ***anonymous function*** is an expression that represents an "in-line" method definition. An anonymous function does not have a value or type in and of itself, but is convertible to a compatible delegate or expression-tree type. The evaluation of an anonymous-function conversion depends on the target type of the conversion: If it is a delegate type, the conversion evaluates to a delegate value referencing the method that the anonymous function defines. If it is an expression-tree type, the conversion evaluates to an expression tree that represents the structure of the method as an object structure.

> *Note*: For historical reasons, there are two syntactic flavors of anonymous functions, namely *lambda_expression*s and *anonymous_method_expression*s. For almost all purposes, *lambda_expression*s are more concise and expressive than *anonymous_method_expression*s, which remain in the language for backwards compatibility. *end note*

ANTLR

```
lambda_expression
    : 'async'? anonymous_function_signature '=>' anonymous_function_body
    ;

anonymous_method_expression
    : 'async'? 'delegate' explicit_anonymous_function_signature? block
    ;

anonymous_function_signature
    : explicit_anonymous_function_signature
    | implicit_anonymous_function_signature
    ;

explicit_anonymous_function_signature
    : '(' explicit_anonymous_function_parameter_list? ')'
    ;

explicit_anonymous_function_parameter_list
    : explicit_anonymous_function_parameter
      (',' explicit_anonymous_function_parameter)*
    ;

explicit_anonymous_function_parameter
    : anonymous_function_parameter_modifier? type identifier
    ;

anonymous_function_parameter_modifier
```

```
    : 'ref'
    | 'out'
    | 'in'
    ;

implicit_anonymous_function_signature
    : '(' implicit_anonymous_function_parameter_list? ')'
    | implicit_anonymous_function_parameter
    ;

implicit_anonymous_function_parameter_list
    : implicit_anonymous_function_parameter
      (',' implicit_anonymous_function_parameter)*
    ;

implicit_anonymous_function_parameter
    : identifier
    ;

anonymous_function_body
    : null_conditional_invocation_expression
    | expression
    | 'ref' variable_reference
    | block
    ;
```

When recognising an *anonymous_function_body* if both the *null_conditional_invocation_expression* and *expression* alternatives are applicable then the former shall be chosen.

> *Note*: The overlapping of, and priority between, alternatives here is solely for descriptive convenience; the grammar rules could be elaborated to remove the overlap. ANTLR, and other grammar systems, adopt the same convenience and so *anonymous_function_body* has the specified semantics automatically. *end note*

> *Note*: When treated as an *expression*, a syntactic form such as `x?.M()` would be an error if the result type of `M` is `void` (§12.8.12). But when treated as a *null_conditional_invocation_expression*, the result type is permitted to be `void`. *end note*

> *Example*: The result type of `List<T>.Reverse` is `void`. In the following code, the body of the anonymous expression is a *null_conditional_invocation_expression*, so it is not an error.

```
C#
```

```
Action<List<int>> a = x => x?.Reverse();
```

end example

The `=>` operator has the same precedence as assignment (`=`) and is right-associative.

An anonymous function with the `async` modifier is an async function and follows the rules described in §15.15.

The parameters of an anonymous function in the form of a *lambda_expression* can be explicitly or implicitly typed. In an explicitly typed parameter list, the type of each parameter is explicitly stated. In an implicitly typed parameter list, the types of the parameters are inferred from the context in which the anonymous function occurs—specifically, when the anonymous function is converted to a compatible delegate type or expression tree type, that type provides the parameter types (§10.7).

In a *lambda_expression* with a single, implicitly typed parameter, the parentheses may be omitted from the parameter list. In other words, an anonymous function of the form

C#

```
( «param» ) => «expr»
```

can be abbreviated to

C#

```
«param» => «expr»
```

The parameter list of an anonymous function in the form of an *anonymous_method_expression* is optional. If given, the parameters shall be explicitly typed. If not, the anonymous function is convertible to a delegate with any parameter list not containing out parameters.

A *block* body of an anonymous function is always reachable (§13.2).

Example: Some examples of anonymous functions follow below:

C#

```
x => x + 1                              // Implicitly typed, expression
body
x => { return x + 1; }                  // Implicitly typed, block body
```

```
(int x) => x + 1                        // Explicitly typed, expression
body
(int x) => { return x + 1; }            // Explicitly typed, block body
(x, y) => x * y                         // Multiple parameters
() => Console.WriteLine()               // No parameters
async (t1,t2) => await t1 + await t2    // Async
delegate (int x) { return x + 1; }      // Anonymous method expression
delegate { return 1 + 1; }              // Parameter list omitted
```

end example

The behavior of *lambda_expression*s and *anonymous_method_expression*s is the same except for the following points:

- *anonymous_method_expression*s permit the parameter list to be omitted entirely, yielding convertibility to delegate types of any list of value parameters.
- *lambda_expression*s permit parameter types to be omitted and inferred whereas *anonymous_method_expression*s require parameter types to be explicitly stated.
- The body of a *lambda_expression* can be an expression or a block whereas the body of an *anonymous_method_expression* shall be a block.
- Only *lambda_expression*s have conversions to compatible expression tree types (§8.6).

12.19.2 Anonymous function signatures

The *anonymous_function_signature* of an anonymous function defines the names and optionally the types of the formal parameters for the anonymous function. The scope of the parameters of the anonymous function is the *anonymous_function_body* (§7.7). Together with the parameter list (if given) the anonymous-method-body constitutes a declaration space (§7.3). It is thus a compile-time error for the name of a parameter of the anonymous function to match the name of a local variable, local constant or parameter whose scope includes the *anonymous_method_expression* or *lambda_expression*.

If an anonymous function has an *explicit_anonymous_function_signature*, then the set of compatible delegate types and expression tree types is restricted to those that have the same parameter types and modifiers in the same order (§10.7). In contrast to method group conversions (§10.8), contra-variance of anonymous function parameter types is not supported. If an anonymous function does not have an *anonymous_function_signature*, then the set of compatible delegate types and expression tree types is restricted to those that have no out parameters.

Note that an *anonymous_function_signature* cannot include attributes or a parameter array. Nevertheless, an *anonymous_function_signature* may be compatible with a delegate type whose parameter list contains a parameter array.

Note also that conversion to an expression tree type, even if compatible, may still fail at compile-time (§8.6).

12.19.3 Anonymous function bodies

The body (*expression* or *block*) of an anonymous function is subject to the following rules:

- If the anonymous function includes a signature, the parameters specified in the signature are available in the body. If the anonymous function has no signature it can be converted to a delegate type or expression type having parameters (§10.7), but the parameters cannot be accessed in the body.
- Except for `in`, `out`, or `ref` parameters specified in the signature (if any) of the nearest enclosing anonymous function, it is a compile-time error for the body to access an `in`, `out`, or `ref` parameter.
- Except for parameters specified in the signature (if any) of the nearest enclosing anonymous function, it is a compile-time error for the body to access a parameter of a `ref struct` type.
- When the type of `this` is a struct type, it is a compile-time error for the body to access `this`. This is true whether the access is explicit (as in `this.x`) or implicit (as in `x` where `x` is an instance member of the struct). This rule simply prohibits such access and does not affect whether member lookup results in a member of the struct.
- The body has access to the outer variables (§12.19.6) of the anonymous function. Access of an outer variable will reference the instance of the variable that is active at the time the *lambda_expression* or *anonymous_method_expression* is evaluated (§12.19.7).
- It is a compile-time error for the body to contain a `goto` statement, a `break` statement, or a `continue` statement whose target is outside the body or within the body of a contained anonymous function.
- A `return` statement in the body returns control from an invocation of the nearest enclosing anonymous function, not from the enclosing function member.

It is explicitly unspecified whether there is any way to execute the block of an anonymous function other than through evaluation and invocation of the *lambda_expression* or *anonymous_method_expression*. In particular, the compiler may choose to implement an anonymous function by synthesizing one or more named

methods or types. The names of any such synthesized elements shall be of a form reserved for compiler use (§6.4.3).

12.19.4 Overload resolution

Anonymous functions in an argument list participate in type inference and overload resolution. Refer to §12.6.3 and §12.6.4 for the exact rules.

> *Example*: The following example illustrates the effect of anonymous functions on overload resolution.

```
C#
```

```
class ItemList<T> : List<T>
{
    public int Sum(Func<T, int> selector)
    {
        int sum = 0;
        foreach (T item in this)
        {
            sum += selector(item);
        }
        return sum;
    }

    public double Sum(Func<T, double> selector)
    {
        double sum = 0;
        foreach (T item in this)
        {
            sum += selector(item);
        }
        return sum;
    }
}
```

The `ItemList<T>` class has two `Sum` methods. Each takes a `selector` argument, which extracts the value to sum over from a list item. The extracted value can be either an `int` or a `double` and the resulting sum is likewise either an `int` or a `double`.

The `Sum` methods could for example be used to compute sums from a list of detail lines in an order.

```
C#
```

```
class Detail
{
```

```
    public int UnitCount;
    public double UnitPrice;
    ...
}

class A
{
    void ComputeSums()
    {
        ItemList<Detail> orderDetails = GetOrderDetails( ... );
        int totalUnits = orderDetails.Sum(d => d.UnitCount);
        double orderTotal = orderDetails.Sum(d => d.UnitPrice *
d.UnitCount);
        ...
    }

    ItemList<Detail> GetOrderDetails( ... )
    {
        ...
    }
}
```

In the first invocation of `orderDetails.Sum`, both `Sum` methods are applicable because the anonymous function `d => d.UnitCount` is compatible with both `Func<Detail,int>` and `Func<Detail,double>`. However, overload resolution picks the first `Sum` method because the conversion to `Func<Detail,int>` is better than the conversion to `Func<Detail,double>`.

In the second invocation of `orderDetails.Sum`, only the second `Sum` method is applicable because the anonymous function `d => d.UnitPrice * d.UnitCount` produces a value of type `double`. Thus, overload resolution picks the second `Sum` method for that invocation.

end example

12.19.5 Anonymous functions and dynamic binding

An anonymous function cannot be a receiver, argument, or operand of a dynamically bound operation.

12.19.6 Outer variables

12.19.6.1 General

Any local variable, value parameter, or parameter array whose scope includes the *lambda_expression* or *anonymous_method_expression* is called an ***outer variable*** of the anonymous function. In an instance function member of a class, the this value is considered a value parameter and is an outer variable of any anonymous function contained within the function member.

12.19.6.2 Captured outer variables

When an outer variable is referenced by an anonymous function, the outer variable is said to have been ***captured*** by the anonymous function. Ordinarily, the lifetime of a local variable is limited to execution of the block or statement with which it is associated (§9.2.9). However, the lifetime of a captured outer variable is extended at least until the delegate or expression tree created from the anonymous function becomes eligible for garbage collection.

> *Example*: In the example
>
> C#
>
> ```
> delegate int D();
>
> class Test
> {
>     static D F()
>     {
>         int x = 0;
>         D result = () => ++x;
>         return result;
>     }
>
>     static void Main()
>     {
>         D d = F();
>         Console.WriteLine(d());
>         Console.WriteLine(d());
>         Console.WriteLine(d());
>     }
> }
> ```
>
> the local variable `x` is captured by the anonymous function, and the lifetime of `x` is extended at least until the delegate returned from `F` becomes eligible for garbage collection. Since each invocation of the anonymous function operates on the same instance of `x`, the output of the example is:
>
> Console

```
1
2
3
```

end example

When a local variable or a value parameter is captured by an anonymous function, the local variable or parameter is no longer considered to be a fixed variable (§23.4), but is instead considered to be a moveable variable. However, captured outer variables cannot be used in a `fixed` statement (§23.7), so the address of a captured outer variable cannot be taken.

> *Note*: Unlike an uncaptured variable, a captured local variable can be simultaneously exposed to multiple threads of execution. *end note*

12.19.6.3 Instantiation of local variables

A local variable is considered to be ***instantiated*** when execution enters the scope of the variable.

> *Example*: For example, when the following method is invoked, the local variable `x` is instantiated and initialized three times—once for each iteration of the loop.

C#

```
static void F()
{
    for (int i = 0; i < 3; i++)
    {
        int x = i * 2 + 1;
        ...
    }
}
```

However, moving the declaration of `x` outside the loop results in a single instantiation of `x`:

C#

```
static void F()
{
    int x;
    for (int i = 0; i < 3; i++)
    {
```

```
            x = i * 2 + 1;
            ...
        }
    }
```

end example

When not captured, there is no way to observe exactly how often a local variable is instantiated—because the lifetimes of the instantiations are disjoint, it is possible for each instantation to simply use the same storage location. However, when an anonymous function captures a local variable, the effects of instantiation become apparent.

Example: The example

C#

```
delegate void D();
class Test
{
    static D[] F()
    {
        D[] result = new D[3];
        for (int i = 0; i < 3; i++)
        {
            int x = i * 2 + 1;
            result[i] = () => Console.WriteLine(x);
        }
        return result;
    }

    static void Main()
    {
        foreach (D d in F())
        {
            d();
        }
    }
}
```

produces the output:

Console

```
1
3
5
```

However, when the declaration of `x` is moved outside the loop:

C#

```
delegate void D();

class Test
{
    static D[] F()
    {
        D[] result = new D[3];
        int x;
        for (int i = 0; i < 3; i++)
        {
            x = i * 2 + 1;
            result[i] = () => Console.WriteLine(x);
        }
        return result;
   }

   static void Main()
   {
       foreach (D d in F())
       {
           d();
       }
   }
}
```

the output is:

Console

```
5
5
5
```

Note that the compiler is permitted (but not required) to optimize the three instantiations into a single delegate instance (§10.7.2).

end example

If a for-loop declares an iteration variable, that variable itself is considered to be declared outside of the loop.

Example: Thus, if the example is changed to capture the iteration variable itself:

C#

```
delegate void D();

class Test
{
    static D[] F()
    {
        D[] result = new D[3];
        for (int i = 0; i < 3; i++)
        {
            result[i] = () => Console.WriteLine(i);
        }
        return result;
   }

   static void Main()
   {
       foreach (D d in F())
       {
           d();
       }
   }
}
```

only one instance of the iteration variable is captured, which produces the output:

Console

```
3
3
3
```

end example

It is possible for anonymous function delegates to share some captured variables yet have separate instances of others.

Example: For example, if `F` is changed to

C#

```
static D[] F()
{
    D[] result = new D[3];
    int x = 0;
    for (int i = 0; i < 3; i++)
    {
        int y = 0;
        result[i] = () => Console.WriteLine($"{++x} {++y}");
    }
```

```
    return result;
}
```

the three delegates capture the same instance of `x` but separate instances of `y`, and the output is:

```
Console
```

```
1 1
2 1
3 1
```

end example

Separate anonymous functions can capture the same instance of an outer variable.

Example: In the example:

```
C#
```

```
delegate void Setter(int value);
delegate int Getter();

class Test
{
    static void Main()
    {
        int x = 0;
        Setter s = (int value) => x = value;
        Getter g = () => x;
        s(5);
        Console.WriteLine(g());
        s(10);
        Console.WriteLine(g());
    }
}
```

the two anonymous functions capture the same instance of the local variable `x`, and they can thus "communicate" through that variable. The output of the example is:

```
Console
```

```
5
10
```

end example

12.19.7 Evaluation of anonymous function expressions

An anonymous function `F` shall always be converted to a delegate type `D` or an expression-tree type `E`, either directly or through the execution of a delegate creation expression `new D(F)`. This conversion determines the result of the anonymous function, as described in §10.7.

12.19.8 Implementation Example

This subclause is informative.

This subclause describes a possible implementation of anonymous function conversions in terms of other C# constructs. The implementation described here is based on the same principles used by a commercial C# compiler, but it is by no means a mandated implementation, nor is it the only one possible. It only briefly mentions conversions to expression trees, as their exact semantics are outside the scope of this specification.

The remainder of this subclause gives several examples of code that contains anonymous functions with different characteristics. For each example, a corresponding translation to code that uses only other C# constructs is provided. In the examples, the identifier `D` is assumed by represent the following delegate type:

C#

```
public delegate void D();
```

The simplest form of an anonymous function is one that captures no outer variables:

C#

```
delegate void D();

class Test
{
    static void F()
    {
        D d = () => Console.WriteLine("test");
    }
}
```

This can be translated to a delegate instantiation that references a compiler generated static method in which the code of the anonymous function is placed:

C#

```
delegate void D();

class Test
{
    static void F()
    {
        D d = new D(__Method1);
    }

    static void __Method1()
    {
        Console.WriteLine("test");
    }
}
```

In the following example, the anonymous function references instance members of `this`:

C#

```
delegate void D();

class Test
{
    int x;

    void F()
    {
        D d = () => Console.WriteLine(x);
    }
}
```

This can be translated to a compiler generated instance method containing the code of the anonymous function:

C#

```
delegate void D();

class Test
{
   int x;

   void F()
   {
       D d = new D(__Method1);
   }

   void __Method1()
   {
```

```
        Console.WriteLine(x);
    }
}
```

In this example, the anonymous function captures a local variable:

C#

```
delegate void D();

class Test
{
    void F()
    {
        int y = 123;
        D d = () => Console.WriteLine(y);
    }
}
```

The lifetime of the local variable must now be extended to at least the lifetime of the anonymous function delegate. This can be achieved by "hoisting" the local variable into a field of a compiler-generated class. Instantiation of the local variable (§12.19.6.3) then corresponds to creating an instance of the compiler generated class, and accessing the local variable corresponds to accessing a field in the instance of the compiler generated class. Furthermore, the anonymous function becomes an instance method of the compiler-generated class:

C#

```
delegate void D();

class Test
{
    void F()
    {
        __Locals1 __locals1 = new __Locals1();
        __locals1.y = 123;
        D d = new D(__locals1.__Method1);
    }

    class __Locals1
    {
        public int y;

        public void __Method1()
        {
            Console.WriteLine(y);
        }
```

```
    }
}
```

Finally, the following anonymous function captures `this` as well as two local variables with different lifetimes:

C#

```
delegate void D();

class Test
{
   int x;

   void F()
   {
       int y = 123;
       for (int i = 0; i < 10; i++)
       {
           int z = i * 2;
           D d = () => Console.WriteLine(x + y + z);
       }
   }
}
```

Here, a compiler-generated class is created for each block in which locals are captured such that the locals in the different blocks can have independent lifetimes. An instance of `__Locals2`, the compiler generated class for the inner block, contains the local variable `z` and a field that references an instance of `__Locals1`. An instance of `__Locals1`, the compiler generated class for the outer block, contains the local variable `y` and a field that references `this` of the enclosing function member. With these data structures, it is possible to reach all captured outer variables through an instance of `__Local2`, and the code of the anonymous function can thus be implemented as an instance method of that class.

C#

```
delegate void D();

class Test
{
    int x;

    void F()
    {
        __Locals1 __locals1 = new __Locals1();
        __locals1.__this = this;
        __locals1.y = 123;
```

```
        for (int i = 0; i < 10; i++)
        {
            __Locals2 __locals2 = new __Locals2();
            __locals2.__locals1 = __locals1;
            __locals2.z = i * 2;
            D d = new D(__locals2.__Method1);
        }
    }

    class __Locals1
    {
        public Test __this;
        public int y;
    }

    class __Locals2
    {
        public __Locals1 __locals1;
        public int z;

        public void __Method1()
        {
            Console.WriteLine(__locals1.__this.x + __locals1.y + z);
        }
    }
}
```

The same technique applied here to capture local variables can also be used when converting anonymous functions to expression trees: references to the compiler-generated objects can be stored in the expression tree, and access to the local variables can be represented as field accesses on these objects. The advantage of this approach is that it allows the "lifted" local variables to be shared between delegates and expression trees.

End of informative text.

12.20 Query expressions

12.20.1 General

Query expressions provide a language-integrated syntax for queries that is similar to relational and hierarchical query languages such as SQL and XQuery.

ANTLR

```
query_expression
    : from_clause query_body
    ;
```

```
from_clause
    : 'from' type? identifier 'in' expression
    ;

query_body
    : query_body_clauses? select_or_group_clause query_continuation?
    ;

query_body_clauses
    : query_body_clause
    | query_body_clauses query_body_clause
    ;

query_body_clause
    : from_clause
    | let_clause
    | where_clause
    | join_clause
    | join_into_clause
    | orderby_clause
    ;

let_clause
    : 'let' identifier '=' expression
    ;

where_clause
    : 'where' boolean_expression
    ;

join_clause
    : 'join' type? identifier 'in' expression 'on' expression
      'equals' expression
    ;

join_into_clause
    : 'join' type? identifier 'in' expression 'on' expression
      'equals' expression 'into' identifier
    ;

orderby_clause
    : 'orderby' orderings
    ;

orderings
    : ordering (',' ordering)*
    ;

ordering
    : expression ordering_direction?
    ;

ordering_direction
    : 'ascending'
```

```
    | 'descending'
    ;

select_or_group_clause
    : select_clause
    | group_clause
    ;

select_clause
    : 'select' expression
    ;

group_clause
    : 'group' expression 'by' expression
    ;

query_continuation
    : 'into' identifier query_body
    ;
```

A query expression begins with a `from` clause and ends with either a `select` or `group` clause. The initial `from` clause may be followed by zero or more `from`, `let`, `where`, `join` or `orderby` clauses. Each `from` clause is a generator introducing a ***range variable*** that ranges over the elements of a ***sequence***. Each `let` clause introduces a range variable representing a value computed by means of previous range variables. Each `where` clause is a filter that excludes items from the result. Each `join` clause compares specified keys of the source sequence with keys of another sequence, yielding matching pairs. Each `orderby` clause reorders items according to specified criteria.The final `select` or `group` clause specifies the shape of the result in terms of the range variables. Finally, an `into` clause can be used to "splice" queries by treating the results of one query as a generator in a subsequent query.

12.20.2 Ambiguities in query expressions

Query expressions use a number of contextual keywords (§6.4.4): `ascending`, `by`, `descending`, `equals`, `from`, `group`, `into`, `join`, `let`, `on`, `orderby`, `select` and `where`.

To avoid ambiguities that could arise from the use of these identifiers both as keywords and simple names these identifiers are considered keywords anywhere within a query expression, unless they are prefixed with "`@`" (§6.4.4) in which case they are considered identifiers. For this purpose, a query expression is any expression that starts with "`from` *identifier*" followed by any token except "`;`", "`=`" or "`,`".

12.20.3 Query expression translation

12.20.3.1 General

The C# language does not specify the execution semantics of query expressions. Rather, query expressions are translated into invocations of methods that adhere to the query-expression pattern (§12.20.4). Specifically, query expressions are translated into invocations of methods named `Where`, `Select`, `SelectMany`, `Join`, `GroupJoin`, `OrderBy`, `OrderByDescending`, `ThenBy`, `ThenByDescending`, `GroupBy`, and `Cast`. These methods are expected to have particular signatures and return types, as described in §12.20.4. These methods may be instance methods of the object being queried or extension methods that are external to the object. These methods implement the actual execution of the query.

The translation from query expressions to method invocations is a syntactic mapping that occurs before any type binding or overload resolution has been performed. Following translation of query expressions, the resulting method invocations are processed as regular method invocations, and this may in turn uncover compile time errors. These error conditions include, but are not limited to, methods that do not exist, arguments of the wrong types, and generic methods where type inference fails.

A query expression is processed by repeatedly applying the following translations until no further reductions are possible. The translations are listed in order of application: each section assumes that the translations in the preceding sections have been performed exhaustively, and once exhausted, a section will not later be revisited in the processing of the same query expression.

It is a compile time error for a query expression to include an assignment to a range variable, or the use of a range variable as an argument for a `ref` or `out` parameter.

Certain translations inject range variables with *transparent identifiers* denoted by *. These are described further in §12.20.3.8.

12.20.3.2 Query expressions with continuations

A query expression with a continuation following its query body

```
C#
```

```
from «x1» in «e1» «b1» into «x2» «b2»
```

is translated into

```
C#
```

```
from «x2» in ( from «x1» in «e1» «b1» ) «b2»
```

The translations in the following sections assume that queries have no continuations.

> *Example*: The example:
>
> ```
> C#
> from c in customers
> group c by c.Country into g
> select new { Country = g.Key, CustCount = g.Count() }
> ```
>
> is translated into:
>
> ```
> C#
> from g in
>    (from c in customers
>    group c by c.Country)
> select new { Country = g.Key, CustCount = g.Count() }
> ```
>
> the final translation of which is:
>
> ```
> C#
> customers.
> GroupBy(c => c.Country).
> Select(g => new { Country = g.Key, CustCount = g.Count() })
> ```
>
> *end example*

12.20.3.3 Explicit range variable types

A `from` clause that explicitly specifies a range variable type

```
C#
from «T» «x» in «e»
```

is translated into

```
C#
```

```
from «x» in ( «e» ) . Cast < «T» > ( )
```

A `join` clause that explicitly specifies a range variable type

```
C#

join «T» «x» in «e» on «k1» equals «k2»
```

is translated into

```
C#

join «x» in ( «e» ) . Cast < «T» > ( ) on «k1» equals «k2»
```

The translations in the following sections assume that queries have no explicit range variable types.

> *Example*: The example
>
> ```
> C#
>
> from Customer c in customers
> where c.City == "London"
> select c
> ```
>
> is translated into
>
> ```
> C#
>
> from c in (customers).Cast<Customer>()
> where c.City == "London"
> select c
> ```
>
> the final translation of which is
>
> ```
> C#
>
> customers.
> Cast<Customer>().
> Where(c => c.City == "London")
> ```
>
> *end example*

> *Note*: Explicit range variable types are useful for querying collections that implement the non-generic `IEnumerable` interface, but not the generic `IEnumerable<T>` interface. In the example above, this would be the case if customers were of type `ArrayList`. *end note*

12.20.3.4 Degenerate query expressions

A query expression of the form

```
C#

from «x» in «e» select «x»
```

is translated into

```
C#

( «e» ) . Select ( «x» => «x» )
```

> *Example*: The example
>
> ```
> C#
>
> from c in customers
> select c
> ```
>
> is translated into
>
> ```
> C#
>
> (customers).Select(c => c)
> ```
>
> *end example*

A degenerate query expression is one that trivially selects the elements of the source.

> *Note*: Later phases of the translation (§12.20.3.6 and §12.20.3.7) remove degenerate queries introduced by other translation steps by replacing them with their source. It is important, however, to ensure that the result of a query expression is never the source object itself. Otherwise, returning the result of such a query might inadvertently expose private data (e.g., an element array) to a caller. Therefore this step protects degenerate queries written directly in source code by explicitly calling

> `Select` on the source. It is then up to the implementers of `Select` and other query operators to ensure that these methods never return the source object itself. *end note*

12.20.3.5 From, let, where, join and orderby clauses

A query expression with a second `from` clause followed by a `select` clause

```
C#

from «x1» in «e1»
from «x2» in «e2»
select «v»
```

is translated into

```
C#

( «e1» ) . SelectMany( «x1» => «e2» , ( «x1» , «x2» ) => «v» )
```

> *Example*: The example
>
> ```
> C#
>
> from c in customers
> from o in c.Orders
> select new { c.Name, o.OrderID, o.Total }
> ```
>
> is translated into
>
> ```
> C#
>
> (customers).
> SelectMany(c => c.Orders,
> (c,o) => new { c.Name, o.OrderID, o.Total }
> )
> ```
>
> *end example*

A query expression with a second `from` clause followed by a query body `Q` containing a non-empty set of query body clauses:

```
C#
```

```
from «x1» in «e1»
from «x2» in «e2»
Q
```

is translated into

C#

```
from * in («e1») . SelectMany( «x1» => «e2» ,
                               ( «x1» , «x2» ) => new { «x1» , «x2» } )
Q
```

Example: The example

C#

```
from c in customers
from o in c.Orders
orderby o.Total descending
select new { c.Name, o.OrderID, o.Total }
```

is translated into

C#

```
from * in (customers).
   SelectMany(c => c.Orders, (c,o) => new { c, o })
orderby o.Total descending
select new { c.Name, o.OrderID, o.Total }
```

the final translation of which is

C#

```
customers.
SelectMany(c => c.Orders, (c,o) => new { c, o }).
OrderByDescending(x => x.o.Total).
Select(x => new { x.c.Name, x.o.OrderID, x.o.Total })
```

where `x` is a compiler generated identifier that is otherwise invisible and inaccessible.

end example

A `let` expression along with its preceding `from` clause:

```
C#

from «x» in «e»
let «y» = «f»
...
```

is translated into

```
C#

from * in ( «e» ) . Select ( «x» => new { «x» , «y» = «f» } )
...
```

Example: The example

```
C#

from o in orders
let t = o.Details.Sum(d => d.UnitPrice * d.Quantity)
where t >= 1000
select new { o.OrderID, Total = t }
```

is translated into

```
C#

from * in (orders).Select(
    o => new { o, t = o.Details.Sum(d => d.UnitPrice * d.Quantity) })
where t >= 1000
select new { o.OrderID, Total = t }
```

the final translation of which is

```
C#

orders
    .Select(o => new { o, t = o.Details.Sum(d => d.UnitPrice *
d.Quantity) })
    .Where(x => x.t >= 1000)
    .Select(x => new { x.o.OrderID, Total = x.t })
```

where `x` is a compiler generated identifier that is otherwise invisible and inaccessible.

end example

A `where` expression along with its preceding `from` clause:

```
C#

from «x» in «e»
where «f»
...
```

is translated into

```
C#

from «x» in ( «e» ) . Where ( «x» => «f» )
...
```

A `join` clause immediately followed by a `select` clause

```
C#

from «x1» in «e1»
join «x2» in «e2» on «k1» equals «k2»
select «v»
```

is translated into

```
C#

( «e1» ) . Join( «e2» , «x1» => «k1» , «x2» => «k2» , ( «x1» , «x2» ) => «v»
)
```

> *Example*: The example
>
> ```
> C#
>
> from c in customersh
> join o in orders on c.CustomerID equals o.CustomerID
> select new { c.Name, o.OrderDate, o.Total }
> ```
>
> is translated into
>
> ```
> C#
>
> (customers).Join(
>    orders,
> ```

```
    c => c.CustomerID, o => o.CustomerID,
    (c, o) => new { c.Name, o.OrderDate, o.Total })
```

end example

A `join` clause followed by a query body clause:

```
C#

from «x1» in «e1»
join «x2» in «e2» on «k1» equals «k2»
...
```

is translated into

```
C#

from * in ( «e1» ) . Join(
«e2» , «x1» => «k1» , «x2» => «k2» ,
( «x1» , «x2» ) => new { «x1» , «x2» })
...
```

A `join`-`into` clause immediately followed by a `select` clause

```
C#

from «x1» in «e1»
join «x2» in «e2» on «k1» equals «k2» into «g»
select «v»
```

is translated into

```
C#

( «e1» ) . GroupJoin( «e2» , «x1» => «k1» , «x2» => «k2» ,
                      ( «x1» , «g» ) => «v» )
```

A `join into` clause followed by a query body clause

```
C#

from «x1» in «e1»
join «x2» in «e2» on «k1» equals «k2» into *g»
...
```

is translated into

```
C#

from * in ( «e1» ) . GroupJoin(
   «e2» , «x1» => «k1» , «x2» => «k2» , ( «x1» , «g» ) => new { «x1» , «g»
})
...
```

Example: The example

```
C#

from c in customers
join o in orders on c.CustomerID equals o.CustomerID into co
let n = co.Count()
where n >= 10
select new { c.Name, OrderCount = n }
```

is translated into

```
C#

from * in (customers).GroupJoin(
    orders,
    c => c.CustomerID,
    o => o.CustomerID,
    (c, co) => new { c, co })
let n = co.Count()
where n >= 10
select new { c.Name, OrderCount = n }
```

the final translation of which is

```
C#

customers
    .GroupJoin(
        orders,
        c => c.CustomerID,
        o => o.CustomerID,
        (c, co) => new { c, co })
    .Select(x => new { x, n = x.co.Count() })
    .Where(y => y.n >= 10)
    .Select(y => new { y.x.c.Name, OrderCount = y.n })
```

where `x` and `y` are compiler generated identifiers that are otherwise invisible and inaccessible.

end example

An `orderby` clause and its preceding `from` clause:

```
C#

from «x» in «e»
orderby «k1» , «k2» , ... , «kn»
...
```

is translated into

```
C#

from «x» in ( «e» ) .
OrderBy ( «x» => «k1» ) .
ThenBy ( «x» => «k2» ) .
... .
ThenBy ( «x» => «kn» )
...
```

If an `ordering` clause specifies a descending direction indicator, an invocation of `OrderByDescending` or `ThenByDescending` is produced instead.

Example: The example

```
C#

from o in orders
orderby o.Customer.Name, o.Total descending
select o
```

has the final translation

```
C#

(orders)
    .OrderBy(o => o.Customer.Name)
    .ThenByDescending(o => o.Total)
```

end example

The following translations assume that there are no `let`, `where`, `join` or `orderby` clauses, and no more than the one initial `from` clause in each query expression.

12.20.3.6 Select clauses

A query expression of the form

```
C#

from «x» in «e» select «v»
```

is translated into

```
C#

( «e» ) . Select ( «x» => «v» )
```

except when `«v»` is the identifier `«x»`, the translation is simply

```
C#

( «e» )
```

> *Example*: The example
>
> ```
> C#
>
> from c in customers.Where(c => c.City == "London")
> select c
> ```
>
> is simply translated into
>
> ```
> C#
>
> (customers).Where(c => c.City == "London")
> ```
>
> *end example*

12.20.3.7 Group clauses

A `group` clause

```
C#

from «x» in «e» group «v» by «k»
```

is translated into

```
C#

( «e» ) . GroupBy ( «x» => «k» , «x» => «v» )
```

except when `«v»` is the identifier `«x»`, the translation is

```
C#

( «e» ) . GroupBy ( «x» => «k» )
```

> *Example*: The example
>
> ```
> C#
>
> from c in customers
> group c.Name by c.Country
> ```
>
> is translated into
>
> ```
> C#
>
> (customers).GroupBy(c => c.Country, c => c.Name)
> ```
>
> *end example*

12.20.3.8 Transparent identifiers

Certain translations inject range variables with ***transparent identifiers*** denoted by `*`. Transparent identifiers exist only as an intermediate step in the query-expression translation process.

When a query translation injects a transparent identifier, further translation steps propagate the transparent identifier into anonymous functions and anonymous object initializers. In those contexts, transparent identifiers have the following behavior:

- When a transparent identifier occurs as a parameter in an anonymous function, the members of the associated anonymous type are automatically in scope in the body of the anonymous function.
- When a member with a transparent identifier is in scope, the members of that member are in scope as well.

- When a transparent identifier occurs as a member declarator in an anonymous object initializer, it introduces a member with a transparent identifier.

In the translation steps described above, transparent identifiers are always introduced together with anonymous types, with the intent of capturing multiple range variables as members of a single object. An implementation of C# is permitted to use a different mechanism than anonymous types to group together multiple range variables. The following translation examples assume that anonymous types are used, and shows one possible translation of transparent identifiers.

> *Example*: The example
>
> C#
>
> ```
> from c in customers
> from o in c.Orders
> orderby o.Total descending
> select new { c.Name, o.Total }
> ```
>
> is translated into
>
> C#
>
> ```
> from * in (customers).SelectMany(c => c.Orders, (c,o) => new { c, o })
> orderby o.Total descending
> select new { c.Name, o.Total }
> ```
>
> which is further translated into
>
> C#
>
> ```
> customers
>     .SelectMany(c => c.Orders, (c,o) => new { c, o })
>     .OrderByDescending(* => o.Total)
>     .Select(* => new { c.Name, o.Total })
> ```
>
> which, when transparent identifiers are erased, is equivalent to
>
> C#
>
> ```
> customers
>     .SelectMany(c => c.Orders, (c,o) => new { c, o })
>     .OrderByDescending(x => x.o.Total)
>     .Select(x => new { x.c.Name, x.o.Total })
> ```

where `x` is a compiler generated identifier that is otherwise invisible and inaccessible.

The example

C#

```
from c in customers
join o in orders on c.CustomerID equals o.CustomerID
join d in details on o.OrderID equals d.OrderID
join p in products on d.ProductID equals p.ProductID
select new { c.Name, o.OrderDate, p.ProductName }
```

is translated into

C#

```
from * in (customers).Join(
    orders,
    c => c.CustomerID,
    o => o.CustomerID,
    (c, o) => new { c, o })
join d in details on o.OrderID equals d.OrderID
join p in products on d.ProductID equals p.ProductID
select new { c.Name, o.OrderDate, p.ProductName }
```

which is further reduced to

C#

```
customers
    .Join(orders, c => c.CustomerID,
        o => o.CustomerID, (c, o) => new { c, o })
    .Join(details, * => o.OrderID, d => d.OrderID, (*, d) => new { *, d
})
    .Join(products, * => d.ProductID, p => p.ProductID,
        (*, p) => new { c.Name, o.OrderDate, p.ProductName })
```

the final translation of which is

C#

```
customers
    .Join(orders, c => c.CustomerID,
        o => o.CustomerID, (c, o) => new { c, o })
    .Join(details, x => x.o.OrderID, d => d.OrderID, (x, d) => new { x,
d })
```

```
    .Join(products, y => y.d.ProductID, p => p.ProductID,
        (y, p) => new { y.x.c.Name, y.x.o.OrderDate, p.ProductName })
```

where `x` and `y` are compiler-generated identifiers that are otherwise invisible and inaccessible. *end example*

12.20.4 The query-expression pattern

The ***Query-expression pattern*** establishes a pattern of methods that types can implement to support query expressions.

A generic type `C<T>` supports the query-expression-pattern if its public member methods and the publicly accessible extension methods could be replaced by the following class definition. The members and accessible extenson methods is referred to as the "shape" of a generic type `C<T>`. A generic type is used in order to illustrate the proper relationships between parameter and return types, but it is possible to implement the pattern for non-generic types as well.

C#

```
delegate R Func<T1,R>(T1 arg1);
delegate R Func<T1,T2,R>(T1 arg1, T2 arg2);

class C
{
    public C<T> Cast<T>() { ... }
}

class C<T> : C
{
    public C<T> Where(Func<T,bool> predicate) { ... }
    public C<U> Select<U>(Func<T,U> selector) { ... }
    public C<V> SelectMany<U,V>(Func<T,C<U>> selector,
        Func<T,U,V> resultSelector) { ... }
    public C<V> Join<U,K,V>(C<U> inner, Func<T,K> outerKeySelector,
        Func<U,K> innerKeySelector, Func<T,U,V> resultSelector) { ... }
    public C<V> GroupJoin<U,K,V>(C<U> inner, Func<T,K> outerKeySelector,
        Func<U,K> innerKeySelector, Func<T,C<U>,V> resultSelector) { ... }
    public O<T> OrderBy<K>(Func<T,K> keySelector) { ... }
    public O<T> OrderByDescending<K>(Func<T,K> keySelector) { ... }
    public C<G<K,T>> GroupBy<K>(Func<T,K> keySelector) { ... }
    public C<G<K,E>> GroupBy<K,E>(Func<T,K> keySelector,
        Func<T,E> elementSelector) { ... }
}

class O<T> : C<T>
{
    public O<T> ThenBy<K>(Func<T,K> keySelector) { ... }
    public O<T> ThenByDescending<K>(Func<T,K> keySelector) { ... }
```

```
}

class G<K,T> : C<T>
{
    public K Key { get; }
}
```

The methods above use the generic delegate types `Func<T1, R>` and `Func<T1, T2, R>`, but they could equally well have used other delegate or expression-tree types with the same relationships in parameter and return types.

> *Note*: The recommended relationship between `C<T>` and `O<T>` that ensures that the `ThenBy` and `ThenByDescending` methods are available only on the result of an `OrderBy` or `OrderByDescending`. *end note*

> *Note*: The recommended shape of the result of `GroupBy`—a sequence of sequences, where each inner sequence has an additional `Key` property. *end note*

> *Note*: Because query expressions are translated to method invocations by means of a syntactic mapping, types have considerable flexibility in how they implement any or all of the query-expression pattern. For example, the methods of the pattern can be implemented as instance methods or as extension methods because the two have the same invocation syntax, and the methods can request delegates or expression trees because anonymous functions are convertible to both. Types implementing only some of the query expression pattern support only query expression translations that map to the methods that type supports. *end note*

> *Note*: The `System.Linq` namespace provides an implementation of the query-expression pattern for any type that implements the `System.Collections.Generic.IEnumerable<T>` interface. *end note*

12.21 Assignment operators

12.21.1 General

All but one of the assignment operators assigns a new value to a variable, a property, an event, or an indexer element. The exception, `= ref`, assigns a variable reference (§9.5) to a reference variable (§9.7).

```
ANTLR
```

```
assignment
    : unary_expression assignment_operator expression
    ;

assignment_operator
    : '=' 'ref'? | '+=' | '-=' | '*=' | '/=' | '%=' | '&=' | '|=' | '^=' |
'<<='
    | right_shift_assignment
    ;
```

The left operand of an assignment shall be an expression classified as a variable, or, except for `= ref`, a property access, an indexer access, an event access or a tuple. A declaration expression is not directly permitted as a left operand, but may occur as a step in the evaluation of a deconstructing assignment.

The `=` operator is called the ***simple assignment operator***. It assigns the value or values of the right operand to the variable, property, indexer element or tuple elements given by the left operand. The left operand of the simple assignment operator shall not be an event access (except as described in §15.8.2). The simple assignment operator is described in §12.21.2.

The operator `= ref` is called the ***ref assignment operator***. It makes the right operand, which shall be a *variable_reference* (§9.5), the referent of the reference variable designated by the left operand. The ref assignment operator is described in §12.21.3.

The assignment operators other than the `=` and `= ref` operators are called the ***compound assignment operators***. These operators perform the indicated operation on the two operands, and then assign the resulting value to the variable, property, or indexer element given by the left operand. The compound assignment operators are described in §12.21.4.

The `+=` and `-=` operators with an event access expression as the left operand are called the ***event assignment operators***. No other assignment operator is valid with an event access as the left operand. The event assignment operators are described in §12.21.5.

The assignment operators are right-associative, meaning that operations are grouped from right to left.

> *Example*: An expression of the form `a = b = c` is evaluated as `a = (b = c)`. *end example*

12.21.2 Simple assignment

The `=` operator is called the simple assignment operator.

If the left operand of a simple assignment is of the form `E.P` or `E[Ei]` where `E` has the compile-time type `dynamic`, then the assignment is dynamically bound (§12.3.3). In this case, the compile-time type of the assignment expression is `dynamic`, and the resolution described below will take place at run-time based on the run-time type of `E`. If the left operand is of the form `E[Ei]` where at least one element of `Ei` has the compile-time type `dynamic`, and the compile-time type of `E` is not an array, the resulting indexer access is dynamically bound, but with limited compile-time checking (§12.6.5).

A simple assignment where the left operand is classified as a tuple is also called a ***deconstructing assignment***. If any of the tuple elements of the left operand has an element name, a compile-time error occurs. If any of the tuple elements of the left operand is a *declaration_expression* and any other element is not a *declaration_expression* or a simple discard, a compile-time error occurs.

The type of a simple assignment `x = y` is the type of an assignment to `x` of `y`, which is recursively determined as follows:

- If `x` is a tuple expression `(x1, ..., xn)`, and `y` can be deconstructed to a tuple expression `(y1, ..., yn)` with `n` elements (§12.7), and each assignment to `xi` of `yi` has the type `Ti`, then the assignment has the type `(T1, ..., Tn)`.
- Otherwise, if `x` is classified as a variable, the variable is not `readonly`, `x` has a type `T`, and `y` has an implicit conversion to `T`, then the assignment has the type `T`.
- Otherwise, if `x` is classified as an implicitly typed variable (i.e. an implicitly typed declaration expression) and `y` has a type `T`, then the inferred type of the variable is `T`, and the assignment has the type `T`.
- Otherwise, if `x` is classified as a property or indexer access, the property or indexer has an accessible set accessor, `x` has a type `T`, and `y` has an implicit conversion to `T`, then the assignment has the type `T`.
- Otherwise the assignment is not valid and a binding-time error occurs.

The run-time processing of a simple assignment of the form `x = y` with type `T` is performed as an assignment to `x` of `y` with type `T`, which consists of the following recursive steps:

- `x` is evaluated if it wasn't already.
- If `x` is classified as a variable, `y` is evaluated and, if required, converted to `T` through an implicit conversion (§10.2).
  - If the variable given by `x` is an array element of a *reference_type*, a run-time check is performed to ensure that the value computed for `y` is compatible with

the array instance of which `x` is an element. The check succeeds if `y` is `null`, or if an implicit reference conversion (§10.2.8) exists from the type of the instance referenced by `y` to the actual element type of the array instance containing `x`. Otherwise, a `System.ArrayTypeMismatchException` is thrown.
    - The value resulting from the evaluation and conversion of `y` is stored into the location given by the evaluation of `x`, and is yielded as a result of the assignment.
- If `x` is classified as a property or indexer access:
    - `y` is evaluated and, if required, converted to `T` through an implicit conversion (§10.2).
    - The set accessor of `x` is invoked with the value resulting from the evaluation and conversion of `y` as its value argument.
    - The value resulting from the evaluation and conversion of `y` is yielded as the result of the assignment.
- If `x` is classified as a tuple `(x1, ..., xn)` with arity `n`:
    - `y` is deconstructed with `n` elements to a tuple expression `e`.
    - a result tuple `t` is created by converting `e` to `T` using an implicit tuple conversion.
    - for each `xi` in order from left to right, an assignment to `xi` of `t.Itemi` is performed, except that the `xi` are not evaluated again.
    - `t` is yielded as the result of the assignment.

> *Note*: if the compile time type of `x` is `dynamic` and there is an implicit conversion from the compile time type of `y` to `dynamic`, no runtime resolution is required. *end note*

> *Note*: The array co-variance rules (§17.6) permit a value of an array type `A[]` to be a reference to an instance of an array type `B[]`, provided an implicit reference conversion exists from `B` to `A`. Because of these rules, assignment to an array element of a *reference_type* requires a run-time check to ensure that the value being assigned is compatible with the array instance. In the example

C#

```
string[] sa = new string[10];
object[] oa = sa;
oa[0] = null;               // OK
oa[1] = "Hello";            // OK
oa[2] = new ArrayList();    // ArrayTypeMismatchException
```

> the last assignment causes a `System.ArrayTypeMismatchException` to be thrown because a reference to an `ArrayList` cannot be stored in an element of a `string[]`.
>
> *end note*

When a property or indexer declared in a *struct_type* is the target of an assignment, the instance expression associated with the property or indexer access shall be classified as a variable. If the instance expression is classified as a value, a binding-time error occurs.

> *Note*: Because of §12.8.7, the same rule also applies to fields. *end note*

> *Example*: Given the declarations:

```C#
struct Point
{
   int x, y;

   public Point(int x, int y)
   {
      this.x = x;
      this.y = y;
   }

   public int X
   {
      get { return x; }
      set { x = value; }
   }

   public int Y {
      get { return y; }
      set { y = value; }
   }
}

struct Rectangle
{
    Point a, b;

    public Rectangle(Point a, Point b)
    {
        this.a = a;
        this.b = b;
    }

    public Point A
    {
        get { return a; }
```

```
        set { a = value; }
    }

    public Point B
    {
        get { return b; }
        set { b = value; }
    }
}
```

in the example

```
C#

Point p = new Point();
p.X = 100;
p.Y = 100;
Rectangle r = new Rectangle();
r.A = new Point(10, 10);
r.B = p;
```

the assignments to `p.X`, `p.Y`, `r.A`, and `r.B` are permitted because `p` and `r` are variables. However, in the example

```
C#

Rectangle r = new Rectangle();
r.A.X = 10;
r.A.Y = 10;
r.B.X = 100;
r.B.Y = 100;
```

the assignments are all invalid, since `r.A` and `r.B` are not variables.

end example

12.21.3 Ref assignment

The `= ref` operator is known as the *ref assignment* operator.

The left operand shall be an expression that binds to a reference variable (§9.7), a reference parameter (other than `this`), an output parameter, or an input parameter. The right operand shall be an expression that yields a *variable_reference* designating a value of the same type as the left operand.

It is a compile time error if the ref-safe-context (§9.7.2) of the left operand is wider than the ref-safe-context of the right operand.

The right operand shall be definitely assigned at the point of the ref assignment.

When the left operand binds to an `out` parameter, it is an error if that `out` parameter has not been definitely assigned at the beginning of the ref assignment operator.

If the left operand is a writeable ref (i.e., it designates anything other than a `ref readonly` local or `in` parameter), then the right operand shall be a writeable *variable_reference*. If the right operand variable is writeable, the left operand may be a writeable or read-only ref.

The operation makes the left operand an alias of the right operand variable. The alias may be made read-only even if the right operand variable is writeable.

The ref assignment operator yields a *variable_reference* of the assigned type. It is writeable if the left operand is writeable.

The ref assignment operator shall not read the storage location referenced by the right operand.

> *Example*: Here are some examples of using `= ref`:
>
> C#
>
> ```
> public static int M1() { ... }
> public static ref int M2() { ... }
> public static ref uint M2u() { ... }
> public static ref readonly int M3() { ... }
> public static void Test()
> {
> int v = 42;
> ref int r1 = ref v; // OK, r1 refers to v, which has value 42
> r1 = ref M1();      // Error; M1 returns a value, not a reference
> r1 = ref M2();      // OK; makes an alias
> r1 = ref M2u();     // Error; lhs and rhs have different types
> r1 = ref M3();    // error; M3 returns a ref readonly, which r1 cannot
> honor
> ref readonly int r2 = ref v; // OK; make readonly alias to ref
> r2 = ref M2();      // OK; makes an alias, adding read-only protection
> r2 = ref M3();      // OK; makes an alias and honors the read-only
> r2 = ref (r1 = ref M2());  // OK; r1 is an alias to a writable variable,
>              // r2 is an alias (with read-only access) to the same
> variable
> }
> ```
>
> *end example*

> *Note*: When reading code using an `= ref` operator, it can be tempting to read the `ref` part as being part of the operand. This is particularly confusing when the operand is a conditional `?:` expression. For example, when reading `ref int a = ref b ? ref x : ref y;` it's important to read this as `= ref` being the operator, and `b ? ref x : ref y` being the right operand: `ref int a = ref (b ? ref x : ref y);`. Importantly, the expression `ref b` is *not* part of that statement, even though it might appear so at first glance. *end note*

12.21.4 Compound assignment

If the left operand of a compound assignment is of the form `E.P` or `E[Ei]` where `E` has the compile-time type `dynamic`, then the assignment is dynamically bound (§12.3.3). In this case, the compile-time type of the assignment expression is `dynamic`, and the resolution described below will take place at run-time based on the run-time type of `E`. If the left operand is of the form `E[Ei]` where at least one element of `Ei` has the compile-time type `dynamic`, and the compile-time type of `E` is not an array, the resulting indexer access is dynamically bound, but with limited compile-time checking (§12.6.5).

An operation of the form `x «op»= y` is processed by applying binary operator overload resolution (§12.4.5) as if the operation was written `x «op» y`. Then,

- If the return type of the selected operator is implicitly convertible to the type of `x`, the operation is evaluated as `x = x «op» y`, except that `x` is evaluated only once.
- Otherwise, if the selected operator is a predefined operator, if the return type of the selected operator is explicitly convertible to the type of `x` , and if `y` is implicitly convertible to the type of `x` or the operator is a shift operator, then the operation is evaluated as `x = (T)(x «op» y)`, where `T` is the type of `x`, except that `x` is evaluated only once.
- Otherwise, the compound assignment is invalid, and a binding-time error occurs.

The term "evaluated only once" means that in the evaluation of `x «op» y`, the results of any constituent expressions of `x` are temporarily saved and then reused when performing the assignment to `x`.

> *Example*: In the assignment `A()[B()] += C()`, where `A` is a method returning `int[]`, and `B` and `C` are methods returning `int`, the methods are invoked only once, in the order `A`, `B`, `C`. *end example*

When the left operand of a compound assignment is a property access or indexer access, the property or indexer shall have both a get accessor and a set accessor. If this

is not the case, a binding-time error occurs.

The second rule above permits `x «op»= y` to be evaluated as `x = (T)(x «op» y)` in certain contexts. The rule exists such that the predefined operators can be used as compound operators when the left operand is of type `sbyte`, `byte`, `short`, `ushort`, or `char`. Even when both arguments are of one of those types, the predefined operators produce a result of type `int`, as described in §12.4.7.3. Thus, without a cast it would not be possible to assign the result to the left operand.

The intuitive effect of the rule for predefined operators is simply that `x «op»= y` is permitted if both of `x «op» y` and `x = y` are permitted.

> *Example*: In the following code
>
> C#
>
> ```
> byte b = 0;
> char ch = '\0';
> int i = 0;
> b += 1;            // OK
> b += 1000;         // Error, b = 1000 not permitted
> b += i;            // Error, b = i not permitted
> b += (byte)i;      // OK
> ch += 1;           // Error, ch = 1 not permitted
> ch += (char)1;     // OK
> ```
>
> the intuitive reason for each error is that a corresponding simple assignment would also have been an error.
>
> *end example*

> *Note*: This also means that compound assignment operations support lifted operators. Since a compound assignment `x «op»= y` is evaluated as either `x = x «op» y` or `x = (T)(x «op» y)`, the rules of evaluation implicitly cover lifted operators. *end note*

12.21.5 Event assignment

If the left operand of `a += or -=` operator is classified as an event access, then the expression is evaluated as follows:

- The instance expression, if any, of the event access is evaluated.
- The right operand of the `+=` or `-=` operator is evaluated, and, if required, converted to the type of the left operand through an implicit conversion (§10.2).

- An event accessor of the event is invoked, with an argument list consisting of the value computed in the previous step. If the operator was `+=`, the add accessor is invoked; if the operator was `-=`, the remove accessor is invoked.

An event assignment expression does not yield a value. Thus, an event assignment expression is valid only in the context of a *statement_expression* (§13.7).

12.22 Expression

An *expression* is either a *non_assignment_expression* or an *assignment*.

ANTLR

```
expression
    : non_assignment_expression
    | assignment
    ;

non_assignment_expression
    : declaration_expression
    | conditional_expression
    | lambda_expression
    | query_expression
    ;
```

12.23 Constant expressions

A constant expression is an expression that shall be fully evaluated at compile-time.

ANTLR

```
constant_expression
    : expression
    ;
```

A constant expression shall either have the value `null` or one of the following types:

- `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, `double`, `decimal`, `bool`, `string`;
- an enumeration type; or
- a default value expression (§12.8.20) for a reference type.

Only the following constructs are permitted in constant expressions:

- Literals (including the `null` literal).
- References to `const` members of class and struct types.
- References to members of enumeration types.
- References to local constants.
- Parenthesized subexpressions, which are themselves constant expressions.
- Cast expressions.
- `checked` and `unchecked` expressions.
- `nameof` expressions.
- The predefined `+`, `-`, `!`, and `~` unary operators.
- The predefined `+`, `-`, `*`, `/`, `%`, `<<`, `>>`, `&`, `|`, `^`, `&&`, `||`, `==`, `!=`, `<`, `>`, `<=`, and `>=` binary operators.
- The `?:` conditional operator.
- `sizeof` expressions, provided the unmanaged-type is one of the types specified in §23.6.9 for which `sizeof` returns a constant value.
- Default value expressions, provided the type is one of the types listed above.

The following conversions are permitted in constant expressions:

- Identity conversions
- Numeric conversions
- Enumeration conversions
- Constant expression conversions
- Implicit and explicit reference conversions, provided the source of the conversions is a constant expression that evaluates to the `null` value.

> *Note*: Other conversions including boxing, unboxing, and implicit reference conversions of non-`null` values are not permitted in constant expressions. *end note*

> *Example*: In the following code
>
> C#
>
> ```
> class C
> {
>     const object i = 5;          // error: boxing conversion not
> permitted
>     const object str = "hello"; // error: implicit reference conversion
> }
> ```
>
> the initialization of `i` is an error because a boxing conversion is required. The initialization of `str` is an error because an implicit reference conversion from a non-`null` value is required.

> *end example*

Whenever an expression fulfills the requirements listed above, the expression is evaluated at compile-time. This is true even if the expression is a subexpression of a larger expression that contains non-constant constructs.

The compile-time evaluation of constant expressions uses the same rules as run-time evaluation of non-constant expressions, except that where run-time evaluation would have thrown an exception, compile-time evaluation causes a compile-time error to occur.

Unless a constant expression is explicitly placed in an `unchecked` context, overflows that occur in integral-type arithmetic operations and conversions during the compile-time evaluation of the expression always cause compile-time errors (§12.8.19).

Constant expressions are required in the contexts listed below and this is indicated in the grammar by using *constant_expression*. In these contexts, a compile-time error occurs if an expression cannot be fully evaluated at compile-time.

- Constant declarations (§15.4)
- Enumeration member declarations (§19.4)
- Default arguments of formal parameter lists (§15.6.2)
- `case` labels of a `switch` statement (§13.8.3).
- `goto case` statements (§13.10.4)
- Dimension lengths in an array creation expression (§12.8.16.5) that includes an initializer.
- Attributes (§22)
- In a *constant_pattern* (§11.2.3)

An implicit constant expression conversion (§10.2.11) permits a constant expression of type `int` to be converted to `sbyte`, `byte`, `short`, `ushort`, `uint`, or `ulong`, provided the value of the constant expression is within the range of the destination type.

12.24 Boolean expressions

A *boolean_expression* is an expression that yields a result of type `bool`; either directly or through application of `operator true` in certain contexts as specified in the following:

ANTLR

```
boolean_expression
    : expression
```

```
    ;
```

The controlling conditional expression of an *if_statement* (§13.8.2), *while_statement* (§13.9.2), *do_statement* (§13.9.3), or *for_statement* (§13.9.4) is a *boolean_expression*. The controlling conditional expression of the `?:` operator (§12.18) follows the same rules as a *boolean_expression*, but for reasons of operator precedence is classified as a *null_coalescing_expression*.

A *boolean_expression* `E` is required to be able to produce a value of type `bool`, as follows:

- If E is implicitly convertible to `bool` then at run-time that implicit conversion is applied.
- Otherwise, unary operator overload resolution (§12.4.4) is used to find a unique best implementation of `operator true` on `E`, and that implementation is applied at run-time.
- If no such operator is found, a binding-time error occurs.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

13 Statements

Article • 02/07/2024

13.1 General

C# provides a variety of statements.

> *Note*: Most of these statements will be familiar to developers who have programmed in C and C++. *end note*

ANTLR

```
statement
    : labeled_statement
    | declaration_statement
    | embedded_statement
    ;

embedded_statement
    : block
    | empty_statement
    | expression_statement
    | selection_statement
    | iteration_statement
    | jump_statement
    | try_statement
    | checked_statement
    | unchecked_statement
    | lock_statement
    | using_statement
    | yield_statement
    | unsafe_statement   // unsafe code support
    | fixed_statement    // unsafe code support
    ;
```

unsafe_statement (§23.2) and *fixed_statement* (§23.7) are only available in unsafe code (§23).

The *embedded_statement* nonterminal is used for statements that appear within other statements. The use of *embedded_statement* rather than *statement* excludes the use of declaration statements and labeled statements in these contexts.

> *Example*: The code

C#

```
void F(bool b)
{
   if (b)
      int i = 44;
}
```

results in a compile-time error because an `if` statement requires an *embedded_statement* rather than a *statement* for its `if` branch. If this code were permitted, then the variable `i` would be declared, but it could never be used. Note, however, that by placing `i`'s declaration in a block, the example is valid.

end example

13.2 End points and reachability

Every statement has an ***end point***. In intuitive terms, the end point of a statement is the location that immediately follows the statement. The execution rules for composite statements (statements that contain embedded statements) specify the action that is taken when control reaches the end point of an embedded statement.

Example: When control reaches the end point of a statement in a block, control is transferred to the next statement in the block. *end example*

If a statement can possibly be reached by execution, the statement is said to be ***reachable***. Conversely, if there is no possibility that a statement will be executed, the statement is said to be ***unreachable***.

Example: In the following code

C#

```
void F()
{
    Console.WriteLine("reachable");
    goto Label;
    Console.WriteLine("unreachable");
  Label:
    Console.WriteLine("reachable");
}
```

the second invocation of Console.WriteLine is unreachable because there is no possibility that the statement will be executed.

end example

A warning is reported if a statement other than *throw_statement*, *block*, or *empty_statement* is unreachable. It is specifically not an error for a statement to be unreachable.

Note: To determine whether a particular statement or end point is reachable, the compiler performs flow analysis according to the reachability rules defined for each statement. The flow analysis takes into account the values of constant expressions (§12.23) that control the behavior of statements, but the possible values of non-constant expressions are not considered. In other words, for purposes of control flow analysis, a non-constant expression of a given type is considered to have any possible value of that type.

In the example

C#

```
void F()
{
    const int i = 1;
    if (i == 2)
        Console.WriteLine("unreachable");
}
```

the Boolean expression of the `if` statement is a constant expression because both operands of the `==` operator are constants. As the constant expression is evaluated at compile-time, producing the value `false`, the `Console.WriteLine` invocation is considered unreachable. However, if `i` is changed to be a local variable

C#

```
void F()
{
    int i = 1;
    if (i == 2)
        Console.WriteLine("reachable");
}
```

the `Console.WriteLine` invocation is considered reachable, even though, in reality, it will never be executed.

end note

The *block* of a function member or an anonymous function is always considered reachable. By successively evaluating the reachability rules of each statement in a block, the reachability of any given statement can be determined.

> *Example*: In the following code
>
> ```C#
> void F(int x)
> {
>     Console.WriteLine("start");
>     if (x < 0)
>         Console.WriteLine("negative");
> }
> ```
>
> the reachability of the second `Console.WriteLine` is determined as follows:
>
> - The first `Console.WriteLine` expression statement is reachable because the block of the `F` method is reachable (§13.3).
> - The end point of the first `Console.WriteLine` expression statement is reachable because that statement is reachable (§13.7 and §13.3).
> - The `if` statement is reachable because the end point of the first `Console.WriteLine` expression statement is reachable (§13.7 and §13.3).
> - The second `Console.WriteLine` expression statement is reachable because the Boolean expression of the `if` statement does not have the constant value `false`.
>
> *end example*

There are two situations in which it is a compile-time error for the end point of a statement to be reachable:

- Because the `switch` statement does not permit a switch section to "fall through" to the next switch section, it is a compile-time error for the end point of the statement list of a switch section to be reachable. If this error occurs, it is typically an indication that a `break` statement is missing.

- It is a compile-time error for the end point of the block of a function member or an anonymous function that computes a value to be reachable. If this error occurs, it typically is an indication that a `return` statement is missing (§13.10.5).

13.3 Blocks

13.3.1 General

A *block* permits multiple statements to be written in contexts where a single statement is allowed.

ANTLR

```
block
    : '{' statement_list? '}'
    ;
```

A *block* consists of an optional *statement_list* (§13.3.2), enclosed in braces. If the statement list is omitted, the block is said to be empty.

A block may contain declaration statements (§13.6). The scope of a local variable or constant declared in a block is the block.

A block is executed as follows:

- If the block is empty, control is transferred to the end point of the block.
- If the block is not empty, control is transferred to the statement list. When and if control reaches the end point of the statement list, control is transferred to the end point of the block.

The statement list of a block is reachable if the block itself is reachable.

The end point of a block is reachable if the block is empty or if the end point of the statement list is reachable.

A *block* that contains one or more `yield` statements (§13.15) is called an iterator block. Iterator blocks are used to implement function members as iterators (§15.14). Some additional restrictions apply to iterator blocks:

- It is a compile-time error for a `return` statement to appear in an iterator block (but `yield return` statements are permitted).
- It is a compile-time error for an iterator block to contain an unsafe context (§23.2). An iterator block always defines a safe context, even when its declaration is nested in an unsafe context.

13.3.2 Statement lists

A ***statement list*** consists of one or more statements written in sequence. Statement lists occur in *block*s (§13.3) and in *switch_block*s (§13.8.3).

ANTLR

```
statement_list
    : statement+
    ;
```

A statement list is executed by transferring control to the first statement. When and if control reaches the end point of a statement, control is transferred to the next statement. When and if control reaches the end point of the last statement, control is transferred to the end point of the statement list.

A statement in a statement list is reachable if at least one of the following is true:

- The statement is the first statement and the statement list itself is reachable.
- The end point of the preceding statement is reachable.
- The statement is a labeled statement and the label is referenced by a reachable `goto` statement.

The end point of a statement list is reachable if the end point of the last statement in the list is reachable.

13.4 The empty statement

An *empty_statement* does nothing.

ANTLR

```
empty_statement
    : ';'
    ;
```

An empty statement is used when there are no operations to perform in a context where a statement is required.

Execution of an empty statement simply transfers control to the end point of the statement. Thus, the end point of an empty statement is reachable if the empty statement is reachable.

> *Example*: An empty statement can be used when writing a `while` statement with a null body:
>
> C#

```
bool ProcessMessage() {...}
void ProcessMessages()
{
    while (ProcessMessage())
        ;
}
```

Also, an empty statement can be used to declare a label just before the closing "`}`" of a block:

C#

```
void F(bool done)
{
    ...
    if (done)
    {
        goto exit;
    }
    ...
  exit:
    ;
}
```

end example

13.5 Labeled statements

A *labeled_statement* permits a statement to be prefixed by a label. Labeled statements are permitted in blocks, but are not permitted as embedded statements.

ANTLR

```
labeled_statement
    : identifier ':' statement
    ;
```

A labeled statement declares a label with the name given by the *identifier*. The scope of a label is the whole block in which the label is declared, including any nested blocks. It is a compile-time error for two labels with the same name to have overlapping scopes.

A label can be referenced from `goto` statements (§13.10.4) within the scope of the label.

> *Note*: This means that `goto` statements can transfer control within blocks and out of blocks, but never into blocks. *end note*

Labels have their own declaration space and do not interfere with other identifiers.

> *Example*: The example
>
> C#
>
> ```
> int F(int x)
> {
>     if (x >= 0)
>     {
>         goto x;
>     }
>     x = -x;
>   x:
>     return x;
> }
> ```
>
> is valid and uses the name x as both a parameter and a label.
>
> *end example*

Execution of a labeled statement corresponds exactly to execution of the statement following the label.

In addition to the reachability provided by normal flow of control, a labeled statement is reachable if the label is referenced by a reachable `goto` statement, unless the `goto` statement is inside the `try` block or a `catch` block of a *try_statement* that includes a `finally` block whose end point is unreachable, and the labeled statement is outside the *try_statement*.

13.6 Declaration statements

13.6.1 General

A *declaration_statement* declares one or more local variables, one or more local constants, or a local function. Declaration statements are permitted in blocks and switch blocks, but are not permitted as embedded statements.

ANTLR

```
declaration_statement
    : local_variable_declaration ';'
    | local_constant_declaration ';'
```

```
    | local_function_declaration
    ;
```

A local variable is declared using a *local_variable_declaration* (§13.6.2). A local constant is declared using a *local_constant_declaration* (§13.6.3). A local function is declared using a *local_function_declaration* (§13.6.4).

The declared names are introduced into the nearest enclosing declaration space (§7.3).

13.6.2 Local variable declarations

13.6.2.1 General

A *local_variable_declaration* declares one or more local variables.

ANTLR

```
local_variable_declaration
    : implicitly_typed_local_variable_declaration
    | explicitly_typed_local_variable_declaration
    | ref_local_variable_declaration
    ;
```

Local variable declarations fall into one of the three categories: implicitly typed, explicitly typed, and ref local.

Implicitly typed declarations contain the contextual keyword (§6.4.4) `var` resulting in a syntactic ambiguity between the three categories which is resolved as follows:

- If there is no type named `var` in scope and the input matches *implicitly_typed_local_variable_declaration* then it is chosen;
- Otherwise if a type named `var` is in scope then *implicitly_typed_local_variable_declaration* is not considered as a possible match.

Within a *local_variable_declaration* each variable is introduced by a ***declarator***, which is one of *implicitly_typed_local_variable_declarator*, *explicitly_typed_local_variable_declarator* or *ref_local_variable_declarator* for implicitly typed, explicitly typed and ref local variables respectively. The declarator defines the name (*identifier*) and initial value, if any, of the introduced variable.

If there are multiple declarators in a declaration then they are processed, including any initializing expressions, in order left to right (§9.4.4.5).

Note: For a *local_variable_declaration* not occuring as a *for_initializer* (§13.9.4) or *resource_acquisition* (§13.14) this left to right order is equivalent to each declarator being within a separate *local_variable_declaration*. For example:

```
C#

void F()
{
    int x = 1, y, z = x * 2;
}
```

is equivalent to:

```
C#

void F()
{
    int x = 1;
    int y;
    int z = x * 2;
}
```

end note

The value of a local variable is obtained in an expression using a *simple_name* (§12.8.4). A local variable shall be definitely assigned (§9.4) at each location where its value is obtained. Each local variable introduced by a *local_variable_declaration* is *initially unassigned* (§9.4.3). If a declarator has an initializing expression then the introduced local variable is classified as *assigned* at the end of the declarator (§9.4.4.5).

The scope of a local variable introduced by a *local_variable_declaration* is defined as follows (§7.7):

- If the declaration occurs as a *for_initializer* then the scope is the *for_initializer*, *for_condition*, *for_iterator*, and *embedded_statement* (§13.9.4);
- If the declaration occurs as a *resource_acquisition* then the scope is the outermost block of the semantically equivalent expansion of the *using_statement* (§13.14);
- Otherwise the scope is the block in which the declaration occurs.

It is an error to refer to a local variable by name in a textual position that precedes its declarator, or within any initializing expression within its declarator. Within the scope of a local variable, it is a compile-time error to declare another local variable, local function or constant with the same name.

The ref-safe-context (§9.7.2) of a ref local variable is the ref-safe-context of its initializing *variable_reference*. The ref-safe-context of non-ref local variables is *declaration-block*.

13.6.2.2 Implicitly typed local variable declarations

ANTLR

```
implicitly_typed_local_variable_declaration
    : 'var' implicitly_typed_local_variable_declarator
    | ref_kind 'var' ref_local_variable_declarator
    ;

implicitly_typed_local_variable_declarator
    : identifier '=' expression
    ;
```

An *implicity_typed_local_variable_declaration* introduces a single local variable, *identifier*. The *expression* or *variable_reference* shall have a compile-time type, `T`. The first alternative declares a variable with type `T` and an initial value of *expression*. The second alternative declares a ref variable with type `ref T` and an initial value of `ref` *variable_reference*.

> *Example*:
>
> C#
>
> ```
> var i = 5;
> var s = "Hello";
> var d = 1.0;
> var numbers = new int[] {1, 2, 3};
> var orders = new Dictionary<int,Order>();
> ref var j = ref i;
> ref readonly var k = ref i;
> ```
>
> The implicitly typed local variable declarations above are precisely equivalent to the following explicitly typed declarations:
>
> C#
>
> ```
> int i = 5;
> string s = "Hello";
> double d = 1.0;
> int[] numbers = new int[] {1, 2, 3};
> Dictionary<int,Order> orders = new Dictionary<int,Order>();
> ref int j = ref i;
> ref readonly int k = ref i;
> ```

The following are incorrect implicitly typed local variable declarations:

```
C#

var x;                  // Error, no initializer to infer type from
var y = {1, 2, 3};      // Error, array initializer not permitted
var z = null;           // Error, null does not have a type
var u = x => x + 1;     // Error, anonymous functions do not have a type
var v = v++;            // Error, initializer cannot refer to v itself
```

end example

13.6.2.3 Explicitly typed local variable declarations

```
ANTLR

explicitly_typed_local_variable_declaration
    : type explicitly_typed_local_variable_declarators
    ;

explicitly_typed_local_variable_declarators
    : explicitly_typed_local_variable_declarator
      (',' explicitly_typed_local_variable_declarator)*
    ;

explicitly_typed_local_variable_declarator
    : identifier ('=' local_variable_initializer)?
    ;

local_variable_initializer
    : expression
    | array_initializer
    ;
```

An *explicity_typed_local_variable_declaration* introduces one or more local variables with the specified *type*.

If a *local_variable_initializer* is present then its type shall be appropriate according to the rules of simple assignment (§12.21.2) or array initialization (§17.7) and its value is assigned as the initial value of the variable.

13.6.2.4 Ref local variable declarations

```
ANTLR

ref_local_variable_declaration
    : ref_kind type ref_local_variable_declarators
```

```
    ;

ref_local_variable_declarators
    : ref_local_variable_declarator (',' ref_local_variable_declarator)*
    ;

ref_local_variable_declarator
    : identifier '=' 'ref' variable_reference
    ;
```

The initializing *variable_reference* shall have type *type* and meet the same requirements as for a *ref assignment* (§12.21.3).

If *ref_kind* is `ref readonly`, the *identifier*(s) being declared are references to variables that are treated as read-only. Otherwise, if *ref_kind* is `ref`, the *identifier*(s) being declared are references to variables that shall be writable.

It is a compile-time error to declare a ref local variable, or a variable of a `ref struct` type, within a method declared with the *method_modifier* `async`, or within an iterator (§15.14).

13.6.3 Local constant declarations

A *local_constant_declaration* declares one or more local constants.

ANTLR

```
local_constant_declaration
    : 'const' type constant_declarators
    ;

constant_declarators
    : constant_declarator (',' constant_declarator)*
    ;

constant_declarator
    : identifier '=' constant_expression
    ;
```

The *type* of a *local_constant_declaration* specifies the type of the constants introduced by the declaration. The type is followed by a list of *constant_declarators*, each of which introduces a new constant. A *constant_declarator* consists of an *identifier* that names the constant, followed by an "`=`" token, followed by a *constant_expression* (§12.23) that gives the value of the constant.

The *type* and *constant_expression* of a local constant declaration shall follow the same rules as those of a constant member declaration (§15.4).

The value of a local constant is obtained in an expression using a *simple_name* (§12.8.4).

The scope of a local constant is the block in which the declaration occurs. It is an error to refer to a local constant in a textual position that precedes the end of its *constant_declarator*. Within the scope of a local constant, it is a compile-time error to declare another local variable, local function or constant with the same name.

A local constant declaration that declares multiple constants is equivalent to multiple declarations of single constants with the same type.

13.6.4 Local function declarations

A *local_function_declaration* declares a local function.

ANTLR

```
local_function_declaration
    : local_function_modifier* return_type local_function_header
      local_function_body
    | ref_local_function_modifier* ref_kind ref_return_type
      local_function_header ref_local_function_body
    ;

local_function_header
    : identifier '(' formal_parameter_list? ')'
    | identifier type_parameter_list '(' formal_parameter_list? ')'
      type_parameter_constraints_clause*
    ;

local_function_modifier
    : ref_local_function_modifier
    | 'async'
    ;

ref_local_function_modifier
    : 'static'
    | unsafe_modifier   // unsafe code support
    ;

local_function_body
    : block
    | '=>' null_conditional_invocation_expression ';'
    | '=>' expression ';'
    ;

ref_local_function_body
    : block
```

```
    | '=>' 'ref' variable_reference ';'
    ;
```

Grammar note: When recognising a *local_function_body* if both the *null_conditional_invocation_expression* and *expression* alternatives are applicable then the former shall be chosen. (§15.6.1)

> *Example*: There are two common use cases for local functions: iterator methods and async methods. In iterator methods, any exceptions are observed only when calling code that enumerates the returned sequence. In async methods, any exceptions are only observed when the returned Task is awaited. The following example demonstrates separating parameter validation from the iterator implementation using a local function:
>
> C#
>
> ```
> public static IEnumerable<char> AlphabetSubset(char start, char end)
> {
>     if (start < 'a' || start > 'z')
>     {
>         throw new ArgumentOutOfRangeException(paramName: nameof(start),
>             message: "start must be a letter");
>     }
>     if (end < 'a' || end > 'z')
>     {
>         throw new ArgumentOutOfRangeException(paramName: nameof(end),
>             message: "end must be a letter");
>     }
>     if (end <= start)
>     {
>         throw new ArgumentException(
>             $"{nameof(end)} must be greater than {nameof(start)}");
>     }
>     return AlphabetSubsetImplementation();
>
>     IEnumerable<char> AlphabetSubsetImplementation()
>     {
>         for (var c = start; c < end; c++)
>         {
>             yield return c;
>         }
>     }
> }
> ```
>
> *end example*

Unless specified otherwise below, the semantics of all grammar elements is the same as for *method_declaration* (§15.6.1), read in the context of a local function instead of a

method.

The *identifier* of a *local_function_declaration* shall be unique in its declared block scope, including any enclosing local variable declaration spaces. One consequence of this is that overloaded *local_function_declaration*s are not allowed.

A *local_function_declaration* may include one `async` (§15.15) modifier and one `unsafe` (§23.1) modifier. If the declaration includes the `async` modifier then the return type shall be `void` or a `«TaskType»` type (§15.15.1). If the declaration includes the `static` modifier, the function is a ***static local function***; otherwise, it is a ***non-static local function***. It is a compile-time error for *type_parameter_list* or *formal_parameter_list* to contain *attributes*. If the local function is declared in an unsafe context (§23.2), the local function may include unsafe code, even if the local function declaration doesn't include the `unsafe` modifier.

A local function is declared at block scope. A non-static local function may capture variables from the enclosing scope while a static local function shall not (so it has no access to enclosing locals, parameters, non-static local functions, or `this`). It is a compile-time error if a captured variable is read by the body of a non-static local function but is not definitely assigned before each call to the function. The compiler shall determine which variables are definitely assigned on return (§9.4.4.33).

When the type of `this` is a struct type, it is a compile-time error for the body of a local function to access `this`. This is true whether the access is explicit (as in `this.x`) or implicit (as in `x` where `x` is an instance member of the struct). This rule only prohibits such access and does not affect whether member lookup results in a member of the struct.

It is a compile-time error for the body of the local function to contain a `goto` statement, a `break` statement, or a `continue` statement whose target is outside the body of the local function.

> *Note*: the above rules for `this` and `goto` mirror the rules for anonymous functions in §12.19.3. *end note*

A local function may be called from a lexical point prior to its declaration. However, it is a compile-time error for the function to be declared lexically prior to the declaration of a variable used in the local function (§7.7).

It is a compile-time error for a local function to declare a parameter, type parameter or local variable with the same name as one declared in any enclosing local variable declaration space.

Local function bodies are always reachable. The endpoint of a local function declaration is reachable if the beginning point of the local function declaration is reachable.

> *Example*: In the following example, the body of `L` is reachable even though the beginning point of `L` is not reachable. Because the beginning point of `L` isn't reachable, the statement following the endpoint of `L` is not reachable:
>
> ```
> C#
> ```
>
> ```
> class C
> {
>     int M()
>     {
>         L();
>         return 1;
>
>         // Beginning of L is not reachable
>         int L()
>         {
>             // The body of L is reachable
>             return 2;
>         }
>         // Not reachable, because beginning point of L is not reachable
>         return 3;
>     }
> }
> ```
>
> In other words, the location of a local function declaration doesn't affect the reachability of any statements in the containing function. *end example*

If the type of the argument to a local function is `dynamic`, the function to be called shall be resolved at compile time, not runtime.

A local function shall not be used in an expression tree.

A static local function

- May reference static members, type parameters, constant definitions and static local functions from the enclosing scope.
- Shall not reference `this` or `base` nor instance members from an implicit `this` reference, nor local variables, parameters, or non-static local functions from the enclosing scope. However, all these are permitted in a `nameof()` expression.

13.7 Expression statements

An *expression_statement* evaluates a given expression. The value computed by the expression, if any, is discarded.

ANTLR

```
expression_statement
    : statement_expression ';'
    ;

statement_expression
    : null_conditional_invocation_expression
    | invocation_expression
    | object_creation_expression
    | assignment
    | post_increment_expression
    | post_decrement_expression
    | pre_increment_expression
    | pre_decrement_expression
    | await_expression
    ;
```

Not all expressions are permitted as statements.

> *Note*: In particular, expressions such as `x + y` and `x == 1`, that merely compute a value (which will be discarded), are not permitted as statements. *end note*

Execution of an *expression_statement* evaluates the contained expression and then transfers control to the end point of the *expression_statement*. The end point of an *expression_statement* is reachable if that *expression_statement* is reachable.

13.8 Selection statements

13.8.1 General

Selection statements select one of a number of possible statements for execution based on the value of some expression.

ANTLR

```
selection_statement
    : if_statement
    | switch_statement
    ;
```

13.8.2 The if statement

The `if` statement selects a statement for execution based on the value of a Boolean expression.

ANTLR

```
if_statement
    : 'if' '(' boolean_expression ')' embedded_statement
    | 'if' '(' boolean_expression ')' embedded_statement
      'else' embedded_statement
    ;
```

An `else` part is associated with the lexically nearest preceding `if` that is allowed by the syntax.

> *Example*: Thus, an `if` statement of the form
>
> C#
>
> ```
> if (x) if (y) F(); else G();
> ```
>
> is equivalent to
>
> C#
>
> ```
> if (x)
> {
>     if (y)
>     {
>         F();
>     }
>     else
>     {
>         G();
>     }
> }
> ```
>
> *end example*

An `if` statement is executed as follows:

- The *boolean_expression* (§12.24) is evaluated.
- If the Boolean expression yields `true`, control is transferred to the first embedded statement. When and if control reaches the end point of that statement, control is

transferred to the end point of the `if` statement.

- If the Boolean expression yields `false` and if an `else` part is present, control is transferred to the second embedded statement. When and if control reaches the end point of that statement, control is transferred to the end point of the `if` statement.
- If the Boolean expression yields `false` and if an `else` part is not present, control is transferred to the end point of the `if` statement.

The first embedded statement of an `if` statement is reachable if the `if` statement is reachable and the Boolean expression does not have the constant value `false`.

The second embedded statement of an `if` statement, if present, is reachable if the `if` statement is reachable and the Boolean expression does not have the constant value `true`.

The end point of an `if` statement is reachable if the end point of at least one of its embedded statements is reachable. In addition, the end point of an `if` statement with no `else` part is reachable if the `if` statement is reachable and the Boolean expression does not have the constant value `true`.

13.8.3 The switch statement

The `switch` statement selects for execution a statement list having an associated switch label that corresponds to the value of the switch expression.

ANTLR

```
switch_statement
    : 'switch' '(' expression ')' switch_block
    ;

switch_block
    : '{' switch_section* '}'
    ;

switch_section
    : switch_label+ statement_list
    ;

switch_label
    : 'case' pattern case_guard?  ':'
    | 'default' ':'
    ;

case_guard
```

```
    : 'when' expression
    ;
```

A *switch_statement* consists of the keyword `switch`, followed by a parenthesized expression (called the ***switch expression***), followed by a *switch_block*. The *switch_block* consists of zero or more *switch_sections*, enclosed in braces. Each *switch_section* consists of one or more *switch_labels* followed by a *statement_list* (§13.3.2). Each *switch_label* containing `case` has an associated pattern (§11) against which the value of the switch expression is tested. If *case_guard* is present, its expression shall be implicitly convertible to the type `bool` and that expression is evaluated as an additional condition for the case to be considered satisfied.

The ***governing type*** of a `switch` statement is established by the switch expression.

- If the type of the switch expression is `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `bool`, `string`, or an *enum_type*, or if it is the nullable value type corresponding to one of these types, then that is the governing type of the `switch` statement.
- Otherwise, if exactly one user-defined implicit conversion exists from the type of the switch expression to one of the following possible governing types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `string`, or, a nullable value type corresponding to one of those types, then the converted type is the governing type of the `switch` statement.
- Otherwise, the governing type of the `switch` statement is the type of the switch expression. It is an error if no such type exists.

There can be at most one `default` label in a `switch` statement.

It is an error if the pattern of any switch label is not *applicable* (§11.2.1) to the type of the input expression.

It is an error if the pattern of any switch label is *subsumed* by (§11.3) the set of patterns of earlier switch labels of the switch statement that do not have a case guard or whose case guard is a constant expression with the value true.

> *Example*:
>
> C#
>
> ```
> switch (shape)
> {
>     case var x:
>         break;
>     case var _: // error: pattern subsumed, as previous case always
> ```

```
matches
        break;
    default:
        break;  // warning: unreachable, all possible values already
handled.
}
```

end example

A `switch` statement is executed as follows:

- The switch expression is evaluated and converted to the governing type.
- Control is transferred according to the value of the converted switch expression:
  - The lexically first pattern in the set of `case` labels in the same `switch` statement that matches the value of the switch expression, and for which the guard expression is either absent or evaluates to true, causes control to be transferred to the statement list following the matched `case` label.
  - Otherwise, if a `default` label is present, control is transferred to the statement list following the `default` label.
  - Otherwise, control is transferred to the end point of the `switch` statement.

> *Note*: The order in which patterns are matched at runtime is not defined. A compiler is permitted (but not required) to match patterns out of order, and to reuse the results of already matched patterns to compute the result of matching of other patterns. Nevertheless, the compiler is required to determine the lexically first pattern that matches the expression and for which the guard clause is either absent or evaluates to `true`. *end note*

If the end point of the statement list of a switch section is reachable, a compile-time error occurs. This is known as the “no fall through” rule.

> *Example*: The example

C#

```
switch (i)
{
    case 0:
        CaseZero();
        break;
    case 1:
        CaseOne();
        break;
    default:
        CaseOthers();
```

```
        break;
}
```

is valid because no switch section has a reachable end point. Unlike C and C++, execution of a switch section is not permitted to "fall through" to the next switch section, and the example

C#

```
switch (i)
{
    case 0:
        CaseZero();
    case 1:
        CaseZeroOrOne();
    default:
        CaseAny();
}
```

results in a compile-time error. When execution of a switch section is to be followed by execution of another switch section, an explicit `goto case` or `goto default` statement shall be used:

C#

```
switch (i)
{
    case 0:
        CaseZero();
        goto case 1;
    case 1:
        CaseZeroOrOne();
        goto default;
    default:
        CaseAny();
        break;
}
```

end example

Multiple labels are permitted in a *switch_section*.

Example: The example

C#

```
switch (i)
{
    case 0:
        CaseZero();
        break;
    case 1:
        CaseOne();
        break;
    case 2:
    default:
        CaseTwo();
        break;
}
```

is valid. The example does not violate the "no fall through" rule because the labels `case 2:` and `default:` are part of the same *switch_section*.

end example

Note: The "no fall through" rule prevents a common class of bugs that occur in C and C++ when `break` statements are accidentally omitted. For example, the sections of the `switch` statement above can be reversed without affecting the behavior of the statement:

C#

```
switch (i)
{
    default:
        CaseAny();
        break;
    case 1:
        CaseZeroOrOne();
        goto default;
    case 0:
        CaseZero();
        goto case 1;
}
```

end note

Note: The statement list of a switch section typically ends in a `break`, `goto case`, or `goto default` statement, but any construct that renders the end point of the statement list unreachable is permitted. For example, a `while` statement controlled by the Boolean expression `true` is known to never reach its end point. Likewise, a

`throw` or `return` statement always transfers control elsewhere and never reaches its end point. Thus, the following example is valid:

```C#
switch (i)
{
     case 0:
         while (true)
         {
             F();
         }
     case 1:
         throw new ArgumentException();
     case 2:
         return;
}
```

end note

Example: The governing type of a `switch` statement can be the type `string`. For example:

```C#
void DoCommand(string command)
{
    switch (command.ToLower())
    {
        case "run":
            DoRun();
            break;
        case "save":
            DoSave();
            break;
        case "quit":
            DoQuit();
            break;
        default:
            InvalidCommand(command);
            break;
    }
}
```

end example

Note: Like the string equality operators (§12.12.8), the `switch` statement is case sensitive and will execute a given switch section only if the switch expression string

> exactly matches a `case` label constant. *end note* When the governing type of a `switch` statement is `string` or a nullable value type, the value `null` is permitted as a `case` label constant.

The *statement_lists* of a *switch_block* may contain declaration statements (§13.6). The scope of a local variable or constant declared in a switch block is the switch block.

A switch label is reachable if at least one of the following is true:

- The switch expression is a constant value and either
  - the label is a `case` whose pattern *would match* (§11.2.1) that value, and label's guard is either absent or not a constant expression with the value false; or
  - it is a `default` label, and no switch section contains a case label whose pattern would match that value, and whose guard is either absent or a constant expression with the value true.
- The switch expression is not a constant value and either
  - the label is a `case` without a guard or with a guard whose value is not the constant false; or
  - it is a `default` label and
    - the set of patterns appearing among the cases of the switch statement that do not have guards or have guards whose value is the constant true, is not *exhaustive* (§11.4) for the switch governing type; or
    - the switch governing type is a nullable type and the set of patterns appearing among the cases of the switch statement that do not have guards or have guards whose value is the constant true does not contain a pattern that would match the value `null`.
- The switch label is referenced by a reachable `goto case` or `goto default` statement.

The statement list of a given switch section is reachable if the `switch` statement is reachable and the switch section contains a reachable switch label.

The end point of a `switch` statement is reachable if the switch statement is reachable and at least one of the following is true:

- The `switch` statement contains a reachable `break` statement that exits the `switch` statement.
- No `default` label is present and either
  - The switch expression is a non-constant value, and the set of patterns appearing among the cases of the switch statement that do not have guards or have

guards whose value is the constant true, is not *exhaustive* (§11.4) for the switch governing type.

- The switch expression is a non-constant value of a nullable type, and no pattern appearing among the cases of the switch statement that do not have guards or have guards whose value is the constant true would match the value `null`.
- The switch expression is a constant value and no `case` label without a guard or whose guard is the constant true would match that value.

Example: The following code shows a succinct use of the `when` clause:

```C#
static object CreateShape(string shapeDescription)
{
   switch (shapeDescription)
   {
        case "circle":
            return new Circle(2);
        …
        case var o when string.IsNullOrWhiteSpace(o):
            return null;
        default:
            return "invalid shape description";
    }
}
```

The var case matches `null`, the empty string, or any string that contains only white space. *end example*

13.9 Iteration statements

13.9.1 General

Iteration statements repeatedly execute an embedded statement.

```ANTLR
iteration_statement
    : while_statement
    | do_statement
    | for_statement
    | foreach_statement
    ;
```

13.9.2 The while statement

The `while` statement conditionally executes an embedded statement zero or more times.

ANTLR

```
while_statement
    : 'while' '(' boolean_expression ')' embedded_statement
    ;
```

A `while` statement is executed as follows:

- The *boolean_expression* (§12.24) is evaluated.
- If the Boolean expression yields `true`, control is transferred to the embedded statement. When and if control reaches the end point of the embedded statement (possibly from execution of a `continue` statement), control is transferred to the beginning of the `while` statement.
- If the Boolean expression yields `false`, control is transferred to the end point of the `while` statement.

Within the embedded statement of a `while` statement, a `break` statement (§13.10.2) may be used to transfer control to the end point of the `while` statement (thus ending iteration of the embedded statement), and a `continue` statement (§13.10.3) may be used to transfer control to the end point of the embedded statement (thus performing another iteration of the `while` statement).

The embedded statement of a `while` statement is reachable if the `while` statement is reachable and the Boolean expression does not have the constant value `false`.

The end point of a `while` statement is reachable if at least one of the following is true:

- The `while` statement contains a reachable `break` statement that exits the `while` statement.
- The `while` statement is reachable and the Boolean expression does not have the constant value `true`.

13.9.3 The do statement

The `do` statement conditionally executes an embedded statement one or more times.

ANTLR

```
do_statement
    : 'do' embedded_statement 'while' '(' boolean_expression ')' ';'
    ;
```

A `do` statement is executed as follows:

- Control is transferred to the embedded statement.
- When and if control reaches the end point of the embedded statement (possibly from execution of a `continue` statement), the *boolean_expression* (§12.24) is evaluated. If the Boolean expression yields `true`, control is transferred to the beginning of the `do` statement. Otherwise, control is transferred to the end point of the `do` statement.

Within the embedded statement of a `do` statement, a `break` statement (§13.10.2) may be used to transfer control to the end point of the `do` statement (thus ending iteration of the embedded statement), and a `continue` statement (§13.10.3) may be used to transfer control to the end point of the embedded statement (thus performing another iteration of the `do` statement).

The embedded statement of a `do` statement is reachable if the `do` statement is reachable.

The end point of a `do` statement is reachable if at least one of the following is true:

- The `do` statement contains a reachable `break` statement that exits the `do` statement.
- The end point of the embedded statement is reachable and the Boolean expression does not have the constant value `true`.

13.9.4 The for statement

The `for` statement evaluates a sequence of initialization expressions and then, while a condition is true, repeatedly executes an embedded statement and evaluates a sequence of iteration expressions.

ANTLR

```
for_statement
    : 'for' '(' for_initializer? ';' for_condition? ';' for_iterator? ')'
      embedded_statement
    ;

for_initializer
```

```
    : local_variable_declaration
    | statement_expression_list
    ;

for_condition
    : boolean_expression
    ;

for_iterator
    : statement_expression_list
    ;

statement_expression_list
    : statement_expression (',' statement_expression)*
    ;
```

The *for_initializer*, if present, consists of either a *local_variable_declaration* (§13.6.2) or a list of *statement_expression*s (§13.7) separated by commas. The scope of a local variable declared by a *for_initializer* is the *for_initializer*, *for_condition*, *for_iterator*, and *embedded_statement*.

The *for_condition*, if present, shall be a *boolean_expression* (§12.24).

The *for_iterator*, if present, consists of a list of *statement_expression*s (§13.7) separated by commas.

A `for` statement is executed as follows:

- If a *for_initializer* is present, the variable initializers or statement expressions are executed in the order they are written. This step is only performed once.
- If a *for_condition* is present, it is evaluated.
- If the *for_condition* is not present or if the evaluation yields `true`, control is transferred to the embedded statement. When and if control reaches the end point of the embedded statement (possibly from execution of a `continue` statement), the expressions of the *for_iterator*, if any, are evaluated in sequence, and then another iteration is performed, starting with evaluation of the *for_condition* in the step above.
- If the *for_condition* is present and the evaluation yields `false`, control is transferred to the end point of the `for` statement.

Within the embedded statement of a `for` statement, a `break` statement (§13.10.2) may be used to transfer control to the end point of the `for` statement (thus ending iteration of the embedded statement), and a `continue` statement (§13.10.3) may be used to transfer control to the end point of the embedded statement (thus executing the

for_iterator and performing another iteration of the `for` statement, starting with the *for_condition*).

The embedded statement of a `for` statement is reachable if one of the following is true:

- The `for` statement is reachable and no *for_condition* is present.
- The `for` statement is reachable and a *for_condition* is present and does not have the constant value `false`.

The end point of a `for` statement is reachable if at least one of the following is true:

- The `for` statement contains a reachable `break` statement that exits the `for` statement.
- The `for` statement is reachable and a *for_condition* is present and does not have the constant value `true`.

13.9.5 The foreach statement

The `foreach` statement enumerates the elements of a collection, executing an embedded statement for each element of the collection.

ANTLR

```
foreach_statement
    : 'foreach' '(' ref_kind? local_variable_type identifier 'in'
      expression ')' embedded_statement
    ;
```

The *local_variable_type* and *identifier* of a foreach statement declare the ***iteration variable*** of the statement. If the `var` identifier is given as the *local_variable_type*, and no type named `var` is in scope, the iteration variable is said to be an ***implicitly typed iteration variable***, and its type is taken to be the element type of the `foreach` statement, as specified below.

If the *foreach_statement* contains both or neither `ref` and `readonly`, the iteration variable denotes a variable that is treated as read-only. Otherwise, if *foreach_statement* contains `ref` without `readonly`, the iteration variable denotes a variable that shall be writable.

The iteration variable corresponds to a local variable with a scope that extends over the embedded statement. During execution of a `foreach` statement, the iteration variable represents the collection element for which an iteration is currently being performed. If the iteration variable denotes a read-only variable, a compile-time error occurs if the

embedded statement attempts to modify it (via assignment or the `++` and `--` operators) or pass it as a `ref` or `out` parameter.

In the following, for brevity, `IEnumerable`, `IEnumerator`, `IEnumerable<T>` and `IEnumerator<T>` refer to the corresponding types in the namespaces `System.Collections` and `System.Collections.Generic`.

The compile-time processing of a `foreach` statement first determines the ***collection type***, ***enumerator type*** and ***iteration type*** of the expression. This determination proceeds as follows:

- If the type `X` of *expression* is an array type then there is an implicit reference conversion from X to the `IEnumerable` interface (since `System.Array` implements this interface). The collection type is the `IEnumerable` interface, the enumerator type is the `IEnumerator` interface and the iteration type is the element type of the array type `X`.
- If the type `X` of *expression* is `dynamic` then there is an implicit conversion from *expression* to the `IEnumerable` interface (§10.2.10). The collection type is the `IEnumerable` interface and the enumerator type is the `IEnumerator` interface. If the `var` identifier is given as the *local_variable_type* then the iteration type is `dynamic`, otherwise it is `object`.
- Otherwise, determine whether the type `X` has an appropriate `GetEnumerator` method:
  - Perform member lookup on the type `X` with identifier `GetEnumerator` and no type arguments. If the member lookup does not produce a match, or it produces an ambiguity, or produces a match that is not a method group, check for an enumerable interface as described below. It is recommended that a warning be issued if member lookup produces anything except a method group or no match.
  - Perform overload resolution using the resulting method group and an empty argument list. If overload resolution results in no applicable methods, results in an ambiguity, or results in a single best method but that method is either static or not public, check for an enumerable interface as described below. It is recommended that a warning be issued if overload resolution produces anything except an unambiguous public instance method or no applicable methods.
  - If the return type `E` of the `GetEnumerator` method is not a class, struct or interface type, an error is produced and no further steps are taken.
  - Member lookup is performed on `E` with the identifier `Current` and no type arguments. If the member lookup produces no match, the result is an error, or

the result is anything except a public instance property that permits reading, an error is produced and no further steps are taken.
    - Member lookup is performed on `E` with the identifier `MoveNext` and no type arguments. If the member lookup produces no match, the result is an error, or the result is anything except a method group, an error is produced and no further steps are taken.
    - Overload resolution is performed on the method group with an empty argument list. If overload resolution results in no applicable methods, results in an ambiguity, or results in a single best method but that method is either static or not public, or its return type is not `bool`, an error is produced and no further steps are taken.
    - The collection type is `X`, the enumerator type is `E`, and the iteration type is the type of the `Current` property. The `Current` property may include the `ref` modifier, in which case, the expression returned is a *variable_reference* (§9.5) that is optionally read-only.
- Otherwise, check for an enumerable interface:
    - If among all the types T_i for which there is an implicit conversion from `X` to `IEnumerable<`T_i`>`, there is a unique type `T` such that `T` is not `dynamic` and for all the other T_i there is an implicit conversion from `IEnumerable<T>` to `IEnumerable<`T_i`>`, then the collection type is the interface `IEnumerable<T>`, the enumerator type is the interface `IEnumerator<T>`, and the iteration type is `T`.
    - Otherwise, if there is more than one such type `T`, then an error is produced and no further steps are taken.
    - Otherwise, if there is an implicit conversion from `X` to the `System.Collections.IEnumerable` interface, then the collection type is this interface, the enumerator type is the interface `System.Collections.IEnumerator`, and the iteration type is `object`.
    - Otherwise, an error is produced and no further steps are taken.

The above steps, if successful, unambiguously produce a collection type `C`, enumerator type `E` and iteration type `T`, `ref T`, or `ref readonly T`. A `foreach` statement of the form

```
C#
```

```
foreach (V v in x) «embedded_statement»
```

is then equivalent to:

```
C#
```

```
{
    E e = ((C)(x)).GetEnumerator();
    try
    {
        while (e.MoveNext())
        {
            V v = (V)(T)e.Current;
            «embedded_statement»
        }
    }
    finally
    {
        ... // Dispose e
    }
}
```

The variable `e` is not visible to or accessible to the expression `x` or the embedded statement or any other source code of the program. The variable `v` is read-only in the embedded statement. If there is not an explicit conversion (§10.3) from `T` (the iteration type) to `V` (the *local_variable_type* in the `foreach` statement), an error is produced and no further steps are taken.

When the iteration variable is a reference variable (§9.7), a `foreach` statement of the form

C#

```
foreach (ref V v in x) «embedded_statement»
```

is then equivalent to:

C#

```
{
    E e = ((C)(x)).GetEnumerator();
    try
    {
        while (e.MoveNext())
        {
            ref V v = ref e.Current;
            «embedded_statement»
        }
    }
    finally
    {
        ... // Dispose e
```

```
    }
}
```

The variable `e` is not visible or accessible to the expression `x` or the embedded statement or any other source code of the program. The reference variable `v` is read-write in the embedded statement, but `v` shall not be ref-reassigned (§12.21.3). If there is not an identity conversion (§10.2.2) from `T` (the iteration type) to `V` (the *local_variable_type* in the `foreach` statement), an error is produced and no further steps are taken.

A `foreach` statement of the form `foreach (ref readonly V v in x) «embedded_statement»` has a similar equivalent form, but the reference variable `v` is `ref readonly` in the embedded statement, and therefore cannot be ref-reassigned or reassigned.

> *Note*: If `x` has the value `null`, a `System.NullReferenceException` is thrown at run-time. *end note*

An implementation is permitted to implement a given *foreach_statement* differently; e.g., for performance reasons, as long as the behavior is consistent with the above expansion.

The placement of `v` inside the `while` loop is important for how it is captured (§12.19.6.2) by any anonymous function occurring in the *embedded_statement*.

> *Example*:
>
> C#
>
> ```
> int[] values = { 7, 9, 13 };
> Action f = null;
> foreach (var value in values)
> {
>     if (f == null)
>     {
>         f = () => Console.WriteLine("First value: " + value);
>     }
> }
> f();
> ```
>
> If `v` in the expanded form were declared outside of the `while` loop, it would be shared among all iterations, and its value after the `for` loop would be the final value, `13`, which is what the invocation of `f` would print. Instead, because each iteration has its own variable `v`, the one captured by `f` in the first iteration will

> continue to hold the value `7`, which is what will be printed. (Note that earlier versions of C# declared `v` outside of the `while` loop.)
>
> *end example*

The body of the `finally` block is constructed according to the following steps:

- If there is an implicit conversion from `E` to the `System.IDisposable` interface, then
  - If `E` is a non-nullable value type then the `finally` clause is expanded to the semantic equivalent of:

    ```
    C#

    finally
    {
        ((System.IDisposable)e).Dispose();
    }
    ```

  - Otherwise the `finally` clause is expanded to the semantic equivalent of:

    ```
    C#

    finally
    {
        System.IDisposable d = e as System.IDisposable;
        if (d != null)
        {
            d.Dispose();
        }
    }
    ```

    except that if `E` is a value type, or a type parameter instantiated to a value type, then the conversion of `e` to `System.IDisposable` shall not cause boxing to occur.

- Otherwise, if `E` is a sealed type, the `finally` clause is expanded to an empty block:

  ```
  C#

  finally {}
  ```

- Otherwise, the `finally` clause is expanded to:

  ```
  C#
  ```

```
finally
{
    System.IDisposable d = e as System.IDisposable;
    if (d != null)
    {
        d.Dispose();
    }
}
```

The local variable `d` is not visible to or accessible to any user code. In particular, it does not conflict with any other variable whose scope includes the `finally` block.

The order in which `foreach` traverses the elements of an array, is as follows: For single-dimensional arrays elements are traversed in increasing index order, starting with index 0 and ending with index `Length - 1`. For multi-dimensional arrays, elements are traversed such that the indices of the rightmost dimension are increased first, then the next left dimension, and so on to the left.

> *Example*: The following example prints out each value in a two-dimensional array, in element order:
>
> C#
>
> ```
> class Test
> {
>     static void Main()
>     {
>         double[,] values =
>         {
>             {1.2, 2.3, 3.4, 4.5},
>             {5.6, 6.7, 7.8, 8.9}
>         };
>         foreach (double elementValue in values)
>         {
>             Console.Write($"{elementValue} ");
>         }
>         Console.WriteLine();
>     }
> }
> ```
>
> The output produced is as follows:
>
> Console
>
> ```
> 1.2 2.3 3.4 4.5 5.6 6.7 7.8 8.9
> ```
>
> *end example*

Example: In the following example

```C#
int[] numbers = { 1, 3, 5, 7, 9 };
foreach (var n in numbers)
{
    Console.WriteLine(n);
}
```

the type of `n` is inferred to be `int`, the iteration type of `numbers`.

end example

13.10 Jump statements

13.10.1 General

Jump statements unconditionally transfer control.

```ANTLR
jump_statement
    : break_statement
    | continue_statement
    | goto_statement
    | return_statement
    | throw_statement
    ;
```

The location to which a jump statement transfers control is called the ***target*** of the jump statement.

When a jump statement occurs within a block, and the target of that jump statement is outside that block, the jump statement is said to ***exit*** the block. While a jump statement can transfer control out of a block, it can never transfer control into a block.

Execution of jump statements is complicated by the presence of intervening `try` statements. In the absence of such `try` statements, a jump statement unconditionally transfers control from the jump statement to its target. In the presence of such intervening `try` statements, execution is more complex. If the jump statement exits one or more `try` blocks with associated `finally` blocks, control is initially transferred to the `finally` block of the innermost `try` statement. When and if control reaches the end point of a `finally` block, control is transferred to the `finally` block of the next

enclosing `try` statement. This process is repeated until the `finally` blocks of all intervening `try` statements have been executed.

> *Example*: In the following code
>
> C#
>
> ```
> class Test
> {
>     static void Main()
>     {
>         while (true)
>         {
>             try
>             {
>                 try
>                 {
>                     Console.WriteLine("Before break");
>                     break;
>                 }
>                 finally
>                 {
>                     Console.WriteLine("Innermost finally block");
>                 }
>             }
>             finally
>             {
>                 Console.WriteLine("Outermost finally block");
>             }
>         }
>         Console.WriteLine("After break");
>     }
> }
> ```
>
> the `finally` blocks associated with two `try` statements are executed before control is transferred to the target of the jump statement. The output produced is as follows:
>
> Console
>
> ```
> Before break
> Innermost finally block
> Outermost finally block
> After break
> ```
>
> *end example*

13.10.2 The break statement

The `break` statement exits the nearest enclosing `switch`, `while`, `do`, `for`, or `foreach` statement.

ANTLR

```
break_statement
    : 'break' ';'
    ;
```

The target of a `break` statement is the end point of the nearest enclosing `switch`, `while`, `do`, `for`, or `foreach` statement. If a `break` statement is not enclosed by a `switch`, `while`, `do`, `for`, or `foreach` statement, a compile-time error occurs.

When multiple `switch`, `while`, `do`, `for`, or `foreach` statements are nested within each other, a `break` statement applies only to the innermost statement. To transfer control across multiple nesting levels, a `goto` statement (§13.10.4) shall be used.

A `break` statement cannot exit a `finally` block (§13.11). When a `break` statement occurs within a `finally` block, the target of the `break` statement shall be within the same `finally` block; otherwise a compile-time error occurs.

A `break` statement is executed as follows:

- If the `break` statement exits one or more `try` blocks with associated `finally` blocks, control is initially transferred to the `finally` block of the innermost `try` statement. When and if control reaches the end point of a `finally` block, control is transferred to the `finally` block of the next enclosing `try` statement. This process is repeated until the `finally` blocks of all intervening `try` statements have been executed.
- Control is transferred to the target of the `break` statement.

Because a `break` statement unconditionally transfers control elsewhere, the end point of a `break` statement is never reachable.

13.10.3 The continue statement

The `continue` statement starts a new iteration of the nearest enclosing `while`, `do`, `for`, or `foreach` statement.

ANTLR

```
continue_statement
    : 'continue' ';'
```

```
    ;
```

The target of a `continue` statement is the end point of the embedded statement of the nearest enclosing `while`, `do`, `for`, or `foreach` statement. If a `continue` statement is not enclosed by a `while`, `do`, `for`, or `foreach` statement, a compile-time error occurs.

When multiple `while`, `do`, `for`, or `foreach` statements are nested within each other, a `continue` statement applies only to the innermost statement. To transfer control across multiple nesting levels, a `goto` statement (§13.10.4) shall be used.

A `continue` statement cannot exit a `finally` block (§13.11). When a `continue` statement occurs within a `finally` block, the target of the `continue` statement shall be within the same `finally` block; otherwise a compile-time error occurs.

A `continue` statement is executed as follows:

- If the `continue` statement exits one or more `try` blocks with associated `finally` blocks, control is initially transferred to the `finally` block of the innermost `try` statement. When and if control reaches the end point of a `finally` block, control is transferred to the `finally` block of the next enclosing `try` statement. This process is repeated until the `finally` blocks of all intervening `try` statements have been executed.
- Control is transferred to the target of the `continue` statement.

Because a `continue` statement unconditionally transfers control elsewhere, the end point of a `continue` statement is never reachable.

13.10.4 The goto statement

The `goto` statement transfers control to a statement that is marked by a label.

ANTLR

```
goto_statement
    : 'goto' identifier ';'
    | 'goto' 'case' constant_expression ';'
    | 'goto' 'default' ';'
    ;
```

The target of a `goto` *identifier* statement is the labeled statement with the given label. If a label with the given name does not exist in the current function member, or if the `goto` statement is not within the scope of the label, a compile-time error occurs.

Note: This rule permits the use of a `goto` statement to transfer control *out of* a nested scope, but not *into* a nested scope. In the example

C#

```
class Test
{
    static void Main(string[] args)
    {
        string[,] table =
        {
            {"Red", "Blue", "Green"},
            {"Monday", "Wednesday", "Friday"}
        };
        foreach (string str in args)
        {
            int row, colm;
            for (row = 0; row <= 1; ++row)
            {
                for (colm = 0; colm <= 2; ++colm)
                {
                    if (str == table[row,colm])
                    {
                        goto done;
                    }
                }
            }
            Console.WriteLine($"{str} not found");
            continue;
          done:
            Console.WriteLine($"Found {str} at [{row}][{colm}]");
        }
    }
}
```

a `goto` statement is used to transfer control out of a nested scope.

end note

The target of a `goto case` statement is the statement list in the immediately enclosing `switch` statement (§13.8.3) which contains a `case` label with a constant pattern of the given constant value and no guard. If the `goto case` statement is not enclosed by a `switch` statement, if the nearest enclosing `switch` statement does not contain such a `case`, or if the *constant_expression* is not implicitly convertible (§10.2) to the governing type of the nearest enclosing `switch` statement, a compile-time error occurs.

The target of a `goto default` statement is the statement list in the immediately enclosing `switch` statement (§13.8.3), which contains a `default` label. If the `goto`

`default` statement is not enclosed by a `switch` statement, or if the nearest enclosing `switch` statement does not contain a `default` label, a compile-time error occurs.

A `goto` statement cannot exit a `finally` block (§13.11). When a `goto` statement occurs within a `finally` block, the target of the `goto` statement shall be within the same `finally` block, or otherwise a compile-time error occurs.

A `goto` statement is executed as follows:

- If the `goto` statement exits one or more `try` blocks with associated `finally` blocks, control is initially transferred to the `finally` block of the innermost `try` statement. When and if control reaches the end point of a `finally` block, control is transferred to the `finally` block of the next enclosing `try` statement. This process is repeated until the `finally` blocks of all intervening `try` statements have been executed.
- Control is transferred to the target of the `goto` statement.

Because a `goto` statement unconditionally transfers control elsewhere, the end point of a `goto` statement is never reachable.

13.10.5 The return statement

The `return` statement returns control to the current caller of the function member in which the return statement appears, optionally returning a value or a *variable_reference* (§9.5).

ANTLR

```
return_statement
    : 'return' ';'
    | 'return' expression ';'
    | 'return' 'ref' variable_reference ';'
    ;
```

A *return_statement* without *expression* is called a ***return-no-value***; one containing `ref` *expression* is called a ***return-by-ref***; and one containing only *expression* is called a ***return-by-value***.

It is a compile-time error to use a return-no-value from a method declared as being returns-by-value or returns-by-ref (§15.6.1).

It is a compile-time error to use a return-by-ref from a method declared as being returns-no-value or returns-by-value.

It is a compile-time error to use a return-by-value from a method declared as being returns-no-value or returns-by-ref.

It is a compile-time error to use a return-by-ref if *expression* is not a *variable_reference* or is a reference to a variable whose ref-safe-context is not caller-context (§9.7.2).

It is a compile-time error to use a return-by-ref from a method declared with the *method_modifier* `async`.

A function member is said to ***compute a value*** if it is a method with a returns-by-value method (§15.6.11), a returns-by-value get accessor of a property or indexer, or a user-defined operator. Function members that are returns-no-value do not compute a value and are methods with the effective return type `void`, set accessors of properties and indexers, add and remove accessors of events, instance constructors, static constructors and finalizers. Function members that are returns-by-ref do not compute a value.

For a return-by-value, an implicit conversion (§10.2) shall exist from the type of *expression* to the effective return type (§15.6.11) of the containing function member. For a return-by-ref, an identity conversion (§10.2.2) shall exist between the type of *expression* and the effective return type of the containing function member.

`return` statements can also be used in the body of anonymous function expressions (§12.19), and participate in determining which conversions exist for those functions (§10.7.1).

It is a compile-time error for a `return` statement to appear in a `finally` block (§13.11).

A `return` statement is executed as follows:

- For a return-by-value, *expression* is evaluated and its value is converted to the effective return type of the containing function by an implicit conversion. The result of the conversion becomes the result value produced by the function. For a return-by-ref, the *expression* is evaluated, and the result shall be classified as a variable. If the enclosing method's return-by-ref includes `readonly`, the resulting variable is read-only.
- If the `return` statement is enclosed by one or more `try` or `catch` blocks with associated `finally` blocks, control is initially transferred to the `finally` block of the innermost `try` statement. When and if control reaches the end point of a `finally` block, control is transferred to the `finally` block of the next enclosing `try` statement. This process is repeated until the `finally` blocks of all enclosing `try` statements have been executed.
- If the containing function is not an async function, control is returned to the caller of the containing function along with the result value, if any.

- If the containing function is an async function, control is returned to the current caller, and the result value, if any, is recorded in the return task as described in (§15.15.3).

Because a `return` statement unconditionally transfers control elsewhere, the end point of a `return` statement is never reachable.

13.10.6 The throw statement

The `throw` statement throws an exception.

ANTLR

```
throw_statement
    : 'throw' expression? ';'
    ;
```

A `throw` statement with an expression throws an exception produced by evaluating the expression. The expression shall be implicitly convertible to `System.Exception`, and the result of evaluating the expression is converted to `System.Exception` before being thrown. If the result of the conversion is `null`, a `System.NullReferenceException` is thrown instead.

A `throw` statement with no expression can be used only in a `catch` block, in which case, that statement re-throws the exception that is currently being handled by that `catch` block.

Because a `throw` statement unconditionally transfers control elsewhere, the end point of a `throw` statement is never reachable.

When an exception is thrown, control is transferred to the first `catch` clause in an enclosing `try` statement that can handle the exception. The process that takes place from the point of the exception being thrown to the point of transferring control to a suitable exception handler is known as ***exception propagation***. Propagation of an exception consists of repeatedly evaluating the following steps until a `catch` clause that matches the exception is found. In this description, the ***throw point*** is initially the location at which the exception is thrown. This behavior is specified in (§21.4).

- In the current function member, each `try` statement that encloses the throw point is examined. For each statement `S`, starting with the innermost `try` statement and ending with the outermost `try` statement, the following steps are evaluated:

    - If the `try` block of `S` encloses the throw point and if `S` has one or more `catch` clauses, the `catch` clauses are examined in order of appearance to locate a suitable handler for the exception. The first `catch` clause that specifies an exception type `T` (or a type parameter that at run-time denotes an exception type `T`) such that the run-time type of `E` derives from `T` is considered a match. If the clause contains an exception filter, the exception object is assigned to the exception variable, and the exception filter is evaluated. When a `catch` clause contains an exception filter, that `catch` clause is considered a match if the exception filter evaluates to `true`. A general `catch` (§13.11) clause is considered a match for any exception type. If a matching `catch` clause is located, the exception propagation is completed by transferring control to the block of that `catch` clause.
    - Otherwise, if the `try` block or a `catch` block of `S` encloses the throw point and if `S` has a `finally` block, control is transferred to the `finally` block. If the `finally` block throws another exception, processing of the current exception is terminated. Otherwise, when control reaches the end point of the `finally` block, processing of the current exception is continued.

- If an exception handler was not located in the current function invocation, the function invocation is terminated, and one of the following occurs:

    - If the current function is non-async, the steps above are repeated for the caller of the function with a throw point corresponding to the statement from which the function member was invoked.

    - If the current function is async and task-returning, the exception is recorded in the return task, which is put into a faulted or canceled state as described in §15.15.3.

    - If the current function is async and `void`-returning, the synchronization context of the current thread is notified as described in §15.15.4.

- If the exception processing terminates all function member invocations in the current thread, indicating that the thread has no handler for the exception, then the thread is itself terminated. The impact of such termination is implementation-defined.

13.11 The try statement

The `try` statement provides a mechanism for catching exceptions that occur during execution of a block. Furthermore, the `try` statement provides the ability to specify a

block of code that is always executed when control leaves the `try` statement.

ANTLR

```
try_statement
    : 'try' block catch_clauses
    | 'try' block catch_clauses? finally_clause
    ;

catch_clauses
    : specific_catch_clause+
    | specific_catch_clause* general_catch_clause
    ;

specific_catch_clause
    : 'catch' exception_specifier exception_filter? block
    | 'catch' exception_filter block
    ;

exception_specifier
    : '(' type identifier? ')'
    ;

exception_filter
    : 'when' '(' boolean_expression ')'
    ;

general_catch_clause
    : 'catch' block
    ;

finally_clause
    : 'finally' block
    ;
```

A *try_statement* consists of the keyword `try` followed by a *block*, then zero or more *catch_clauses*, then an optional *finally_clause*. There shall be at least one *catch_clause* or a *finally_clause*.

In an *exception_specifier* the *type*, or its effective base class if it is a *type_parameter*, shall be `System.Exception` or a type that derives from it.

When a `catch` clause specifies both a *class_type* and an *identifier*, an ***exception variable*** of the given name and type is declared. The exception variable is introduced into the declaration space of the *specific_catch_clause* (§7.3). During execution of the *exception_filter* and `catch` block, the exception variable represents the exception currently being handled. For purposes of definite assignment checking, the exception variable is considered definitely assigned in its entire scope.

Unless a `catch` clause includes an exception variable name, it is impossible to access the exception object in the filter and `catch` block.

A `catch` clause that specifies neither an exception type nor an exception variable name is called a general `catch` clause. A `try` statement can only have one general `catch` clause, and, if one is present, it shall be the last `catch` clause.

> *Note*: Some programming languages might support exceptions that are not representable as an object derived from `System.Exception`, although such exceptions could never be generated by C# code. A general `catch` clause might be used to catch such exceptions. Thus, a general `catch` clause is semantically different from one that specifies the type `System.Exception`, in that the former might also catch exceptions from other languages. *end note*

In order to locate a handler for an exception, `catch` clauses are examined in lexical order. If a `catch` clause specifies a type but no exception filter, it is a compile-time error for a later `catch` clause of the same `try` statement to specify a type that is the same as, or is derived from, that type.

> *Note*: Without this restriction, it would be possible to write unreachable `catch` clauses. *end note*

Within a `catch` block, a `throw` statement (§13.10.6) with no expression can be used to re-throw the exception that was caught by the `catch` block. Assignments to an exception variable do not alter the exception that is re-thrown.

> *Example*: In the following code

C#

```
class Test
{
    static void F()
    {
        try
        {
            G();
        }
        catch (Exception e)
        {
            Console.WriteLine("Exception in F: " + e.Message);
            e = new Exception("F");
            throw; // re-throw
        }
```

```
    }

    static void G() => throw new Exception("G");

    static void Main()
    {
        try
        {
            F();
        }
        catch (Exception e)
        {
            Console.WriteLine("Exception in Main: " + e.Message);
        }
    }
}
```

the method `F` catches an exception, writes some diagnostic information to the console, alters the exception variable, and re-throws the exception. The exception that is re-thrown is the original exception, so the output produced is:

Console

```
Exception in F: G
Exception in Main: G
```

If the first `catch` block had thrown `e` instead of rethrowing the current exception, the output produced would be as follows:

Console

```
Exception in F: G
Exception in Main: F
```

end example

It is a compile-time error for a `break`, `continue`, or `goto` statement to transfer control out of a `finally` block. When a `break`, `continue`, or `goto` statement occurs in a `finally` block, the target of the statement shall be within the same `finally` block, or otherwise a compile-time error occurs.

It is a compile-time error for a `return` statement to occur in a `finally` block.

When execution reaches a `try` statement, control is transferred to the `try` block. If control reaches the end point of the `try` block without an exception being propagated,

control is transferred to the `finally` block if one exists. If no `finally` block exists, control is transferred to the end point of the `try` statement.

If an exception has been propagated, the `catch` clauses, if any, are examined in lexical order seeking the first match for the exception. The search for a matching `catch` clause continues with all enclosing blocks as described in §13.10.6. A `catch` clause is a match if the exception type matches any *exception_specifier* and any *exception_filter* is true. A `catch` clause without an *exception_specifier* matches any exception type. The exception type matches the *exception_specifier* when the *exception_specifier* specifies the exception type or a base type of the exception type. If the clause contains an exception filter, the exception object is assigned to the exception variable, and the exception filter is evaluated.

If an exception has been propagated and a matching `catch` clause is found, control is transferred to the first matching `catch` block. If control reaches the end point of the `catch` block without an exception being propagated, control is transferred to the `finally` block if one exists. If no `finally` block exists, control is transferred to the end point of the `try` statement. If an exception has been propagated from the `catch` block, control transfers to the `finally` block if one exists. The exception is propagated to the next enclosing `try` statement.

If an exception has been propagated, and no matching `catch` clause is found, control transfers to the `finally` block, if it exists. The exception is propagated to the next enclosing `try` statement.

The statements of a `finally` block are always executed when control leaves a `try` statement. This is true whether the control transfer occurs as a result of normal execution, as a result of executing a `break`, `continue`, `goto`, or `return` statement, or as a result of propagating an exception out of the `try` statement. If control reaches the end point of the `finally` block without an exception being propagated, control is transferred to the end point of the `try` statement.

If an exception is thrown during execution of a `finally` block, and is not caught within the same `finally` block, the exception is propagated to the next enclosing `try` statement. If another exception was in the process of being propagated, that exception is lost. The process of propagating an exception is discussed further in the description of the `throw` statement (§13.10.6).

> *Example*: In the following code
>
> ```
> C#
> ```

```
public class Test
{
    static void Main()
    {
        try
        {
            Method();
        }
        catch (Exception ex) when (ExceptionFilter(ex))
        {
            Console.WriteLine("Catch");
        }

        bool ExceptionFilter(Exception ex)
        {
            Console.WriteLine("Filter");
            return true;
        }
    }

    static void Method()
    {
        try
        {
            throw new ArgumentException();
        }
        finally
        {
            Console.WriteLine("Finally");
        }
    }
}
```

the method `Method` throws an exception. The first action is to examine the enclosing `catch` clauses, executing any *exception filters*. Then, the `finally` clause in `Method` executes before control transfers to the enclosing matching `catch` clause. The resulting output is:

Console

```
Filter
Finally
Catch
```

end example

The `try` block of a `try` statement is reachable if the `try` statement is reachable.

A `catch` block of a `try` statement is reachable if the `try` statement is reachable.

The `finally` block of a `try` statement is reachable if the `try` statement is reachable.

The end point of a `try` statement is reachable if both of the following are true:

- The end point of the `try` block is reachable or the end point of at least one `catch` block is reachable.
- If a `finally` block is present, the end point of the `finally` block is reachable.

13.12 The checked and unchecked statements

The `checked` and `unchecked` statements are used to control the ***overflow-checking context*** for integral-type arithmetic operations and conversions.

ANTLR

```
checked_statement
    : 'checked' block
    ;

unchecked_statement
    : 'unchecked' block
    ;
```

The `checked` statement causes all expressions in the *block* to be evaluated in a checked context, and the `unchecked` statement causes all expressions in the *block* to be evaluated in an unchecked context.

The `checked` and `unchecked` statements are precisely equivalent to the `checked` and `unchecked` operators (§12.8.19), except that they operate on blocks instead of expressions.

13.13 The lock statement

The `lock` statement obtains the mutual-exclusion lock for a given object, executes a statement, and then releases the lock.

ANTLR

```
lock_statement
    : 'lock' '(' expression ')' embedded_statement
    ;
```

The *expression* of a `lock` statement shall denote a value of a type known to be a *reference*. No implicit boxing conversion (§10.2.9) is ever performed for the *expression* of a `lock` statement, and thus it is a compile-time error for the expression to denote a value of a *value_type*.

A `lock` statement of the form

`lock (x)` ...

where `x` is an expression of a *reference_type*, is precisely equivalent to:

C#

```
bool __lockWasTaken = false;
try
{
    System.Threading.Monitor.Enter(x, ref __lockWasTaken);
    ...
}
finally
{
    if (__lockWasTaken)
    {
        System.Threading.Monitor.Exit(x);
    }
}
```

except that `x` is only evaluated once.

While a mutual-exclusion lock is held, code executing in the same execution thread can also obtain and release the lock. However, code executing in other threads is blocked from obtaining the lock until the lock is released.

13.14 The using statement

The `using` statement obtains one or more resources, executes a statement, and then disposes of the resource.

ANTLR

```
using_statement
    : 'using' '(' resource_acquisition ')' embedded_statement
    ;

resource_acquisition
    : local_variable_declaration
```

```
    | expression
    ;
```

A ***resource*** is a class or struct that implements the `System.IDisposable` interface, which includes a single parameterless method named `Dispose`. Code that is using a resource can call `Dispose` to indicate that the resource is no longer needed.

If the form of *resource_acquisition* is *local_variable_declaration* then the type of the *local_variable_declaration* shall be either `dynamic` or a type that can be implicitly converted to `System.IDisposable`. If the form of *resource_acquisition* is *expression* then this expression shall be implicitly convertible to `System.IDisposable`.

Local variables declared in a *resource_acquisition* are read-only, and shall include an initializer. A compile-time error occurs if the embedded statement attempts to modify these local variables (via assignment or the `++` and `--` operators), take the address of them, or pass them as `ref` or `out` parameters.

A `using` statement is translated into three parts: acquisition, usage, and disposal. Usage of the resource is implicitly enclosed in a `try` statement that includes a `finally` clause. This `finally` clause disposes of the resource. If a `null` resource is acquired, then no call to `Dispose` is made, and no exception is thrown. If the resource is of type `dynamic` it is dynamically converted through an implicit dynamic conversion (§10.2.10) to `IDisposable` during acquisition in order to ensure that the conversion is successful before the usage and disposal.

A `using` statement of the form

C#

```
using (ResourceType resource = «expression» ) «statement»
```

corresponds to one of three possible expansions. When `ResourceType` is a non-nullable value type or a type parameter with the value type constraint (§15.2.5), the expansion is semantically equivalent to

C#

```
{
    ResourceType resource = «expression»;
    try
    {
        «statement»;
    }
    finally
```

```
    {
        ((IDisposable)resource).Dispose();
    }
}
```

except that the cast of `resource` to `System.IDisposable` shall not cause boxing to occur.

Otherwise, when `ResourceType` is `dynamic`, the expansion is

C#

```
{
    ResourceType resource = «expression»;
    IDisposable d = resource;
    try
    {
        «statement»;
    }
    finally
    {
        if (d != null)
        {
            d.Dispose();
        }
    }
}
```

Otherwise, the expansion is

C#

```
{
    ResourceType resource = «expression»;
    try
    {
        «statement»;
    }
    finally
    {
        IDisposable d = (IDisposable)resource;
        if (d != null)
        {
            d.Dispose();
        }
    }
}
```

In any expansion, the `resource` variable is read-only in the embedded statement, and the `d` variable is inaccessible in, and invisible to, the embedded statement.

An implementation is permitted to implement a given *using_statement* differently, e.g., for performance reasons, as long as the behavior is consistent with the above expansion.

A `using` statement of the form:

```
C#

using («expression») «statement»
```

has the same three possible expansions. In this case `ResourceType` is implicitly the compile-time type of the *expression*, if it has one. Otherwise the interface `IDisposable` itself is used as the `ResourceType`. The `resource` variable is inaccessible in, and invisible to, the embedded *statement*.

When a *resource_acquisition* takes the form of a *local_variable_declaration*, it is possible to acquire multiple resources of a given type. A `using` statement of the form

```
C#

using (ResourceType r1 = e1, r2 = e2, ..., rN = eN) «statement»
```

is precisely equivalent to a sequence of nested `using` statements:

```
C#

using (ResourceType r1 = e1)
using (ResourceType r2 = e2)
...
using (ResourceType rN = eN)
«statement»
```

> *Example*: The example below creates a file named log.txt and writes two lines of text to the file. The example then opens that same file for reading and copies the contained lines of text to the console.
>
> ```
> C#
>
> class Test
> {
>     static void Main()
>     {
>         using (TextWriter w = File.CreateText("log.txt"))
>         {
>             w.WriteLine("This is line one");
>             w.WriteLine("This is line two");
>         }
> ```

```
        using (TextReader r = File.OpenText("log.txt"))
        {
            string s;
            while ((s = r.ReadLine()) != null)
            {
                Console.WriteLine(s);
            }
        }
    }
}
```

Since the `TextWriter` and `TextReader` classes implement the `IDisposable` interface, the example can use `using` statements to ensure that the underlying file is properly closed following the write or read operations.

end example

13.15 The yield statement

The `yield` statement is used in an iterator block (§13.3) to yield a value to the enumerator object (§15.14.5) or enumerable object (§15.14.6) of an iterator or to signal the end of the iteration.

ANTLR

```
yield_statement
    : 'yield' 'return' expression ';'
    | 'yield' 'break' ';'
    ;
```

`yield` is a contextual keyword (§6.4.4) and has special meaning only when used immediately before a `return` or `break` keyword.

There are several restrictions on where a `yield` statement can appear, as described in the following.

- It is a compile-time error for a `yield` statement (of either form) to appear outside a *method_body*, *operator_body*, or *accessor_body*.
- It is a compile-time error for a `yield` statement (of either form) to appear inside an anonymous function.
- It is a compile-time error for a `yield` statement (of either form) to appear in the `finally` clause of a `try` statement.
- It is a compile-time error for a `yield return` statement to appear anywhere in a `try` statement that contains any *catch_clauses*.

Example: The following example shows some valid and invalid uses of `yield` statements.

C#

```
delegate IEnumerable<int> D();

IEnumerator<int> GetEnumerator()
{
    try
    {
        yield return 1; // Ok
        yield break;    // Ok
    }
    finally
    {
        yield return 2; // Error, yield in finally
        yield break;    // Error, yield in finally
    }
    try
    {
        yield return 3; // Error, yield return in try/catch
        yield break;    // Ok
    }
    catch
    {
        yield return 4; // Error, yield return in try/catch
        yield break;    // Ok
    }
    D d = delegate
    {
        yield return 5; // Error, yield in an anonymous function
    };
}

int MyMethod()
{
    yield return 1;     // Error, wrong return type for an iterator
block
}
```

end example

An implicit conversion (§10.2) shall exist from the type of the expression in the `yield return` statement to the yield type (§15.14.4) of the iterator.

A `yield return` statement is executed as follows:

- The expression given in the statement is evaluated, implicitly converted to the yield type, and assigned to the `Current` property of the enumerator object.

- Execution of the iterator block is suspended. If the `yield return` statement is within one or more `try` blocks, the associated `finally` blocks are *not* executed at this time.
- The `MoveNext` method of the enumerator object returns `true` to its caller, indicating that the enumerator object successfully advanced to the next item.

The next call to the enumerator object's `MoveNext` method resumes execution of the iterator block from where it was last suspended.

A `yield break` statement is executed as follows:

- If the `yield break` statement is enclosed by one or more `try` blocks with associated `finally` blocks, control is initially transferred to the `finally` block of the innermost `try` statement. When and if control reaches the end point of a `finally` block, control is transferred to the `finally` block of the next enclosing `try` statement. This process is repeated until the `finally` blocks of all enclosing `try` statements have been executed.
- Control is returned to the caller of the iterator block. This is either the `MoveNext` method or `Dispose` method of the enumerator object.

Because a `yield break` statement unconditionally transfers control elsewhere, the end point of a `yield break` statement is never reachable.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

14 Namespaces

Article • 02/07/2024

14.1 General

C# programs are organized using namespaces. Namespaces are used both as an "internal" organization system for a program, and as an "external" organization system —a way of presenting program elements that are exposed to other programs.

Using directives (§14.5) are provided to facilitate the use of namespaces.

14.2 Compilation units

A *compilation_unit* consists of zero or more *extern_alias_directive*s followed by zero or more *using_directive*s followed by zero or one *global_attributes* followed by zero or more *namespace_member_declaration*s. The *compilation_unit* defines the overall structure of the input.

ANTLR

```
compilation_unit
    : extern_alias_directive* using_directive* global_attributes?
      namespace_member_declaration*
    ;
```

A C# program consists of one or more compilation units. When a C# program is compiled, all of the compilation units are processed together. Thus, compilation units can depend on each other, possibly in a circular fashion.

The *extern_alias_directive*s of a compilation unit affect the *using_directive*s, *global_attributes* and *namespace_member_declaration*s of that compilation unit, but have no effect on other compilation units.

The *using_directive*s of a compilation unit affect the *global_attributes* and *namespace_member_declaration*s of that compilation unit, but have no effect on other compilation units.

The *global_attributes* (§22.3) of a compilation unit permit the specification of attributes for the target assembly and module. Assemblies and modules act as physical containers for types. An assembly may consist of several physically separate modules.

The *namespace_member_declarations* of each compilation unit of a program contribute members to a single declaration space called the global namespace.

> *Example*:
>
> C#
>
> ```
> // File A.cs:
> class A {}
> // File B.cs:
> class B {}
> ```
>
> The two compilation units contribute to the single global namespace, in this case declaring two classes with the fully qualified names `A` and `B`. Because the two compilation units contribute to the same declaration space, it would have been an error if each contained a declaration of a member with the same name.
>
> *end example*

14.3 Namespace declarations

A *namespace_declaration* consists of the keyword namespace, followed by a namespace name and body, optionally followed by a semicolon.

ANTLR

```
namespace_declaration
    : 'namespace' qualified_identifier namespace_body ';'?
    ;

qualified_identifier
    : identifier ('.' identifier)*
    ;

namespace_body
    : '{' extern_alias_directive* using_directive*
      namespace_member_declaration* '}'
    ;
```

A *namespace_declaration* may occur as a top-level declaration in a *compilation_unit* or as a member declaration within another *namespace_declaration*. When a *namespace_declaration* occurs as a top-level declaration in a *compilation_unit*, the namespace becomes a member of the global namespace. When a *namespace_declaration* occurs within another *namespace_declaration*, the inner

namespace becomes a member of the outer namespace. In either case, the name of a namespace shall be unique within the containing namespace.

Namespaces are implicitly `public` and the declaration of a namespace cannot include any access modifiers.

Within a *namespace_body*, the optional *using_directive*s import the names of other namespaces, types and members, allowing them to be referenced directly instead of through qualified names. The optional *namespace_member_declaration*s contribute members to the declaration space of the namespace. Note that all *using_directive*s shall appear before any member declarations.

The *qualified_identifier* of a *namespace_declaration* may be a single identifier or a sequence of identifiers separated by "`.`" tokens. The latter form permits a program to define a nested namespace without lexically nesting several namespace declarations.

> *Example*:
>
> ```C#
> namespace N1.N2
> {
>     class A {}
>     class B {}
> }
> ```
>
> is semantically equivalent to
>
> ```C#
> namespace N1
> {
>     namespace N2
>     {
>         class A {}
>         class B {}
>     }
> }
> ```
>
> *end example*

Namespaces are open-ended, and two namespace declarations with the same fully qualified name (§7.8.2) contribute to the same declaration space (§7.3).

> *Example*: In the following code

C#

```
namespace N1.N2
{
    class A {}
}

namespace N1.N2
{
    class B {}
}
```

the two namespace declarations above contribute to the same declaration space, in this case declaring two classes with the fully qualified names `N1.N2.A` and `N1.N2.B`. Because the two declarations contribute to the same declaration space, it would have been an error if each contained a declaration of a member with the same name.

end example

14.4 Extern alias directives

An *extern_alias_directive* introduces an identifier that serves as an alias for a namespace. The specification of the aliased namespace is external to the source code of the program and applies also to nested namespaces of the aliased namespace.

ANTLR

```
extern_alias_directive
    : 'extern' 'alias' identifier ';'
    ;
```

The scope of an *extern_alias_directive* extends over the *using_directive*s, *global_attributes* and *namespace_member_declaration*s of its immediately containing *compilation_unit* or *namespace_body*.

Within a compilation unit or namespace body that contains an *extern_alias_directive*, the identifier introduced by the *extern_alias_directive* can be used to reference the aliased namespace. It is a compile-time error for the *identifier* to be the word `global`.

The alias introduced by an *extern_alias_directive* is very similar to the alias introduced by a *using_alias_directive*. See §14.5.2 for more detailed discussion of *extern_alias_directive*s and *using_alias_directive*s.

`alias` is a contextual keyword (§6.4.4) and only has special meaning when it immediately follows the `extern` keyword in an *extern_alias_directive*.

An error occurs if a program declares an extern alias for which no external definition is provided.

> *Example*: The following program declares and uses two extern aliases, `X` and `Y`, each of which represent the root of a distinct namespace hierarchy:
>
> C#
>
> ```
> extern alias X;
> extern alias Y;
>
> class Test
> {
>     X::N.A a;
>     X::N.B b1;
>     Y::N.B b2;
>     Y::N.C c;
> }
> ```
>
> The program declares the existence of the extern aliases `X` and `Y`, but the actual definitions of the aliases are external to the program. The identically named `N.B` classes can now be referenced as `X.N.B` and `Y.N.B`, or, using the namespace alias qualifier, `X::N.B` and `Y::N.B`. *end example*

14.5 Using directives

14.5.1 General

Using directives facilitate the use of namespaces and types defined in other namespaces. Using directives impact the name resolution process of *namespace_or_type_name*s (§7.8) and *simple_name*s (§12.8.4), but unlike declarations, *using_directive*s do not contribute new members to the underlying declaration spaces of the compilation units or namespaces within which they are used.

ANTLR

```
using_directive
    : using_alias_directive
    | using_namespace_directive
    | using_static_directive
    ;
```

A *using_alias_directive* (§14.5.2) introduces an alias for a namespace or type.

A *using_namespace_directive* (§14.5.3) imports the type members of a namespace.

A *using_static_directive* (§14.5.4) imports the nested types and static members of a type.

The scope of a *using_directive* extends over the *namespace_member_declarations* of its immediately containing compilation unit or namespace body. The scope of a *using_directive* specifically does not include its peer *using_directive*s. Thus, peer *using_directive*s do not affect each other, and the order in which they are written is insignificant. In contrast, the scope of an *extern_alias_directive* includes the *using_directive*s defined in the same compilation unit or namespace body.

14.5.2 Using alias directives

A *using_alias_directive* introduces an identifier that serves as an alias for a namespace or type within the immediately enclosing compilation unit or namespace body.

ANTLR

```
using_alias_directive
    : 'using' identifier '=' namespace_or_type_name ';'
    ;
```

Within global attributes and member declarations in a compilation unit or namespace body that contains a *using_alias_directive*, the identifier introduced by the *using_alias_directive* can be used to reference the given namespace or type.

> *Example*:
>
> C#
>
> ```
> namespace N1.N2
> {
>     class A {}
> }
> namespace N3
> {
>     using A = N1.N2.A;
>
>     class B: A {}
> }
> ```
>
> Above, within member declarations in the `N3` namespace, `A` is an alias for `N1.N2.A`, and thus class `N3.B` derives from class `N1.N2.A`. The same effect can be obtained by

creating an alias `R` for `N1.N2` and then referencing `R.A`:

```
C#

namespace N3
{
    using R = N1.N2;

    class B : R.A {}
}
```

end example

Within using directives, global attributes and member declarations in a compilation unit or namespace body that contains an *extern_alias_directive*, the identifier introduced by the *extern_alias_directive* can be used to reference the associated namespace.

Example: For example:

```
C#

namespace N1
{
    extern alias N2;

    class B : N2::A {}
}
```

Above, within member declarations in the `N1` namespace, `N2` is an alias for some namespace whose definition is external to the source code of the program. Class `N1.B` derives from class `N2.A`. The same effect can be obtained by creating an alias `A` for `N2.A` and then referencing `A`:

```
C#

namespace N1
{
    extern alias N2;

    using A = N2::A;

    class B : A {}
}
```

end example

An *extern_alias_directive* or *using_alias_directive* makes an alias available within a particular compilation unit or namespace body, but it does not contribute any new members to the underlying declaration space. In other words, an alias directive is not transitive, but, rather, affects only the compilation unit or namespace body in which it occurs.

> *Example*: In the following code
>
> ```
> C#
> ```
>
> ```
> namespace N3
> {
>     extern alias R1;
>
>     using R2 = N1.N2;
> }
>
> namespace N3
> {
>     class B : R1::A, R2.I {} // Error, R1 and R2 unknown
> }
> ```
>
> the scopes of the alias directives that introduce `R1` and `R2` only extend to member declarations in the namespace body in which they are contained, so `R1` and `R2` are unknown in the second namespace declaration. However, placing the alias directives in the containing compilation unit causes the alias to become available within both namespace declarations:
>
> ```
> C#
> ```
>
> ```
> extern alias R1;
>
> using R2 = N1.N2;
>
> namespace N3
> {
>     class B : R1::A, R2.I {}
> }
>
> namespace N3
> {
>     class C : R1::A, R2.I {}
> }
> ```
>
> *end example*

Each *extern_alias_directive* or *using_alias_directive* in a *compilation_unit* or *namespace_body* contributes a name to the alias declaration space (§7.3) of the immediately enclosing *compilation_unit* or *namespace_body*. The *identifier* of the alias directive shall be unique within the corresponding alias declaration space. The alias identifier need not be unique within the global declaration space or the declaration space of the corresponding namespace.

> *Example*:
>
> C#
>
> ```
> extern alias X;
> extern alias Y;
>
> using X = N1.N2; // Error: alias X already exists
>
> class Y {} // Ok
> ```
>
> The using alias named `X` causes an error since there is already an alias named `X` in the same compilation unit. The class named `Y` does not conflict with the extern alias named `Y` since these names are added to distinct declaration spaces. The former is added to the global declaration space and the latter is added to the alias declaration space for this compilation unit.
>
> When an alias name matches the name of a member of a namespace, usage of either shall be appropriately qualified:
>
> C#
>
> ```
> namespace N1.N2
> {
>     class B {}
> }
>
> namespace N3
> {
>     class A {}
>     class B : A {}
> }
>
> namespace N3
> {
>     using A = N1.N2;
>     using B = N1.N2.B;
>
>     class W : B {} // Error: B is ambiguous
>     class X : A.B {} // Error: A is ambiguous
>     class Y : A::B {} // Ok: uses N1.N2.B
> ```

```
    class Z : N3.B {} // Ok: uses N3.B
}
```

In the second namespace body for `N3`, unqualified use of `B` results in an error, since `N3` contains a member named `B` and the namespace body that also declares an alias with name `B`; likewise for `A`. The class `N3.B` can be referenced as `N3.B` or `global::N3.B`. The alias `A` can be used in a *qualified-alias-member* (§14.8), such as `A::B`. The alias `B` is essentially useless. It cannot be used in a *qualified_alias_member* since only namespace aliases can be used in a *qualified_alias_member* and `B` aliases a type.

end example

Just like regular members, names introduced by *alias_directives* are hidden by similarly named members in nested scopes.

Example: In the following code

```
C#

using R = N1.N2;

namespace N3
{
    class R {}
    class B: R.A {} // Error, R has no member A
}
```

the reference to `R.A` in the declaration of `B` causes a compile-time error because `R` refers to `N3.R`, not `N1.N2`.

end example

The order in which *extern_alias_directive*s are written has no significance. Likewise, the order in which *using_alias_directive*s are written has no significance, but all *using_alias_directives* shall come after all *extern_alias_directive*s in the same compilation unit or namespace body. Resolution of the *namespace_or_type_name* referenced by a *using_alias_directive* is not affected by the *using_alias_directive* itself or by other *using_directive*s in the immediately containing compilation unit or namespace body, but may be affected by *extern_alias_directive*s in the immediately containing compilation unit or namespace body. In other words, the *namespace_or_type_name* of a *using_alias_directive* is resolved as if the immediately containing compilation unit or

namespace body had no *using_directive*s but has the correct set of *extern_alias_directive*s.

> *Example*: In the following code
>
> C#
>
> ```
> namespace N1.N2 {}
>
> namespace N3
> {
>     extern alias X;
>
>     using R1 = X::N; // OK
>     using R2 = N1; // OK
>     using R3 = N1.N2; // OK
>     using R4 = R2.N2; // Error, R2 unknown
> }
> ```
>
> the last *using_alias_directive* results in a compile-time error because it is not affected by the previous *using_alias_directive*. The first *using_alias_directive* does not result in an error since the scope of the extern alias X includes the *using_alias_directive*.
>
> *end example*

A *using_alias_directive* can create an alias for any namespace or type, including the namespace within which it appears and any namespace or type nested within that namespace.

Accessing a namespace or type through an alias yields exactly the same result as accessing that namespace or type through its declared name.

> *Example*: Given
>
> C#
>
> ```
> namespace N1.N2
> {
>     class A {}
> }
>
> namespace N3
> {
>     using R1 = N1;
>     using R2 = N1.N2;
>
>     class B
>     {
> ```

```
        N1.N2.A a; // refers to N1.N2.A
        R1.N2.A b; // refers to N1.N2.A
        R2.A c; // refers to N1.N2.A
    }
}
```

the names `N1.N2.A`, `R1.N2.A`, and `R2.A` are equivalent and all refer to the class declaration whose fully qualified name is `N1.N2.A`.

end example

Although each part of a partial type (§15.2.7) is declared within the same namespace, the parts are typically written within different namespace declarations. Thus, different *extern_alias_directive*s and *using_directive*s can be present for each part. When interpreting simple names (§12.8.4) within one part, only the *extern_alias_directive*s and *using_directive*s of the namespace bodies and compilation unit enclosing that part are considered. This may result in the same identifier having different meanings in different parts.

Example:

C#

```
namespace N
{
    using List = System.Collections.ArrayList;

    partial class A
    {
        List x; // x has type System.Collections.ArrayList
    }
}

namespace N
{
    using List = Widgets.LinkedList;

    partial class A
    {
        List y; // y has type Widgets.LinkedList
    }
}
```

end example

Using aliases can name a closed constructed type, but cannot name an unbound generic type declaration without supplying type arguments.

> *Example*:
>
> C#
>
> ```
> namespace N1
> {
>     class A<T>
>     {
>         class B {}
>     }
> }
>
> namespace N2
> {
>     using W = N1.A;       // Error, cannot name unbound generic type
>     using X = N1.A.B;     // Error, cannot name unbound generic type
>     using Y = N1.A<int>;  // Ok, can name closed constructed type
>     using Z<T> = N1.A<T>; // Error, using alias cannot have type
> parameters
> }
> ```
>
> *end example*

14.5.3 Using namespace directives

A *using_namespace_directive* imports the types contained in a namespace into the immediately enclosing compilation unit or namespace body, enabling the identifier of each type to be used without qualification.

ANTLR

```
using_namespace_directive
    : 'using' namespace_name ';'
    ;
```

Within member declarations in a compilation unit or namespace body that contains a *using_namespace_directive*, the types contained in the given namespace can be referenced directly.

> *Example*:
>
> C#

```
namespace N1.N2
{
    class A {}
}

namespace N3
{
    using N1.N2;

    class B : A {}
}
```

Above, within member declarations in the `N3` namespace, the type members of `N1.N2` are directly available, and thus class `N3.B` derives from class `N1.N2.A`.

end example

A *using_namespace_directive* imports the types contained in the given namespace, but specifically does not import nested namespaces.

Example: In the following code

C#

```
namespace N1.N2
{
    class A {}
}

namespace N3
{
    using N1;
    class B : N2.A {} // Error, N2 unknown
}
```

the *using_namespace_directive* imports the types contained in `N1`, but not the namespaces nested in `N1`. Thus, the reference to `N2.A` in the declaration of `B` results in a compile-time error because no members named `N2` are in scope.

end example

Unlike a *using_alias_directive*, a *using_namespace_directive* may import types whose identifiers are already defined within the enclosing compilation unit or namespace body. In effect, names imported by a *using_namespace_directive* are hidden by similarly named members in the enclosing compilation unit or namespace body.

Example:

```
C#

namespace N1.N2
{
    class A {}
    class B {}
}

namespace N3
{
    using N1.N2;
    class A {}
}
```

Here, within member declarations in the `N3` namespace, `A` refers to `N3.A` rather than `N1.N2.A`.

end example

Because names may be ambiguous when more than one imported namespace introduces the same type name, a *using_alias_directive* is useful to disambiguate the reference.

Example: In the following code

```
C#

namespace N1
{
    class A {}
}

namespace N2
{
    class A {}
}

namespace N3
{
    using N1;
    using N2;

    class B : A {} // Error, A is ambiguous
}
```

both `N1` and `N2` contain a member `A`, and because `N3` imports both, referencing `A` in `N3` is a compile-time error. In this situation, the conflict can be resolved either through qualification of references to `A`, or by introducing a *using_alias_directive* that picks a particular `A`. For example:

```
C#

namespace N3
{
    using N1;
    using N2;
    using A = N1.A;

    class B : A {} // A means N1.A
}
```

end example

Furthermore, when more than one namespace or type imported by *using_namespace_directive*s or *using_static_directive*s in the same compilation unit or namespace body contain types or members by the same name, references to that name as a *simple_name* are considered ambiguous.

Example:

```
C#

namespace N1
{
    class A {}
}

class C
{
    public static int A;
}

namespace N2
{
    using N1;
    using static C;

    class B
    {
        void M()
        {
            A a = new A();   // Ok, A is unambiguous as a type-name
            A.Equals(2);     // Error, A is ambiguous as a simple-name
        }
```

```
    }
}
```

`N1` contains a type member `A`, and `C` contains a static field `A`, and because `N2` imports both, referencing `A` as a *simple_name* is ambiguous and a compile-time error.

end example

Like a *using_alias_directive*, a *using_namespace_directive* does not contribute any new members to the underlying declaration space of the compilation unit or namespace, but, rather, affects only the compilation unit or namespace body in which it appears.

The *namespace_name* referenced by a *using_namespace_directive* is resolved in the same way as the *namespace_or_type_name* referenced by a *using_alias_directive*. Thus, *using_namespace_directive*s in the same compilation unit or namespace body do not affect each other and can be written in any order.

14.5.4 Using static directives

A *using_static_directive* imports the nested types and static members contained directly in a type declaration into the immediately enclosing compilation unit or namespace body, enabling the identifier of each member and type to be used without qualification.

ANTLR

```
using_static_directive
    : 'using' 'static' type_name ';'
    ;
```

Within member declarations in a compilation unit or namespace body that contains a *using_static_directive*, the accessible nested types and static members (except extension methods) contained directly in the declaration of the given type can be referenced directly.

Example:

C#

```
namespace N1
{
   class A
   {
        public class B {}
```

```
        public static B M() => new B();
    }
}

namespace N2
{
    using static N1.A;

    class C
    {
        void N()
        {
            B b = M();
        }
    }
}
```

In the preceding code, within member declarations in the `N2` namespace, the static members and nested types of `N1.A` are directly available, and thus the method `N` is able to reference both the `B` and `M` members of `N1.A`.

end example

A *using_static_directive* specifically does not import extension methods directly as static methods, but makes them available for extension method invocation (§12.8.9.3).

Example:

C#

```
namespace N1
{
    static class A
    {
        public static void M(this string s){}
    }
}

namespace N2
{
    using static N1.A;

    class B
    {
        void N()
        {
            M("A");      // Error, M unknown
            "B".M();     // Ok, M known as extension method
            N1.A.M("C"); // Ok, fully qualified
        }
```

```
    }
}
```

the *using_static_directive* imports the extension method `M` contained in `N1.A`, but only as an extension method. Thus, the first reference to `M` in the body of `B.N` results in a compile-time error because no members named `M` are in scope.

end example

A *using_static_directive* only imports members and types declared directly in the given type, not members and types declared in base classes.

Example:

C#

```
namespace N1
{
    class A
    {
        public static void M(string s){}
    }

    class B : A
    {
        public static void M2(string s){}
    }
}

namespace N2
{
    using static N1.B;

    class C
    {
        void N()
        {
            M2("B");      // OK, calls B.M2
            M("C");       // Error. M unknown
        }
    }
}
```

the *using_static_directive* imports the method `M2` contained in `N1.B`, but does not import the method `M` contained in `N1.A`. Thus, the reference to `M` in the body of `C.N` results in a compile-time error because no members named `M` are in scope.

> Developers must add a second `using static` directive to specify that the methods in `N1.A` should also be imported.
>
> *end example*

Ambiguities between multiple *using_namespace_directives* and *using_static_directives* are discussed in §14.5.3.

14.6 Namespace member declarations

A *namespace_member_declaration* is either a *namespace_declaration* (§14.3) or a *type_declaration* (§14.7).

ANTLR

```
namespace_member_declaration
    : namespace_declaration
    | type_declaration
    ;
```

A compilation unit or a namespace body can contain *namespace_member_declaration*s, and such declarations contribute new members to the underlying declaration space of the containing compilation unit or namespace body.

14.7 Type declarations

A *type_declaration* is a *class_declaration* (§15.2), a *struct_declaration* (§16.2), an *interface_declaration* (§18.2), an *enum_declaration* (§19.2), or a *delegate_declaration* (§20.2).

ANTLR

```
type_declaration
    : class_declaration
    | struct_declaration
    | interface_declaration
    | enum_declaration
    | delegate_declaration
    ;
```

A *type_declaration* can occur as a top-level declaration in a compilation unit or as a member declaration within a namespace, class, or struct.

When a type declaration for a type `T` occurs as a top-level declaration in a compilation unit, the fully qualified name (§7.8.2) of the type declaration is the same as the unqualified name of the declaration (§7.8.2). When a type declaration for a type `T` occurs within a namespace, class, or struct declaration, the fully qualified name (§7.8.3) of the type declarationis `S.N`, where `S` is the fully qualified name of the containing namespace, class, or struct declaration, and `N` is the unqualified name of the declaration.

A type declared within a class or struct is called a nested type (§15.3.9).

The permitted access modifiers and the default access for a type declaration depend on the context in which the declaration takes place (§7.5.2):

- Types declared in compilation units or namespaces can have `public` or `internal` access. The default is `internal` access.
- Types declared in classes can have `public`, `protected internal`, `protected`, `private protected`, `internal`, or `private` access. The default is `private` access.
- Types declared in structs can have `public`, `internal`, or `private` access. The default is `private` access.

14.8 Qualified alias member

14.8.1 General

The ***namespace alias qualifier*** `::` makes it possible to guarantee that type name lookups are unaffected by the introduction of new types and members. The namespace alias qualifier always appears between two identifiers referred to as the left-hand and right-hand identifiers. Unlike the regular `.` qualifier, the left-hand identifier of the `::` qualifier is looked up only as an extern or using alias.

A *qualified_alias_member* provides explicit access to the global namespace and to extern or using aliases that are potentially hidden by other entities.

ANTLR

```
qualified_alias_member
    : identifier '::' identifier type_argument_list?
    ;
```

A *qualified_alias_member* can be used as a *namespace_or_type_name* (§7.8) or as the left operand in a *member_access* (§12.8.7).

A *qualified_alias_member* consists of two identifiers, referred to as the left-hand and right-hand identifiers, seperated by the `::` token and optionally followed by a *type_argument_list*. When the left-hand identifier is global then the global namespace is searched for the right-hand identifier. For any other left-hand identifier, that identifier is looked up as an extern or using alias (§14.4 and §14.5.2). A compile-time error occurs if there is no such alias or the alias references a type. If the alias references a namespace then that namespace is searched for the right-hand identifier.

A *qualified_alias_member* has one of two forms:

- `N::I<A₁, ..., Aₑ>`, where `N` and `I` represent identifiers, and `<A₁, ..., Aₑ>` is a type argument list. (`e` is always at least one.)
- `N::I`, where `N` and `I` represent identifiers. (In this case, `e` is considered to be zero.)

Using this notation, the meaning of a *qualified_alias_member* is determined as follows:

- If `N` is the identifier `global`, then the global namespace is searched for `I`:
  - If the global namespace contains a namespace named `I` and `e` is zero, then the *qualified_alias_member* refers to that namespace.
  - Otherwise, if the global namespace contains a non-generic type named `I` and `e` is zero, then the *qualified_alias_member* refers to that type.
  - Otherwise, if the global namespace contains a type named `I` that has `e` type parameters, then the *qualified_alias_member* refers to that type constructed with the given type arguments.
  - Otherwise, the *qualified_alias_member* is undefined and a compile-time error occurs.
- Otherwise, starting with the namespace declaration (§14.3) immediately containing the *qualified_alias_member* (if any), continuing with each enclosing namespace declaration (if any), and ending with the compilation unit containing the *qualified_alias_member*, the following steps are evaluated until an entity is located:
  - If the namespace declaration or compilation unit contains a *using_alias_directive* that associates N with a type, then the *qualified_alias_member* is undefined and a compile-time error occurs.
  - Otherwise, if the namespace declaration or compilation unit contains an *extern_alias_directive* or *using_alias_directive* that associates `N` with a namespace, then:
    - If the namespace associated with `N` contains a namespace named `I` and `e` is zero, then the *qualified_alias_member* refers to that namespace.
    - Otherwise, if the namespace associated with `N` contains a non-generic type named `I` and `e` is zero, then the *qualified_alias_member* refers to that type.

    - Otherwise, if the namespace associated with `N` contains a type named `I` that has `e` type parameters, then the *qualified_alias_member* refers to that type constructed with the given type arguments.
    - Otherwise, the *qualified_alias_member* is undefined and a compile-time error occurs.
- Otherwise, the *qualified_alias_member* is undefined and a compile-time error occurs.

> *Example*: In the code:
>
> ```
> C#
> ```
>
> ```
> using S = System.Net.Sockets;
>
> class A
> {
>     public static int x;
> }
>
> class C
> {
>     public void F(int A, object S)
>     {
>         // Use global::A.x instead of A.x
>         global::A.x += A;
>         // Use S::Socket instead of S.Socket
>         S::Socket s = S as S::Socket;
>     }
> }
> ```
>
> the class `A` is referenced with `global::A` and the type `System.Net.Sockets.Socket` is referenced with `S::Socket`. Using `A.x` and `S.Socket` instead would have caused compile-time errors because `A` and `S` would have resolved to the parameters.
>
> *end example*

> *Note*: The identifier `global` has special meaning only when used as the left-hand identifier of a *qualified_alias_name*. It is not a keyword and it is not itself an alias; it is a contextual keyword (§6.4.4). In the code:
>
> ```
> C#
> ```
>
> ```
> class A { }
>
> class C
> {
>     global.A x; // Error: global is not defined
> ```

```
    global::A y; // Valid: References A in the global namespace
}
```

using `global.A` causes a compile-time error since there is no entity named `global` in scope. If some entity named global were in scope, then `global` in `global.A` would have resolved to that entity.

Using `global` as the left-hand identifier of a *qualified_alias_member* always causes a lookup in the `global` namespace, even if there is a using alias named `global`. In the code:

```C#
using global = MyGlobalTypes;

class A { }

class C
{
    global.A x; // Valid: References MyGlobalTypes.A
    global::A y; // Valid: References A in the global namespace
}
```

`global.A` resolves to `MyGlobalTypes.A` and `global::A` resolves to class `A` in the global namespace.

end note

14.8.2 Uniqueness of aliases

Each compilation unit and namespace body has a separate declaration space for extern aliases and using aliases. Thus, while the name of an extern alias or using alias shall be unique within the set of extern aliases and using aliases declared in the immediately containing compilation unit or namespace body, an alias is permitted to have the same name as a type or namespace as long as it is used only with the `::` qualifier.

Example: In the following:

```C#
namespace N
{
    public class A {}
    public class B {}
}
```

```
namespace N
{
    using A = System.IO;

    class X
    {
        A.Stream s1; // Error, A is ambiguous
        A::Stream s2; // Ok
    }
}
```

the name `A` has two possible meanings in the second namespace body because both the class `A` and the using alias `A` are in scope. For this reason, use of `A` in the qualified name `A.Stream` is ambiguous and causes a compile-time error to occur. However, use of `A` with the `::` qualifier is not an error because `A` is looked up only as a namespace alias.

end example

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

15 Classes

Article • 02/07/2024

15.1 General

A class is a data structure that may contain data members (constants and fields), function members (methods, properties, events, indexers, operators, instance constructors, finalizers, and static constructors), and nested types. Class types support inheritance, a mechanism whereby a ***derived class*** can extend and specialize a ***base class***.

15.2 Class declarations

15.2.1 General

A *class_declaration* is a *type_declaration* (§14.7) that declares a new class.

ANTLR

```
class_declaration
    : attributes? class_modifier* 'partial'? 'class' identifier
        type_parameter_list? class_base? type_parameter_constraints_clause*
        class_body ';'?
    ;
```

A *class_declaration* consists of an optional set of *attributes* (§22), followed by an optional set of *class_modifiers* (§15.2.2), followed by an optional `partial` modifier (§15.2.7), followed by the keyword `class` and an *identifier* that names the class, followed by an optional *type_parameter_list* (§15.2.3), followed by an optional *class_base* specification (§15.2.4), followed by an optional set of *type_parameter_constraints_clauses* (§15.2.5), followed by a *class_body* (§15.2.6), optionally followed by a semicolon.

A class declaration shall not supply a *type_parameter_constraints_clauses* unless it also supplies a *type_parameter_list*.

A class declaration that supplies a *type_parameter_list* is a generic class declaration. Additionally, any class nested inside a generic class declaration or a generic struct declaration is itself a generic class declaration, since type arguments for the containing type shall be supplied to create a constructed type (§8.4).

15.2.2 Class modifiers

15.2.2.1 General

A *class_declaration* may optionally include a sequence of class modifiers:

ANTLR

```
class_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'abstract'
    | 'sealed'
    | 'static'
    | unsafe_modifier   // unsafe code support
    ;
```

unsafe_modifier (§23.2) is only available in unsafe code (§23).

It is a compile-time error for the same modifier to appear multiple times in a class declaration.

The `new` modifier is permitted on nested classes. It specifies that the class hides an inherited member by the same name, as described in §15.3.5. It is a compile-time error for the `new` modifier to appear on a class declaration that is not a nested class declaration.

The `public`, `protected`, `internal`, and `private` modifiers control the accessibility of the class. Depending on the context in which the class declaration occurs, some of these modifiers might not be permitted (§7.5.2).

When a partial type declaration (§15.2.7) includes an accessibility specification (via the `public`, `protected`, `internal`, and `private` modifiers), that specification shall agree with all other parts that include an accessibility specification. If no part of a partial type includes an accessibility specification, the type is given the appropriate default accessibility (§7.5.2).

The `abstract`, `sealed`, and `static` modifiers are discussed in the following subclauses.

15.2.2.2 Abstract classes

The `abstract` modifier is used to indicate that a class is incomplete and that it is intended to be used only as a base class. An ***abstract class*** differs from a ***non-abstract class*** in the following ways:

- An abstract class cannot be instantiated directly, and it is a compile-time error to use the `new` operator on an abstract class. While it is possible to have variables and values whose compile-time types are abstract, such variables and values will necessarily either be `null` or contain references to instances of non-abstract classes derived from the abstract types.
- An abstract class is permitted (but not required) to contain abstract members.
- An abstract class cannot be sealed.

When a non-abstract class is derived from an abstract class, the non-abstract class shall include actual implementations of all inherited abstract members, thereby overriding those abstract members.

> *Example*: In the following code
>
> C#
>
> ```
> abstract class A
> {
>     public abstract void F();
> }
>
> abstract class B : A
> {
>     public void G() {}
> }
>
> class C : B
> {
>     public override void F()
>     {
>         // Actual implementation of F
>     }
> }
> ```
>
> the abstract class `A` introduces an abstract method `F`. Class `B` introduces an additional method `G`, but since it doesn't provide an implementation of `F`, `B` shall also be declared abstract. Class `C` overrides `F` and provides an actual implementation. Since there are no abstract members in `C`, `C` is permitted (but not required) to be non-abstract.
>
> *end example*

If one or more parts of a partial type declaration (§15.2.7) of a class include the `abstract` modifier, the class is abstract. Otherwise, the class is non-abstract.

15.2.2.3 Sealed classes

The `sealed` modifier is used to prevent derivation from a class. A compile-time error occurs if a sealed class is specified as the base class of another class.

A sealed class cannot also be an abstract class.

> *Note*: The `sealed` modifier is primarily used to prevent unintended derivation, but it also enables certain run-time optimizations. In particular, because a sealed class is known to never have any derived classes, it is possible to transform virtual function member invocations on sealed class instances into non-virtual invocations. *end note*

If one or more parts of a partial type declaration (§15.2.7) of a class include the `sealed` modifier, the class is sealed. Otherwise, the class is unsealed.

15.2.2.4 Static classes

15.2.2.4.1 General

The `static` modifier is used to mark the class being declared as a ***static class***. A static class shall not be instantiated, shall not be used as a type and shall contain only static members. Only a static class can contain declarations of extension methods (§15.6.10).

A static class declaration is subject to the following restrictions:

- A static class shall not include a `sealed` or `abstract` modifier. (However, since a static class cannot be instantiated or derived from, it behaves as if it was both sealed and abstract.)
- A static class shall not include a *class_base* specification (§15.2.4) and cannot explicitly specify a base class or a list of implemented interfaces. A static class implicitly inherits from type `object`.
- A static class shall only contain static members (§15.3.8).

  > *Note*: All constants and nested types are classified as static members. *end note*

- A static class shall not have members with `protected`, `private protected`, or `protected internal` declared accessibility.

It is a compile-time error to violate any of these restrictions.

A static class has no instance constructors. It is not possible to declare an instance constructor in a static class, and no default instance constructor (§15.11.5) is provided for a static class.

The members of a static class are not automatically static, and the member declarations shall explicitly include a `static` modifier (except for constants and nested types). When a class is nested within a static outer class, the nested class is not a static class unless it explicitly includes a `static` modifier.

If one or more parts of a partial type declaration (§15.2.7) of a class include the `static` modifier, the class is static. Otherwise, the class is not static.

15.2.2.4.2 Referencing static class types

A *namespace_or_type_name* (§7.8) is permitted to reference a static class if

- The *namespace_or_type_name* is the `T` in a *namespace_or_type_name* of the form `T.I`, or
- The *namespace_or_type-name* is the `T` in a *typeof_expression* (§12.8.17) of the form `typeof(T)`.

A *primary_expression* (§12.8) is permitted to reference a static class if

- The *primary_expression* is the `E` in a *member_access* (§12.8.7) of the form `E.I`.

In any other context, it is a compile-time error to reference a static class.

> *Note*: For example, it is an error for a static class to be used as a base class, a constituent type (§15.3.7) of a member, a generic type argument, or a type parameter constraint. Likewise, a static class cannot be used in an array type, a new expression, a cast expression, an is expression, an as expression, a `sizeof` expression, or a default value expression. *end note*

15.2.3 Type parameters

A type parameter is a simple identifier that denotes a placeholder for a type argument supplied to create a constructed type. By constrast, a type argument (§8.4.2) is the type that is substituted for the type parameter when a constructed type is created.

```
ANTLR
```

```
type_parameter_list
    : '<' type_parameters '>'
  ;

type_parameters
    : attributes? type_parameter
    | type_parameters ',' attributes? type_parameter
    ;
```

type_parameter is defined in §8.5.

Each type parameter in a class declaration defines a name in the declaration space (§7.3) of that class. Thus, it cannot have the same name as another type parameter of that class or a member declared in that class. A type parameter cannot have the same name as the type itself.

Two partial generic type declarations (in the same program) contribute to the same unbound generic type if they have the same fully qualified name (which includes a *generic_dimension_specifier* (§12.8.17) for the number of type parameters) (§7.8.3). Two such partial type declarations shall specify the same name for each type parameter, in order.

15.2.4 Class base specification

15.2.4.1 General

A class declaration may include a *class_base* specification, which defines the direct base class of the class and the interfaces (§18) directly implemented by the class.

ANTLR

```
class_base
    : ':' class_type
    | ':' interface_type_list
    | ':' class_type ',' interface_type_list
    ;

interface_type_list
    : interface_type (',' interface_type)*
    ;
```

15.2.4.2 Base classes

When a *class_type* is included in the *class_base*, it specifies the direct base class of the class being declared. If a non-partial class declaration has no *class_base*, or if the *class_base* lists only interface types, the direct base class is assumed to be `object`. When a partial class declaration includes a base class specification, that base class specification shall reference the same type as all other parts of that partial type that include a base class specification. If no part of a partial class includes a base class specification, the base class is `object`. A class inherits members from its direct base class, as described in §15.3.4.

> *Example*: In the following code
>
> ```
> C#
> ```
>
> ```
> class A {}
> class B : A {}
> ```
>
> Class `A` is said to be the direct base class of `B`, and `B` is said to be derived from `A`. Since `A` does not explicitly specify a direct base class, its direct base class is implicitly `object`.
>
> *end example*

For a constructed class type, including a nested type declared within a generic type declaration (§15.3.9.7), if a base class is specified in the generic class declaration, the base class of the constructed type is obtained by substituting, for each *type_parameter* in the base class declaration, the corresponding *type_argument* of the constructed type.

> *Example*: Given the generic class declarations
>
> ```
> C#
> ```
>
> ```
> class B<U,V> {...}
> class G<T> : B<string,T[]> {...}
> ```
>
> the base class of the constructed type `G<int>` would be `B<string,int[]>`.
>
> *end example*

The base class specified in a class declaration can be a constructed class type (§8.4). A base class cannot be a type parameter on its own (§8.5), though it can involve the type parameters that are in scope.

Example:

```
C#
```

```
class Base<T> {}

// Valid, non-constructed class with constructed base class
class Extend1 : Base<int> {}

// Error, type parameter used as base class
class Extend2<V> : V {}

// Valid, type parameter used as type argument for base class
class Extend3<V> : Base<V> {}
```

end example

The direct base class of a class type shall be at least as accessible as the class type itself (§7.5.5). For example, it is a compile-time error for a public class to derive from a private or internal class.

The direct base class of a class type shall not be any of the following types: `System.Array`, `System.Delegate`, `System.Enum`, `System.ValueType` or the `dynamic` type. Furthermore, a generic class declaration shall not use `System.Attribute` as a direct or indirect base class (§22.2.1).

In determining the meaning of the direct base class specification `A` of a class `B`, the direct base class of `B` is temporarily assumed to be `object`, which ensures that the meaning of a base class specification cannot recursively depend on itself.

Example: The following

```
C#
```

```
class X<T>
{
    public class Y{}
}

class Z : X<Z.Y> {}
```

is in error since in the base class specification `X<Z.Y>` the direct base class of `Z` is considered to be `object`, and hence (by the rules of §7.8) `Z` is not considered to have a member `Y`.

end example

The base classes of a class are the direct base class and its base classes. In other words, the set of base classes is the transitive closure of the direct base class relationship.

Example: In the following:

C#

```
class A {...}
class B<T> : A {...}
class C<T> : B<IComparable<T>> {...}
class D<T> : C<T[]> {...}
```

the base classes of `D<int>` are `C<int[]>`, `B<IComparable<int[]>>`, `A`, and `object`.

end example

Except for class `object`, every class has exactly one direct base class. The `object` class has no direct base class and is the ultimate base class of all other classes.

It is a compile-time error for a class to depend on itself. For the purpose of this rule, a class ***directly depends on*** its direct base class (if any) and *directly depends on* the nearest enclosing class within which it is nested (if any). Given this definition, the complete set of classes upon which a class depends is the transitive closure of the *directly depends on* relationship.

Example: The example

C#

```
class A : A {}
```

is erroneous because the class depends on itself. Likewise, the example

C#

```
class A : B {}
class B : C {}
class C : A {}
```

is in error because the classes circularly depend on themselves. Finally, the example

C#

```
class A : B.C {}
class B : A
{
    public class C {}
}
```

results in a compile-time error because A depends on `B.C` (its direct base class), which depends on `B` (its immediately enclosing class), which circularly depends on `A`.

end example

A class does not depend on the classes that are nested within it.

Example: In the following code

C#

```
class A
{
    class B : A {}
}
```

`B` depends on `A` (because `A` is both its direct base class and its immediately enclosing class), but `A` does not depend on `B` (since `B` is neither a base class nor an enclosing class of `A`). Thus, the example is valid.

end example

It is not possible to derive from a sealed class.

Example: In the following code

C#

```
sealed class A {}
class B : A {} // Error, cannot derive from a sealed class
```

Class `B` is in error because it attempts to derive from the sealed class `A`.

end example

15.2.4.3 Interface implementations

A *class_base* specification may include a list of interface types, in which case the class is said to implement the given interface types. For a constructed class type, including a nested type declared within a generic type declaration (§15.3.9.7), each implemented interface type is obtained by substituting, for each *type_parameter* in the given interface, the corresponding *type_argument* of the constructed type.

The set of interfaces for a type declared in multiple parts (§15.2.7) is the union of the interfaces specified on each part. A particular interface can only be named once on each part, but multiple parts can name the same base interface(s). There shall only be one implementation of each member of any given interface.

> *Example*: In the following:
>
> ```C#
> partial class C : IA, IB {...}
> partial class C : IC {...}
> partial class C : IA, IB {...}
> ```
>
> the set of base interfaces for class `C` is `IA`, `IB`, and `IC`.
>
> *end example*

Typically, each part provides an implementation of the interface(s) declared on that part; however, this is not a requirement. A part can provide the implementation for an interface declared on a different part.

> *Example*:
>
> ```C#
> partial class X
> {
>     int IComparable.CompareTo(object o) {...}
> }
>
> partial class X : IComparable
> {
>     ...
> }
> ```
>
> *end example*

The base interfaces specified in a class declaration can be constructed interface types (§8.4, §18.2). A base interface cannot be a type parameter on its own, though it can

involve the type parameters that are in scope.

> *Example*: The following code illustrates how a class can implement and extend constructed types:
>
> C#
>
> ```
> class C<U, V> {}
> interface I1<V> {}
> class D : C<string, int>, I1<string> {}
> class E<T> : C<int, T>, I1<T> {}
> ```
>
> *end example*

Interface implementations are discussed further in §18.6.

15.2.5 Type parameter constraints

Generic type and method declarations can optionally specify type parameter constraints by including *type_parameter_constraints_clause*s.

ANTLR

```
type_parameter_constraints_clauses
    : type_parameter_constraints_clause
    | type_parameter_constraints_clauses type_parameter_constraints_clause
    ;

type_parameter_constraints_clause
    : 'where' type_parameter ':' type_parameter_constraints
    ;

type_parameter_constraints
    : primary_constraint
    | secondary_constraints
    | constructor_constraint
    | primary_constraint ',' secondary_constraints
    | primary_constraint ',' constructor_constraint
    | secondary_constraints ',' constructor_constraint
    | primary_constraint ',' secondary_constraints ','
constructor_constraint
    ;

primary_constraint
    : class_type
    | 'class'
    | 'struct'
    | 'unmanaged'
    ;
```

```
secondary_constraints
    : interface_type
    | type_parameter
    | secondary_constraints ',' interface_type
    | secondary_constraints ',' type_parameter
    ;

constructor_constraint
    : 'new' '(' ')'
    ;
```

Each *type_parameter_constraints_clause* consists of the token `where`, followed by the name of a type parameter, followed by a colon and the list of constraints for that type parameter. There can be at most one `where` clause for each type parameter, and the `where` clauses can be listed in any order. Like the `get` and `set` tokens in a property accessor, the `where` token is not a keyword.

The list of constraints given in a `where` clause can include any of the following components, in this order: a single primary constraint, one or more secondary constraints, and the constructor constraint, `new()`.

A primary constraint can be a class type, the ***reference type constraint*** `class`, the ***value type constraint*** `struct`, or the ***unmanaged type constraint*** `unmanaged`.

A secondary constraint can be a *type_parameter* or *interface_type*.

The reference type constraint specifies that a type argument used for the type parameter shall be a reference type. All class types, interface types, delegate types, array types, and type parameters known to be a reference type (as defined below) satisfy this constraint.

The value type constraint specifies that a type argument used for the type parameter shall be a non-nullable value type. All non-nullable struct types, enum types, and type parameters having the value type constraint satisfy this constraint. Note that although classified as a value type, a nullable value type (§8.3.12) does not satisfy the value type constraint. A type parameter having the value type constraint shall not also have the *constructor_constraint*, although it may be used as a type argument for another type parameter with a *constructor_constraint*.

> *Note*: The `System.Nullable<T>` type specifies the non-nullable value type constraint for `T`. Thus, recursively constructed types of the forms `T??` and `Nullable<Nullable<T>>` are prohibited. *end note*

Because `unmanaged` is not a keyword, in *primary_constraint* the unmanaged constraint is always syntactically ambiguous with *class_type*. For compatibility reasons, if a name lookup (§12.8.4) of the name `unmanaged` succeeds it is treated as a `class_type`. Otherwise it is treated as the unmanaged constraint.

The unmanaged type constraint specifies that a type argument used for the type parameter shall be a non-nullable unmanaged type (§8.8).

Pointer types are never allowed to be type arguments, and don't satisfy any type constraints, even unmanaged, despite being unmanaged types.

If a constraint is a class type, an interface type, or a type parameter, that type specifies a minimal "base type" that every type argument used for that type parameter shall support. Whenever a constructed type or generic method is used, the type argument is checked against the constraints on the type parameter at compile-time. The type argument supplied shall satisfy the conditions described in §8.4.5.

A *class_type* constraint shall satisfy the following rules:

- The type shall be a class type.
- The type shall not be `sealed`.
- The type shall not be one of the following types: `System.Array` or `System.ValueType`.
- The type shall not be `object`.
- At most one constraint for a given type parameter may be a class type.

A type specified as an *interface_type* constraint shall satisfy the following rules:

- The type shall be an interface type.
- A type shall not be specified more than once in a given `where` clause.

In either case, the constraint may involve any of the type parameters of the associated type or method declaration as part of a constructed type, and may involve the type being declared.

Any class or interface type specified as a type parameter constraint shall be at least as accessible (§7.5.5) as the generic type or method being declared.

A type specified as a *type_parameter* constraint shall satisfy the following rules:

- The type shall be a type parameter.
- A type shall not be specified more than once in a given `where` clause.

In addition there shall be no cycles in the dependency graph of type parameters, where dependency is a transitive relation defined by:

- If a type parameter `T` is used as a constraint for type parameter `S` then `S` ***depends on*** `T`.
- If a type parameter `S` depends on a type parameter `T` and `T` depends on a type parameter `U` then `S` *depends on* `U`.

Given this relation, it is a compile-time error for a type parameter to depend on itself (directly or indirectly).

Any constraints shall be consistent among dependent type parameters. If type parameter `S` depends on type parameter `T` then:

- `T` shall not have the value type constraint. Otherwise, `T` is effectively sealed so `S` would be forced to be the same type as `T`, eliminating the need for two type parameters.
- If `S` has the value type constraint then `T` shall not have a *class_type* constraint.
- If `S` has a *class_type* constraint `A` and `T` has a *class_type* constraint `B` then there shall be an identity conversion or implicit reference conversion from `A` to `B` or an implicit reference conversion from `B` to `A`.
- If `S` also depends on type parameter `U` and `U` has a *class_type* constraint `A` and `T` has a *class_type* constraint `B` then there shall be an identity conversion or implicit reference conversion from `A` to `B` or an implicit reference conversion from `B` to `A`.

It is valid for `S` to have the value type constraint and `T` to have the reference type constraint. Effectively this limits `T` to the types `System.Object`, `System.ValueType`, `System.Enum`, and any interface type.

If the `where` clause for a type parameter includes a constructor constraint (which has the form `new()`), it is possible to use the `new` operator to create instances of the type (§12.8.16.2). Any type argument used for a type parameter with a constructor constraint shall be a value type, a non-abstract class having a public parameterless constructor, or a type parameter having the value type constraint or constructor constraint.

It is a compile-time error for *type_parameter_constraints* having a *primary_constraint* of `struct` or `unmanaged` to also have a *constructor_constraint*.

> *Example*: The following are examples of constraints:
>
> ```
> C#
> ```

```
interface IPrintable
{
    void Print();
}

interface IComparable<T>
{
    int CompareTo(T value);
}

interface IKeyProvider<T>
{
    T GetKey();
}

class Printer<T> where T : IPrintable {...}
class SortedList<T> where T : IComparable<T> {...}

class Dictionary<K,V>
    where K : IComparable<K>
    where V : IPrintable, IKeyProvider<K>, new()
{
    ...
}
```

The following example is in error because it causes a circularity in the dependency graph of the type parameters:

C#

```
class Circular<S,T>
    where S: T
    where T: S // Error, circularity in dependency graph
{
    ...
}
```

The following examples illustrate additional invalid situations:

C#

```
class Sealed<S,T>
    where S : T
    where T : struct // Error, `T` is sealed
{
    ...
}

class A {...}
class B {...}
```

```
class Incompat<S,T>
    where S : A, T
    where T : B // Error, incompatible class-type constraints
{
    ...
}

class StructWithClass<S,T,U>
    where S : struct, T
    where T : U
    where U : A // Error, A incompatible with struct
{
    ...
}
```

end example

The ***dynamic erasure*** of a type `C` is type C_x constructed as follows:

- If `C` is a nested type `Outer.Inner` then C_x is a nested type $\texttt{Outer}_x\texttt{.Inner}_x$.
- If `C` C_x is a constructed type $\texttt{G<A}^1\texttt{, ..., A}^n\texttt{>}$ with type arguments $\texttt{A}^1\texttt{, ..., A}^n$ then C_x is the constructed type $\texttt{G<A}^1{}_x\texttt{, ..., A}^n{}_x\texttt{>}$.
- If `C` is an array type `E[]` then C_x is the array type E_x`[]`.
- If `C` is dynamic then C_x is `object`.
- Otherwise, C_x is `C`.

The ***effective base class*** of a type parameter `T` is defined as follows:

Let `R` be a set of types such that:

- For each constraint of `T` that is a type parameter, `R` contains its effective base class.
- For each constraint of `T` that is a struct type, `R` contains `System.ValueType`.
- For each constraint of `T` that is an enumeration type, `R` contains `System.Enum`.
- For each constraint of `T` that is a delegate type, `R` contains its dynamic erasure.
- For each constraint of `T` that is an array type, `R` contains `System.Array`.
- For each constraint of `T` that is a class type, `R` contains its dynamic erasure.

Then

- If `T` has the value type constraint, its effective base class is `System.ValueType`.
- Otherwise, if `R` is empty then the effective base class is `object`.
- Otherwise, the effective base class of `T` is the most-encompassed type (§10.5.3) of set `R`. If the set has no encompassed type, the effective base class of `T` is `object`.

The consistency rules ensure that the most-encompassed type exists.

If the type parameter is a method type parameter whose constraints are inherited from the base method the effective base class is calculated after type substitution.

These rules ensure that the effective base class is always a *class_type*.

The ***effective interface set*** of a type parameter `T` is defined as follows:

- If `T` has no *secondary_constraints*, its effective interface set is empty.
- If `T` has *interface_type* constraints but no *type_parameter* constraints, its effective interface set is the set of dynamic erasures of its *interface_type* constraints.
- If `T` has no *interface_type* constraints but has *type_parameter* constraints, its effective interface set is the union of the effective interface sets of its *type_parameter* constraints.
- If `T` has both *interface_type* constraints and *type_parameter* constraints, its effective interface set is the union of the set of dynamic erasures of its *interface_type* constraints and the effective interface sets of its *type_parameter* constraints.

A type parameter is *known to be a reference type* if it has the reference type constraint or its effective base class is not `object` or `System.ValueType`.

Values of a constrained type parameter type can be used to access the instance members implied by the constraints.

> *Example*: In the following:
>
> C#
>
> ```
> interface IPrintable
> {
>     void Print();
> }
>
> class Printer<T> where T : IPrintable
> {
>     void PrintOne(T x) => x.Print();
> }
> ```
>
> the methods of `IPrintable` can be invoked directly on `x` because `T` is constrained to always implement `IPrintable`.
>
> *end example*

When a partial generic type declaration includes constraints, the constraints shall agree with all other parts that include constraints. Specifically, each part that includes constraints shall have constraints for the same set of type parameters, and for each type parameter, the sets of primary, secondary, and constructor constraints shall be equivalent. Two sets of constraints are equivalent if they contain the same members. If no part of a partial generic type specifies type parameter constraints, the type parameters are considered unconstrained.

> *Example*:
>
> C#
>
> ```
> partial class Map<K,V>
>     where K : IComparable<K>
>     where V : IKeyProvider<K>, new()
> {
>     ...
> }
>
> partial class Map<K,V>
>     where V : IKeyProvider<K>, new()
>     where K : IComparable<K>
> {
>     ...
> }
>
> partial class Map<K,V>
> {
>     ...
> }
> ```
>
> is correct because those parts that include constraints (the first two) effectively specify the same set of primary, secondary, and constructor constraints for the same set of type parameters, respectively.
>
> *end example*

15.2.6 Class body

The *class_body* of a class defines the members of that class.

ANTLR

```
class_body
    : '{' class_member_declaration* '}'
    ;
```

15.2.7 Partial declarations

The modifier `partial` is used when defining a class, struct, or interface type in multiple parts. The `partial` modifier is a contextual keyword (§6.4.4) and only has special meaning immediately before one of the keywords `class`, `struct`, or `interface`.

Each part of a ***partial type*** declaration shall include a `partial` modifier and shall be declared in the same namespace or containing type as the other parts. The `partial` modifier indicates that additional parts of the type declaration might exist elsewhere, but the existence of such additional parts is not a requirement; it is valid for the only declaration of a type to include the `partial` modifier.

All parts of a partial type shall be compiled together such that the parts can be merged at compile-time. Partial types specifically do not allow already compiled types to be extended.

Nested types can be declared in multiple parts by using the `partial` modifier. Typically, the containing type is declared using `partial` as well, and each part of the nested type is declared in a different part of the containing type.

> *Example*: The following partial class is implemented in two parts, which reside in different compilation units. The first part is machine generated by a database-mapping tool while the second part is manually authored:
>
> C#
>
> ```
> public partial class Customer
> {
>     private int id;
>     private string name;
>     private string address;
>     private List<Order> orders;
>
>     public Customer()
>     {
>         ...
>     }
> }
>
> // File: Customer2.cs
> public partial class Customer
> {
>     public void SubmitOrder(Order orderSubmitted) =>
> orders.Add(orderSubmitted);
>
>     public bool HasOutstandingOrders() => orders.Count > 0;
> }
> ```

When the two parts above are compiled together, the resulting code behaves as if the class had been written as a single unit, as follows:

C#

```
public class Customer
{
    private int id;
    private string name;
    private string address;
    private List<Order> orders;

    public Customer()
    {
        ...
    }

    public void SubmitOrder(Order orderSubmitted) =>
orders.Add(orderSubmitted);

    public bool HasOutstandingOrders() => orders.Count > 0;
}
```

end example

The handling of attributes specified on the type or type parameters of different parts of a partial declaration is discussed in §22.3.

15.3 Class members

15.3.1 General

The members of a class consist of the members introduced by its *class_member_declaration*s and the members inherited from the direct base class.

ANTLR

```
class_member_declaration
    : constant_declaration
    | field_declaration
    | method_declaration
    | property_declaration
    | event_declaration
    | indexer_declaration
    | operator_declaration
    | constructor_declaration
    | finalizer_declaration
```

```
    | static_constructor_declaration
    | type_declaration
    ;
```

The members of a class are divided into the following categories:

- Constants, which represent constant values associated with the class (§15.4).
- Fields, which are the variables of the class (§15.5).
- Methods, which implement the computations and actions that can be performed by the class (§15.6).
- Properties, which define named characteristics and the actions associated with reading and writing those characteristics (§15.7).
- Events, which define notifications that can be generated by the class (§15.8).
- Indexers, which permit instances of the class to be indexed in the same way (syntactically) as arrays (§15.9).
- Operators, which define the expression operators that can be applied to instances of the class (§15.10).
- Instance constructors, which implement the actions required to initialize instances of the class (§15.11)
- Finalizers, which implement the actions to be performed before instances of the class are permanently discarded (§15.13).
- Static constructors, which implement the actions required to initialize the class itself (§15.12).
- Types, which represent the types that are local to the class (§14.7).

A *class_declaration* creates a new declaration space (§7.3), and the *type_parameter*s and the *class_member_declaration*s immediately contained by the *class_declaration* introduce new members into this declaration space. The following rules apply to *class_member_declaration*s:

- Instance constructors, finalizers, and static constructors shall have the same name as the immediately enclosing class. All other members shall have names that differ from the name of the immediately enclosing class.

- The name of a type parameter in the *type_parameter_list* of a class declaration shall differ from the names of all other type parameters in the same *type_parameter_list* and shall differ from the name of the class and the names of all members of the class.

- The name of a type shall differ from the names of all non-type members declared in the same class. If two or more type declarations share the same fully qualified name, the declarations shall have the `partial` modifier (§15.2.7) and these declarations combine to define a single type.

> *Note*: Since the fully qualified name of a type declaration encodes the number of type parameters, two distinct types may share the same name as long as they have different number of type parameters. *end note*

- The name of a constant, field, property, or event shall differ from the names of all other members declared in the same class.
- The name of a method shall differ from the names of all other non-methods declared in the same class. In addition, the signature (§7.6) of a method shall differ from the signatures of all other methods declared in the same class, and two methods declared in the same class shall not have signatures that differ solely by `in`, `out`, and `ref`.
- The signature of an instance constructor shall differ from the signatures of all other instance constructors declared in the same class, and two constructors declared in the same class shall not have signatures that differ solely by `ref` and `out`.
- The signature of an indexer shall differ from the signatures of all other indexers declared in the same class.
- The signature of an operator shall differ from the signatures of all other operators declared in the same class.

The inherited members of a class (§15.3.4) are not part of the declaration space of a class.

> *Note*: Thus, a derived class is allowed to declare a member with the same name or signature as an inherited member (which in effect hides the inherited member). *end note*

The set of members of a type declared in multiple parts (§15.2.7) is the union of the members declared in each part. The bodies of all parts of the type declaration share the same declaration space (§7.3), and the scope of each member (§7.7) extends to the bodies of all the parts. The accessibility domain of any member always includes all the parts of the enclosing type; a private member declared in one part is freely accessible from another part. It is a compile-time error to declare the same member in more than one part of the type, unless that member is a type having the `partial` modifier.

> *Example*:
>
> ```
> C#
> ```

```
partial class A
{
    int x;                    // Error, cannot declare x more than once

    partial class Inner       // Ok, Inner is a partial type
    {
        int y;
    }
}

partial class A
{
    int x;                    // Error, cannot declare x more than once

    partial class Inner       // Ok, Inner is a partial type
    {
        int z;
    }
}
```

end example

Field initialization order can be significant within C# code, and some guarantees are provided, as defined in §15.5.6.1. Otherwise, the ordering of members within a type is rarely significant, but may be significant when interfacing with other languages and environments. In these cases, the ordering of members within a type declared in multiple parts is undefined.

15.3.2 The instance type

Each class declaration has an associated ***instance type***. For a generic class declaration, the instance type is formed by creating a constructed type (§8.4) from the type declaration, with each of the supplied type arguments being the corresponding type parameter. Since the instance type uses the type parameters, it can only be used where the type parameters are in scope; that is, inside the class declaration. The instance type is the type of `this` for code written inside the class declaration. For non-generic classes, the instance type is simply the declared class.

Example: The following shows several class declarations along with their instance types:

C#

```
class A<T>              // instance type: A<T>
{
```

```
    class B {}          // instance type: A<T>.B
    class C<U> {}       // instance type: A<T>.C<U>
}
class D {}              // instance type: D
```

end example

15.3.3 Members of constructed types

The non-inherited members of a constructed type are obtained by substituting, for each *type_parameter* in the member declaration, the corresponding *type_argument* of the constructed type. The substitution process is based on the semantic meaning of type declarations, and is not simply textual substitution.

Example: Given the generic class declaration

C#

```
class Gen<T,U>
{
    public T[,] a;
    public void G(int i, T t, Gen<U,T> gt) {...}
    public U Prop { get {...} set {...} }
    public int H(double d) {...}
}
```

the constructed type `Gen<int[],IComparable<string>>` has the following members:

C#

```
public int[,][] a;
public void G(int i, int[] t, Gen<IComparable<string>,int[]> gt) {...}
public IComparable<string> Prop { get {...} set {...} }
public int H(double d) {...}
```

The type of the member `a` in the generic class declaration `Gen` is "two-dimensional array of `T`", so the type of the member `a` in the constructed type above is "two-dimensional array of single-dimensional array of `int`", or `int[,][]`.

end example

Within instance function members, the type of `this` is the instance type (§15.3.2) of the containing declaration.

All members of a generic class can use type parameters from any enclosing class, either directly or as part of a constructed type. When a particular closed constructed type (§8.4.3) is used at run-time, each use of a type parameter is replaced with the type argument supplied to the constructed type.

> *Example*:
>
> ```
> C#
> ```
>
> ```
> class C<V>
> {
>     public V f1;
>     public C<V> f2 = null;
>
>     public C(V x)
>     {
>         this.f1 = x;
>         this.f2 = this;
>     }
> }
>
> class Application
> {
>     static void Main()
>     {
>         C<int> x1 = new C<int>(1);
>         Console.WriteLine(x1.f1);              // Prints 1
>
>         C<double> x2 = new C<double>(3.1415);
>         Console.WriteLine(x2.f1);              // Prints 3.1415
>     }
> }
> ```
>
> *end example*

15.3.4 Inheritance

A class ***inherits*** the members of its direct base class. Inheritance means that a class implicitly contains all members of its direct base class, except for the instance constructors, finalizers, and static constructors of the base class. Some important aspects of inheritance are:

- Inheritance is transitive. If `C` is derived from `B`, and `B` is derived from `A`, then `C` inherits the members declared in `B` as well as the members declared in `A`.
- A derived class *extends* its direct base class. A derived class can add new members to those it inherits, but it cannot remove the definition of an inherited member.

- Instance constructors, finalizers, and static constructors are not inherited, but all other members are, regardless of their declared accessibility (§7.5). However, depending on their declared accessibility, inherited members might not be accessible in a derived class.
- A derived class can *hide* (§7.7.2.3) inherited members by declaring new members with the same name or signature. However, hiding an inherited member does not remove that member—it merely makes that member inaccessible directly through the derived class.
- An instance of a class contains a set of all instance fields declared in the class and its base classes, and an implicit conversion (§10.2.8) exists from a derived class type to any of its base class types. Thus, a reference to an instance of some derived class can be treated as a reference to an instance of any of its base classes.
- A class can declare virtual methods, properties, indexers, and events, and derived classes can override the implementation of these function members. This enables classes to exhibit polymorphic behavior wherein the actions performed by a function member invocation vary depending on the run-time type of the instance through which that function member is invoked.

The inherited members of a constructed class type are the members of the immediate base class type (§15.2.4.2), which is found by substituting the type arguments of the constructed type for each occurrence of the corresponding type parameters in the *base_class_specification*. These members, in turn, are transformed by substituting, for each *type_parameter* in the member declaration, the corresponding *type_argument* of the *base_class_specification*.

> *Example*:
>
> C#
>
> ```
> class B<U>
> {
>     public U F(long index) {...}
> }
>
> class D<T> : B<T[]>
> {
>     public T G(string s) {...}
> }
> ```
>
> In the code above, the constructed type `D<int>` has a non-inherited member public `int G(string s)` obtained by substituting the type argument `int` for the type

> parameter `T`. `D<int>` also has an inherited member from the class declaration `B`. This inherited member is determined by first determining the base class type `B<int[]>` of `D<int>` by substituting `int` for `T` in the base class specification `B<T[]>`. Then, as a type argument to `B`, `int[]` is substituted for `U` in `public U F(long index)`, yielding the inherited member `public int[] F(long index)`.
>
> *end example*

15.3.5 The new modifier

A *class_member_declaration* is permitted to declare a member with the same name or signature as an inherited member. When this occurs, the derived class member is said to *hide* the base class member. See §7.7.2.3 for a precise specification of when a member hides an inherited member.

An inherited member `M` is considered to be ***available*** if `M` is accessible and there is no other inherited accessible member N that already hides `M`. Implicitly hiding an inherited member is not considered an error, but it does cause the compiler to issue a warning unless the declaration of the derived class member includes a `new` modifier to explicitly indicate that the derived member is intended to hide the base member. If one or more parts of a partial declaration (§15.2.7) of a nested type include the `new` modifier, no warning is issued if the nested type hides an available inherited member.

If a `new` modifier is included in a declaration that doesn't hide an available inherited member, a warning to that effect is issued.

15.3.6 Access modifiers

A *class_member_declaration* can have any one of the permitted kinds of declared accessibility (§7.5.2): `public`, `protected internal`, `protected`, `private protected`, `internal`, or `private`. Except for the `protected internal` and `private protected` combinations, it is a compile-time error to specify more than one access modifier. When a *class_member_declaration* does not include any access modifiers, `private` is assumed.

15.3.7 Constituent types

Types that are used in the declaration of a member are called the ***constituent types*** of that member. Possible constituent types are the type of a constant, field, property, event, or indexer, the return type of a method or operator, and the parameter types of a

method, indexer, operator, or instance constructor. The constituent types of a member shall be at least as accessible as that member itself (§7.5.5).

15.3.8 Static and instance members

Members of a class are either ***static members*** or ***instance members***.

> *Note*: Generally speaking, it is useful to think of static members as belonging to classes and instance members as belonging to objects (instances of classes). *end note*

When a field, method, property, event, operator, or constructor declaration includes a `static` modifier, it declares a static member. In addition, a constant or type declaration implicitly declares a static member. Static members have the following characteristics:

- When a static member `M` is referenced in a *member_access* (§12.8.7) of the form `E.M`, `E` shall denote a type that has a member `M`. It is a compile-time error for `E` to denote an instance.
- A static field in a non-generic class identifies exactly one storage location. No matter how many instances of a non-generic class are created, there is only ever one copy of a static field. Each distinct closed constructed type (§8.4.3) has its own set of static fields, regardless of the number of instances of the closed constructed type.
- A static function member (method, property, event, operator, or constructor) does not operate on a specific instance, and it is a compile-time error to refer to this in such a function member.

When a field, method, property, event, indexer, constructor, or finalizer declaration does not include a static modifier, it declares an instance member. (An instance member is sometimes called a non-static member.) Instance members have the following characteristics:

- When an instance member `M` is referenced in a *member_access* (§12.8.7) of the form `E.M`, `E` shall denote an instance of a type that has a member `M`. It is a binding-time error for E to denote a type.
- Every instance of a class contains a separate set of all instance fields of the class.
- An instance function member (method, property, indexer, instance constructor, or finalizer) operates on a given instance of the class, and this instance can be accessed as `this` (§12.8.13).

Example: The following example illustrates the rules for accessing static and instance members:

```
C#
```

```
class Test
{
    int x;
    static int y;
    void F()
    {
        x = 1;              // Ok, same as this.x = 1
        y = 1;              // Ok, same as Test.y = 1
    }

    static void G()
    {
        x = 1;              // Error, cannot access this.x
        y = 1;              // Ok, same as Test.y = 1
    }

    static void Main()
    {
        Test t = new Test();
        t.x = 1;       // Ok
        t.y = 1;       // Error, cannot access static member through
instance
        Test.x = 1;    // Error, cannot access instance member through
type
        Test.y = 1;    // Ok
    }
}
```

The `F` method shows that in an instance function member, a *simple_name* (§12.8.4) can be used to access both instance members and static members. The `G` method shows that in a static function member, it is a compile-time error to access an instance member through a *simple_name*. The `Main` method shows that in a *member_access* (§12.8.7), instance members shall be accessed through instances, and static members shall be accessed through types.

end example

15.3.9 Nested types

15.3.9.1 General

A type declared within a class or struct is called a ***nested type***. A type that is declared within a compilation unit or namespace is called a ***non-nested type***.

> *Example*: In the following example:
>
> ```
> C#
> ```
>
> ```
> class A
> {
>     class B
>     {
>         static void F()
>         {
>             Console.WriteLine("A.B.F");
>         }
>     }
> }
> ```
>
> class `B` is a nested type because it is declared within class `A`, and class `A` is a non-nested type because it is declared within a compilation unit.
>
> *end example*

15.3.9.2 Fully qualified name

The fully qualified name (§7.8.3) for a nested type declarationis `S.N` where `S` is the fully qualified name of the type declarationin which type `N` is declared and `N` is the unqualified name (§7.8.2) of the nested type declaration (including any *generic_dimension_specifier* (§12.8.17)).

15.3.9.3 Declared accessibility

Non-nested types can have `public` or `internal` declared accessibility and have `internal` declared accessibility by default. Nested types can have these forms of declared accessibility too, plus one or more additional forms of declared accessibility, depending on whether the containing type is a class or struct:

- A nested type that is declared in a class can have any of the permitted kinds of declared accessibility and, like other class members, defaults to `private` declared accessibility.
- A nested type that is declared in a struct can have any of three forms of declared accessibility (`public`, `internal`, or `private`) and, like other struct members, defaults to `private` declared accessibility.

Example: The example

```
C#
```

```
public class List
{
    // Private data structure
    private class Node
    {
        public object Data;
        public Node Next;

        public Node(object data, Node next)
        {
            this.Data = data;
            this.Next = next;
        }
    }

    private Node first = null;
    private Node last = null;

    // Public interface
    public void AddToFront(object o) {...}
    public void AddToBack(object o) {...}
    public object RemoveFromFront() {...}
    public object RemoveFromBack() {...}
    public int Count { get {...} }
}
```

declares a private nested class `Node`.

end example

15.3.9.4 Hiding

A nested type may hide (§7.7.2.2) a base member. The `new` modifier (§15.3.5) is permitted on nested type declarations so that hiding can be expressed explicitly.

Example: The example

```
C#
```

```
class Base
{
    public static void M()
    {
        Console.WriteLine("Base.M");
    }
}
```

```
class Derived: Base
{
    public new class M
    {
        public static void F()
        {
            Console.WriteLine("Derived.M.F");
        }
    }
}

class Test
{
    static void Main()
    {
        Derived.M.F();
    }
}
```

shows a nested class `M` that hides the method `M` defined in `Base`.

end example

15.3.9.5 this access

A nested type and its containing type do not have a special relationship with regard to *this_access* (§12.8.13). Specifically, `this` within a nested type cannot be used to refer to instance members of the containing type. In cases where a nested type needs access to the instance members of its containing type, access can be provided by providing the `this` for the instance of the containing type as a constructor argument for the nested type.

Example: The following example

C#

```
class C
{
    int i = 123;
    public void F()
    {
        Nested n = new Nested(this);
        n.G();
    }

    public class Nested
    {
        C this_c;
```

```
        public Nested(C c)
        {
            this_c = c;
        }

        public void G()
        {
            Console.WriteLine(this_c.i);
        }
    }
}

class Test
{
    static void Main()
    {
        C c = new C();
        c.F();
    }
}
```

shows this technique. An instance of `C` creates an instance of `Nested`, and passes its own this to `Nested`'s constructor in order to provide subsequent access to `C`'s instance members.

end example

15.3.9.6 Access to private and protected members of the containing type

A nested type has access to all of the members that are accessible to its containing type, including members of the containing type that have `private` and `protected` declared accessibility.

Example: The example

C#

```
class C
{
    private static void F() => Console.WriteLine("C.F");

    public class Nested
    {
        public static void G() => F();
    }
}
```

```
class Test
{
    static void Main() => C.Nested.G();
}
```

shows a class `C` that contains a nested class `Nested`. Within `Nested`, the method `G` calls the static method `F` defined in `C`, and `F` has private declared accessibility.

end example

A nested type also may access protected members defined in a base type of its containing type.

Example: In the following code

C#

```
class Base
{
    protected void F() => Console.WriteLine("Base.F");
}

class Derived: Base
{
    public class Nested
    {
        public void G()
        {
            Derived d = new Derived();
            d.F(); // ok
        }
    }
}

class Test
{
    static void Main()
    {
        Derived.Nested n = new Derived.Nested();
        n.G();
    }
}
```

the nested class `Derived.Nested` accesses the protected method `F` defined in `Derived`'s base class, `Base`, by calling through an instance of `Derived`.

end example

15.3.9.7 Nested types in generic classes

A generic class declaration may contain nested type declarations. The type parameters of the enclosing class may be used within the nested types. A nested type declaration may contain additional type parameters that apply only to the nested type.

Every type declaration contained within a generic class declaration is implicitly a generic type declaration. When writing a reference to a type nested within a generic type, the containing constructed type, including its type arguments, shall be named. However, from within the outer class, the nested type may be used without qualification; the instance type of the outer class may be implicitly used when constructing the nested type.

> *Example*: The following shows three different correct ways to refer to a constructed type created from `Inner`; the first two are equivalent:
>
> C#
>
> ```
> class Outer<T>
> {
>     class Inner<U>
>     {
>         public static void F(T t, U u) {...}
>     }
>
>     static void F(T t)
>     {
>         Outer<T>.Inner<string>.F(t, "abc");    // These two statements
> have
>         Inner<string>.F(t, "abc");             // the same effect
>         Outer<int>.Inner<string>.F(3, "abc");  // This type is different
>         Outer.Inner<string>.F(t, "abc");       // Error, Outer needs
> type arg
>     }
> }
> ```
>
> *end example*

Although it is bad programming style, a type parameter in a nested type can hide a member or type parameter declared in the outer type.

> *Example*:
>
> C#
>
> ```
> class Outer<T>
> {
> ```

```
    class Inner<T>                                  // Valid, hides
Outer's T
    {
        public T t;                                 // Refers to Inner's
T
    }
}
```

end example

15.3.10 Reserved member names

15.3.10.1 General

To facilitate the underlying C# run-time implementation, for each source member declaration that is a property, event, or indexer, the implementation shall reserve two method signatures based on the kind of the member declaration, its name, and its type (§15.3.10.2, §15.3.10.3, §15.3.10.4). It is a compile-time error for a program to declare a member whose signature matches a signature reserved by a member declared in the same scope, even if the underlying run-time implementation does not make use of these reservations.

The reserved names do not introduce declarations, thus they do not participate in member lookup. However, a declaration's associated reserved method signatures do participate in inheritance (§15.3.4), and can be hidden with the `new` modifier (§15.3.5).

> *Note*: The reservation of these names serves three purposes:
>
> 1. To allow the underlying implementation to use an ordinary identifier as a method name for get or set access to the C# language feature.
> 2. To allow other languages to interoperate using an ordinary identifier as a method name for get or set access to the C# language feature.
> 3. To help ensure that the source accepted by one conforming compiler is accepted by another, by making the specifics of reserved member names consistent across all C# implementations.
>
> *end note*

The declaration of a finalizer (§15.13) also causes a signature to be reserved (§15.3.10.5).

15.3.10.2 Member names reserved for properties

For a property `P` (§15.7) of type `T`, the following signatures are reserved:

```
C#

T get_P();
void set_P(T value);
```

Both signatures are reserved, even if the property is read-only or write-only.

> *Example*: In the following code
>
> ```
> C#
>
> class A
> {
>     public int P
>     {
>         get => 123;
>     }
> }
>
> class B : A
> {
>     public new int get_P() => 456;
>
>     public new void set_P(int value)
>     {
>     }
> }
>
> class Test
> {
>     static void Main()
>     {
>         B b = new B();
>         A a = b;
>         Console.WriteLine(a.P);
>         Console.WriteLine(b.P);
>         Console.WriteLine(b.get_P());
>     }
> }
> ```
>
> A class `A` defines a read-only property `P`, thus reserving signatures for `get_P` and `set_P` methods. `A` class `B` derives from `A` and hides both of these reserved signatures. The example produces the output:
>
> ```
> Console
>
> 123
> 123
> ```

```
456
```

end example

15.3.10.3 Member names reserved for events

For an event `E` (§15.8) of delegate type `T`, the following signatures are reserved:

```
C#

void add_E(T handler);
void remove_E(T handler);
```

15.3.10.4 Member names reserved for indexers

For an indexer (§15.9) of type `T` with parameter-list `L`, the following signatures are reserved:

```
C#

T get_Item(L);
void set_Item(L, T value);
```

Both signatures are reserved, even if the indexer is read-only or write-only.

Furthermore the member name `Item` is reserved.

15.3.10.5 Member names reserved for finalizers

For a class containing a finalizer (§15.13), the following signature is reserved:

```
C#

void Finalize();
```

15.4 Constants

A ***constant*** is a class member that represents a constant value: a value that can be computed at compile-time. A *constant_declaration* introduces one or more constants of a given type.

```
```

ANTLR

```
constant_declaration
    : attributes? constant_modifier* 'const' type constant_declarators ';'
    ;

constant_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    ;
```

A *constant_declaration* may include a set of *attributes* (§22), a `new` modifier (§15.3.5), and any one of the permitted kinds of declared accessibility (§15.3.6). The attributes and modifiers apply to all of the members declared by the *constant_declaration*. Even though constants are considered static members, a *constant_declaration* neither requires nor allows a `static` modifier. It is an error for the same modifier to appear multiple times in a constant declaration.

The *type* of a *constant_declaration* specifies the type of the members introduced by the declaration. The type is followed by a list of *constant_declarator*s (§13.6.3), each of which introduces a new member. A *constant_declarator* consists of an *identifier* that names the member, followed by an "`=`" token, followed by a *constant_expression* (§12.23) that gives the value of the member.

The *type* specified in a constant declaration shall be `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, `double`, `decimal`, `bool`, `string`, an *enum_type*, or a *reference_type*. Each *constant_expression* shall yield a value of the target type or of a type that can be converted to the target type by an implicit conversion (§10.2).

The *type* of a constant shall be at least as accessible as the constant itself (§7.5.5).

The value of a constant is obtained in an expression using a *simple_name* (§12.8.4) or a *member_access* (§12.8.7).

A constant can itself participate in a *constant_expression*. Thus, a constant may be used in any construct that requires a *constant_expression*.

> *Note*: Examples of such constructs include `case` labels, `goto case` statements, `enum` member declarations, attributes, and other constant declarations. *end note*

> *Note*: As described in §12.23, a *constant_expression* is an expression that can be fully evaluated at compile-time. Since the only way to create a non-null value of a

> *reference_type* other than `string` is to apply the `new` operator, and since the `new` operator is not permitted in a *constant_expression*, the only possible value for constants of *reference_type*s other than `string` is `null`. *end note*

When a symbolic name for a constant value is desired, but when the type of that value is not permitted in a constant declaration, or when the value cannot be computed at compile-time by a *constant_expression*, a readonly field (§15.5.3) may be used instead.

> *Note*: The versioning semantics of `const` and `readonly` differ (§15.5.3.3). *end note*

A constant declaration that declares multiple constants is equivalent to multiple declarations of single constants with the same attributes, modifiers, and type.

> *Example*:
>
> C#
>
> ```
> class A
> {
>     public const double X = 1.0, Y = 2.0, Z = 3.0;
> }
> ```
>
> is equivalent to
>
> C#
>
> ```
> class A
> {
>     public const double X = 1.0;
>     public const double Y = 2.0;
>     public const double Z = 3.0;
> }
> ```
>
> *end example*

Constants are permitted to depend on other constants within the same program as long as the dependencies are not of a circular nature. The compiler automatically arranges to evaluate the constant declarations in the appropriate order.

> *Example*: In the following code
>
> C#

```
class A
{
    public const int X = B.Z + 1;
    public const int Y = 10;
}

class B
{
    public const int Z = A.Y + 1;
}
```

the compiler first evaluates `A.Y`, then evaluates `B.Z`, and finally evaluates `A.X`, producing the values `10`, `11`, and `12`.

end example

Constant declarations may depend on constants from other programs, but such dependencies are only possible in one direction.

> *Example*: Referring to the example above, if `A` and `B` were declared in separate programs, it would be possible for `A.X` to depend on `B.Z`, but `B.Z` could then not simultaneously depend on `A.Y`. *end example*

15.5 Fields

15.5.1 General

A ***field*** is a member that represents a variable associated with an object or class. A *field_declaration* introduces one or more fields of a given type.

ANTLR

```
field_declaration
    : attributes? field_modifier* type variable_declarators ';'
    ;

field_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'static'
    | 'readonly'
    | 'volatile'
```

```
    | unsafe_modifier   // unsafe code support
    ;

variable_declarators
    : variable_declarator (',' variable_declarator)*
    ;

variable_declarator
    : identifier ('=' variable_initializer)?
    ;
```

unsafe_modifier (§23.2) is only available in unsafe code (§23).

A *field_declaration* may include a set of *attributes* (§22), a `new` modifier (§15.3.5), a valid combination of the four access modifiers (§15.3.6), and a `static` modifier (§15.5.2). In addition, a *field_declaration* may include a `readonly` modifier (§15.5.3) or a `volatile` modifier (§15.5.4), but not both. The attributes and modifiers apply to all of the members declared by the *field_declaration*. It is an error for the same modifier to appear multiple times in a *field_declaration*.

The *type* of a *field_declaration* specifies the type of the members introduced by the declaration. The type is followed by a list of *variable_declarator*s, each of which introduces a new member. A *variable_declarator* consists of an *identifier* that names that member, optionally followed by an "`=`" token and a *variable_initializer* (§15.5.6) that gives the initial value of that member.

The *type* of a field shall be at least as accessible as the field itself (§7.5.5).

The value of a field is obtained in an expression using a *simple_name* (§12.8.4), a *member_access* (§12.8.7) or a base_access (§12.8.14). The value of a non-readonly field is modified using an *assignment* (§12.21). The value of a non-readonly field can be both obtained and modified using postfix increment and decrement operators (§12.8.15) and prefix increment and decrement operators (§12.9.6).

A field declaration that declares multiple fields is equivalent to multiple declarations of single fields with the same attributes, modifiers, and type.

> *Example*:
>
> C#
>
> ```
> class A
> {
>     public static int X = 1, Y, Z = 100;
> }
> ```

is equivalent to

```
C#

class A
{
    public static int X = 1;
    public static int Y;
    public static int Z = 100;
}
```

end example

15.5.2 Static and instance fields

When a field declaration includes a `static` modifier, the fields introduced by the declaration are ***static fields***. When no `static` modifier is present, the fields introduced by the declaration are ***instance fields***. Static fields and instance fields are two of the several kinds of variables (§9) supported by C#, and at times they are referred to as ***static variables*** and ***instance variables***, respectively.

As explained in §15.3.8, each instance of a class contains a complete set of the instance fields of the class, while there is only one set of static fields for each non-generic class or closed constructed type, regardless of the number of instances of the class or closed constructed type.

15.5.3 Readonly fields

15.5.3.1 General

When a *field_declaration* includes a `readonly` modifier, the fields introduced by the declaration are ***readonly fields***. Direct assignments to readonly fields can only occur as part of that declaration or in an instance constructor or static constructor in the same class. (A readonly field can be assigned to multiple times in these contexts.) Specifically, direct assignments to a readonly field are permitted only in the following contexts:

- In the *variable_declarator* that introduces the field (by including a *variable_initializer* in the declaration).
- For an instance field, in the instance constructors of the class that contains the field declaration; for a static field, in the static constructor of the class that contains the field declaration. These are also the only contexts in which it is valid to pass a readonly field as an `out` or `ref` parameter.

Attempting to assign to a readonly field or pass it as an `out` or `ref` parameter in any other context is a compile-time error.

15.5.3.2 Using static readonly fields for constants

A static readonly field is useful when a symbolic name for a constant value is desired, but when the type of the value is not permitted in a const declaration, or when the value cannot be computed at compile-time.

> *Example*: In the following code
>
> C#
>
> ```
> public class Color
> {
>     public static readonly Color Black = new Color(0, 0, 0);
>     public static readonly Color White = new Color(255, 255, 255);
>     public static readonly Color Red = new Color(255, 0, 0);
>     public static readonly Color Green = new Color(0, 255, 0);
>     public static readonly Color Blue = new Color(0, 0, 255);
>
>     private byte red, green, blue;
>
>     public Color(byte r, byte g, byte b)
>     {
>         red = r;
>         green = g;
>         blue = b;
>     }
> }
> ```
>
> the `Black`, `White`, `Red`, `Green`, and `Blue` members cannot be declared as const members because their values cannot be computed at compile-time. However, declaring them `static readonly` instead has much the same effect.
>
> *end example*

15.5.3.3 Versioning of constants and static readonly fields

Constants and readonly fields have different binary versioning semantics. When an expression references a constant, the value of the constant is obtained at compile-time, but when an expression references a readonly field, the value of the field is not obtained until run-time.

> *Example*: Consider an application that consists of two separate programs:

```C#
namespace Program1
{
    public class Utils
    {
        public static readonly int x = 1;
    }
}
```

and

```C#
namespace Program2
{
    class Test
    {
        static void Main()
        {
            Console.WriteLine(Program1.Utils.X);
        }
    }
}
```

The `Program1` and `Program2` namespaces denote two programs that are compiled separately. Because `Program1.Utils.X` is declared as a `static readonly` field, the value output by the `Console.WriteLine` statement is not known at compile-time, but rather is obtained at run-time. Thus, if the value of `X` is changed and `Program1` is recompiled, the `Console.WriteLine` statement will output the new value even if `Program2` isn't recompiled. However, had `X` been a constant, the value of `X` would have been obtained at the time `Program2` was compiled, and would remain unaffected by changes in `Program1` until `Program2` is recompiled.

end example

15.5.4 Volatile fields

When a *field_declaration* includes a `volatile` modifier, the fields introduced by that declaration are ***volatile fields***. For non-volatile fields, optimization techniques that reorder instructions can lead to unexpected and unpredictable results in multi-threaded programs that access fields without synchronization such as that provided by the *lock_statement* (§13.13). These optimizations can be performed by the compiler, by the

run-time system, or by hardware. For volatile fields, such reordering optimizations are restricted:

- A read of a volatile field is called a ***volatile read***. A volatile read has "acquire semantics"; that is, it is guaranteed to occur prior to any references to memory that occur after it in the instruction sequence.
- A write of a volatile field is called a ***volatile write***. A volatile write has "release semantics"; that is, it is guaranteed to happen after any memory references prior to the write instruction in the instruction sequence.

These restrictions ensure that all threads will observe volatile writes performed by any other thread in the order in which they were performed. A conforming implementation is not required to provide a single total ordering of volatile writes as seen from all threads of execution. The type of a volatile field shall be one of the following:

- A *reference_type*.
- A *type_parameter* that is known to be a reference type (§15.2.5).
- The type `byte`, `sbyte`, `short`, `ushort`, `int`, `uint`, `char`, `float`, `bool`, `System.IntPtr`, or `System.UIntPtr`.
- An *enum_type* having an *enum_base* type of `byte`, `sbyte`, `short`, `ushort`, `int`, or `uint`.

> *Example*: The example
>
> C#
>
> ```
> class Test
> {
>     public static int result;
>     public static volatile bool finished;
>
>     static void Thread2()
>     {
>         result = 143;
>         finished = true;
>     }
>
>     static void Main()
>     {
>         finished = false;
>
>         // Run Thread2() in a new thread
>         new Thread(new ThreadStart(Thread2)).Start();
>
>         // Wait for Thread2() to signal that it has a result
>         // by setting finished to true.
>         for (;;)
>         {
> ```

```
            if (finished)
            {
                Console.WriteLine($"result = {result}");
                return;
            }
        }
    }
}
```

produces the output:

```
Console

result = 143
```

In this example, the method `Main` starts a new thread that runs the method `Thread2`. This method stores a value into a non-volatile field called `result`, then stores `true` in the volatile field `finished`. The main thread waits for the field `finished` to be set to `true`, then reads the field `result`. Since `finished` has been declared `volatile`, the main thread shall read the value `143` from the field `result`. If the field `finished` had not been declared `volatile`, then it would be permissible for the store to `result` to be visible to the main thread *after* the store to `finished`, and hence for the main thread to read the value 0 from the field `result`. Declaring `finished` as a `volatile` field prevents any such inconsistency.

end example

15.5.5 Field initialization

The initial value of a field, whether it be a static field or an instance field, is the default value (§9.3) of the field's type. It is not possible to observe the value of a field before this default initialization has occurred, and a field is thus never "uninitialized".

Example: The example

```
C#

class Test
{
    static bool b;
    int i;

    static void Main()
    {
```

```
        Test t = new Test();
        Console.WriteLine($"b = {b}, i = {t.i}");
    }
}
```

produces the output

```
Console
```

```
b = False, i = 0
```

because `b` and `i` are both automatically initialized to default values.

end example

15.5.6 Variable initializers

15.5.6.1 General

Field declarations may include *variable_initializer*s. For static fields, variable initializers correspond to assignment statements that are executed during class initialization. For instance fields, variable initializers correspond to assignment statements that are executed when an instance of the class is created.

Example: The example

```
C#
```

```
class Test
{
    static double x = Math.Sqrt(2.0);
    int i = 100;
    string s = "Hello";

    static void Main()
    {
        Test a = new Test();
        Console.WriteLine($"x = {x}, i = {a.i}, s = {a.s}");
    }
}
```

produces the output

```
Console
```

```
x = 1.4142135623730951, i = 100, s = Hello
```

because an assignment to `x` occurs when static field initializers execute and assignments to `i` and `s` occur when the instance field initializers execute.

end example

The default value initialization described in §15.5.5 occurs for all fields, including fields that have variable initializers. Thus, when a class is initialized, all static fields in that class are first initialized to their default values, and then the static field initializers are executed in textual order. Likewise, when an instance of a class is created, all instance fields in that instance are first initialized to their default values, and then the instance field initializers are executed in textual order. When there are field declarations in multiple partial type declarations for the same type, the order of the parts is unspecified. However, within each part the field initializers are executed in order.

It is possible for static fields with variable initializers to be observed in their default value state.

Example: However, this is strongly discouraged as a matter of style. The example

C#

```
class Test
{
    static int a = b + 1;
    static int b = a + 1;

    static void Main()
    {
        Console.WriteLine($"a = {a}, b = {b}");
    }
}
```

exhibits this behavior. Despite the circular definitions of `a` and `b`, the program is valid. It results in the output

Console

```
a = 1, b = 2
```

because the static fields `a` and `b` are initialized to `0` (the default value for `int`) before their initializers are executed. When the initializer for `a` runs, the value of `b`

is zero, and so `a` is initialized to `1`. When the initializer for `b` runs, the value of a is already `1`, and so `b` is initialized to `2`.

end example

15.5.6.2 Static field initialization

The static field variable initializers of a class correspond to a sequence of assignments that are executed in the textual order in which they appear in the class declaration (§15.5.6.1). Within a partial class, the meaning of "textual order" is specified by §15.5.6.1. If a static constructor (§15.12) exists in the class, execution of the static field initializers occurs immediately prior to executing that static constructor. Otherwise, the static field initializers are executed at an implementation-dependent time prior to the first use of a static field of that class.

Example: The example

C#

```
class Test
{
    static void Main()
    {
        Console.WriteLine($"{B.Y} {A.X}");
    }

    public static int F(string s)
    {
        Console.WriteLine(s);
        return 1;
    }
}

class A
{
    public static int X = Test.F("Init A");
}

class B
{
    public static int Y = Test.F("Init B");
}
```

might produce either the output:

Console

```
Init A
Init B
1 1
```

or the output:

Console

```
Init B
Init A
1 1
```

because the execution of `X`'s initializer and `Y`'s initializer could occur in either order; they are only constrained to occur before the references to those fields. However, in the example:

C#

```
class Test
{
    static void Main()
    {
        Console.WriteLine($"{B.Y} {A.X}");
    }

    public static int F(string s)
    {
        Console.WriteLine(s);
        return 1;
    }
}

class A
{
    static A() {}
    public static int X = Test.F("Init A");
}

class B
{
    static B() {}
    public static int Y = Test.F("Init B");
}
```

the output shall be:

Console

```
Init B
Init A
1 1
```

because the rules for when static constructors execute (as defined in §15.12) provide that `B`'s static constructor (and hence `B`'s static field initializers) shall run before `A`'s static constructor and field initializers.

end example

15.5.6.3 Instance field initialization

The instance field variable initializers of a class correspond to a sequence of assignments that are executed immediately upon entry to any one of the instance constructors (§15.11.3) of that class. Within a partial class, the meaning of "textual order" is specified by §15.5.6.1. The variable initializers are executed in the textual order in which they appear in the class declaration (§15.5.6.1). The class instance creation and initialization process is described further in §15.11.

A variable initializer for an instance field cannot reference the instance being created. Thus, it is a compile-time error to reference `this` in a variable initializer, as it is a compile-time error for a variable initializer to reference any instance member through a *simple_name*.

Example: In the following code

C#

```
class A
{
    int x = 1;
    int y = x + 1;     // Error, reference to instance member of this
}
```

the variable initializer for `y` results in a compile-time error because it references a member of the instance being created.

end example

15.6 Methods

15.6.1 General

A ***method*** is a member that implements a computation or action that can be performed by an object or class. Methods are declared using *method_declaration*s:

ANTLR

```
method_declaration
    : attributes? method_modifiers return_type method_header method_body
    | attributes? ref_method_modifiers ref_kind ref_return_type
method_header
      ref_method_body
    ;

method_modifiers
    : method_modifier* 'partial'?
    ;

ref_kind
    : 'ref'
    | 'ref' 'readonly'
    ;

ref_method_modifiers
    : ref_method_modifier*
    ;

method_header
    : member_name '(' formal_parameter_list? ')'
    | member_name type_parameter_list '(' formal_parameter_list? ')'
      type_parameter_constraints_clause*
    ;

method_modifier
    : ref_method_modifier
    | 'async'
    ;

ref_method_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'static'
    | 'virtual'
    | 'sealed'
    | 'override'
    | 'abstract'
    | 'extern'
    | unsafe_modifier   // unsafe code support
    ;

return_type
    : ref_return_type
    | 'void'
```

```
    ;

ref_return_type
    : type
    ;

member_name
    : identifier
    | interface_type '.' identifier
    ;

method_body
    : block
    | '=>' null_conditional_invocation_expression ';'
    | '=>' expression ';'
    | ';'
    ;

ref_method_body
    : block
    | '=>' 'ref' variable_reference ';'
    | ';'
    ;
```

Grammar notes:

- *unsafe_modifier* (§23.2) is only available in unsafe code (§23).
- when recognising a *method_body* if both the *null_conditional_invocation_expression* and *expression* alternatives are applicable then the former shall be chosen.

> *Note*: The overlapping of, and priority between, alternatives here is solely for descriptive convenience; the grammar rules could be elaborated to remove the overlap. ANTLR, and other grammar systems, adopt the same convenience and so *method_body* has the specified semantics automatically. *end note*

A *method_declaration* may include a set of *attributes* (§22) and one of the permitted kinds of declared accessibility (§15.3.6), the `new` (§15.3.5), `static` (§15.6.3), `virtual` (§15.6.4), `override` (§15.6.5), `sealed` (§15.6.6), `abstract` (§15.6.7), `extern` (§15.6.8) and `async` (§15.15) modifiers.

A declaration has a valid combination of modifiers if all of the following are true:

- The declaration includes a valid combination of access modifiers (§15.3.6).
- The declaration does not include the same modifier multiple times.
- The declaration includes at most one of the following modifiers: `static`, `virtual`, and `override`.
- The declaration includes at most one of the following modifiers: `new` and `override`.

- If the declaration includes the `abstract` modifier, then the declaration does not include any of the following modifiers: `static`, `virtual`, `sealed`, or `extern`.
- If the declaration includes the `private` modifier, then the declaration does not include any of the following modifiers: `virtual`, `override`, or `abstract`.
- If the declaration includes the `sealed` modifier, then the declaration also includes the `override` modifier.
- If the declaration includes the `partial` modifier, then it does not include any of the following modifiers: `new`, `public`, `protected`, `internal`, `private`, `virtual`, `sealed`, `override`, `abstract`, or `extern`.

Methods are classified according to what, if anything, they return:

- If `ref` is present, the method is ***returns-by-ref*** and returns a *variable reference*, that is optionally read-only;
- Otherwise, if *return_type* is `void`, the method is ***returns-no-value*** and does not return a value;
- Otherwise, the method is ***returns-by-value*** and returns a value.

The *return_type* of a returns-by-value or returns-no-value method declaration specifies the type of the result, if any, returned by the method. Only a returns-no-value method may include the `partial` modifier (§15.6.9). If the declaration includes the `async` modifier then *return_type* shall be `void` or the method returns-by-value and the return type is a *task type* (§15.15.1).

The *ref_return_type* of a returns-by-ref method declaration specifies the type of the variable referenced by the *variable_reference* returned by the method.

A generic method is a method whose declaration includes a *type_parameter_list*. This specifies the type parameters for the method. The optional *type_parameter_constraints_clause*s specify the constraints for the type parameters.

A generic *method_declaration* for an explicit interface member implementation shall not have any *type_parameter_constraints_clause*s; the declaration inherits any constraints from the constraints on the interface method.

Similarly, a method declaration with the `override` modifier shall not have any *type_parameter_constraints_clause*s and the constraints of the method's type parameters are inherited from the virtual method being overridden.

The *member_name* specifies the name of the method. Unless the method is an explicit interface member implementation (§18.6.2), the *member_name* is simply an *identifier*.

For an explicit interface member implementation, the *member_name* consists of an *interface_type* followed by a "`.`" and an *identifier*. In this case, the declaration shall include no modifiers other than (possibly) `extern` or `async`.

The optional *formal_parameter_list* specifies the parameters of the method (§15.6.2).

The *return_type* or *ref_return_type*, and each of the types referenced in the *formal_parameter_list* of a method, shall be at least as accessible as the method itself (§7.5.5).

The *method_body* of a returns-by-value or returns-no-value method is either a semicolon, a block body or an expression body. A block body consists of a *block*, which specifies the statements to execute when the method is invoked. An expression body consists of `=>`, followed by a *null_conditional_invocation_expression* or *expression*, and a semicolon, and denotes a single expression to perform when the method is invoked.

For abstract and extern methods, the *method_body* consists simply of a semicolon. For partial methods the *method_body* may consist of either a semicolon, a block body or an expression body. For all other methods, the *method_body* is either a block body or an expression body.

If the *method_body* consists of a semicolon, the declaration shall not include the `async` modifier.

The *ref_method_body* of a returns-by-ref method is either a semicolon, a block body or an expression body. A block body consists of a *block*, which specifies the statements to execute when the method is invoked. An expression body consists of `=>`, followed by `ref`, a *variable_reference*, and a semicolon, and denotes a single *variable_reference* to evaluate when the method is invoked.

For abstract and extern methods, the *ref_method_body* consists simply of a semicolon; for all other methods, the *ref_method_body* is either a block body or an expression body.

The name, the number of type parameters, and the formal parameter list of a method define the signature (§7.6) of the method. Specifically, the signature of a method consists of its name, the number of its type parameters, and the number, *parameter_mode_modifier*s (§15.6.2.1), and types of its formal parameters. The return type is not part of a method's signature, nor are the names of the formal parameters, the names of the type parameters, or the constraints. When a formal parameter type references a type parameter of the method, the ordinal position of the type parameter (not the name of the type parameter) is used for type equivalence.

The name of a method shall differ from the names of all other non-methods declared in the same class. In addition, the signature of a method shall differ from the signatures of all other methods declared in the same class, and two methods declared in the same class shall not have signatures that differ solely by `in`, `out`, and `ref`.

The method's *type_parameter*s are in scope throughout the *method_declaration*, and can be used to form types throughout that scope in *return_type* or *ref_return_type*, *method_body* or *ref_method_body*, and *type_parameter_constraints_clause*s but not in *attributes*.

All formal parameters and type parameters shall have different names.

15.6.2 Method parameters

15.6.2.1 General

The parameters of a method, if any, are declared by the method's *formal_parameter_list*.

ANTLR

```
formal_parameter_list
    : fixed_parameters
    | fixed_parameters ',' parameter_array
    | parameter_array
    ;

fixed_parameters
    : fixed_parameter (',' fixed_parameter)*
    ;

fixed_parameter
    : attributes? parameter_modifier? type identifier default_argument?
    ;

default_argument
    : '=' expression
    ;

parameter_modifier
    : parameter_mode_modifier
    | 'this'
    ;

parameter_mode_modifier
    : 'ref'
    | 'out'
    | 'in'
    ;
```

```
parameter_array
    : attributes? 'params' array_type identifier
    ;
```

The formal parameter list consists of one or more comma-separated parameters of which only the last may be a *parameter_array*.

A *fixed_parameter* consists of an optional set of *attributes* (§22); an optional `in`, `out`, `ref`, or `this` modifier; a *type*; an *identifier*; and an optional *default_argument*. Each *fixed_parameter* declares a parameter of the given type with the given name. The `this` modifier designates the method as an extension method and is only allowed on the first parameter of a static method in a non-generic, non-nested static class. If the parameter is a `struct` type or a type parameter constrained to a `struct`, the `this` modifier may be combined with either the `ref` or `in` modifier, but not the `out` modifier. Extension methods are further described in §15.6.10. A *fixed_parameter* with a *default_argument* is known as an ***optional parameter***, whereas a *fixed_parameter* without a *default_argument* is a ***required parameter***. A required parameter shall not appear after an optional parameter in a *formal_parameter_list*.

A parameter with a `ref`, `out` or `this` modifier cannot have a *default_argument*. A parameter with an `in` modifier may have a *default_argument*. The *expression* in a *default_argument* shall be one of the following:

- a *constant_expression*
- an expression of the form `new S()` where `S` is a value type
- an expression of the form `default(S)` where `S` is a value type

The *expression* shall be implicitly convertible by an identity or nullable conversion to the type of the parameter.

If optional parameters occur in an implementing partial method declaration (§15.6.9), an explicit interface member implementation (§18.6.2), a single-parameter indexer declaration (§15.9), or in an operator declaration (§15.10.1) the compiler should give a warning, since these members can never be invoked in a way that permits arguments to be omitted.

A *parameter_array* consists of an optional set of *attributes* (§22), a `params` modifier, an *array_type*, and an *identifier*. A parameter array declares a single parameter of the given array type with the given name. The *array_type* of a parameter array shall be a single-dimensional array type (§17.2). In a method invocation, a parameter array permits either a single argument of the given array type to be specified, or it permits zero or more

arguments of the array element type to be specified. Parameter arrays are described further in §15.6.2.4.

A *parameter_array* may occur after an optional parameter, but cannot have a default value – the omission of arguments for a *parameter_array* would instead result in the creation of an empty array.

> *Example*: The following illustrates different kinds of parameters:
>
> ```
> C#
> ```
>
> ```
> void M<T>(
>     ref int i,
>     decimal d,
>     bool b = false,
>     bool? n = false,
>     string s = "Hello",
>     object o = null,
>     T t = default(T),
>     params int[] a
> ) { }
> ```
>
> In the *formal_parameter_list* for `M`, `i` is a required `ref` parameter, `d` is a required value parameter, `b`, `s`, `o` and `t` are optional value parameters and `a` is a parameter array.
>
> *end example*

A method declaration creates a separate declaration space (§7.3) for parameters and type parameters. Names are introduced into this declaration space by the type parameter list and the formal parameter list of the method. The body of the method, if any, is considered to be nested within this declaration space. It is an error for two members of a method declaration space to have the same name. It is an error for the method declaration space and the local variable declaration space of a nested declaration space to contain elements with the same name.

A method invocation (§12.8.9.2) creates a copy, specific to that invocation, of the formal parameters and local variables of the method, and the argument list of the invocation assigns values or variable references to the newly created formal parameters. Within the *block* of a method, formal parameters can be referenced by their identifiers in *simple_name* expressions (§12.8.4).

The following kinds of formal parameters exist:

- Value parameters, which are declared without any modifiers.

- Input parameters, which are declared with the `in` modifier.
- Output parameters, which are declared with the `out` modifier.
- Reference parameters, which are declared with the `ref` modifier.
- Parameter arrays, which are declared with the `params` modifier.

> *Note*: As described in §7.6, the `in`, `out`, and `ref` modifiers are part of a method's signature, but the `params` modifier is not. *end note*

15.6.2.2 Value parameters

A parameter declared with no modifiers is a value parameter. A value parameter is a local variable that gets its initial value from the corresponding argument supplied in the method invocation.

When a formal parameter is a value parameter, the corresponding argument in a method invocation shall be an expression that is implicitly convertible (§10.2) to the formal parameter type.

A method is permitted to assign new values to a value parameter. Such assignments only affect the local storage location represented by the value parameter—they have no effect on the actual argument given in the method invocation.

15.6.2.3 By-reference parameters

15.6.2.3.1 General

A parameter declared with one of the `in`, `ref` or `out` modifiers is a ***by-reference*** parameter. A by-reference parameter is a local reference variable (§9.7); the initial referent is obtained from the corresponding argument supplied in the method invocation.

> *Note*: The referent of a by-reference parameter can be changed using the ref assignment (`= ref`) operator.

When a formal parameter is a by-reference parameter, the corresponding argument in a method invocation shall consist of the corresponding keyword, `in`, `ref`, or `out`, followed by a *variable_reference* (§9.5) of the same type as the formal parameter. However, when the formal parameter is an `in` parameter, the argument may be an *expression* for which an implicit conversion (§10.2) exists from that argument expression to the type of the corresponding parameter.

By-reference parameters are not allowed on functions declared as an iterator (§15.14) or async function (§15.15).

In a method that takes multiple by-reference parameters, it is possible for multiple names to represent the same storage location.

15.6.2.3.2 Input parameters

A parameter declared with an `in` modifier is a input parameter. The argument corresponding to an `in` parameter is either a variable existing at the point of the method invocation, or one created by the implementation (§12.6.2.3) in the method invocation. A variable shall be definitely assigned before it can be passed as an argument for an input parameter. Within a method, an input parameter is always considered definitely assigned.

It is a compile-time error to modify the value of an input parameter.

> *Note*: The primary purpose of input parameters is for efficiency. When the type of a method parameter is a large struct (in terms of memory requirements), it is useful to be able to avoid copying the whole value of the argument when calling the method. Input parameters allow methods to refer to existing values in memory, while providing protection against unwanted changes to those values. *end note*

15.6.2.3.3 Reference parameters

A parameter declared with a `ref` modifier is a reference parameter. A variable shall be definitely assigned before it can be passed as an argument for a reference parameter. Within a method, a reference parameter is always considered definitely assigned.

> *Example*: The example
>
> C#
>
> ```
> class Test
> {
>     static void Swap(ref int x, ref int y)
>     {
>         int temp = x;
>         x = y;
>         y = temp;
>     }
>
>     static void Main()
>     {
>         int i = 1, j = 2;
> ```

```
        Swap(ref i, ref j);
        Console.WriteLine($"i = {i}, j = {j}");
    }
}
```

produces the output

```
Console
```

```
i = 2, j = 1
```

For the invocation of `Swap` in `Main`, `x` represents `i` and `y` represents `j`. Thus, the invocation has the effect of swapping the values of `i` and `j`.

end example

Example: In the following code

```
C#
```

```
class A
{
    string s;
    void F(ref string a, ref string b)
    {
        s = "One";
        a = "Two";
        b = "Three";
    }

    void G()
    {
        F(ref s, ref s);
    }
}
```

the invocation of `F` in `G` passes a reference to `s` for both `a` and `b`. Thus, for that invocation, the names `s`, `a`, and `b` all refer to the same storage location, and the three assignments all modify the instance field `s`.

end example

15.6.2.3.4 Output parameters

A parameter declared with an `out` modifier is an output parameter. Within a method, just like a local variable, an output parameter is initially considered unassigned and shall

be definitely assigned before its value is used. Every output parameter of a method shall be definitely assigned before the method returns.

A method declared as a partial method (§15.6.9) shall not have output parameters.

> *Note: Output parameters are typically used in methods that produce multiple return values. *end note*

> *Example*:

C#

```
class Test
{
    static void SplitPath(string path, out string dir, out string name)
    {
        int i = path.Length;
        while (i > 0)
        {
            char ch = path[i - 1];
            if (ch == '\\' || ch == '/' || ch == ':')
            {
                break;
            }
            i--;
        }
        dir = path.Substring(0, i);
        name = path.Substring(i);
    }

    static void Main()
    {
        string dir, name;
        SplitPath(@"c:\Windows\System\hello.txt", out dir, out name);
        Console.WriteLine(dir);
        Console.WriteLine(name);
    }
}
```

The example produces the output:

Console

```
c:\Windows\System\
hello.txt
```

Note that the `dir` and `name` variables can be unassigned before they are passed to `SplitPath`, and that they are considered definitely assigned following the call.

> *end example*

15.6.2.4 Parameter arrays

A parameter declared with a `params` modifier is a parameter array. If a formal parameter list includes a parameter array, it shall be the last parameter in the list and it shall be of a single-dimensional array type.

> *Example*: The types `string[]` and `string[][]` can be used as the type of a parameter array, but the type `string[,]` can not. *end example*

> *Note*: It is not possible to combine the `params` modifier with the modifiers `in`, `out`, or `ref`. *end note*

A parameter array permits arguments to be specified in one of two ways in a method invocation:

- The argument given for a parameter array can be a single expression that is implicitly convertible (§10.2) to the parameter array type. In this case, the parameter array acts precisely like a value parameter.
- Alternatively, the invocation can specify zero or more arguments for the parameter array, where each argument is an expression that is implicitly convertible (§10.2) to the element type of the parameter array. In this case, the invocation creates an instance of the parameter array type with a length corresponding to the number of arguments, initializes the elements of the array instance with the given argument values, and uses the newly created array instance as the actual argument.

Except for allowing a variable number of arguments in an invocation, a parameter array is precisely equivalent to a value parameter (§15.6.2.2) of the same type.

> *Example*: The example

```C#
class Test
{
    static void F(params int[] args)
    {
        Console.Write($"Array contains {args.Length} elements:");
        foreach (int i in args)
        {
            Console.Write($" {i}");
        }
        Console.WriteLine();
```

```
    }

    static void Main()
    {
        int[] arr = {1, 2, 3};
        F(arr);
        F(10, 20, 30, 40);
        F();
    }
}
```

produces the output

```
Console

Array contains 3 elements: 1 2 3
Array contains 4 elements: 10 20 30 40
Array contains 0 elements:
```

The first invocation of `F` simply passes the array `arr` as a value parameter. The second invocation of F automatically creates a four-element `int[]` with the given element values and passes that array instance as a value parameter. Likewise, the third invocation of `F` creates a zero-element `int[]` and passes that instance as a value parameter. The second and third invocations are precisely equivalent to writing:

```
C#

F(new int[] {10, 20, 30, 40});
F(new int[] {});
```

end example

When performing overload resolution, a method with a parameter array might be applicable, either in its normal form or in its expanded form (§12.6.4.2). The expanded form of a method is available only if the normal form of the method is not applicable and only if an applicable method with the same signature as the expanded form is not already declared in the same type.

Example: The example

```
C#

class Test
{
```

```
    static void F(params object[] a) =>
        Console.WriteLine("F(object[])");

    static void F() =>
        Console.WriteLine("F()");

    static void F(object a0, object a1) =>
        Console.WriteLine("F(object,object)");

    static void Main()
    {
        F();
        F(1);
        F(1, 2);
        F(1, 2, 3);
        F(1, 2, 3, 4);
    }
}
```

produces the output

Console

```
F()
F(object[])
F(object,object)
F(object[])
F(object[])
```

In the example, two of the possible expanded forms of the method with a parameter array are already included in the class as regular methods. These expanded forms are therefore not considered when performing overload resolution, and the first and third method invocations thus select the regular methods. When a class declares a method with a parameter array, it is not uncommon to also include some of the expanded forms as regular methods. By doing so, it is possible to avoid the allocation of an array instance that occurs when an expanded form of a method with a parameter array is invoked.

end example

An array is a reference type, so the value passed for a parameter array can be `null`.

Example: The example:

C#

```
class Test
{
```

```
    static void F(params string[] array) =>
        Console.WriteLine(array == null);

    static void Main()
    {
        F(null);
        F((string) null);
    }
}
```

produces the output:

```
Console

True
False
```

The second invocation produces `False` as it is equivalent to `F(new string[] { null })` and passes an array containing a single null reference.

end example

When the type of a parameter array is `object[]`, a potential ambiguity arises between the normal form of the method and the expanded form for a single `object` parameter. The reason for the ambiguity is that an `object[]` is itself implicitly convertible to type `object`. The ambiguity presents no problem, however, since it can be resolved by inserting a cast if needed.

Example: The example

```
C#

class Test
{
    static void F(params object[] args)
    {
        foreach (object o in args)
        {
            Console.Write(o.GetType().FullName);
            Console.Write(" ");
        }
        Console.WriteLine();
    }

    static void Main()
    {
        object[] a = {1, "Hello", 123.456};
        object o = a;
```

```
            F(a);
            F((object)a);
            F(o);
            F((object[])o);
        }
    }
```

produces the output

```
Console

System.Int32 System.String System.Double
System.Object[]
System.Object[]
System.Int32 System.String System.Double
```

In the first and last invocations of `F`, the normal form of `F` is applicable because an implicit conversion exists from the argument type to the parameter type (both are of type `object[]`). Thus, overload resolution selects the normal form of `F`, and the argument is passed as a regular value parameter. In the second and third invocations, the normal form of `F` is not applicable because no implicit conversion exists from the argument type to the parameter type (type `object` cannot be implicitly converted to type `object[]`). However, the expanded form of `F` is applicable, so it is selected by overload resolution. As a result, a one-element `object[]` is created by the invocation, and the single element of the array is initialized with the given argument value (which itself is a reference to an `object[]`).

end example

15.6.3 Static and instance methods

When a method declaration includes a `static` modifier, that method is said to be a static method. When no `static` modifier is present, the method is said to be an instance method.

A static method does not operate on a specific instance, and it is a compile-time error to refer to `this` in a static method.

An instance method operates on a given instance of a class, and that instance can be accessed as `this` (§12.8.13).

The differences between static and instance members are discussed further in §15.3.8.

15.6.4 Virtual methods

When an instance method declaration includes a virtual modifier, that method is said to be a ***virtual method***. When no virtual modifier is present, the method is said to be a ***non-virtual method***.

The implementation of a non-virtual method is invariant: The implementation is the same whether the method is invoked on an instance of the class in which it is declared or an instance of a derived class. In contrast, the implementation of a virtual method can be superseded by derived classes. The process of superseding the implementation of an inherited virtual method is known as ***overriding*** that method (§15.6.5).

In a virtual method invocation, the ***run-time type*** of the instance for which that invocation takes place determines the actual method implementation to invoke. In a non-virtual method invocation, the ***compile-time type*** of the instance is the determining factor. In precise terms, when a method named `N` is invoked with an argument list `A` on an instance with a compile-time type `C` and a run-time type `R` (where `R` is either `C` or a class derived from `C`), the invocation is processed as follows:

- At binding-time, overload resolution is applied to `C`, `N`, and `A`, to select a specific method `M` from the set of methods declared in and inherited by `C`. This is described in §12.8.9.2.
- Then at run-time:
  - If `M` is a non-virtual method, `M` is invoked.
  - Otherwise, `M` is a virtual method, and the most derived implementation of `M` with respect to `R` is invoked.

For every virtual method declared in or inherited by a class, there exists a ***most derived implementation*** of the method with respect to that class. The most derived implementation of a virtual method `M` with respect to a class `R` is determined as follows:

- If `R` contains the introducing virtual declaration of `M`, then this is the most derived implementation of `M` with respect to `R`.
- Otherwise, if `R` contains an override of `M`, then this is the most derived implementation of `M` with respect to `R`.
- Otherwise, the most derived implementation of `M` with respect to `R` is the same as the most derived implementation of `M` with respect to the direct base class of `R`.

> *Example*: The following example illustrates the differences between virtual and non-virtual methods:
>
> ```
> C#
> ```

```
class A
{
    public void F() => Console.WriteLine("A.F");
    public virtual void G() => Console.WriteLine("A.G");
}

class B : A
{
    public new void F() => Console.WriteLine("B.F");
    public override void G() => Console.WriteLine("B.G");
}

class Test
{
    static void Main()
    {
        B b = new B();
        A a = b;
        a.F();
        b.F();
        a.G();
        b.G();
    }
}
```

In the example, `A` introduces a non-virtual method `F` and a virtual method `G`. The class `B` introduces a *new* non-virtual method `F`, thus *hiding* the inherited `F`, and also *overrides* the inherited method `G`. The example produces the output:

Console

```
A.F
B.F
B.G
B.G
```

Notice that the statement `a.G()` invokes `B.G`, not `A.G`. This is because the run-time type of the instance (which is `B`), not the compile-time type of the instance (which is `A`), determines the actual method implementation to invoke.

end example

Because methods are allowed to hide inherited methods, it is possible for a class to contain several virtual methods with the same signature. This does not present an ambiguity problem, since all but the most derived method are hidden.

Example: In the following code

C#

```
class A
{
    public virtual void F() => Console.WriteLine("A.F");
}

class B : A
{
    public override void F() => Console.WriteLine("B.F");
}

class C : B
{
    public new virtual void F() => Console.WriteLine("C.F");
}

class D : C
{
    public override void F() => Console.WriteLine("D.F");
}

class Test
{
    static void Main()
    {
        D d = new D();
        A a = d;
        B b = d;
        C c = d;
        a.F();
        b.F();
        c.F();
        d.F();
    }
}
```

the `C` and `D` classes contain two virtual methods with the same signature: The one introduced by `A` and the one introduced by `C`. The method introduced by `C` hides the method inherited from `A`. Thus, the override declaration in `D` overrides the method introduced by `C`, and it is not possible for `D` to override the method introduced by `A`. The example produces the output:

Console

```
B.F
B.F
D.F
D.F
```

> Note that it is possible to invoke the hidden virtual method by accessing an instance of `D` through a less derived type in which the method is not hidden.
>
> *end example*

15.6.5 Override methods

When an instance method declaration includes an `override` modifier, the method is said to be an ***override method***. An override method overrides an inherited virtual method with the same signature. Whereas a virtual method declaration *introduces* a new method, an override method declaration *specializes* an existing inherited virtual method by providing a new implementation of that method.

The method overridden by an override declaration is known as the ***overridden base method*** For an override method `M` declared in a class `C`, the overridden base method is determined by examining each base class of `C`, starting with the direct base class of `C` and continuing with each successive direct base class, until in a given base class type at least one accessible method is located which has the same signature as `M` after substitution of type arguments. For the purposes of locating the overridden base method, a method is considered accessible if it is `public`, if it is `protected`, if it is `protected internal`, or if it is either `internal` or `private protected` and declared in the same program as `C`.

A compile-time error occurs unless all of the following are true for an override declaration:

- An overridden base method can be located as described above.
- There is exactly one such overridden base method. This restriction has effect only if the base class type is a constructed type where the substitution of type arguments makes the signature of two methods the same.
- The overridden base method is a virtual, abstract, or override method. In other words, the overridden base method cannot be static or non-virtual.
- The overridden base method is not a sealed method.
- There is an identity conversion between the return type of the overridden base method and the override method.
- The override declaration and the overridden base method have the same declared accessibility. In other words, an override declaration cannot change the accessibility of the virtual method. However, if the overridden base method is protected internal and it is declared in a different assembly than the assembly containing the override declaration then the override declaration's declared accessibility shall be protected.

- The override declaration does not specify any *type_parameter_constraints_clause*s. Instead, the constraints are inherited from the overridden base method. Constraints that are type parameters in the overridden method may be replaced by type arguments in the inherited constraint. This can lead to constraints that are not valid when explicitly specified, such as value types or sealed types.

> *Example*: The following demonstrates how the overriding rules work for generic classes:
>
> C#
>
> ```
> abstract class C<T>
> {
>     public virtual T F() {...}
>     public virtual C<T> G() {...}
>     public virtual void H(C<T> x) {...}
> }
>
> class D : C<string>
> {
>     public override string F() {...}           // Ok
>     public override C<string> G() {...}        // Ok
>     public override void H(C<T> x) {...}       // Error, should be
> C<string>
> }
>
> class E<T,U> : C<U>
> {
>     public override U F() {...}                // Ok
>     public override C<U> G() {...}             // Ok
>     public override void H(C<T> x) {...}       // Error, should be C<U>
> }
> ```
>
> *end example*

An override declaration can access the overridden base method using a *base_access* (§12.8.14).

> *Example*: In the following code
>
> C#
>
> ```
> class A
> {
>     int x;
>
>     public virtual void PrintFields() => Console.WriteLine($"x = {x}");
> }
> ```

```
class B : A
{
    int y;

    public override void PrintFields()
    {
        base.PrintFields();
        Console.WriteLine($"y = {y}");
    }
}
```

the `base.PrintFields()` invocation in `B` invokes the PrintFields method declared in `A`. A *base_access* disables the virtual invocation mechanism and simply treats the base method as a non-`virtual` method. Had the invocation in `B` been written `((A)this).PrintFields()`, it would recursively invoke the `PrintFields` method declared in `B`, not the one declared in `A`, since `PrintFields` is virtual and the run-time type of `((A)this)` is `B`.

end example

Only by including an `override` modifier can a method override another method. In all other cases, a method with the same signature as an inherited method simply hides the inherited method.

Example: In the following code

C#

```
class A
{
    public virtual void F() {}
}

class B : A
{
    public virtual void F() {} // Warning, hiding inherited F()
}
```

the `F` method in `B` does not include an `override` modifier and therefore does not override the `F` method in `A`. Rather, the `F` method in `B` hides the method in `A`, and a warning is reported because the declaration does not include a new modifier.

end example

Example: In the following code

C#

```
class A
{
    public virtual void F() {}
}

class B : A
{
    private new void F() {} // Hides A.F within body of B
}

class C : B
{
    public override void F() {} // Ok, overrides A.F
}
```

the `F` method in `B` hides the virtual `F` method inherited from `A`. Since the new `F` in `B` has private access, its scope only includes the class body of `B` and does not extend to `C`. Therefore, the declaration of `F` in `C` is permitted to override the `F` inherited from `A`.

end example

15.6.6 Sealed methods

When an instance method declaration includes a `sealed` modifier, that method is said to be a ***sealed method***. A sealed method overrides an inherited virtual method with the same signature. A sealed method shall also be marked with the `override` modifier. Use of the `sealed` modifier prevents a derived class from further overriding the method.

Example: The example

C#

```
class A
{
    public virtual void F() => Console.WriteLine("A.F");
    public virtual void G() => Console.WriteLine("A.G");
}

class B : A
{
    public sealed override void F() => Console.WriteLine("B.F");
    public override void G()        => Console.WriteLine("B.G");
}

class C : B
```

```
{
    public override void G() => Console.WriteLine("C.G");
}
```

the class `B` provides two override methods: an `F` method that has the `sealed` modifier and a `G` method that does not. `B`'s use of the `sealed` modifier prevents `C` from further overriding `F`.

end example

15.6.7 Abstract methods

When an instance method declaration includes an `abstract` modifier, that method is said to be an ***abstract method***. Although an abstract method is implicitly also a virtual method, it cannot have the modifier `virtual`.

An abstract method declaration introduces a new virtual method but does not provide an implementation of that method. Instead, non-abstract derived classes are required to provide their own implementation by overriding that method. Because an abstract method provides no actual implementation, the method body of an abstract method simply consists of a semicolon.

Abstract method declarations are only permitted in abstract classes (§15.2.2.2).

Example: In the following code

C#

```
public abstract class Shape
{
    public abstract void Paint(Graphics g, Rectangle r);
}

public class Ellipse : Shape
{
    public override void Paint(Graphics g, Rectangle r) =>
g.DrawEllipse(r);
}

public class Box : Shape
{
    public override void Paint(Graphics g, Rectangle r) =>
g.DrawRect(r);
}
```

> the `Shape` class defines the abstract notion of a geometrical shape object that can paint itself. The `Paint` method is abstract because there is no meaningful default implementation. The `Ellipse` and `Box` classes are concrete `Shape` implementations. Because these classes are non-abstract, they are required to override the `Paint` method and provide an actual implementation.
>
> *end example*

It is a compile-time error for a *base_access* (§12.8.14) to reference an abstract method.

> *Example*: In the following code
>
> C#
>
> ```
> abstract class A
> {
>     public abstract void F();
> }
>
> class B : A
> {
>     // Error, base.F is abstract
>     public override void F() => base.F();
> }
> ```
>
> a compile-time error is reported for the `base.F()` invocation because it references an abstract method.
>
> *end example*

An abstract method declaration is permitted to override a virtual method. This allows an abstract class to force re-implementation of the method in derived classes, and makes the original implementation of the method unavailable.

> *Example*: In the following code
>
> C#
>
> ```
> class A
> {
>     public virtual void F() => Console.WriteLine("A.F");
> }
>
> abstract class B: A
> {
>     public abstract override void F();
> ```

```
}

class C : B
{
    public override void F() => Console.WriteLine("C.F");
}
```

class `A` declares a virtual method, class `B` overrides this method with an abstract method, and class `C` overrides the abstract method to provide its own implementation.

end example

15.6.8 External methods

When a method declaration includes an `extern` modifier, the method is said to be an ***external method***. External methods are implemented externally, typically using a language other than C#. Because an external method declaration provides no actual implementation, the method body of an external method simply consists of a semicolon. An external method shall not be generic.

The mechanism by which linkage to an external method is achieved, is implementation-defined.

Example: The following example demonstrates the use of the `extern` modifier and the `DllImport` attribute:

C#

```
class Path
{
    [DllImport("kernel32", SetLastError=true)]
    static extern bool CreateDirectory(string name, SecurityAttribute
sa);

    [DllImport("kernel32", SetLastError=true)]
    static extern bool RemoveDirectory(string name);

    [DllImport("kernel32", SetLastError=true)]
    static extern int GetCurrentDirectory(int bufSize, StringBuilder
buf);

    [DllImport("kernel32", SetLastError=true)]
    static extern bool SetCurrentDirectory(string name);
}
```

> *end example*

15.6.9 Partial methods

When a method declaration includes a `partial` modifier, that method is said to be a ***partial method***. Partial methods may only be declared as members of partial types (§15.2.7), and are subject to a number of restrictions.

Partial methods may be defined in one part of a type declaration and implemented in another. The implementation is optional; if no part implements the partial method, the partial method declaration and all calls to it are removed from the type declaration resulting from the combination of the parts.

Partial methods shall not define access modifiers; they are implicitly private. Their return type shall be `void`, and their parameters shall not have the `out` modifier. The identifier partial is recognized as a contextual keyword (§6.4.4) in a method declaration only if it appears immediately before the `void` keyword. A partial method cannot explicitly implement interface methods.

There are two kinds of partial method declarations: If the body of the method declaration is a semicolon, the declaration is said to be a ***defining partial method declaration***. If the body is other than a semicolon, the declaration is said to be an ***implementing partial method declaration***. Across the parts of a type declaration, there may be only one defining partial method declaration with a given signature, and there may be only one implementing partial method declaration with a given signature. If an implementing partial method declaration is given, a corresponding defining partial method declaration shall exist, and the declarations shall match as specified in the following:

- The declarations shall have the same modifiers (although not necessarily in the same order), method name, number of type parameters and number of parameters.
- Corresponding parameters in the declarations shall have the same modifiers (although not necessarily in the same order) and the same types, or identity convertible types (modulo differences in type parameter names).
- Corresponding type parameters in the declarations shall have the same constraints (modulo differences in type parameter names).

An implementing partial method declaration can appear in the same part as the corresponding defining partial method declaration.

Only a defining partial method participates in overload resolution. Thus, whether or not an implementing declaration is given, invocation expressions may resolve to invocations of the partial method. Because a partial method always returns `void`, such invocation expressions will always be expression statements. Furthermore, because a partial method is implicitly `private`, such statements will always occur within one of the parts of the type declaration within which the partial method is declared.

> *Note*: The definition of matching defining and implementing partial method declarations does not require parameter names to match. This can produce *surprising*, albeit *well defined*, behaviour when named arguments (§12.6.2.1) are used. For example, given the defining partial method declaration for `M` in one file, and the implementing partial method declaration in another file:
>
> ```
> C#
> ```
>
> ```
> // File P1.cs:
> partial class P
> {
>     static partial void M(int x);
> }
>
> // File P2.cs:
> partial class P
> {
>     static void Caller() => M(y: 0);
>     static partial void M(int y) {}
> }
> ```
>
> is **invalid** as the invocation uses the argument name from the implementing and not the defining partial method declaration.
>
> *end note*

If no part of a partial type declaration contains an implementing declaration for a given partial method, any expression statement invoking it is simply removed from the combined type declaration. Thus the invocation expression, including any subexpressions, has no effect at run-time. The partial method itself is also removed and will not be a member of the combined type declaration.

If an implementing declaration exists for a given partial method, the invocations of the partial methods are retained. The partial method gives rise to a method declaration similar to the implementing partial method declaration except for the following:

- The `partial` modifier is not included.

- The attributes in the resulting method declaration are the combined attributes of the defining and the implementing partial method declaration in unspecified order. Duplicates are not removed.
- The attributes on the parameters of the resulting method declaration are the combined attributes of the corresponding parameters of the defining and the implementing partial method declaration in unspecified order. Duplicates are not removed.

If a defining declaration but not an implementing declaration is given for a partial method `M`, the following restrictions apply:

- It is a compile-time error to create a delegate from `M` (§12.8.16.6).
- It is a compile-time error to refer to `M` inside an anonymous function that is converted to an expression tree type (§8.6).
- Expressions occurring as part of an invocation of `M` do not affect the definite assignment state (§9.4), which can potentially lead to compile-time errors.
- `M` cannot be the entry point for an application (§7.1).

Partial methods are useful for allowing one part of a type declaration to customize the behavior of another part, e.g., one that is generated by a tool. Consider the following partial class declaration:

C#

```
partial class Customer
{
    string name;

    public string Name
    {
        get => name;
        set
        {
            OnNameChanging(value);
            name = value;
            OnNameChanged();
        }
    }

    partial void OnNameChanging(string newName);
    partial void OnNameChanged();
}
```

If this class is compiled without any other parts, the defining partial method declarations and their invocations will be removed, and the resulting combined class declaration will be equivalent to the following:

C#

```
class Customer
{
    string name;

    public string Name
    {
        get => name;
        set => name = value;
    }
}
```

Assume that another part is given, however, which provides implementing declarations of the partial methods:

C#

```
partial class Customer
{
    partial void OnNameChanging(string newName) =>
        Console.WriteLine($"Changing {name} to {newName}");

    partial void OnNameChanged() =>
        Console.WriteLine($"Changed to {name}");
}
```

Then the resulting combined class declaration will be equivalent to the following:

C#

```
class Customer
{
    string name;

    public string Name
    {
        get => name;
        set
        {
            OnNameChanging(value);
            name = value;
            OnNameChanged();
        }
    }
```

```
    void OnNameChanging(string newName) =>
        Console.WriteLine($"Changing {name} to {newName}");

    void OnNameChanged() =>
        Console.WriteLine($"Changed to {name}");
}
```

15.6.10 Extension methods

When the first parameter of a method includes the `this` modifier, that method is said to be an ***extension method***. Extension methods shall only be declared in non-generic, non-nested static classes. The first parameter of an extension method is restricted, as follows:

- It may have the parameter modifier `in` only if the parameter has a value type
- It may have the parameter modifier `ref` only if the parameter has a value type or is a generic type constrained to struct
- It shall not be a pointer type.

> *Example*: The following is an example of a static class that declares two extension methods:
>
> C#
>
> ```
> public static class Extensions
> {
>     public static int ToInt32(this string s) => Int32.Parse(s);
>
>     public static T[] Slice<T>(this T[] source, int index, int count)
>     {
>         if (index < 0 || count < 0 || source.Length - index < count)
>         {
>             throw new ArgumentException();
>         }
>         T[] result = new T[count];
>         Array.Copy(source, index, result, 0, count);
>         return result;
>     }
> }
> ```
>
> *end example*

An extension method is a regular static method. In addition, where its enclosing static class is in scope, an extension method may be invoked using instance method invocation syntax (§12.8.9.3), using the receiver expression as the first argument.

> *Example*: The following program uses the extension methods declared above:

C#

```
static class Program
{
    static void Main()
    {
        string[] strings = { "1", "22", "333", "4444" };
        foreach (string s in strings.Slice(1, 2))
        {
            Console.WriteLine(s.ToInt32());
        }
    }
}
```

The `Slice` method is available on the `string[]`, and the `ToInt32` method is available on `string`, because they have been declared as extension methods. The meaning of the program is the same as the following, using ordinary static method calls:

C#

```
static class Program
{
    static void Main()
    {
        string[] strings = { "1", "22", "333", "4444" };
        foreach (string s in Extensions.Slice(strings, 1, 2))
        {
            Console.WriteLine(Extensions.ToInt32(s));
        }
    }
}
```

end example

15.6.11 Method body

The method body of a method declaration consists of either a block body, an expression body or a semicolon.

Abstract and external method declarations do not provide a method implementation, so their method bodies simply consist of a semicolon. For any other method, the method body is a block (§13.3) that contains the statements to execute when that method is invoked.

The ***effective return type*** of a method is `void` if the return type is `void`, or if the method is async and the return type is `«TaskType»` (§15.15.1). Otherwise, the effective return type of a non-async method is its return type, and the effective return type of an async method with return type `«TaskType»<T>` (§15.15.1) is `T`.

When the effective return type of a method is `void` and the method has a block body, `return` statements (§13.10.5) in the block shall not specify an expression. If execution of the block of a void method completes normally (that is, control flows off the end of the method body), that method simply returns to its caller.

When the effective return type of a method is `void` and the method has an expression body, the expression `E` shall be a *statement_expression*, and the body is exactly equivalent to a block body of the form `{ E; }`.

For a returns-by-value method (§15.6.1), each return statement in that method's body shall specify an expression that is implicitly convertible to the effective return type.

For a returns-by-ref method (§15.6.1), each return statement in that method's body shall specify an expression whose type is that of the effective return type, and has a *ref-safe-context* of *caller-context* (§9.7.2).

For returns-by-value and returns-by-ref methods the endpoint of the method body shall not be reachable. In other words, control is not permitted to flow off the end of the method body.

> *Example*: In the following code

```C#
class A
{
    public int F() {} // Error, return value required

    public int G()
    {
        return 1;
    }

    public int H(bool b)
    {
        if (b)
        {
            return 1;
        }
        else
        {
            return 0;
```

```
        }
    }

    public int I(bool b) => b ? 1 : 0;
}
```

the value-returning `F` method results in a compile-time error because control can flow off the end of the method body. The `G` and `H` methods are correct because all possible execution paths end in a return statement that specifies a return value. The `I` method is correct, because its body is equivalent to a block with just a single return statement in it.

end example

15.7 Properties

15.7.1 General

A ***property*** is a member that provides access to a characteristic of an object or a class. Examples of properties include the length of a string, the size of a font, the caption of a window, the name of a customer, and so on. Properties are a natural extension of fields—both are named members with associated types, and the syntax for accessing fields and properties is the same. However, unlike fields, properties do not denote storage locations. Instead, properties have ***accessors*** that specify the statements to be executed when their values are read or written. Properties thus provide a mechanism for associating actions with the reading and writing of an object's characteristics; furthermore, they permit such characteristics to be computed.

Properties are declared using *property_declarations*:

ANTLR

```
property_declaration
    : attributes? property_modifier* type member_name property_body
    | attributes? property_modifier* ref_kind type member_name
ref_property_body
    ;

property_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'static'
```

```
    | 'virtual'
    | 'sealed'
    | 'override'
    | 'abstract'
    | 'extern'
    | unsafe_modifier   // unsafe code support
    ;

property_body
    : '{' accessor_declarations '}' property_initializer?
    | '=>' expression ';'
    ;

property_initializer
    : '=' variable_initializer ';'
    ;

ref_property_body
    : '{' ref_get_accessor_declaration '}'
    | '=>' 'ref' variable_reference ';'
    ;
```

unsafe_modifier (§23.2) is only available in unsafe code (§23).

There are two kinds of *property_declaration*:

- The first declares a non-ref-valued property. Its value has type *type*. This kind of property may be readable and/or writeable.
- The second declares a ref-valued property. Its value is a *variable_reference* (§9.5), that may be `readonly`, to a variable of type *type*. This kind of property is only readable.

A *property_declaration* may include a set of *attributes* (§22) and any one of the permitted kinds of declared accessibility (§15.3.6), the `new` (§15.3.5), `static` (§15.7.2), `virtual` (§15.6.4, §15.7.6), `override` (§15.6.5, §15.7.6), `sealed` (§15.6.6), `abstract` (§15.6.7, §15.7.6), and `extern` (§15.6.8) modifiers.

Property declarations are subject to the same rules as method declarations (§15.6) with regard to valid combinations of modifiers.

The *member_name* (§15.6.1) specifies the name of the property. Unless the property is an explicit interface member implementation, the *member_name* is simply an *identifier*. For an explicit interface member implementation (§18.6.2), the *member_name* consists of an *interface_type* followed by a "`.`" and an *identifier*.

The *type* of a property shall be at least as accessible as the property itself (§7.5.5).

A *property_body* may either consist of a statement body or an expression body. In a statement body, *accessor_declarations*, which shall be enclosed in "`{`" and "`}`" tokens, declare the accessors (§15.7.3) of the property. The accessors specify the executable statements associated with reading and writing the property.

In a *property_body* an expression body consisting of `=>` followed by an *expression* `E` and a semicolon is exactly equivalent to the statement body `{ get { return E; } }`, and can therefore only be used to specify read-only properties where the result of the get accessor is given by a single expression.

A *property_initializer* may only be given for an automatically implemented property (§15.7.4), and causes the initialization of the underlying field of such properties with the value given by the *expression*.

A *ref_property_body* may either consist of a statement body or an expression body. In a statement body a *get_accessor_declaration* declares the get accessor (§15.7.3) of the property. The accessor specifies the executable statements associated with reading the property.

In a *ref_property_body* an expression body consisting of `=>` followed by `ref`, a *variable_reference* `V` and a semicolon is exactly equivalent to the statement body `{ get { return ref V; } }`.

> *Note*: Even though the syntax for accessing a property is the same as that for a field, a property is not classified as a variable. Thus, it is not possible to pass a property as an `in`, `out`, or `ref` argument unless the property is ref-valued and therefore returns a variable reference (§9.7). *end note*

When a property declaration includes an `extern` modifier, the property is said to be an ***external property***. Because an external property declaration provides no actual implementation, each of its *accessor_declarations* consists of a semicolon.

15.7.2 Static and instance properties

When a property declaration includes a `static` modifier, the property is said to be a ***static property***. When no `static` modifier is present, the property is said to be an ***instance property***.

A static property is not associated with a specific instance, and it is a compile-time error to refer to `this` in the accessors of a static property.

An instance property is associated with a given instance of a class, and that instance can be accessed as `this` (§12.8.13) in the accessors of that property.

The differences between static and instance members are discussed further in §15.3.8.

15.7.3 Accessors

Note: This clause applies to both properties (§15.7) and indexers (§15.9). The clause is written in terms of properties, when reading for indexers substitute indexer/indexers for property/properties and consult the list of differences between properties and indexers given in §15.9.2. *end note*

The *accessor_declarations* of a property specify the executable statements associated with writing and/or reading that property.

ANTLR

```
accessor_declarations
    : get_accessor_declaration set_accessor_declaration?
    | set_accessor_declaration get_accessor_declaration?
    ;

get_accessor_declaration
    : attributes? accessor_modifier? 'get' accessor_body
    ;

set_accessor_declaration
    : attributes? accessor_modifier? 'set' accessor_body
    ;

accessor_modifier
    : 'protected'
    | 'internal'
    | 'private'
    | 'protected' 'internal'
    | 'internal' 'protected'
    | 'protected' 'private'
    | 'private' 'protected'
    ;

accessor_body
    : block
    | '=>' expression ';'
    | ';'
    ;

ref_get_accessor_declaration
    : attributes? accessor_modifier? 'get' ref_accessor_body
    ;
```

```
ref_accessor_body
    : block
    | '=>' 'ref' variable_reference ';'
    | ';'
    ;
```

The *accessor_declarations* consist of a *get_accessor_declaration*, a *set_accessor_declaration*, or both. Each accessor declaration consists of optional attributes, an optional *accessor_modifier*, the token `get` or `set`, followed by an *accessor_body*.

For a ref-valued property the *ref_get_accessor_declaration* consists optional attributes, an optional *accessor_modifier*, the token `get`, followed by an *ref_accessor_body*.

The use of *accessor_modifier*s is governed by the following restrictions:

- An *accessor_modifier* shall not be used in an interface or in an explicit interface member implementation.
- For a property or indexer that has no `override` modifier, an *accessor_modifier* is permitted only if the property or indexer has both a get and set accessor, and then is permitted only on one of those accessors.
- For a property or indexer that includes an `override` modifier, an accessor shall match the *accessor_modifier*, if any, of the accessor being overridden.
- The *accessor_modifier* shall declare an accessibility that is strictly more restrictive than the declared accessibility of the property or indexer itself. To be precise:
  - If the property or indexer has a declared accessibility of `public`, the accessibility declared by *accessor_modifier* may be either `private protected`, `protected internal`, `internal`, `protected`, or `private`.
  - If the property or indexer has a declared accessibility of `protected internal`, the accessibility declared by *accessor_modifier* may be either `private protected`, `protected private`, `internal`, `protected`, or `private`.
  - If the property or indexer has a declared accessibility of `internal` or `protected`, the accessibility declared by *accessor_modifier* shall be either `private protected` or `private`.
  - If the property or indexer has a declared accessibility of `private protected`, the accessibility declared by *accessor_modifier* shall be `private`.
  - If the property or indexer has a declared accessibility of `private`, no *accessor_modifier* may be used.

For `abstract` and `extern` non-ref-valued properties, any *accessor_body* for each accessor specified is simply a semicolon. A non-abstract, non-extern property, but not an indexer, may also have the *accessor_body* for all accessors specified be a semicolon, in

which case it is an ***automatically implemented property*** (§15.7.4). An automatically implemented property shall have at least a get accessor. For the accessors of any other non-abstract, non-extern property, the *accessor_body* is either:

- a *block* that specifies the statements to be executed when the corresponding accessor is invoked; or
- an expression body, which consists of `=>` followed by an *expression* and a semicolon, and denotes a single expression to be executed when the corresponding accessor is invoked.

For `abstract` and `extern` ref-valued properties the *ref_accessor_body* is simply a semicolon. For the accessor of any other non-abstract, non-extern property, the *ref_accessor_body* is either:

- a *block* that specifies the statements to be executed when the get accessor is invoked; or
- an expression body, which consists of `=>` followed by `ref`, a *variable_reference* and a semicolon. The variable reference is evaluated when the get accessor is invoked.

A get accessor for a non-ref-valued property corresponds to a parameterless method with a return value of the property type. Except as the target of an assignment, when such a property is referenced in an expression its get accessor is invoked to compute the value of the property (§12.2.2).

The body of a get accessor for a non-ref-valued property shall conform to the rules for value-returning methods described in §15.6.11. In particular, all `return` statements in the body of a get accessor shall specify an expression that is implicitly convertible to the property type. Furthermore, the endpoint of a get accessor shall not be reachable.

A get accessor for a ref-valued property corresponds to a parameterless method with a return value of a *variable_reference* to a variable of the property type. When such a property is referenced in an expression its get accessor is invoked to compute the *variable_reference* value of the property. That *variable reference*, like any other, is then used to read or, for non-readonly *variable_reference*s, write the referenced variable as required by the context.

> *Example*: The following example illustrates a ref-valued property as the target of an assignment:

```
C#
```

```
class Program
{
    static int field;
```

```
    static ref int Property => ref field;

    static void Main()
    {
        field = 10;
        Console.WriteLine(Property); // Prints 10
        Property = 20;               // This invokes the get accessor,
then assigns
                                     // via the resulting variable
reference
        Console.WriteLine(field);    // Prints 20
    }
}
```

end example

The body of a get accessor for a ref-valued property shall conform to the rules for ref-valued methods described in §15.6.11.

A set accessor corresponds to a method with a single value parameter of the property type and a `void` return type. The implicit parameter of a set accessor is always named `value`. When a property is referenced as the target of an assignment (§12.21), or as the operand of `++` or `--` (§12.8.15, §12.9.6), the set accessor is invoked with an argument that provides the new value (§12.21.2). The body of a set accessor shall conform to the rules for `void` methods described in §15.6.11. In particular, return statements in the set accessor body are not permitted to specify an expression. Since a set accessor implicitly has a parameter named `value`, it is a compile-time error for a local variable or constant declaration in a set accessor to have that name.

Based on the presence or absence of the get and set accessors, a property is classified as follows:

- A property that includes both a get accessor and a set accessor is said to be a ***read-write property***.
- A property that has only a get accessor is said to be a ***read-only property***. It is a compile-time error for a read-only property to be the target of an assignment.
- A property that has only a set accessor is said to be a ***write-only property***. Except as the target of an assignment, it is a compile-time error to reference a write-only property in an expression.

Note: The pre- and postfix `++` and `--` operators and compound assignment operators cannot be applied to write-only properties, since these operators read the old value of their operand before they write the new one. *end note*

Example: In the following code

```C#
public class Button : Control
{
    private string caption;

    public string Caption
    {
        get => caption;
        set
        {
            if (caption != value)
            {
                caption = value;
                Repaint();
            }
        }
    }

    public override void Paint(Graphics g, Rectangle r)
    {
        // Painting code goes here
    }
}
```

the `Button` control declares a public `Caption` property. The get accessor of the Caption property returns the `string` stored in the private `caption` field. The set accessor checks if the new value is different from the current value, and if so, it stores the new value and repaints the control. Properties often follow the pattern shown above: The get accessor simply returns a value stored in a `private` field, and the set accessor modifies that `private` field and then performs any additional actions required to update fully the state of the object. Given the `Button` class above, the following is an example of use of the `Caption` property:

```C#
Button okButton = new Button();
okButton.Caption = "OK"; // Invokes set accessor
string s = okButton.Caption; // Invokes get accessor
```

Here, the set accessor is invoked by assigning a value to the property, and the get accessor is invoked by referencing the property in an expression.

end example

The get and set accessors of a property are not distinct members, and it is not possible to declare the accessors of a property separately.

> *Example*: The example
>
> C#
>
> ```
> class A
> {
>     private string name;
>
>     // Error, duplicate member name
>     public string Name
>     {
>         get => name;
>     }
>
>     // Error, duplicate member name
>     public string Name
>     {
>         set => name = value;
>     }
> }
> ```
>
> does not declare a single read-write property. Rather, it declares two properties with the same name, one read-only and one write-only. Since two members declared in the same class cannot have the same name, the example causes a compile-time error to occur.
>
> *end example*

When a derived class declares a property by the same name as an inherited property, the derived property hides the inherited property with respect to both reading and writing.

> *Example*: In the following code
>
> C#
>
> ```
> class A
> {
>     public int P
>     {
>         set {...}
>     }
> }
>
> class B : A
> ```

```
{
    public new int P
    {
        get {...}
    }
}
```

the `P` property in `B` hides the `P` property in `A` with respect to both reading and writing. Thus, in the statements

C#

```
B b = new B();
b.P = 1;       // Error, B.P is read-only
((A)b).P = 1;  // Ok, reference to A.P
```

the assignment to `b.P` causes a compile-time error to be reported, since the read-only `P` property in `B` hides the write-only `P` property in `A`. Note, however, that a cast can be used to access the hidden `P` property.

end example

Unlike public fields, properties provide a separation between an object's internal state and its public interface.

Example: Consider the following code, which uses a `Point` struct to represent a location:

C#

```
class Label
{
    private int x, y;
    private string caption;

    public Label(int x, int y, string caption)
    {
        this.x = x;
        this.y = y;
        this.caption = caption;
    }

    public int X => x;
    public int Y => y;
    public Point Location => new Point(x, y);
    public string Caption => caption;
}
```

Here, the `Label` class uses two `int` fields, `x` and `y`, to store its location. The location is publicly exposed both as an `X` and a `Y` property and as a `Location` property of type `Point`. If, in a future version of `Label`, it becomes more convenient to store the location as a `Point` internally, the change can be made without affecting the public interface of the class:

C#

```
class Label
{
    private Point location;
    private string caption;

    public Label(int x, int y, string caption)
    {
        this.location = new Point(x, y);
        this.caption = caption;
    }

    public int X => location.X;
    public int Y => location.Y;
    public Point Location => location;
    public string Caption => caption;
}
```

Had `x` and `y` instead been `public readonly` fields, it would have been impossible to make such a change to the `Label` class.

end example

Note: Exposing state through properties is not necessarily any less efficient than exposing fields directly. In particular, when a property is non-virtual and contains only a small amount of code, the execution environment might replace calls to accessors with the actual code of the accessors. This process is known as ***inlining***, and it makes property access as efficient as field access, yet preserves the increased flexibility of properties. *end note*

Example: Since invoking a get accessor is conceptually equivalent to reading the value of a field, it is considered bad programming style for get accessors to have observable side-effects. In the example

C#

```
class Counter
{
```

```
    private int next;

    public int Next => next++;
}
```

the value of the `Next` property depends on the number of times the property has previously been accessed. Thus, accessing the property produces an observable side effect, and the property should be implemented as a method instead.

The “no side-effects” convention for get accessors doesn’t mean that get accessors should always be written simply to return values stored in fields. Indeed, get accessors often compute the value of a property by accessing multiple fields or invoking methods. However, a properly designed get accessor performs no actions that cause observable changes in the state of the object.

end example

Properties can be used to delay initialization of a resource until the moment it is first referenced.

Example:

C#

```
public class Console
{
    private static TextReader reader;
    private static TextWriter writer;
    private static TextWriter error;

    public static TextReader In
    {
        get
        {
            if (reader == null)
            {
                reader = new StreamReader(Console.OpenStandardInput());
            }
            return reader;
        }
    }

    public static TextWriter Out
    {
        get
        {
            if (writer == null)
            {
                writer = new StreamWriter(Console.OpenStandardOutput());
```

```
            }
            return writer;
        }
    }

    public static TextWriter Error
    {
        get
        {
            if (error == null)
            {
                error = new StreamWriter(Console.OpenStandardError());
            }
            return error;
        }
    }
...
}
```

The `Console` class contains three properties, `In`, `Out`, and `Error`, that represent the standard input, output, and error devices, respectively. By exposing these members as properties, the `Console` class can delay their initialization until they are actually used. For example, upon first referencing the `Out` property, as in

```
C#

Console.Out.WriteLine("hello, world");
```

the underlying `TextWriter` for the output device is created. However, if the application makes no reference to the `In` and `Error` properties, then no objects are created for those devices.

end example

15.7.4 Automatically implemented properties

An automatically implemented property (or auto-property for short), is a non-abstract, non-extern, non-ref-valued property with semicolon-only accessor bodies. Auto-properties shall have a get accessor and may optionally have a set accessor.

When a property is specified as an automatically implemented property, a hidden backing field is automatically available for the property, and the accessors are implemented to read from and write to that backing field. The hidden backing field is inaccessible, it can be read and written only through the automatically implemented property accessors, even within the containing type. If the auto-property has no set

accessor, the backing field is considered `readonly` (§15.5.3). Just like a `readonly` field, a read-only auto-property may also be assigned to in the body of a constructor of the enclosing class. Such an assignment assigns directly to the read-only backing field of the property.

An auto-property may optionally have a *property_initializer*, which is applied directly to the backing field as a *variable_initializer* (§17.7).

> *Example*:
>
> ```C#
> public class Point
> {
>     public int X { get; set; } // Automatically implemented
>     public int Y { get; set; } // Automatically implemented
> }
> ```
>
> is equivalent to the following declaration:
>
> ```C#
> public class Point
> {
>     private int x;
>     private int y;
>
>     public int X { get { return x; } set { x = value; } }
>     public int Y { get { return y; } set { y = value; } }
> }
> ```
>
> *end example*

> *Example*: In the following
>
> ```C#
> public class ReadOnlyPoint
> {
>     public int X { get; }
>     public int Y { get; }
>
>     public ReadOnlyPoint(int x, int y)
>     {
>         X = x;
>         Y = y;
> ```

```
    }
}
```

is equivalent to the following declaration:

C#

```
public class ReadOnlyPoint
{
    private readonly int __x;
    private readonly int __y;
    public int X { get { return __x; } }
    public int Y { get { return __y; } }

    public ReadOnlyPoint(int x, int y)
    {
        __x = x;
        __y = y;
    }
}
```

The assignments to the read-only field are valid, because they occur within the constructor.

end example

Although the backing field is hidden, that field may have field-targeted attributes applied directly to it via the automatically implemented property's *property_declaration* (§15.7.1).

Example: The following code

C#

```
[Serializable]
public class Foo
{
    [field: NonSerialized]
    public string MySecret { get; set; }
}
```

results in the field-targeted attribute `NonSerialized` being applied to the compiler-generated backing field, as if the code had been written as follows:

C#

```
[Serializable]
public class Foo
{
    [NonSerialized]
    private string _mySecretBackingField;
    public string MySecret
    {
        get { return _mySecretBackingField; }
        set { _mySecretBackingField = value; }
    }
}
```

end example

15.7.5 Accessibility

If an accessor has an *accessor_modifier*, the accessibility domain (§7.5.3) of the accessor is determined using the declared accessibility of the *accessor_modifier*. If an accessor does not have an *accessor_modifier*, the accessibility domain of the accessor is determined from the declared accessibility of the property or indexer.

The presence of an *accessor_modifier* never affects member lookup (§12.5) or overload resolution (§12.6.4). The modifiers on the property or indexer always determine which property or indexer is bound to, regardless of the context of the access.

Once a particular non-ref-valued property or non-ref-valued indexer has been selected, the accessibility domains of the specific accessors involved are used to determine if that usage is valid:

- If the usage is as a value (§12.2.2), the get accessor shall exist and be accessible.
- If the usage is as the target of a simple assignment (§12.21.2), the set accessor shall exist and be accessible.
- If the usage is as the target of compound assignment (§12.21.4), or as the target of the `++` or `--` operators (§12.8.15, §12.9.6), both the get accessors and the set accessor shall exist and be accessible.

Example: In the following example, the property `A.Text` is hidden by the property `B.Text`, even in contexts where only the set accessor is called. In contrast, the property `B.Count` is not accessible to class `M`, so the accessible property `A.Count` is used instead.

```
C#
```

```
class A
{
    public string Text
    {
        get => "hello";
        set { }
    }

    public int Count
    {
        get => 5;
        set { }
    }
}

class B : A
{
    private string text = "goodbye";
    private int count = 0;

    public new string Text
    {
        get => text;
        protected set => text = value;
    }

    protected new int Count
    {
        get => count;
        set => count = value;
    }
}

class M
{
    static void Main()
    {
        B b = new B();
        b.Count = 12;       // Calls A.Count set accessor
        int i = b.Count;    // Calls A.Count get accessor
        b.Text = "howdy";   // Error, B.Text set accessor not accessible
        string s = b.Text;  // Calls B.Text get accessor
    }
}
```

end example

Once a particular ref-valued property or ref-valued indexer has been selected; whether the usage is as a value, the target of a simple assignment, or the target of a compound assignment; the accessibility domain of the get accessor involved is used to determine if that usage is valid.

An accessor that is used to implement an interface shall not have an *accessor_modifier*. If only one accessor is used to implement an interface, the other accessor may be declared with an *accessor_modifier*:

> *Example*:
>
> ```
> C#
> ```
>
> ```
> public interface I
> {
>     string Prop { get; }
> }
>
> public class C : I
> {
>     public string Prop
>     {
>         get => "April";     // Must not have a modifier here
>         internal set {...}  // Ok, because I.Prop has no set accessor
>     }
> }
> ```
>
> *end example*

15.7.6 Virtual, sealed, override, and abstract accessors

Note: This clause applies to both properties (§15.7) and indexers (§15.9). The clause is written in terms of properties, when reading for indexers substitute indexer/indexers for property/properties and consult the list of differences between properties and indexers given in §15.9.2. *end note*

A virtual property declaration specifies that the accessors of the property are virtual. The `virtual` modifier applies to all non-private accessors of a property. When an accessor of a virtual property has the private *accessor_modifier*, the `private` accessor is implicitly not virtual.

An abstract property declaration specifies that the accessors of the property are virtual, but does not provide an actual implementation of the accessors. Instead, non-abstract derived classes are required to provide their own implementation for the accessors by overriding the property. Because an accessor for an abstract property declaration provides no actual implementation, its *accessor_body* simply consists of a semicolon. An abstract property shall not have a `private` accessor.

A property declaration that includes both the `abstract` and `override` modifiers specifies that the property is abstract and overrides a base property. The accessors of such a

property are also abstract.

Abstract property declarations are only permitted in abstract classes (§15.2.2.2). The accessors of an inherited virtual property can be overridden in a derived class by including a property declaration that specifies an `override` directive. This is known as an ***overriding property declaration***. An overriding property declaration does not declare a new property. Instead, it simply specializes the implementations of the accessors of an existing virtual property.

The override declaration and the overridden base property are required to have the same declared accessibility. In other words, an override declaration shall not change the accessibility of the base property. However, if the overridden base property is protected internal and it is declared in a different assembly than the assembly containing the override declaration then the override declaration's declared accessibility shall be protected. If the inherited property has only a single accessor (i.e., if the inherited property is read-only or write-only), the overriding property shall include only that accessor. If the inherited property includes both accessors (i.e., if the inherited property is read-write), the overriding property can include either a single accessor or both accessors. There shall be an identity conversion between the type of the overriding and the inherited property.

An overriding property declaration may include the `sealed` modifier. Use of this modifier prevents a derived class from further overriding the property. The accessors of a sealed property are also sealed.

Except for differences in declaration and invocation syntax, virtual, sealed, override, and abstract accessors behave exactly like virtual, sealed, override and abstract methods. Specifically, the rules described in §15.6.4, §15.6.5, §15.6.6, and §15.6.7 apply as if accessors were methods of a corresponding form:

- A get accessor corresponds to a parameterless method with a return value of the property type and the same modifiers as the containing property.
- A set accessor corresponds to a method with a single value parameter of the property type, a void return type, and the same modifiers as the containing property.

> *Example*: In the following code
>
> C#
>
> ```
> abstract class A
> {
>     int y;
> ```

```
    public virtual int X
    {
        get => 0;
    }

    public virtual int Y
    {
        get => y;
        set => y = value;
    }

    public abstract int Z { get; set; }
}
```

`X` is a virtual read-only property, `Y` is a virtual read-write property, and `Z` is an abstract read-write property. Because `Z` is abstract, the containing class A shall also be declared abstract.

A class that derives from `A` is shown below:

C#

```
class B : A
{
    int z;

    public override int X
    {
        get => base.X + 1;
    }

    public override int Y
    {
        set => base.Y = value < 0 ? 0: value;
    }

    public override int Z
    {
        get => z;
        set => z = value;
    }
}
```

Here, the declarations of `X`, `Y`, and `Z` are overriding property declarations. Each property declaration exactly matches the accessibility modifiers, type, and name of the corresponding inherited property. The get accessor of `X` and the set accessor of `Y` use the base keyword to access the inherited accessors. The declaration of `Z` overrides both abstract accessors—thus, there are no outstanding `abstract` function members in `B`, and `B` is permitted to be a non-abstract class.

end example

When a property is declared as an override, any overridden accessors shall be accessible to the overriding code. In addition, the declared accessibility of both the property or indexer itself, and of the accessors, shall match that of the overridden member and accessors.

Example:

C#

```
public class B
{
    public virtual int P
    {
        get {...}
        protected set {...}
    }
}

public class D: B
{
    public override int P
    {
        get {...}            // Must not have a modifier here
        protected set {...}  // Must specify protected here
    }
}
```

end example

15.8 Events

15.8.1 General

An ***event*** is a member that enables an object or class to provide notifications. Clients can attach executable code for events by supplying ***event handlers***.

Events are declared using *event_declarations*:

ANTLR

```
event_declaration
    : attributes? event_modifier* 'event' type variable_declarators ';'
    | attributes? event_modifier* 'event' type member_name
        '{' event_accessor_declarations '}'
```

```
    ;

event_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'static'
    | 'virtual'
    | 'sealed'
    | 'override'
    | 'abstract'
    | 'extern'
    | unsafe_modifier   // unsafe code support
    ;

event_accessor_declarations
    : add_accessor_declaration remove_accessor_declaration
    | remove_accessor_declaration add_accessor_declaration
    ;

add_accessor_declaration
    : attributes? 'add' block
    ;

remove_accessor_declaration
    : attributes? 'remove' block
    ;
```

unsafe_modifier (§23.2) is only available in unsafe code (§23).

An *event_declaration* may include a set of *attributes* (§22) and any one of the permitted kinds of declared accessibility (§15.3.6), the `new` (§15.3.5), `static` (§15.6.3, §15.8.4), `virtual` (§15.6.4, §15.8.5), `override` (§15.6.5, §15.8.5), `sealed` (§15.6.6), `abstract` (§15.6.7, §15.8.5), and `extern` (§15.6.8) modifiers.

Event declarations are subject to the same rules as method declarations (§15.6) with regard to valid combinations of modifiers.

The *type* of an event declaration shall be a *delegate_type* (§8.2.8), and that *delegate_type* shall be at least as accessible as the event itself (§7.5.5).

An event declaration can include *event_accessor_declaration*s. However, if it does not, for non-extern, non-abstract events, the compiler shall supply them automatically (§15.8.2); for `extern` events, the accessors are provided externally.

An event declaration that omits *event_accessor_declaration*s defines one or more events —one for each of the *variable_declarator*s. The attributes and modifiers apply to all of

the members declared by such an *event_declaration*.

It is a compile-time error for an *event_declaration* to include both the `abstract` modifier and *event_accessor_declaration*s.

When an event declaration includes an `extern` modifier, the event is said to be an ***external event***. Because an external event declaration provides no actual implementation, it is an error for it to include both the `extern` modifier and *event_accessor_declaration*s.

It is a compile-time error for a *variable_declarator* of an event declaration with an `abstract` or `external` modifier to include a *variable_initializer*.

An event can be used as the left operand of the `+=` and `-=` operators. These operators are used, respectively, to attach event handlers to, or to remove event handlers from an event, and the access modifiers of the event control the contexts in which such operations are permitted.

The only operations that are permitted on an event by code that is outside the type in which that event is declared, are `+=` and `-=`. Therefore, while such code can add and remove handlers for an event, it cannot directly obtain or modify the underlying list of event handlers.

In an operation of the form `x += y` or `x -= y`, when `x` is an event the result of the operation has type `void` (§12.21.5) (as opposed to having the type of `x`, with the value of `x` after the assignment, as for other the `+=` and `-=` operators defined on non-event types). This prevents external code from indirectly examining the underlying delegate of an event.

> *Example*: The following example shows how event handlers are attached to instances of the `Button` class:

```C#
public delegate void EventHandler(object sender, EventArgs e);

public class Button : Control
{
    public event EventHandler Click;
}

public class LoginDialog : Form
{
    Button okButton;
    Button cancelButton;
```

```
    public LoginDialog()
    {
        okButton = new Button(...);
        okButton.Click += new EventHandler(OkButtonClick);
        cancelButton = new Button(...);
        cancelButton.Click += new EventHandler(CancelButtonClick);
    }

    void OkButtonClick(object sender, EventArgs e)
    {
        // Handle okButton.Click event
    }

    void CancelButtonClick(object sender, EventArgs e)
    {
        // Handle cancelButton.Click event
    }
}
```

Here, the `LoginDialog` instance constructor creates two `Button` instances and attaches event handlers to the `Click` events.

end example

15.8.2 Field-like events

Within the program text of the class or struct that contains the declaration of an event, certain events can be used like fields. To be used in this way, an event shall not be abstract or extern, and shall not explicitly include *event_accessor_declaration*s. Such an event can be used in any context that permits a field. The field contains a delegate (§20), which refers to the list of event handlers that have been added to the event. If no event handlers have been added, the field contains `null`.

Example: In the following code

C#

```
public delegate void EventHandler(object sender, EventArgs e);

public class Button : Control
{
    public event EventHandler Click;

    protected void OnClick(EventArgs e)
    {
        EventHandler handler = Click;
        if (handler != null)
        {
            handler(this, e);
```

```
        }
    }

    public void Reset() => Click = null;
}
```

`Click` is used as a field within the `Button` class. As the example demonstrates, the field can be examined, modified, and used in delegate invocation expressions. The `OnClick` method in the `Button` class "raises" the `Click` event. The notion of raising an event is precisely equivalent to invoking the delegate represented by the event—thus, there are no special language constructs for raising events. Note that the delegate invocation is preceded by a check that ensures the delegate is non-null and that the check is made on a local copy to ensure thread safety.

Outside the declaration of the `Button` class, the `Click` member can only be used on the left-hand side of the `+=` and `-=` operators, as in

C#

```
b.Click += new EventHandler(...);
```

which appends a delegate to the invocation list of the `Click` event, and

C#

```
Click -= new EventHandler(...);
```

which removes a delegate from the invocation list of the `Click` event.

end example

When compiling a field-like event, the compiler automatically creates storage to hold the delegate, and creates accessors for the event that add or remove event handlers to the delegate field. The addition and removal operations are thread safe, and may (but are not required to) be done while holding the lock (§13.13) on the containing object for an instance event, or the `System.Type` object (§12.8.17) for a static event.

Note: Thus, an instance event declaration of the form:

C#

```
class X
{
```

```
    public event D Ev;
}
```

shall be compiled to something equivalent to:

```C#
class X
{
    private D __Ev; // field to hold the delegate

    public event D Ev
    {
        add
        {
            /* Add the delegate in a thread safe way */
        }
        remove
        {
            /* Remove the delegate in a thread safe way */
        }
    }
}
```

Within the class `X`, references to `Ev` on the left-hand side of the `+=` and `-=` operators cause the add and remove accessors to be invoked. All other references to `Ev` are compiled to reference the hidden field `__Ev` instead (§12.8.7). The name "`__Ev`" is arbitrary; the hidden field could have any name or no name at all.

end note

15.8.3 Event accessors

Note: Event declarations typically omit *event_accessor_declarations*, as in the `Button` example above. For example, they might be included if the storage cost of one field per event is not acceptable. In such cases, a class can include *event_accessor_declarations* and use a private mechanism for storing the list of event handlers. *end note*

The *event_accessor_declarations* of an event specify the executable statements associated with adding and removing event handlers.

The accessor declarations consist of an *add_accessor_declaration* and a *remove_accessor_declaration*. Each accessor declaration consists of the token add or remove followed by a *block*. The *block* associated with an *add_accessor_declaration*

specifies the statements to execute when an event handler is added, and the *block* associated with a *remove_accessor_declaration* specifies the statements to execute when an event handler is removed.

Each *add_accessor_declaration* and *remove_accessor_declaration* corresponds to a method with a single value parameter of the event type, and a `void` return type. The implicit parameter of an event accessor is named `value`. When an event is used in an event assignment, the appropriate event accessor is used. Specifically, if the assignment operator is `+=` then the add accessor is used, and if the assignment operator is `-=` then the remove accessor is used. In either case, the right operand of the assignment operator is used as the argument to the event accessor. The block of an *add_accessor_declaration* or a *remove_accessor_declaration* shall conform to the rules for `void` methods described in §15.6.9. In particular, `return` statements in such a block are not permitted to specify an expression.

Since an event accessor implicitly has a parameter named `value`, it is a compile-time error for a local variable or constant declared in an event accessor to have that name.

> *Example*: In the following code

C#

```
class Control : Component
{
    // Unique keys for events
    static readonly object mouseDownEventKey = new object();
    static readonly object mouseUpEventKey = new object();

    // Return event handler associated with key
    protected Delegate GetEventHandler(object key) {...}

    // Add event handler associated with key
    protected void AddEventHandler(object key, Delegate handler) {...}

    // Remove event handler associated with key
    protected void RemoveEventHandler(object key, Delegate handler)
{...}

    // MouseDown event
    public event MouseEventHandler MouseDown
    {
        add { AddEventHandler(mouseDownEventKey, value); }
        remove { RemoveEventHandler(mouseDownEventKey, value); }
    }

    // MouseUp event
    public event MouseEventHandler MouseUp
    {
```

```
        add { AddEventHandler(mouseUpEventKey, value); }
        remove { RemoveEventHandler(mouseUpEventKey, value); }
    }

    // Invoke the MouseUp event
    protected void OnMouseUp(MouseEventArgs args)
    {
        MouseEventHandler handler;
        handler = (MouseEventHandler)GetEventHandler(mouseUpEventKey);
        if (handler != null)
        {
            handler(this, args);
        }
    }
}
```

the `Control` class implements an internal storage mechanism for events. The `AddEventHandler` method associates a delegate value with a key, the `GetEventHandler` method returns the delegate currently associated with a key, and the `RemoveEventHandler` method removes a delegate as an event handler for the specified event. Presumably, the underlying storage mechanism is designed such that there is no cost for associating a null delegate value with a key, and thus unhandled events consume no storage.

end example

15.8.4 Static and instance events

When an event declaration includes a `static` modifier, the event is said to be a ***static event***. When no `static` modifier is present, the event is said to be an ***instance event***.

A static event is not associated with a specific instance, and it is a compile-time error to refer to `this` in the accessors of a static event.

An instance event is associated with a given instance of a class, and this instance can be accessed as `this` (§12.8.13) in the accessors of that event.

The differences between static and instance members are discussed further in §15.3.8.

15.8.5 Virtual, sealed, override, and abstract accessors

A virtual event declaration specifies that the accessors of that event are virtual. The `virtual` modifier applies to both accessors of an event.

An abstract event declaration specifies that the accessors of the event are virtual, but does not provide an actual implementation of the accessors. Instead, non-abstract derived classes are required to provide their own implementation for the accessors by overriding the event. Because an accessor for an abstract event declaration provides no actual implementation, it shall not provide *event_accessor_declaration*s.

An event declaration that includes both the `abstract` and `override` modifiers specifies that the event is abstract and overrides a base event. The accessors of such an event are also abstract.

Abstract event declarations are only permitted in abstract classes (§15.2.2.2).

The accessors of an inherited virtual event can be overridden in a derived class by including an event declaration that specifies an `override` modifier. This is known as an ***overriding event declaration***. An overriding event declaration does not declare a new event. Instead, it simply specializes the implementations of the accessors of an existing virtual event.

An overriding event declaration shall specify the exact same accessibility modifiers and name as the overridden event, there shall be an identity conversion between the type of the overriding and the overridden event, and both the add and remove accessors shall be specified within the declaration.

An overriding event declaration can include the `sealed` modifier. Use of `this` modifier prevents a derived class from further overriding the event. The accessors of a sealed event are also sealed.

It is a compile-time error for an overriding event declaration to include a `new` modifier.

Except for differences in declaration and invocation syntax, virtual, sealed, override, and abstract accessors behave exactly like virtual, sealed, override and abstract methods. Specifically, the rules described in §15.6.4, §15.6.5, §15.6.6, and §15.6.7 apply as if accessors were methods of a corresponding form. Each accessor corresponds to a method with a single value parameter of the event type, a `void` return type, and the same modifiers as the containing event.

15.9 Indexers

15.9.1 General

An ***indexer*** is a member that enables an object to be indexed in the same way as an array. Indexers are declared using *indexer_declaration*s:

ANTLR

```
indexer_declaration
    : attributes? indexer_modifier* indexer_declarator indexer_body
    | attributes? indexer_modifier* ref_kind indexer_declarator
ref_indexer_body
    ;

indexer_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'virtual'
    | 'sealed'
    | 'override'
    | 'abstract'
    | 'extern'
    | unsafe_modifier   // unsafe code support
    ;

indexer_declarator
    : type 'this' '[' formal_parameter_list ']'
    | type interface_type '.' 'this' '[' formal_parameter_list ']'
    ;

indexer_body
    : '{' accessor_declarations '}'
    | '=>' expression ';'
    ;

ref_indexer_body
    : '{' ref_get_accessor_declaration '}'
    | '=>' 'ref' variable_reference ';'
    ;
```

unsafe_modifier (§23.2) is only available in unsafe code (§23).

There are two kinds of *indexer_declaration*:

- The first declares a non-ref-valued indexer. Its value has type *type*. This kind of indexer may be readable and/or writeable.
- The second declares a ref-valued indexer. Its value is a *variable_reference* (§9.5), that may be `readonly`, to a variable of type *type*. This kind of indexer is only readable.

An *indexer_declaration* may include a set of *attributes* (§22) and any one of the permitted kinds of declared accessibility (§15.3.6), the `new` (§15.3.5), `virtual` (§15.6.4), `override` (§15.6.5), `sealed` (§15.6.6), `abstract` (§15.6.7), and `extern` (§15.6.8) modifiers.

Indexer declarations are subject to the same rules as method declarations (§15.6) with regard to valid combinations of modifiers, with the one exception being that the `static` modifier is not permitted on an indexer declaration.

The *type* of an indexer declaration specifies the element type of the indexer introduced by the declaration.

> *Note*: As indexers are designed to be used in array element-like contexts, the term *element type* as defined for an array is also used with an indexer. *end note*

Unless the indexer is an explicit interface member implementation, the *type* is followed by the keyword `this`. For an explicit interface member implementation, the *type* is followed by an *interface_type*, a "`.`", and the keyword `this`. Unlike other members, indexers do not have user-defined names.

The *formal_parameter_list* specifies the parameters of the indexer. The formal parameter list of an indexer corresponds to that of a method (§15.6.2), except that at least one parameter shall be specified, and that the `this`, `ref`, and `out` parameter modifiers are not permitted.

The *type* of an indexer and each of the types referenced in the *formal_parameter_list* shall be at least as accessible as the indexer itself (§7.5.5).

An *indexer_body* may either consist of a statement body (§15.7.1) or an expression body (§15.6.1). In a statement body, *accessor_declarations*, which shall be enclosed in "`{`" and "`}`" tokens, declare the accessors (§15.7.3) of the indexer. The accessors specify the executable statements associated with reading and writing indexer elements.

In a *indexer_body* an expression body consisting of "`=>`" followed by an expression `E` and a semicolon is exactly equivalent to the statement body `{ get { return E; } }`, and can therefore only be used to specify read-only indexers where the result of the get accessor is given by a single expression.

A *ref_indexer_body* may either consist of a statement body or an expression body. In a statement body a *get_accessor_declaration* declares the get accessor (§15.7.3) of the property. The accessor specifies the executable statements associated with reading the property.

In a *ref_indexer_body* an expression body consisting of `=>` followed by `ref`, a *variable_reference* `V` and a semicolon is exactly equivalent to the statement body `{ get { return ref V; } }`.

> *Note*: Even though the syntax for accessing an indexer element is the same as that for an array element, an indexer element is not classified as a variable. Thus, it is not possible to pass an indexer element as an `in`, `out`, or `ref` argument unless the indexer is ref-valued and therefore returns a reference (§9.7). *end note*

The *formal_parameter_list* of an indexer defines the signature (§7.6) of the indexer. Specifically, the signature of an indexer consists of the number and types of its formal parameters. The element type and names of the formal parameters are not part of an indexer's signature.

The signature of an indexer shall differ from the signatures of all other indexers declared in the same class.

When an indexer declaration includes an `extern` modifier, the indexer is said to be an ***external indexer***. Because an external indexer declaration provides no actual implementation, each of its *accessor_declarations* consists of a semicolon.

> *Example*: The example below declares a `BitArray` class that implements an indexer for accessing the individual bits in the bit array.

C#

```
class BitArray
{
    int[] bits;
    int length;

    public BitArray(int length)
    {
        if (length < 0)
        {
            throw new ArgumentException();
        }
        bits = new int[((length - 1) >> 5) + 1];
        this.length = length;
    }

    public int Length => length;

    public bool this[int index]
    {
        get
        {
            if (index < 0 || index >= length)
            {
                throw new IndexOutOfRangeException();
            }
            return (bits[index >> 5] & 1 << index) != 0;
```

```
        }
        set
        {
            if (index < 0 || index >= length)
            {
                throw new IndexOutOfRangeException();
            }
            if (value)
            {
                bits[index >> 5] |= 1 << index;
            }
            else
            {
                bits[index >> 5] &= ~(1 << index);
            }
        }
    }
}
```

An instance of the `BitArray` class consumes substantially less memory than a corresponding `bool[]` (since each value of the former occupies only one bit instead of the latter's one `byte`), but it permits the same operations as a `bool[]`.

The following `CountPrimes` class uses a `BitArray` and the classical "sieve" algorithm to compute the number of primes between 2 and a given maximum:

C#

```
class CountPrimes
{
    static int Count(int max)
    {
        BitArray flags = new BitArray(max + 1);
        int count = 0;
        for (int i = 2; i <= max; i++)
        {
            if (!flags[i])
            {
                for (int j = i * 2; j <= max; j += i)
                {
                    flags[j] = true;
                }
                count++;
            }
        }
        return count;
    }

    static void Main(string[] args)
    {
        int max = int.Parse(args[0]);
        int count = Count(max);
```

```
        Console.WriteLine($"Found {count} primes between 2 and {max}");
    }
}
```

Note that the syntax for accessing elements of the `BitArray` is precisely the same as for a `bool[]`.

The following example shows a 26×10 grid class that has an indexer with two parameters. The first parameter is required to be an upper- or lowercase letter in the range A–Z, and the second is required to be an integer in the range 0–9.

C#

```
class Grid
{
    const int NumRows = 26;
    const int NumCols = 10;
    int[,] cells = new int[NumRows, NumCols];

    public int this[char row, int col]
    {
        get
        {
            row = Char.ToUpper(row);
            if (row < 'A' || row > 'Z')
            {
                throw new ArgumentOutOfRangeException("row");
            }
            if (col < 0 || col >= NumCols)
            {
                throw new ArgumentOutOfRangeException ("col");
            }
            return cells[row - 'A', col];
        }
        set
        {
            row = Char.ToUpper(row);
            if (row < 'A' || row > 'Z')
            {
                throw new ArgumentOutOfRangeException ("row");
            }
            if (col < 0 || col >= NumCols)
            {
                throw new ArgumentOutOfRangeException ("col");
            }
            cells[row - 'A', col] = value;
        }
    }
}
```

end example

15.9.2 Indexer and Property Differences

Indexers and properties are very similar in concept, but differ in the following ways:

- A property is identified by its name, whereas an indexer is identified by its signature.
- A property is accessed through a *simple_name* (§12.8.4) or a *member_access* (§12.8.7), whereas an indexer element is accessed through an *element_access* (§12.8.11.3).
- A property can be a static member, whereas an indexer is always an instance member.
- A get accessor of a property corresponds to a method with no parameters, whereas a get accessor of an indexer corresponds to a method with the same formal parameter list as the indexer.
- A set accessor of a property corresponds to a method with a single parameter named `value`, whereas a set accessor of an indexer corresponds to a method with the same formal parameter list as the indexer, plus an additional parameter named `value`.
- It is a compile-time error for an indexer accessor to declare a local variable or local constant with the same name as an indexer parameter.
- In an overriding property declaration, the inherited property is accessed using the syntax `base.P`, where `P` is the property name. In an overriding indexer declaration, the inherited indexer is accessed using the syntax `base[E]`, where `E` is a comma-separated list of expressions.
- There is no concept of an "automatically implemented indexer". It is an error to have a non-abstract, non-external indexer with semicolon accessors.

Aside from these differences, all rules defined in §15.7.3, §15.7.5 and §15.7.6 apply to indexer accessors as well as to property accessors.

This replacing of property/properties with indexer/indexers when reading §15.7.3, §15.7.5 and §15.7.6 applies to defined terms as well. Specifically, *read-write property* becomes ***read-write indexer***, *read-only property* becomes ***read-only indexer***, and *write-only property* becomes ***write-only indexer***.

15.10 Operators

15.10.1 General

An ***operator*** is a member that defines the meaning of an expression operator that can be applied to instances of the class. Operators are declared using *operator_declarations*:

ANTLR

```
operator_declaration
    : attributes? operator_modifier+ operator_declarator operator_body
    ;

operator_modifier
    : 'public'
    | 'static'
    | 'extern'
    | unsafe_modifier   // unsafe code support
    ;

operator_declarator
    : unary_operator_declarator
    | binary_operator_declarator
    | conversion_operator_declarator
    ;

unary_operator_declarator
    : type 'operator' overloadable_unary_operator '(' fixed_parameter ')'
    ;

overloadable_unary_operator
    : '+' | '-' | '!' | '~' | '++' | '--' | 'true' | 'false'
    ;

binary_operator_declarator
    : type 'operator' overloadable_binary_operator
        '(' fixed_parameter ',' fixed_parameter ')'
    ;

overloadable_binary_operator
    : '+'  | '-'  | '*'  | '/'  | '%'  | '&' | '|' | '^'  | '<<'
    | right_shift | '==' | '!=' | '>' | '<' | '>=' | '<='
    ;

conversion_operator_declarator
    : 'implicit' 'operator' type '(' fixed_parameter ')'
    | 'explicit' 'operator' type '(' fixed_parameter ')'
    ;

operator_body
    : block
    | '=>' expression ';'
    | ';'
    ;
```

unsafe_modifier (§23.2) is only available in unsafe code (§23).

There are three categories of overloadable operators: Unary operators (§15.10.2), binary operators (§15.10.3), and conversion operators (§15.10.4).

The *operator_body* is either a semicolon, a block body (§15.6.1) or an expression body (§15.6.1). A block body consists of a *block*, which specifies the statements to execute when the operator is invoked. The *block* shall conform to the rules for value-returning methods described in §15.6.11. An expression body consists of `=>` followed by an expression and a semicolon, and denotes a single expression to perform when the operator is invoked.

For `extern` operators, the *operator_body* consists simply of a semicolon. For all other operators, the *operator_body* is either a block body or an expression body.

The following rules apply to all operator declarations:

- An operator declaration shall include both a `public` and a `static` modifier.
- The parameter(s) of an operator shall have no modifiers other than `in`.
- The signature of an operator (§15.10.2, §15.10.3, §15.10.4) shall differ from the signatures of all other operators declared in the same class.
- All types referenced in an operator declaration shall be at least as accessible as the operator itself (§7.5.5).
- It is an error for the same modifier to appear multiple times in an operator declaration.

Each operator category imposes additional restrictions, as described in the following subclauses.

Like other members, operators declared in a base class are inherited by derived classes. Because operator declarations always require the class or struct in which the operator is declared to participate in the signature of the operator, it is not possible for an operator declared in a derived class to hide an operator declared in a base class. Thus, the `new` modifier is never required, and therefore never permitted, in an operator declaration.

Additional information on unary and binary operators can be found in §12.4.

Additional information on conversion operators can be found in §10.5.

15.10.2 Unary operators

The following rules apply to unary operator declarations, where `T` denotes the instance type of the class or struct that contains the operator declaration:

- A unary `+`, `-`, `!`, or `~` operator shall take a single parameter of type `T` or `T?` and can return any type.
- A unary `++` or `--` operator shall take a single parameter of type `T` or `T?` and shall return that same type or a type derived from it.

- A unary `true` or `false` operator shall take a single parameter of type `T` or `T?` and shall return type `bool`.

The signature of a unary operator consists of the operator token (`+`, `-`, `!`, `~`, `++`, `--`, `true`, or `false`) and the type of the single formal parameter. The return type is not part of a unary operator's signature, nor is the name of the formal parameter.

The `true` and `false` unary operators require pair-wise declaration. A compile-time error occurs if a class declares one of these operators without also declaring the other. The `true` and `false` operators are described further in §12.24.

> *Example*: The following example shows an implementation and subsequent usage of operator++ for an integer vector class:

C#

```
public class IntVector
{
    public IntVector(int length) {...}
    public int Length { get { ... } }                       // Read-only
property
    public int this[int index] { get { ... } set { ... } } // Read-write
indexer

    public static IntVector operator++(IntVector iv)
    {
        IntVector temp = new IntVector(iv.Length);
        for (int i = 0; i < iv.Length; i++)
        {
            temp[i] = iv[i] + 1;
        }
        return temp;
    }
}

class Test
{
    static void Main()
    {
        IntVector iv1 = new IntVector(4); // Vector of 4 x 0
        IntVector iv2;
        iv2 = iv1++;              // iv2 contains 4 x 0, iv1 contains 4
x 1
        iv2 = ++iv1;              // iv2 contains 4 x 2, iv1 contains 4
x 2
    }
}
```

> Note how the operator method returns the value produced by adding 1 to the operand, just like the postfix increment and decrement operators (§12.8.15), and the prefix increment and decrement operators (§12.9.6). Unlike in C++, this method should not modify the value of its operand directly as this would violate the standard semantics of the postfix increment operator (§12.8.15).
>
> *end example*

15.10.3 Binary operators

The following rules apply to binary operator declarations, where `T` denotes the instance type of the class or struct that contains the operator declaration:

- A binary non-shift operator shall take two parameters, at least one of which shall have type `T` or `T?`, and can return any type.
- A binary `<<` or `>>` operator (§12.11) shall take two parameters, the first of which shall have type `T` or T? and the second of which shall have type `int` or `int?`, and can return any type.

The signature of a binary operator consists of the operator token (`+`, `-`, `*`, `/`, `%`, `&`, `|`, `^`, `<<`, `>>`, `==`, `!=`, `>`, `<`, `>=`, or `<=`) and the types of the two formal parameters. The return type and the names of the formal parameters are not part of a binary operator's signature.

Certain binary operators require pair-wise declaration. For every declaration of either operator of a pair, there shall be a matching declaration of the other operator of the pair. Two operator declarations match if identity conversions exist between their return types and their corresponding parameter types. The following operators require pair-wise declaration:

- operator `==` and operator `!=`
- operator `>` and operator `<`
- operator `>=` and operator `<=`

15.10.4 Conversion operators

A conversion operator declaration introduces a ***user-defined conversion*** (§10.5), which augments the pre-defined implicit and explicit conversions.

A conversion operator declaration that includes the `implicit` keyword introduces a user-defined implicit conversion. Implicit conversions can occur in a variety of situations,

including function member invocations, cast expressions, and assignments. This is described further in §10.2.

A conversion operator declaration that includes the `explicit` keyword introduces a user-defined explicit conversion. Explicit conversions can occur in cast expressions, and are described further in §10.3.

A conversion operator converts from a source type, indicated by the parameter type of the conversion operator, to a target type, indicated by the return type of the conversion operator.

For a given source type `S` and target type `T`, if `S` or `T` are nullable value types, let S_0 and T_0 refer to their underlying types; otherwise, S_0 and T_0 are equal to `S` and `T` respectively. A class or struct is permitted to declare a conversion from a source type `S` to a target type `T` only if all of the following are true:

- S_0 and T_0 are different types.
- Either S_0 or T_0 is the instance type of the class or struct that contains the operator declaration.
- Neither S_0 nor T_0 is an *interface_type*.
- Excluding user-defined conversions, a conversion does not exist from `S` to `T` or from `T` to `S`.

For the purposes of these rules, any type parameters associated with `S` or `T` are considered to be unique types that have no inheritance relationship with other types, and any constraints on those type parameters are ignored.

> *Example*: In the following:
>
> C#
>
> ```
> class C<T> {...}
>
> class D<T> : C<T>
> {
>     public static implicit operator C<int>(D<T> value) {...}     // Ok
>     public static implicit operator C<string>(D<T> value) {...}  // Ok
>     public static implicit operator C<T>(D<T> value) {...}       //
> Error
> }
> ```

> the first two operator declarations are permitted because `T` and `int` and `string`, respectively are considered unique types with no relationship. However, the third operator is an error because `C<T>` is the base class of `D<T>`.
>
> *end example*

From the second rule, it follows that a conversion operator shall convert either to or from the class or struct type in which the operator is declared.

> *Example*: It is possible for a class or struct type `C` to define a conversion from `C` to `int` and from `int` to `C`, but not from `int` to `bool`. *end example*

It is not possible to directly redefine a pre-defined conversion. Thus, conversion operators are not allowed to convert from or to `object` because implicit and explicit conversions already exist between `object` and all other types. Likewise, neither the source nor the target types of a conversion can be a base type of the other, since a conversion would then already exist. However, it *is* possible to declare operators on generic types that, for particular type arguments, specify conversions that already exist as pre-defined conversions.

> *Example*:
>
> C#
>
> ```
> struct Convertible<T>
> {
>     public static implicit operator Convertible<T>(T value) {...}
>     public static explicit operator T(Convertible<T> value) {...}
> }
> ```
>
> when type `object` is specified as a type argument for `T`, the second operator declares a conversion that already exists (an implicit, and therefore also an explicit, conversion exists from any type to type object).
>
> *end example*

In cases where a pre-defined conversion exists between two types, any user-defined conversions between those types are ignored. Specifically:

- If a pre-defined implicit conversion (§10.2) exists from type `S` to type `T`, all user-defined conversions (implicit or explicit) from `S` to `T` are ignored.

- If a pre-defined explicit conversion (§10.3) exists from type `S` to type `T`, any user-defined explicit conversions from `S` to `T` are ignored. Furthermore:
  - If either `S` or `T` is an interface type, user-defined implicit conversions from `S` to `T` are ignored.
  - Otherwise, user-defined implicit conversions from `S` to `T` are still considered.

For all types but `object`, the operators declared by the `Convertible<T>` type above do not conflict with pre-defined conversions.

> *Example*:
>
> ```
> C#
> ```
>
> ```
> void F(int i, Convertible<int> n)
> {
>     i = n;                    // Error
>     i = (int)n;               // User-defined explicit conversion
>     n = i;                    // User-defined implicit conversion
>     n = (Convertible<int>)i;  // User-defined implicit conversion
> }
> ```
>
> However, for type `object`, pre-defined conversions hide the user-defined conversions in all cases but one:
>
> ```
> C#
> ```
>
> ```
> void F(object o, Convertible<object> n)
> {
>     o = n;                        // Pre-defined boxing conversion
>     o = (object)n;                // Pre-defined boxing conversion
>     n = o;                        // User-defined implicit conversion
>     n = (Convertible<object>)o;   // Pre-defined unboxing conversion
> }
> ```
>
> *end example*

User-defined conversions are not allowed to convert from or to *interface_type*s. In particular, this restriction ensures that no user-defined transformations occur when converting to an *interface_type*, and that a conversion to an *interface_type* succeeds only if the `object` being converted actually implements the specified *interface_type*.

The signature of a conversion operator consists of the source type and the target type. (This is the only form of member for which the return type participates in the signature.) The implicit or explicit classification of a conversion operator is not part of the

operator's signature. Thus, a class or struct cannot declare both an implicit and an explicit conversion operator with the same source and target types.

> *Note*: In general, user-defined implicit conversions should be designed to never throw exceptions and never lose information. If a user-defined conversion can give rise to exceptions (for example, because the source argument is out of range) or loss of information (such as discarding high-order bits), then that conversion should be defined as an explicit conversion. *end note*

> *Example*: In the following code

```
C#
```

```
public struct Digit
{
    byte value;

    public Digit(byte value)
    {
        if (value < 0 || value > 9)
        {
            throw new ArgumentException();
        }
        this.value = value;
    }

    public static implicit operator byte(Digit d) => d.value;
    public static explicit operator Digit(byte b) => new Digit(b);
}
```

> the conversion from `Digit` to `byte` is implicit because it never throws exceptions or loses information, but the conversion from `byte` to `Digit` is explicit since `Digit` can only represent a subset of the possible values of a `byte`.
>
> *end example*

15.11 Instance constructors

15.11.1 General

An ***instance constructor*** is a member that implements the actions required to initialize an instance of a class. Instance constructors are declared using *constructor_declarations*:

```
ANTLR
```

```
constructor_declaration
    : attributes? constructor_modifier* constructor_declarator
constructor_body
    ;

constructor_modifier
    : 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'extern'
    | unsafe_modifier   // unsafe code support
    ;

constructor_declarator
    : identifier '(' formal_parameter_list? ')' constructor_initializer?
    ;

constructor_initializer
    : ':' 'base' '(' argument_list? ')'
    | ':' 'this' '(' argument_list? ')'
    ;

constructor_body
    : block
    | '=>' expression ';'
    | ';'
    ;
```

unsafe_modifier (§23.2) is only available in unsafe code (§23).

A *constructor_declaration* may include a set of *attributes* (§22), any one of the permitted kinds of declared accessibility (§15.3.6), and an `extern` (§15.6.8) modifier. A constructor declaration is not permitted to include the same modifier multiple times.

The *identifier* of a *constructor_declarator* shall name the class in which the instance constructor is declared. If any other name is specified, a compile-time error occurs.

The optional *formal_parameter_list* of an instance constructor is subject to the same rules as the *formal_parameter_list* of a method (§15.6). As the `this` modifier for parameters only applies to extension methods (§15.6.10), no parameter in a constructor's *formal_parameter_list* shall contain the `this` modifier. The formal parameter list defines the signature (§7.6) of an instance constructor and governs the process whereby overload resolution (§12.6.4) selects a particular instance constructor in an invocation.

Each of the types referenced in the *formal_parameter_list* of an instance constructor shall be at least as accessible as the constructor itself (§7.5.5).

The optional *constructor_initializer* specifies another instance constructor to invoke before executing the statements given in the *constructor_body* of this instance constructor. This is described further in §15.11.2.

When a constructor declaration includes an `extern` modifier, the constructor is said to be an ***external constructor***. Because an external constructor declaration provides no actual implementation, its *constructor_body* consists of a semicolon. For all other constructors, the *constructor_body* consists of either

- a *block*, which specifies the statements to initialize a new instance of the class; or
- an expression body, which consists of `=>` followed by an *expression* and a semicolon, and denotes a single expression to initialize a new instance of the class.

A *constructor_body* that is a *block* or expression body corresponds exactly to the *block* of an instance method with a `void` return type (§15.6.11).

Instance constructors are not inherited. Thus, a class has no instance constructors other than those actually declared in the class, with the exception that if a class contains no instance constructor declarations, a default instance constructor is automatically provided (§15.11.5).

Instance constructors are invoked by *object_creation_expression*s (§12.8.16.2) and through *constructor_initializer*s.

15.11.2 Constructor initializers

All instance constructors (except those for class `object`) implicitly include an invocation of another instance constructor immediately before the *constructor_body*. The constructor to implicitly invoke is determined by the *constructor_initializer*:

- An instance constructor initializer of the form `base(` *argument_list* `)` (where *argument_list* is optional) causes an instance constructor from the direct base class to be invoked. That constructor is selected using *argument_list* and the overload resolution rules of §12.6.4. The set of candidate instance constructors consists of all the accessible instance constructors of the direct base class. If this set is empty, or if a single best instance constructor cannot be identified, a compile-time error occurs.
- An instance constructor initializer of the form `this(` *argument_list* `)` (where *argument_list* is optional) invokes another instance constructor from the same class. The constructor is selected using *argument_list* and the overload resolution rules of §12.6.4. The set of candidate instance constructors consists of all instance constructors declared in the class itself. If the resulting set of applicable instance

constructors is empty, or if a single best instance constructor cannot be identified, a compile-time error occurs. If an instance constructor declaration invokes itself through a chain of one or more constructor initializers, a compile-time error occurs.

If an instance constructor has no constructor initializer, a constructor initializer of the form `base()` is implicitly provided.

> *Note*: Thus, an instance constructor declaration of the form
>
> ```C#
> C(...) {...}
> ```
>
> is exactly equivalent to
>
> ```C#
> C(...) : base() {...}
> ```
>
> *end note*

The scope of the parameters given by the *formal_parameter_list* of an instance constructor declaration includes the constructor initializer of that declaration. Thus, a constructor initializer is permitted to access the parameters of the constructor.

> *Example*:
>
> ```C#
> class A
> {
>     public A(int x, int y) {}
> }
>
> class B: A
> {
>     public B(int x, int y) : base(x + y, x - y) {}
> }
> ```
>
> *end example*

An instance constructor initializer cannot access the instance being created. Therefore it is a compile-time error to reference this in an argument expression of the constructor

initializer, as it is a compile-time error for an argument expression to reference any instance member through a *simple_name*.

15.11.3 Instance variable initializers

When an instance constructor has no constructor initializer, or it has a constructor initializer of the form `base(...)`, that constructor implicitly performs the initializations specified by the *variable_initializer*s of the instance fields declared in its class. This corresponds to a sequence of assignments that are executed immediately upon entry to the constructor and before the implicit invocation of the direct base class constructor. The variable initializers are executed in the textual order in which they appear in the class declaration (§15.5.6).

15.11.4 Constructor execution

Variable initializers are transformed into assignment statements, and these assignment statements are executed *before* the invocation of the base class instance constructor. This ordering ensures that all instance fields are initialized by their variable initializers before *any* statements that have access to that instance are executed.

> *Example*: Given the following:

```C#
class A
{
    public A()
    {
        PrintFields();
    }

    public virtual void PrintFields() {}
}
class B: A
{
    int x = 1;
    int y;

    public B()
    {
        y = -1;
    }

    public override void PrintFields() =>
        Console.WriteLine($"x = {x}, y = {y}");
}
```

when new `B()` is used to create an instance of `B`, the following output is produced:

Console

```
x = 1, y = 0
```

The value of `x` is 1 because the variable initializer is executed before the base class instance constructor is invoked. However, the value of `y` is 0 (the default value of an `int`) because the assignment to `y` is not executed until after the base class constructor returns. It is useful to think of instance variable initializers and constructor initializers as statements that are automatically inserted before the *constructor_body*. The example

C#

```
class A
{
    int x = 1, y = -1, count;

    public A()
    {
        count = 0;
    }

    public A(int n)
    {
        count = n;
    }
}

class B : A
{
    double sqrt2 = Math.Sqrt(2.0);
    ArrayList items = new ArrayList(100);
    int max;

    public B(): this(100)
    {
        items.Add("default");
    }

    public B(int n) : base(n - 1)
    {
        max = n;
    }
}
```

contains several variable initializers; it also contains constructor initializers of both forms (`base` and `this`). The example corresponds to the code shown below, where

each comment indicates an automatically inserted statement (the syntax used for the automatically inserted constructor invocations isn't valid, but merely serves to illustrate the mechanism).

C#

```
class A
{
    int x, y, count;
    public A()
    {
        x = 1;      // Variable initializer
        y = -1;     // Variable initializer
        object();   // Invoke object() constructor
        count = 0;
    }

    public A(int n)
    {
        x = 1;      // Variable initializer
        y = -1;     // Variable initializer
        object();   // Invoke object() constructor
        count = n;
    }
}

class B : A
{
    double sqrt2;
    ArrayList items;
    int max;
    public B() : this(100)
    {
        B(100);                     // Invoke B(int) constructor
        items.Add("default");
    }

    public B(int n) : base(n - 1)
    {
        sqrt2 = Math.Sqrt(2.0);     // Variable initializer
        items = new ArrayList(100); // Variable initializer
        A(n - 1);                   // Invoke A(int) constructor
        max = n;
    }
}
```

end example

15.11.5 Default constructors

If a class contains no instance constructor declarations, a default instance constructor is automatically provided. That default constructor simply invokes a constructor of the direct base class, as if it had a constructor initializer of the form `base()`. If the class is abstract then the declared accessibility for the default constructor is protected. Otherwise, the declared accessibility for the default constructor is public.

> *Note*: Thus, the default constructor is always of the form
>
> ```C#
> protected C(): base() {}
> ```
>
> or
>
> ```C#
> public C(): base() {}
> ```
>
> where `C` is the name of the class.
>
> *end note*

If overload resolution is unable to determine a unique best candidate for the base-class constructor initializer then a compile-time error occurs.

> *Example*: In the following code
>
> ```C#
> class Message
> {
>     object sender;
>     string text;
> }
> ```
>
> a default constructor is provided because the class contains no instance constructor declarations. Thus, the example is precisely equivalent to
>
> ```C#
> class Message
> {
>     object sender;
>     string text;
> ```

```
    public Message() : base() {}
}
```

end example

15.12 Static constructors

A ***static constructor*** is a member that implements the actions required to initialize a closed class. Static constructors are declared using *static_constructor_declarations*:

ANTLR

```
static_constructor_declaration
    : attributes? static_constructor_modifiers identifier '(' ')'
        static_constructor_body
    ;

static_constructor_modifiers
    : 'static'
    | 'static' 'extern' unsafe_modifier?
    | 'static' unsafe_modifier 'extern'?
    | 'extern' 'static' unsafe_modifier?
    | 'extern' unsafe_modifier 'static'
    | unsafe_modifier 'static' 'extern'?
    | unsafe_modifier 'extern' 'static'
    ;

static_constructor_body
    : block
    | '=>' expression ';'
    | ';'
    ;
```

unsafe_modifier (§23.2) is only available in unsafe code (§23).

A *static_constructor_declaration* may include a set of *attributes* (§22) and an `extern` modifier (§15.6.8).

The *identifier* of a *static_constructor_declaration* shall name the class in which the static constructor is declared. If any other name is specified, a compile-time error occurs.

When a static constructor declaration includes an `extern` modifier, the static constructor is said to be an ***external static constructor***. Because an external static constructor declaration provides no actual implementation, its *static_constructor_body* consists of a semicolon. For all other static constructor declarations, the *static_constructor_body* consists of either

- a *block*, which specifies the statements to execute in order to initialize the class; or
- an expression body, which consists of `=>` followed by an *expression* and a semicolon, and denotes a single expression to execute in order to initialize the class.

A *static_constructor_body* that is a *block* or expression body corresponds exactly to the *method_body* of a static method with a `void` return type (§15.6.11).

Static constructors are not inherited, and cannot be called directly.

The static constructor for a closed class executes at most once in a given application domain. The execution of a static constructor is triggered by the first of the following events to occur within an application domain:

- An instance of the class is created.
- Any of the static members of the class are referenced.

If a class contains the `Main` method (§7.1) in which execution begins, the static constructor for that class executes before the `Main` method is called.

To initialize a new closed class type, first a new set of static fields (§15.5.2) for that particular closed type is created. Each of the static fields is initialized to its default value (§15.5.5). Next, the static field initializers (§15.5.6.2) are executed for those static fields. Finally, the static constructor is executed.

> *Example*: The example
>
> C#
>
> ```
> class Test
> {
>     static void Main()
>     {
>         A.F();
>         B.F();
>     }
> }
>
> class A
> {
>     static A()
>     {
>         Console.WriteLine("Init A");
>     }
>
>     public static void F()
>     {
>         Console.WriteLine("A.F");
> ```

```
    }
}

class B
{
    static B()
    {
        Console.WriteLine("Init B");
    }

    public static void F()
    {
        Console.WriteLine("B.F");
    }
}
```

must produce the output:

```
Console

Init A
A.F
Init B
B.F
```

because the execution of `A`'s static constructor is triggered by the call to `A.F`, and the execution of `B`'s static constructor is triggered by the call to `B.F`.

end example

It is possible to construct circular dependencies that allow static fields with variable initializers to be observed in their default value state.

Example: The example

```
C#

class A
{
    public static int X;

    static A()
    {
        X = B.Y + 1;
    }
}

class B
{
```

```
    public static int Y = A.X + 1;

    static B() {}

    static void Main()
    {
        Console.WriteLine($"X = {A.X}, Y = {B.Y}");
    }
}
```

produces the output

```
Console
X = 1, Y = 2
```

To execute the `Main` method, the system first runs the initializer for `B.Y`, prior to class `B`'s static constructor. `Y`'s initializer causes `A`'s `static` constructor to be run because the value of `A.X` is referenced. The static constructor of `A` in turn proceeds to compute the value of `X`, and in doing so fetches the default value of `Y`, which is zero. `A.X` is thus initialized to 1. The process of running `A`'s static field initializers and static constructor then completes, returning to the calculation of the initial value of `Y`, the result of which becomes 2.

end example

Because the static constructor is executed exactly once for each closed constructed class type, it is a convenient place to enforce run-time checks on the type parameter that cannot be checked at compile-time via constraints (§15.2.5).

Example: The following type uses a static constructor to enforce that the type argument is an enum:

```
C#
class Gen<T> where T : struct
{
    static Gen()
    {
        if (!typeof(T).IsEnum)
        {
            throw new ArgumentException("T must be an enum");
        }
    }
}
```

> *end example*

15.13 Finalizers

> *Note*: In an earlier version of this specification, what is now referred to as a "finalizer" was called a "destructor". Experience has shown that the term "destructor" caused confusion and often resulted to incorrect expectations, especially to programmers knowing C++. In C++, a destructor is called in a determinate manner, whereas, in C#, a finalizer is not. To get determinate behavior from C#, one should use `Dispose`. *end note*

A ***finalizer*** is a member that implements the actions required to finalize an instance of a class. A finalizer is declared using a *finalizer_declaration*:

ANTLR

```
finalizer_declaration
    : attributes? '~' identifier '(' ')' finalizer_body
    | attributes? 'extern' unsafe_modifier? '~' identifier '(' ')'
      finalizer_body
    | attributes? unsafe_modifier 'extern'? '~' identifier '(' ')'
      finalizer_body
    ;

finalizer_body
    : block
    | '=>' expression ';'
    | ';'
    ;
```

unsafe_modifier (§23.2) is only available in unsafe code (§23).

A *finalizer_declaration* may include a set of *attributes* (§22).

The *identifier* of a *finalizer_declarator* shall name the class in which the finalizer is declared. If any other name is specified, a compile-time error occurs.

When a finalizer declaration includes an `extern` modifier, the finalizer is said to be an ***external finalizer***. Because an external finalizer declaration provides no actual implementation, its *finalizer_body* consists of a semicolon. For all other finalizers, the *finalizer_body* consists of either

- a *block*, which specifies the statements to execute in order to finalize an instance of the class.

- or an expression body, which consists of `=>` followed by an *expression* and a semicolon, and denotes a single expression to execute in order to finalize an instance of the class.

A *finalizer_body* that is a *block* or expression body corresponds exactly to the *method_body* of an instance method with a `void` return type (§15.6.11).

Finalizers are not inherited. Thus, a class has no finalizers other than the one that may be declared in that class.

> *Note*: Since a finalizer is required to have no parameters, it cannot be overloaded, so a class can have, at most, one finalizer. *end note*

Finalizers are invoked automatically, and cannot be invoked explicitly. An instance becomes eligible for finalization when it is no longer possible for any code to use that instance. Execution of the finalizer for the instance may occur at any time after the instance becomes eligible for finalization (§7.9). When an instance is finalized, the finalizers in that instance's inheritance chain are called, in order, from most derived to least derived. A finalizer may be executed on any thread. For further discussion of the rules that govern when and how a finalizer is executed, see §7.9.

> *Example*: The output of the example
>
> C#
>
> ```
> class A
> {
>     ~A()
>     {
>         Console.WriteLine("A's finalizer");
>     }
> }
>
> class B : A
> {
>     ~B()
>     {
>         Console.WriteLine("B's finalizer");
>     }
> }
>
> class Test
> {
>     static void Main()
>     {
>         B b = new B();
>         b = null;
>         GC.Collect();
> ```

```
        GC.WaitForPendingFinalizers();
    }
}
```

is

```
Console
```

```
B's finalizer
A's finalizer
```

since finalizers in an inheritance chain are called in order, from most derived to least derived.

end example

Finalizers are implemented by overriding the virtual method `Finalize` on `System.Object`. C# programs are not permitted to override this method or call it (or overrides of it) directly.

Example: For instance, the program

```
C#
```

```
class A
{
    override protected void Finalize() {}  // Error
    public void F()
    {
        this.Finalize();                   // Error
    }
}
```

contains two errors.

end example

The compiler behaves as if this method, and overrides of it, do not exist at all.

Example: Thus, this program:

```
C#
```

```
class A
{
```

```
    void Finalize() {}  // Permitted
}
```

is valid and the method shown hides `System.Object`'s `Finalize` method.

end example

For a discussion of the behavior when an exception is thrown from a finalizer, see §21.4.

15.14 Iterators

15.14.1 General

A function member (§12.6) implemented using an iterator block (§13.3) is called an ***iterator***.

An iterator block may be used as the body of a function member as long as the return type of the corresponding function member is one of the enumerator interfaces (§15.14.2) or one of the enumerable interfaces (§15.14.3). It may occur as a *method_body*, *operator_body* or *accessor_body*, whereas events, instance constructors, static constructors and finalizer shall not be implemented as iterators.

When a function member is implemented using an iterator block, it is a compile-time error for the formal parameter list of the function member to specify any `in`, `out`, or `ref` parameters, or an parameter of a `ref struct` type.

15.14.2 Enumerator interfaces

The ***enumerator interfaces*** are the non-generic interface `System.Collections.IEnumerator` and all instantiations of the generic interface `System.Collections.Generic.IEnumerator<T>`. For the sake of brevity, in this subclause and its siblings these interfaces are referenced as `IEnumerator` and `IEnumerator<T>`, respectively.

15.14.3 Enumerable interfaces

The ***enumerable interfaces*** are the non-generic interface `System.Collections.IEnumerable` and all instantiations of the generic interface `System.Collections.Generic.IEnumerable<T>`. For the sake of brevity, in this subclause

and its siblings these interfaces are referenced as `IEnumerable` and `IEnumerable<T>`, respectively.

15.14.4 Yield type

An iterator produces a sequence of values, all of the same type. This type is called the ***yield type*** of the iterator.

- The yield type of an iterator that returns `IEnumerator` or `IEnumerable` is `object`.
- The yield type of an iterator that returns `IEnumerator<T>` or `IEnumerable<T>` is `T`.

15.14.5 Enumerator objects

15.14.5.1 General

When a function member returning an enumerator interface type is implemented using an iterator block, invoking the function member does not immediately execute the code in the iterator block. Instead, an ***enumerator object*** is created and returned. This object encapsulates the code specified in the iterator block, and execution of the code in the iterator block occurs when the enumerator object's `MoveNext` method is invoked. An enumerator object has the following characteristics:

- It implements `IEnumerator` and `IEnumerator<T>`, where `T` is the yield type of the iterator.
- It implements `System.IDisposable`.
- It is initialized with a copy of the argument values (if any) and instance value passed to the function member.
- It has four potential states, **before**, **running**, **suspended**, and **after**, and is initially in the **before** state.

An enumerator object is typically an instance of a compiler-generated enumerator class that encapsulates the code in the iterator block and implements the enumerator interfaces, but other methods of implementation are possible. If an enumerator class is generated by the compiler, that class will be nested, directly or indirectly, in the class containing the function member, it will have private accessibility, and it will have a name reserved for compiler use (§6.4.3).

An enumerator object may implement more interfaces than those specified above.

The following subclauses describe the required behavior of the `MoveNext`, `Current`, and `Dispose` members of the `IEnumerator` and `IEnumerator<T>` interface implementations

provided by an enumerator object.

Enumerator objects do not support the `IEnumerator.Reset` method. Invoking this method causes a `System.NotSupportedException` to be thrown.

15.14.5.2 The MoveNext method

The `MoveNext` method of an enumerator object encapsulates the code of an iterator block. Invoking the `MoveNext` method executes code in the iterator block and sets the `Current` property of the enumerator object as appropriate. The precise action performed by `MoveNext` depends on the state of the enumerator object when `MoveNext` is invoked:

- If the state of the enumerator object is **before**, invoking `MoveNext`:
  - Changes the state to **running**.
  - Initializes the parameters (including `this`) of the iterator block to the argument values and instance value saved when the enumerator object was initialized.
  - Executes the iterator block from the beginning until execution is interrupted (as described below).
- If the state of the enumerator object is **running**, the result of invoking `MoveNext` is unspecified.
- If the state of the enumerator object is **suspended**, invoking MoveNext:
  - Changes the state to **running**.
  - Restores the values of all local variables and parameters (including `this`) to the values saved when execution of the iterator block was last suspended.

    > *Note*: The contents of any objects referenced by these variables may have changed since the previous call to `MoveNext`. *end note*

  - Resumes execution of the iterator block immediately following the yield return statement that caused the suspension of execution and continues until execution is interrupted (as described below).
- If the state of the enumerator object is **after**, invoking `MoveNext` returns false.

When `MoveNext` executes the iterator block, execution can be interrupted in four ways: By a `yield return` statement, by a `yield break` statement, by encountering the end of the iterator block, and by an exception being thrown and propagated out of the iterator block.

- When a `yield return` statement is encountered (§9.4.4.20):

    - The expression given in the statement is evaluated, implicitly converted to the yield type, and assigned to the `Current` property of the enumerator object.
    - Execution of the iterator body is suspended. The values of all local variables and parameters (including `this`) are saved, as is the location of this `yield return` statement. If the `yield return` statement is within one or more `try` blocks, the associated finally blocks are *not* executed at this time.
    - The state of the enumerator object is changed to **suspended**.
    - The `MoveNext` method returns `true` to its caller, indicating that the iteration successfully advanced to the next value.
- When a `yield break` statement is encountered (§9.4.4.20):
    - If the `yield break` statement is within one or more `try` blocks, the associated `finally` blocks are executed.
    - The state of the enumerator object is changed to **after**.
    - The `MoveNext` method returns `false` to its caller, indicating that the iteration is complete.
- When the end of the iterator body is encountered:
    - The state of the enumerator object is changed to **after**.
    - The `MoveNext` method returns `false` to its caller, indicating that the iteration is complete.
- When an exception is thrown and propagated out of the iterator block:
    - Appropriate `finally` blocks in the iterator body will have been executed by the exception propagation.
    - The state of the enumerator object is changed to **after**.
    - The exception propagation continues to the caller of the `MoveNext` method.

15.14.5.3 The Current property

An enumerator object's `Current` property is affected by `yield return` statements in the iterator block.

When an enumerator object is in the **suspended** state, the value of `Current` is the value set by the previous call to `MoveNext`. When an enumerator object is in the **before**, **running**, or **after** states, the result of accessing `Current` is unspecified.

For an iterator with a yield type other than `object`, the result of accessing `Current` through the enumerator object's `IEnumerable` implementation corresponds to accessing `Current` through the enumerator object's `IEnumerator<T>` implementation and casting the result to `object`.

15.14.5.4 The Dispose method

The `Dispose` method is used to clean up the iteration by bringing the enumerator object to the **after** state.

- If the state of the enumerator object is **before**, invoking `Dispose` changes the state to **after**.
- If the state of the enumerator object is **running**, the result of invoking `Dispose` is unspecified.
- If the state of the enumerator object is **suspended**, invoking `Dispose`:
  - Changes the state to **running**.
  - Executes any finally blocks as if the last executed `yield return` statement were a `yield break` statement. If this causes an exception to be thrown and propagated out of the iterator body, the state of the enumerator object is set to **after** and the exception is propagated to the caller of the `Dispose` method.
  - Changes the state to **after**.
- If the state of the enumerator object is **after**, invoking `Dispose` has no affect.

15.14.6 Enumerable objects

15.14.6.1 General

When a function member returning an enumerable interface type is implemented using an iterator block, invoking the function member does not immediately execute the code in the iterator block. Instead, an ***enumerable object*** is created and returned. The enumerable object's `GetEnumerator` method returns an enumerator object that encapsulates the code specified in the iterator block, and execution of the code in the iterator block occurs when the enumerator object's `MoveNext` method is invoked. An enumerable object has the following characteristics:

- It implements `IEnumerable` and `IEnumerable<T>`, where `T` is the yield type of the iterator.
- It is initialized with a copy of the argument values (if any) and instance value passed to the function member.

An enumerable object is typically an instance of a compiler-generated enumerable class that encapsulates the code in the iterator block and implements the enumerable interfaces, but other methods of implementation are possible. If an enumerable class is generated by the compiler, that class will be nested, directly or indirectly, in the class containing the function member, it will have private accessibility, and it will have a name reserved for compiler use (§6.4.3).

An enumerable object may implement more interfaces than those specified above.

> *Note*: For example, an enumerable object may also implement `IEnumerator` and `IEnumerator<T>`, enabling it to serve as both an enumerable and an enumerator. Typically, such an implementation would return its own instance (to save allocations) from the first call to `GetEnumerator`. Subsequent invocations of `GetEnumerator`, if any, would return a new class instance, typically of the same class, so that calls to different enumerator instances will not affect each other. It cannot return the same instance even if the previous enumerator has already enumerated past the end of the sequence, since all future calls to an exhausted enumerator must throw exceptions. *end note*

15.14.6.2 The GetEnumerator method

An enumerable object provides an implementation of the `GetEnumerator` methods of the `IEnumerable` and `IEnumerable<T>` interfaces. The two `GetEnumerator` methods share a common implementation that acquires and returns an available enumerator object. The enumerator object is initialized with the argument values and instance value saved when the enumerable object was initialized, but otherwise the enumerator object functions as described in §15.14.5.

15.15 Async Functions

15.15.1 General

A method (§15.6) or anonymous function (§12.19) with the `async` modifier is called an ***async function***. In general, the term ***async*** is used to describe any kind of function that has the `async` modifier.

It is a compile-time error for the formal parameter list of an async function to specify any `in`, `out`, or `ref` parameters, or any parameter of a `ref struct` type.

The *return_type* of an async method shall be either `void` or a ***task type***. For an async method that produces a result value, a task type shall be generic. For an async method that does not produce a result value, a task type shall not be generic. Such types are referred to in this specification as `«TaskType»<T>` and `«TaskType»`, respectively. The Standard library type `System.Threading.Tasks.Task` and types constructed from `System.Threading.Tasks.Task<TResult>` are task types, as well as a class, struct or interface type that is associated with a ***task builder type*** via the attribute `System.Runtime.CompilerServices.AsyncMethodBuilderAttribute`. Such types are referred

to in this specification as `«TaskBuilderType»<T>` and `«TaskBuilderType»`. A task type can have at most one type parameter and cannot be nested in a generic type.

An async method returning a task type is said to be ***task-returning***.

Task types can vary in their exact definition, but from the language's point of view, a task type is in one of the states *incomplete*, *succeeded* or *faulted*. A *faulted* task records a pertinent exception. A *succeeded* `«TaskType»<T>` records a result of type `T`. Task types are awaitable, and tasks can therefore be the operands of await expressions (§12.9.8).

> *Example*: The task type `MyTask<T>` is associated with the task builder type `MyTaskMethodBuilder<T>` and the awaiter type `Awaiter<T>`:
>
> C#
>
> ```
> using System.Runtime.CompilerServices;
> [AsyncMethodBuilder(typeof(MyTaskMethodBuilder<>))]
> class MyTask<T>
> {
>     public Awaiter<T> GetAwaiter() { ... }
> }
>
> class Awaiter<T> : INotifyCompletion
> {
>     public void OnCompleted(Action completion) { ... }
>     public bool IsCompleted { get; }
>     public T GetResult() { ... }
> }
> ```
>
> *end example*

A task builder type is a class or struct type that corresponds to a specific task type (§15.15.2). The task builder type shall exactly match the declared accessibility of its corresponding task type.

> *Note:* If the task type is declared `internal`, the the corresponding builder type must also be declared `internal` and be defined in the same assembly. If the task type is nested inside another type, the task buider type must also be nested in that same type. *end note*

An async function has the ability to suspend evaluation by means of await expressions (§12.9.8) in its body. Evaluation may later be resumed at the point of the suspending await expression by means of a ***resumption delegate***. The resumption delegate is of type `System.Action`, and when it is invoked, evaluation of the async function invocation

will resume from the await expression where it left off. The ***current caller*** of an async function invocation is the original caller if the function invocation has never been suspended or the most recent caller of the resumption delegate otherwise.

15.15.2 Task-type builder pattern

A task builder type can have at most one type parameter and cannot be nested in a generic type. A task builder type shall have the following members (for non-generic task builder types, `SetResult` has no parameters) with declared `public` accessibility:

C#

```
class «TaskBuilderType»<T>
{
    public static «TaskBuilderType»<T> Create();
    public void Start<TStateMachine>(ref TStateMachine stateMachine)
                where TStateMachine : IAsyncStateMachine;
    public void SetStateMachine(IAsyncStateMachine stateMachine);
    public void SetException(Exception exception);
    public void SetResult(T result);
    public void AwaitOnCompleted<TAwaiter, TStateMachine>(
        ref TAwaiter awaiter, ref TStateMachine stateMachine)
        where TAwaiter : INotifyCompletion
        where TStateMachine : IAsyncStateMachine;
    public void AwaitUnsafeOnCompleted<TAwaiter, TStateMachine>(
        ref TAwaiter awaiter, ref TStateMachine stateMachine)
        where TAwaiter : ICriticalNotifyCompletion
        where TStateMachine : IAsyncStateMachine;
    public «TaskType»<T> Task { get; }
}
```

The compiler generates code that uses the «TaskBuilderType» to implement the semantics of suspending and resuming the evaluation of the async function. The compiler uses the «TaskBuilderType» as follows:

- `«TaskBuilderType».Create()` is invoked to create an instance of the «TaskBuilderType», named `builder` in this list.
- `builder.Start(ref stateMachine)` is invoked to associate the builder with a compiler-generated state machine instance, `stateMachine`.
  - The builder shall call `stateMachine.MoveNext()` either in `Start()` or after `Start()` has returned to advance the state machine.
- After `Start()` returns, the `async` method invokes `builder.Task` for the task to return from the async method.
- Each call to `stateMachine.MoveNext()` will advance the state machine.

- If the state machine completes successfully, `builder.SetResult()` is called, with the method return value, if any.
- Otherwise, if an exception, `e` is thrown in the state machine, `builder.SetException(e)` is called.
- If the state machine reaches an `await expr` expression, `expr.GetAwaiter()` is invoked.
- If the awaiter implements `ICriticalNotifyCompletion` and `IsCompleted` is false, the state machine invokes `builder.AwaitUnsafeOnCompleted(ref awaiter, ref stateMachine)`.
  - `AwaitUnsafeOnCompleted()` should call `awaiter.UnsafeOnCompleted(action)` with an `Action` that calls `stateMachine.MoveNext()` when the awaiter completes.
- Otherwise, the state machine invokes `builder.AwaitOnCompleted(ref awaiter, ref stateMachine)`.
  - `AwaitOnCompleted()` should call `awaiter.OnCompleted(action)` with an `Action` that calls `stateMachine.MoveNext()` when the awaiter completes.
- `SetStateMachine(IAsyncStateMachine)` may be called by the compiler-generated `IAsyncStateMachine` implementation to identify the instance of the builder associated with a state machine instance, particularly for cases where the state machine is implemented as a value type.
  - If the builder calls `stateMachine.SetStateMachine(stateMachine)`, the `stateMachine` will call `builder.SetStateMachine(stateMachine)` on the *builder instance associated with* `stateMachine`.

> *Note*: For both `SetResult(T result)` and `«TaskType»<T> Task { get; }`, the parameter and argument respectively must be identity convertible to `T`. This allows a task-type builder to support types such as tuples, where two types that aren't the same are identity convertible. *end note*

15.15.3 Evaluation of a task-returning async function

Invocation of a task-returning async function causes an instance of the returned task type to be generated. This is called the ***return task*** of the async function. The task is initially in an *incomplete* state.

The async function body is then evaluated until it is either suspended (by reaching an await expression) or terminates, at which point control is returned to the caller, along with the return task.

When the body of the async function terminates, the return task is moved out of the incomplete state:

- If the function body terminates as the result of reaching a return statement or the end of the body, any result value is recorded in the return task, which is put into a *succeeded* state.
- If the function body terminates because of an uncaught `OperationCanceledException`, the exception is recorded in the return task which is put into the *canceled* state.
- If the function body terminates as the result of any other uncaught exception (§13.10.6) the exception is recorded in the return task which is put into a *faulted* state.

15.15.4 Evaluation of a void-returning async function

If the return type of the async function is `void`, evaluation differs from the above in the following way: Because no task is returned, the function instead communicates completion and exceptions to the current thread's ***synchronization context***. The exact definition of synchronization context is implementation-dependent, but is a representation of "where" the current thread is running. The synchronization context is notified when evaluation of a `void`-returning async function commences, completes successfully, or causes an uncaught exception to be thrown.

This allows the context to keep track of how many `void`-returning async functions are running under it, and to decide how to propagate exceptions coming out of them.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

16 Structs

Article • 02/07/2024

16.1 General

Structs are similar to classes in that they represent data structures that can contain data members and function members. However, unlike classes, structs are value types and do not require heap allocation. A variable of a `struct` type directly contains the data of the `struct`, whereas a variable of a class type contains a reference to the data, the latter known as an object.

> *Note*: Structs are particularly useful for small data structures that have value semantics. Complex numbers, points in a coordinate system, or key-value pairs in a dictionary are all good examples of structs. Key to these data structures is that they have few data members, that they do not require use of inheritance or reference semantics, rather they can be conveniently implemented using value semantics where assignment copies the value instead of the reference. *end note*

As described in §8.3.5, the simple types provided by C#, such as `int`, `double`, and `bool`, are, in fact, all struct types.

16.2 Struct declarations

16.2.1 General

A *struct_declaration* is a *type_declaration* (§14.7) that declares a new struct:

ANTLR

```
struct_declaration
    : attributes? struct_modifier* 'ref'? 'partial'? 'struct'
      identifier type_parameter_list? struct_interfaces?
      type_parameter_constraints_clause* struct_body ';'?
    ;
```

A *struct_declaration* consists of an optional set of *attributes* (§22), followed by an optional set of *struct_modifiers* (§16.2.2), followed by an optional `ref` modifier (§16.2.3), followed by an optional partial modifier (§15.2.7), followed by the keyword `struct` and an *identifier* that names the struct, followed by an optional *type_parameter_list*

specification (§15.2.3), followed by an optional *struct_interfaces* specification (§16.2.5), followed by an optional *type_parameter_constraints-clauses* specification (§15.2.5), followed by a *struct_body* (§16.2.6), optionally followed by a semicolon.

A struct declaration shall not supply a *type_parameter_constraints_clauses* unless it also supplies a *type_parameter_list*.

A struct declaration that supplies a *type_parameter_list* is a generic struct declaration. Additionally, any struct nested inside a generic class declaration or a generic struct declaration is itself a generic struct declaration, since type arguments for the containing type shall be supplied to create a constructed type (§8.4).

A struct declaration that includes a `ref` keyword shall not have a *struct_interfaces* part.

16.2.2 Struct modifiers

A *struct_declaration* may optionally include a sequence of *struct_modifiers*:

ANTLR

```
struct_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'readonly'
    | unsafe_modifier   // unsafe code support
    ;
```

unsafe_modifier (§23.2) is only available in unsafe code (§23).

It is a compile-time error for the same modifier to appear multiple times in a struct declaration.

Except for `readonly`, the modifiers of a struct declaration have the same meaning as those of a class declaration (§15.2.2).

The `readonly` modifier indicates that the *struct_declaration* declares a type whose instances are immutable.

A readonly struct has the following constraints:

- Each of its instance fields shall also be declared `readonly`.
- None of its instance properties shall have a *set_accessor_declaration* (§15.7.3).
- It shall not declare any field-like events (§15.8.2).

When an instance of a readonly struct is passed to a method, its `this` is treated like an `in` argument/parameter, which disallows write access to any instance fields (except by constructors).

16.2.3 Ref modifier

The `ref` modifier indicates that the *struct_declaration* declares a type whose instances are allocated on the execution stack. These types are called ***ref struct*** types. The `ref` modifier declares that instances may contain ref-like fields, and shall not be copied out of its safe-context (§16.4.12). The rules for determining the safe context of a ref struct are described in §16.4.12.

It is a compile-time error if a ref struct type is used in any of the following contexts:

- As the element type of an array.
- As the declared type of a field of a class or a struct that does not have the `ref` modifier.
- Being boxed to `System.ValueType` or `System.Object`.
- As a type argument.
- As the type of a tuple element.
- An async method.
- An iterator.
- There is no conversion from a `ref struct` type to the type `object` or the type `System.ValueType`.
- A `ref struct` type shall not be declared to implement any interface.
- An instance method declared in `object` or in `System.ValueType` but not overridden in a `ref struct` type shall not be called with a receiver of that `ref struct` type.
- An instance method of a `ref struct` type shall not be captured by method group conversion to a delegate type.
- A ref struct shall not be captured by a lambda expression or a local function.

> *Note*: A `ref struct` shall not declare `async` instance methods nor use a `yield return` or `yield break` statement within an instance method, because the implicit `this` parameter cannot be used in those contexts. *end note*

These constraints ensure that a variable of `ref struct` type does not refer to stack memory that is no longer valid, or to variables that are no longer valid.

16.2.4 Partial modifier

The `partial` modifier indicates that this *struct_declaration* is a partial type declaration. Multiple partial struct declarations with the same name within an enclosing namespace or type declaration combine to form one struct declaration, following the rules specified in §15.2.7.

16.2.5 Struct interfaces

A struct declaration may include a *struct_interfaces* specification, in which case the struct is said to directly implement the given interface types. For a constructed struct type, including a nested type declared within a generic type declaration (§15.3.9.7), each implemented interface type is obtained by substituting, for each *type_parameter* in the given interface, the corresponding *type_argument* of the constructed type.

ANTLR

```
struct_interfaces
    : ':' interface_type_list
    ;
```

The handling of interfaces on multiple parts of a partial struct declaration (§15.2.7) are discussed further in §15.2.4.3.

Interface implementations are discussed further in §18.6.

16.2.6 Struct body

The *struct_body* of a struct defines the members of the struct.

ANTLR

```
struct_body
    : '{' struct_member_declaration* '}'
    ;
```

16.3 Struct members

The members of a struct consist of the members introduced by its *struct_member_declaration*s and the members inherited from the type `System.ValueType`.

ANTLR

```
struct_member_declaration
    : constant_declaration
```

```
    | field_declaration
    | method_declaration
    | property_declaration
    | event_declaration
    | indexer_declaration
    | operator_declaration
    | constructor_declaration
    | static_constructor_declaration
    | type_declaration
    | fixed_size_buffer_declaration   // unsafe code support
    ;
```

fixed_size_buffer_declaration (§23.8.2) is only available in unsafe code (§23).

> *Note*: All kinds of *class_member_declaration*s except *finalizer_declaration* are also *struct_member_declaration*s. *end note*

Except for the differences noted in §16.4, the descriptions of class members provided in §15.3 through §15.12 apply to struct members as well.

16.4 Class and struct differences

16.4.1 General

Structs differ from classes in several important ways:

- Structs are value types (§16.4.2).
- All struct types implicitly inherit from the class `System.ValueType` (§16.4.3).
- Assignment to a variable of a struct type creates a *copy* of the value being assigned (§16.4.4).
- The default value of a struct is the value produced by setting all fields to their default value (§16.4.5).
- Boxing and unboxing operations are used to convert between a struct type and certain reference types (§16.4.6).
- The meaning of `this` is different within struct members (§16.4.7).
- Instance field declarations for a struct are not permitted to include variable initializers (§16.4.8).
- A struct is not permitted to declare a parameterless instance constructor (§16.4.9).
- A struct is not permitted to declare a finalizer.

16.4.2 Value semantics

Structs are value types (§8.3) and are said to have value semantics. Classes, on the other hand, are reference types (§8.2) and are said to have reference semantics.

A variable of a struct type directly contains the data of the struct, whereas a variable of a class type contains a reference to an object that contains the data. When a struct `B` contains an instance field of type `A` and `A` is a struct type, it is a compile-time error for `A` to depend on `B` or a type constructed from `B`. `A struct X` *directly depends on* a struct `Y` if `X` contains an instance field of type `Y`. Given this definition, the complete set of structs upon which a struct depends is the transitive closure of the *directly depends on* relationship.

> *Example*:
>
> C#
>
> ```
> struct Node
> {
>     int data;
>     Node next; // error, Node directly depends on itself
> }
> ```
>
> is an error because `Node` contains an instance field of its own type. Another example
>
> C#
>
> ```
> struct A { B b; }
> struct B { C c; }
> struct C { A a; }
> ```
>
> is an error because each of the types `A`, `B`, and `C` depend on each other.
>
> *end example*

With classes, it is possible for two variables to reference the same object, and thus possible for operations on one variable to affect the object referenced by the other variable. With structs, the variables each have their own copy of the data (except in the case of `in`, `out` and `ref` parameter variables), and it is not possible for operations on one to affect the other. Furthermore, except when explicitly nullable (§8.3.12), it is not possible for values of a struct type to be `null`.

> *Note*: If a struct contains a field of reference type then the contents of the object referenced can be altered by other operations. However the value of the field itself,

> i.e., which object it references, cannot be changed through a mutation of a different struct value. *end note*

> *Example*: Given the following

```C#
struct Point
{
    public int x, y;

    public Point(int x, int y)
    {
        this.x = x;
        this.y = y;
    }
}

class A
{
    static void Main()
    {
        Point a = new Point(10, 10);
        Point b = a;
        a.x = 100;
        Console.WriteLine(b.x);
    }
}
```

> the output is `10`. The assignment of `a` to `b` creates a copy of the value, and `b` is thus unaffected by the assignment to `a.x`. Had `Point` instead been declared as a class, the output would be `100` because `a` and `b` would reference the same object.

> *end example*

16.4.3 Inheritance

All struct types implicitly inherit from the class `System.ValueType`, which, in turn, inherits from class `object`. A struct declaration may specify a list of implemented interfaces, but it is not possible for a struct declaration to specify a base class.

Struct types are never abstract and are always implicitly sealed. The `abstract` and `sealed` modifiers are therefore not permitted in a struct declaration.

Since inheritance isn't supported for structs, the declared accessibility of a struct member cannot be `protected`, `private protected`, or `protected internal`.

Function members in a struct cannot be abstract or virtual, and the `override` modifier is allowed only to override methods inherited from `System.ValueType`.

16.4.4 Assignment

Assignment to a variable of a struct type creates a *copy* of the value being assigned. This differs from assignment to a variable of a class type, which copies the reference but not the object identified by the reference.

Similar to an assignment, when a struct is passed as a value parameter or returned as the result of a function member, a copy of the struct is created. A struct may be passed by reference to a function member using an `in`, `out`, or `ref` parameter.

When a property or indexer of a struct is the target of an assignment, the instance expression associated with the property or indexer access shall be classified as a variable. If the `instance` expression is classified as a value, a compile-time error occurs. This is described in further detail in §12.21.2.

16.4.5 Default values

As described in §9.3, several kinds of variables are automatically initialized to their default value when they are created. For variables of class types and other reference types, this default value is `null`. However, since structs are value types that cannot be `null`, the default value of a struct is the value produced by setting all value type fields to their default value and all reference type fields to `null`.

> *Example*: Referring to the `Point` struct declared above, the example
>
> ```
> C#
> Point[] a = new Point[100];
> ```
>
> initializes each `Point` in the array to the value produced by setting the `x` and `y` fields to zero.
>
> *end example*

The default value of a struct corresponds to the value returned by the default constructor of the struct (§8.3.3). Unlike a class, a struct is not permitted to declare a parameterless instance constructor. Instead, every struct implicitly has a parameterless

instance constructor, which always returns the value that results from setting all fields to their default values.

> *Note*: Structs should be designed to consider the default initialization state a valid state. In the example
>
> C#
>
> ```
> struct KeyValuePair
> {
>     string key;
>     string value;
>
>     public KeyValuePair(string key, string value)
>     {
>         if (key == null || value == null)
>         {
>             throw new ArgumentException();
>         }
>
>         this.key = key;
>         this.value = value;
>     }
> }
> ```
>
> the user-defined instance constructor protects against `null` values only where it is explicitly called. In cases where a `KeyValuePair` variable is subject to default value initialization, the `key` and `value` fields will be `null`, and the struct should be prepared to handle this state.
>
> *end note*

16.4.6 Boxing and unboxing

A value of a class type can be converted to type `object` or to an interface type that is implemented by the class simply by treating the reference as another type at compile-time. Likewise, a value of type `object` or a value of an interface type can be converted back to a class type without changing the reference (but, of course, a run-time type check is required in this case).

Since structs are not reference types, these operations are implemented differently for struct types. When a value of a struct type is converted to certain reference types (as defined in §10.2.9), a boxing operation takes place. Likewise, when a value of certain reference types (as defined in §10.3.7) is converted back to a struct type, an unboxing

operation takes place. A key difference from the same operations on class types is that boxing and unboxing *copies* the struct value either into or out of the boxed instance.

> *Note*: Thus, following a boxing or unboxing operation, changes made to the unboxed `struct` are not reflected in the boxed `struct`. *end note*

For further details on boxing and unboxing, see §10.2.9 and §10.3.7.

16.4.7 Meaning of this

The meaning of `this` in a struct differs from the meaning of `this` in a class, as described in §12.8.13. When a struct type overrides a virtual method inherited from `System.ValueType` (such as `Equals`, `GetHashCode`, or `ToString`), invocation of the virtual method through an instance of the struct type does not cause boxing to occur. This is true even when the struct is used as a type parameter and the invocation occurs through an instance of the type parameter type.

> *Example*:

C#

```
struct Counter
{
    int value;
    public override string ToString()
    {
        value++;
        return value.ToString();
    }
}

class Program
{
    static void Test<T>() where T : new()
    {
        T x = new T();
        Console.WriteLine(x.ToString());
        Console.WriteLine(x.ToString());
        Console.WriteLine(x.ToString());
    }

    static void Main() => Test<Counter>();
}
```

> The output of the program is:

Console

```
1
2
3
```

Although it is bad style for `ToString` to have side effects, the example demonstrates that no boxing occurred for the three invocations of `x.ToString()`.

end example

Similarly, boxing never implicitly occurs when accessing a member on a constrained type parameter when the member is implemented within the value type. For example, suppose an interface `ICounter` contains a method `Increment`, which can be used to modify a value. If `ICounter` is used as a constraint, the implementation of the `Increment` method is called with a reference to the variable that `Increment` was called on, never a boxed copy.

Example:

C#

```
interface ICounter
{
    void Increment();
}

struct Counter : ICounter
{
    int value;

    public override string ToString() => value.ToString();

    void ICounter.Increment() => value++;
}

class Program
{
    static void Test<T>() where T : ICounter, new()
    {
        T x = new T();
        Console.WriteLine(x);
        x.Increment();              // Modify x
        Console.WriteLine(x);
        ((ICounter)x).Increment();  // Modify boxed copy of x
        Console.WriteLine(x);
    }
```

```
    static void Main() => Test<Counter>();
}
```

The first call to `Increment` modifies the value in the variable `x`. This is not equivalent to the second call to `Increment`, which modifies the value in a boxed copy of `x`. Thus, the output of the program is:

Console

```
0
1
1
```

end example

16.4.8 Field initializers

As described in §16.4.5, the default value of a struct consists of the value that results from setting all value type fields to their default value and all reference type fields to `null`. For this reason, a struct does not permit instance field declarations to include variable initializers. This restriction applies only to instance fields. Static fields of a struct are permitted to include variable initializers.

Example: The following

C#

```
struct Point
{
    public int x = 1; // Error, initializer not permitted
    public int y = 1; // Error, initializer not permitted
}
```

is in error because the instance field declarations include variable initializers.

end example

16.4.9 Constructors

Unlike a class, a struct is not permitted to declare a parameterless instance constructor. Instead, every struct implicitly has a parameterless instance constructor, which always returns the value that results from setting all value type fields to their default value and

all reference type fields to `null` (§8.3.3). A struct can declare instance constructors having parameters.

> *Example*: Given the following
>
> C#
>
> ```
> struct Point
> {
>     int x, y;
>
>     public Point(int x, int y)
>     {
>         this.x = x;
>         this.y = y;
>     }
> }
>
> class A
> {
>     static void Main()
>     {
>         Point p1 = new Point();
>         Point p2 = new Point(0, 0);
>     }
> }
> ```
>
> the statements both create a `Point` with `x` and `y` initialized to zero.
>
> *end example*

A struct instance constructor is not permitted to include a constructor initializer of the form `base(`*argument_list*`)`, where *argument_list* is optional.

The `this` parameter of a struct instance constructor corresponds to an `out` parameter of the struct type. As such, `this` shall be definitely assigned (§9.4) at every location where the constructor returns. Similarly, it cannot be read (even implicitly) in the constructor body before being definitely assigned.

If the struct instance constructor specifies a constructor initializer, that initializer is considered a definite assignment to this that occurs prior to the body of the constructor. Therefore, the body itself has no initialization requirements.

> *Example*: Consider the instance constructor implementation below:
>
> C#

```
struct Point
{
    int x, y;

    public int X
    {
        set { x = value; }
    }

    public int Y
    {
        set { y = value; }
    }

    public Point(int x, int y)
    {
        X = x; // error, this is not yet definitely assigned
        Y = y; // error, this is not yet definitely assigned
    }
}
```

No instance function member (including the set accessors for the properties `X` and `Y`) can be called until all fields of the struct being constructed have been definitely assigned. Note, however, that if `Point` were a class instead of a struct, the instance constructor implementation would be permitted. There is one exception to this, and that involves automatically implemented properties (§15.7.4). The definite assignment rules (§12.21.2) specifically exempt assignment to an auto-property of a struct type within an instance constructor of that struct type: such an assignment is considered a definite assignment of the hidden backing field of the auto-property. Thus, the following is allowed:

C#

```
struct Point
{
    public int X { get; set; }
    public int Y { get; set; }

    public Point(int x, int y)
    {
        X = x; // allowed, definitely assigns backing field
        Y = y; // allowed, definitely assigns backing field
   }
}
```

end example]

16.4.10 Static constructors

Static constructors for structs follow most of the same rules as for classes. The execution of a static constructor for a struct type is triggered by the first of the following events to occur within an application domain:

- A static member of the struct type is referenced.
- An explicitly declared constructor of the struct type is called.

> *Note*: The creation of default values (§16.4.5) of struct types does not trigger the static constructor. (An example of this is the initial value of elements in an array.) *end note*

16.4.11 Automatically implemented properties

Automatically implemented properties (§15.7.4) use hidden backing fields, which are only accessible to the property accessors.

> *Note*: This access restriction means that constructors in structs containing automatically implemented properties often need an explicit constructor initializer where they would not otherwise need one, to satisfy the requirement of all fields being definitely assigned before any function member is invoked or the constructor returns. *end note*

16.4.12 Safe context constraint

16.4.12.1 General

At compile-time, each expression is associated with a context where that instance and all its fields can be safely accessed, its ***safe-context***. The safe-context is a context, enclosing an expression, which it is safe for the value to escape to.

Any expression whose compile-time type is not a ref struct has a safe-context of caller-context.

A `default` expression, for any type, has safe-context of caller-context.

For any non-default expression whose compile-time type is a ref struct has a safe-context defined by the following sections.

The safe-context records which context a value may be copied into. Given an assignment from an expression `E1` with a safe-context `S1`, to an expression `E2` with

safe-context `S2`, it is an error if `S2` is a wider context than `S1`.

There are three different safe-context values, the same as the ref-safe-context values defined for reference variables (§9.7.2): **declaration-block**, **function-member**, and **caller-context**. The safe-context of an expression constrains its use as follows:

- For a return statement `return e1`, the safe-context of `e1` shall be caller-context.
- For an assignment `e1 = e2` the safe-context of `e2` shall be at least as wide a context as the safe-context of `e1`.

For a method invocation if there is a `ref` or `out` argument of a `ref struct` type (including the receiver unless the type is `readonly`), with safe-context `S1`, then no argument (including the receiver) may have a narrower safe-context than `S1`.

16.4.12.2 Parameter safe context

A formal parameter of a ref struct type, including the `this` parameter of an instance method, has a safe-context of caller-context.

16.4.12.3 Local variable safe context

A local variable of a ref struct type has a safe-context as follows:

- If the variable is an iteration variable of a `foreach` loop, then the variable's safe-context is the same as the safe-context of the `foreach` loop's expression.
- Otherwise if the variable's declaration has an initializer then the variable's safe-context is the same as the safe-context of that initializer.
- Otherwise the variable is uninitialized at the point of declaration and has a safe-context of caller-context.

16.4.12.4 Field safe context

A reference to a field `e.F`, where the type of `F` is a ref struct type, has a safe-context that is the same as the safe-context of `e`.

16.4.12.5 Operators

The application of a user-defined operator is treated as a method invocation (§16.4.12.6).

For an operator that yields a value, such as `e1 + e2` or `c ? e1 : e2`, the safe-context of the result is the narrowest context among the safe-contexts of the operands of the operator. As a consequence, for a unary operator that yields a value, such as `+e`, the safe-context of the result is the safe-context of the operand.

> *Note*: The first operand of a conditional operator is a `bool`, so its safe-context is caller-context. It follows that the resulting safe-context is the narrowest safe-context of the second and third operand. *end note*

16.4.12.6 Method and property invocation

A value resulting from a method invocation `e1.M(e2, ...)` or property invocation `e.P` has safe-context of the smallest of the following contexts:

- caller-context.
- The safe-context of all argument expressions (including the receiver).

A property invocation (either `get` or `set`) is treated as a method invocation of the underlying method by the above rules.

16.4.12.7 stackalloc

The result of a stackalloc expression has safe-context of function-member.

16.4.12.8 Constructor invocations

A `new` expression that invokes a constructor obeys the same rules as a method invocation that is considered to return the type being constructed.

In addition the safe-context is the smallest of the safe-contexts of all arguments and operands of all object initializer expressions, recursively, if any initializer is present.

> *Note*: These rules rely on `Span<T>` not having a constructor of the following form:
>
> ```
> C#
> ```
>
> ```
> public Span<T>(ref T p)
> ```
>
> Such a constructor makes instances of `Span<T>` used as fields indistinguishable from a `ref` field. The safety rules described in this document depend on `ref` fields not being a valid construct in C# or .NET. *end note*

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

- Open a documentation issue
- Provide product feedback

17 Arrays

Article • 04/07/2023

17.1 General

An array is a data structure that contains a number of variables that are accessed through computed indices. The variables contained in an array, also called the ***elements*** of the array, are all of the same type, and this type is called the ***element type*** of the array.

An array has a rank that determines the number of indices associated with each array element. The rank of an array is also referred to as the dimensions of the array. An array with a rank of one is called a ***single-dimensional array***. An array with a rank greater than one is called a ***multi-dimensional array***. Specific sized multi-dimensional arrays are often referred to as two-dimensional arrays, three-dimensional arrays, and so on. Each dimension of an array has an associated length that is an integral number greater than or equal to zero. The dimension lengths are not part of the type of the array, but rather are established when an instance of the array type is created at run-time. The length of a dimension determines the valid range of indices for that dimension: For a dimension of length `N`, indices can range from `0` to `N - 1` inclusive. The total number of elements in an array is the product of the lengths of each dimension in the array. If one or more of the dimensions of an array have a length of zero, the array is said to be empty.

The element type of an array can itself be an array type (§17.2.1). Such arrays of arrays are distinct from multi-dimensional arrays and can be used to represent "jagged arrays".

> *Example*:
>
> C#
>
> ```
> int[][] pascals =
> {
>     new int[] {1},
>     new int[] {1, 1},
>     new int[] {1, 2, 1},
>     new int[] {1, 3, 3, 1}
> };
> ```
>
> *end example*

Every array type is a reference type (§8.2). The element type of an array can be any type, including value types and array types.

17.2 Array types

17.2.1 General

The grammar productions for array types are provided in §8.2.1.

An array type is written as a *non_array_type* followed by one or more *rank_specifier*s.

A *non_array_type* is any *type* that is not itself an *array_type*.

The rank of an array type is given by the leftmost *rank_specifier* in the *array_type*: A *rank_specifier* indicates that the array is an array with a rank of one plus the number of "`,`" tokens in the *rank_specifier*.

The element type of an array type is the type that results from deleting the leftmost *rank_specifier*:

- An array type of the form `T[R]` is an array with rank `R` and a non-array element type `T`.
- An array type of the form `T[R][R₁]...[Rₓ]` is an array with rank `R` and an element type `T[R₁]...[Rₓ]`.

In effect, the *rank_specifier*s are read from left to right *before* the final non-array element type.

> *Example*: The type in `T[][,,][,]` is a single-dimensional array of three-dimensional arrays of two-dimensional arrays of `int`. *end example*

At run-time, a value of an array type can be `null` or a reference to an instance of that array type.

> *Note*: Following the rules of §17.6, the value may also be a reference to a covariant array type. *end note*

17.2.2 The System.Array type

The type `System.Array` is the abstract base type of all array types. An implicit reference conversion (§10.2.8) exists from any array type to `System.Array` and to any interface type

implemented by `System.Array`. An explicit reference conversion (§10.3.5) exists from `System.Array` and any interface type implemented by `System.Array` to any array type. `System.Array` is not itself an *array_type*. Rather, it is a *class_type* from which all *array_type*s are derived.

At run-time, a value of type `System.Array` can be `null` or a reference to an instance of any array type.

17.2.3 Arrays and the generic collection interfaces

A single-dimensional array `T[]` implements the interface `System.Collections.Generic.IList<T>` (`IList<T>` for short) and its base interfaces. Accordingly, there is an implicit conversion from `T[]` to `IList<T>` and its base interfaces. In addition, if there is an implicit reference conversion from `S` to `T` then `S[]` implements `IList<T>` and there is an implicit reference conversion from `S[]` to `IList<T>` and its base interfaces (§10.2.8). If there is an explicit reference conversion from `S` to `T` then there is an explicit reference conversion from `S[]` to `IList<T>` and its base interfaces (§10.3.5).

Similarly, a single-dimensional array `T[]` also implements the interface `System.Collections.Generic.IReadOnlyList<T>` (`IReadOnlyList<T>` for short) and its base interfaces. Accordingly, there is an implicit conversion from `T[]` to `IReadOnlyList<T>` and its base interfaces. In addition, if there is an implicit reference conversion from `S` to `T` then `S[]` implements `IReadOnlyList<T>` and there is an implicit reference conversion from `S[]` to `IReadOnlyList<T>` and its base interfaces (§10.2.8). If there is an explicit reference conversion from `S` to `T` then there is an explicit reference conversion from `S[]` to `IReadOnlyList<T>` and its base interfaces (§10.3.5).

> *Example*: For example:

C#

```
class Test
{
    static void Main()
    {
        string[] sa = new string[5];
        object[] oa1 = new object[5];
        object[] oa2 = sa;

        IList<string> lst1 = sa;  // Ok
        IList<string> lst2 = oa1; // Error, cast needed
        IList<object> lst3 = sa;  // Ok
```

```
        IList<object> lst4 = oa1; // Ok

        IList<string> lst5 = (IList<string>)oa1; // Exception
        IList<string> lst6 = (IList<string>)oa2; // Ok

        IReadOnlyList<string> lst7 = sa;        // Ok
        IReadOnlyList<string> lst8 = oa1;       // Error, cast needed
        IReadOnlyList<object> lst9 = sa;        // Ok
        IReadOnlyList<object> lst10 = oa1;      // Ok
        IReadOnlyList<string> lst11 = (IReadOnlyList<string>)oa1; //
Exception
        IReadOnlyList<string> lst12 = (IReadOnlyList<string>)oa2; // Ok
    }
}
```

The assignment `lst2 = oa1` generates a compile-time error since the conversion from `object[]` to `IList<string>` is an explicit conversion, not implicit. The cast `(IList<string>)oa1` will cause an exception to be thrown at run-time since `oa1` references an `object[]` and not a `string[]`. However the cast `(IList<string>)oa2` will not cause an exception to be thrown since `oa2` references a `string[]`.

end example

Whenever there is an implicit or explicit reference conversion from `S[]` to `IList<T>`, there is also an explicit reference conversion from `IList<T>` and its base interfaces to `S[]` (§10.3.5).

When an array type `S[]` implements `IList<T>`, some of the members of the implemented interface may throw exceptions. The precise behavior of the implementation of the interface is beyond the scope of this specification.

17.3 Array creation

Array instances are created by *array_creation_expressions* (§12.8.16.5) or by field or local variable declarations that include an *array_initializer* (§17.7). Array instances can also be created implicitly as part of evaluating an argument list involving a parameter array (§15.6.2.4).

When an array instance is created, the rank and length of each dimension are established and then remain constant for the entire lifetime of the instance. In other words, it is not possible to change the rank of an existing array instance, nor is it possible to resize its dimensions.

An array instance is always of an array type. The `System.Array` type is an abstract type that cannot be instantiated.

Elements of arrays created by *array_creation_expressions* are always initialized to their default value (§9.3).

17.4 Array element access

Array elements are accessed using *element_access* expressions (§12.8.11.2) of the form `A[I₁, I₂, ..., Iₓ]`, where `A` is an expression of an array type and each `Iₑ` is an expression of type `int`, `uint`, `long`, `ulong`, or can be implicitly converted to one or more of these types. The result of an array element access is a variable, namely the array element selected by the indices.

The elements of an array can be enumerated using a `foreach` statement (§13.9.5).

17.5 Array members

Every array type inherits the members declared by the `System.Array` type.

17.6 Array covariance

For any two *reference_types* `A` and `B`, if an implicit reference conversion (§10.2.8) or explicit reference conversion (§10.3.5) exists from `A` to `B`, then the same reference conversion also exists from the array type `A[R]` to the array type `B[R]`, where `R` is any given *rank_specifier* (but the same for both array types). This relationship is known as ***array covariance***. Array covariance, in particular, means that a value of an array type `A[R]` might actually be a reference to an instance of an array type `B[R]`, provided an implicit reference conversion exists from `B` to `A`.

Because of array covariance, assignments to elements of reference type arrays include a run-time check which ensures that the value being assigned to the array element is actually of a permitted type (§12.21.2).

> *Example*:
>
> ```
> C#
>
> class Test
> {
>     static void Fill(object[] array, int index, int count, object value)
>     {
> ```

```
            for (int i = index; i < index + count; i++)
            {
                array[i] = value;
            }
        }

        static void Main()
        {
            string[] strings = new string[100];
            Fill(strings, 0, 100, "Undefined");
            Fill(strings, 0, 10, null);
            Fill(strings, 90, 10, 0);
        }
    }
```

The assignment to `array[i]` in the `Fill` method implicitly includes a run-time check, which ensures that `value` is either a `null` reference or a reference to an object of a type that is compatible with the actual element type of `array`. In `Main`, the first two invocations of `Fill` succeed, but the third invocation causes a `System.ArrayTypeMismatchException` to be thrown upon executing the first assignment to `array[i]`. The exception occurs because a boxed `int` cannot be stored in a `string` array.

end example

Array covariance specifically does not extend to arrays of *value_types*. For example, no conversion exists that permits an `int[]` to be treated as an `object[]`.

17.7 Array initializers

Array initializers may be specified in field declarations (§15.5), local variable declarations (§13.6.2), and array creation expressions (§12.8.16.5):

ANTLR

```
array_initializer
    : '{' variable_initializer_list? '}'
    | '{' variable_initializer_list ',' '}'
    ;

variable_initializer_list
    : variable_initializer (',' variable_initializer)*
    ;

variable_initializer
    : expression
```

```
    | array_initializer
    ;
```

An array initializer consists of a sequence of variable initializers, enclosed by "`{`" and "`}`" tokens and separated by "`,`" tokens. Each variable initializer is an expression or, in the case of a multi-dimensional array, a nested array initializer.

The context in which an array initializer is used determines the type of the array being initialized. In an array creation expression, the array type immediately precedes the initializer, or is inferred from the expressions in the array initializer. In a field or variable declaration, the array type is the type of the field or variable being declared. When an array initializer is used in a field or variable declaration,

```
C#

int[] a = {0, 2, 4, 6, 8};
```

it is simply shorthand for an equivalent array creation expression:

```
C#

int[] a = new int[] {0, 2, 4, 6, 8};
```

For a single-dimensional array, the array initializer shall consist of a sequence of expressions, each having an implicit conversion to the element type of the array (§10.2). The expressions initialize array elements in increasing order, starting with the element at index zero. The number of expressions in the array initializer determines the length of the array instance being created.

> *Example*: The array initializer above creates an `int[]` instance of length 5 and then initializes the instance with the following values:
>
> ```
> C#
>
> a[0] = 0; a[1] = 2; a[2] = 4; a[3] = 6; a[4] = 8;
> ```
>
> *end example*

For a multi-dimensional array, the array initializer shall have as many levels of nesting as there are dimensions in the array. The outermost nesting level corresponds to the leftmost dimension and the innermost nesting level corresponds to the rightmost dimension. The length of each dimension of the array is determined by the number of

elements at the corresponding nesting level in the array initializer. For each nested array initializer, the number of elements shall be the same as the other array initializers at the same level.

> *Example*: The example:
>
> ```
> C#
> ```
>
> ```
> int[,] b = {{0, 1}, {2, 3}, {4, 5}, {6, 7}, {8, 9}};
> ```
>
> creates a two-dimensional array with a length of five for the leftmost dimension and a length of two for the rightmost dimension:
>
> ```
> C#
> ```
>
> ```
> int[,] b = new int[5, 2];
> ```
>
> and then initializes the array instance with the following values:
>
> ```
> C#
> ```
>
> ```
> b[0, 0] = 0; b[0, 1] = 1;
> b[1, 0] = 2; b[1, 1] = 3;
> b[2, 0] = 4; b[2, 1] = 5;
> b[3, 0] = 6; b[3, 1] = 7;
> b[4, 0] = 8; b[4, 1] = 9;
> ```
>
> *end example*

If a dimension other than the rightmost is given with length zero, the subsequent dimensions are assumed to also have length zero.

> *Example*:
>
> ```
> C#
> ```
>
> ```
> int[,] c = {};
> ```
>
> creates a two-dimensional array with a length of zero for both the leftmost and the rightmost dimension:
>
> ```
> C#
> ```

```
int[,] c = new int[0, 0];
```

end example

When an array creation expression includes both explicit dimension lengths and an array initializer, the lengths shall be constant expressions and the number of elements at each nesting level shall match the corresponding dimension length.

Example: Here are some examples:

```
C#

int i = 3;
int[] x = new int[3] {0, 1, 2}; // OK
int[] y = new int[i] {0, 1, 2}; // Error, i not a constant
int[] z = new int[3] {0, 1, 2, 3}; // Error, length/initializer mismatch
```

Here, the initializer for `y` results in a compile-time error because the dimension length expression is not a constant, and the initializer for `z` results in a compile-time error because the length and the number of elements in the initializer do not agree.

end example

Note: C# allows a trailing comma at the end of an *array_initializer*. This syntax provides flexibility in adding or deleting members from such a list, and simplifies machine generation of such lists. *end note*

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# Standard documentation feedback

C# Standard documentation is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

18 Interfaces

Article • 02/07/2024

18.1 General

An interface defines a contract. A class or struct that implements an interface shall adhere to its contract. An interface may inherit from multiple base interfaces, and a class or struct may implement multiple interfaces.

Interfaces can contain methods, properties, events, and indexers. The interface itself does not provide implementations for the members that it declares. The interface merely specifies the members that shall be supplied by classes or structs that implement the interface.

18.2 Interface declarations

18.2.1 General

An *interface_declaration* is a *type_declaration* (§14.7) that declares a new interface type.

ANTLR

```
interface_declaration
    : attributes? interface_modifier* 'partial'? 'interface'
      identifier variant_type_parameter_list? interface_base?
      type_parameter_constraints_clause* interface_body ';'?
    ;
```

An *interface_declaration* consists of an optional set of *attributes* (§22), followed by an optional set of *interface_modifier*s (§18.2.2), followed by an optional partial modifier (§15.2.7), followed by the keyword `interface` and an *identifier* that names the interface, followed by an optional *variant_type_parameter_list* specification (§18.2.3), followed by an optional *interface_base* specification (§18.2.4), followed by an optional *type_parameter_constraints_clause*s specification (§15.2.5), followed by an *interface_body* (§18.3), optionally followed by a semicolon.

An interface declaration shall not supply a *type_parameter_constraints_clause*s unless it also supplies a *variant_type_parameter_list*.

An interface declaration that supplies a *variant_type_parameter_list* is a generic interface declaration. Additionally, any interface nested inside a generic class declaration or a

generic struct declaration is itself a generic interface declaration, since type arguments for the containing type shall be supplied to create a constructed type (§8.4).

18.2.2 Interface modifiers

An *interface_declaration* may optionally include a sequence of interface modifiers:

ANTLR

```
interface_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | unsafe_modifier   // unsafe code support
    ;
```

unsafe_modifier (§23.2) is only available in unsafe code (§23).

It is a compile-time error for the same modifier to appear multiple times in an interface declaration.

The `new` modifier is only permitted on interfaces defined within a class. It specifies that the interface hides an inherited member by the same name, as described in §15.3.5.

The `public`, `protected`, `internal`, and `private` modifiers control the accessibility of the interface. Depending on the context in which the interface declaration occurs, only some of these modifiers might be permitted (§7.5.2). When a partial type declaration (§15.2.7) includes an accessibility specification (via the `public`, `protected`, `internal`, and `private` modifiers), the rules in §15.2.2 apply.

18.2.3 Variant type parameter lists

18.2.3.1 General

Variant type parameter lists can only occur on interface and delegate types. The difference from ordinary *type_parameter_lists* is the optional *variance_annotation* on each type parameter.

ANTLR

```
variant_type_parameter_list
    : '<' variant_type_parameters '>'
```

```
    ;
```

ANTLR

```
variant_type_parameters
    : attributes? variance_annotation? type_parameter
    | variant_type_parameters ',' attributes? variance_annotation?
      type_parameter
    ;
```

ANTLR

```
variance_annotation
    : 'in'
    | 'out'
    ;
```

If the variance annotation is `out`, the type parameter is said to be ***covariant***. If the variance annotation is `in`, the type parameter is said to be ***contravariant***. If there is no variance annotation, the type parameter is said to be ***invariant***.

> *Example*: In the following:
>
> C#
>
> ```
> interface C<out X, in Y, Z>
> {
>     X M(Y y);
>     Z P { get; set; }
> }
> ```
>
> `X` is covariant, `Y` is contravariant and `Z` is invariant.
>
> *end example*

If a generic interface is declared in multiple parts (§15.2.3), each partial declaration shall specify the same variance for each type parameter.

18.2.3.2 Variance safety

The occurrence of variance annotations in the type parameter list of a type restricts the places where types can occur within the type declaration.

A type T is ***output-unsafe*** if one of the following holds:

- `T` is a contravariant type parameter
- `T` is an array type with an output-unsafe element type
- `T` is an interface or delegate type $S_i,... A_e$ constructed from a generic type $S<X_i, ... X_e>$ where for at least one A_i one of the following holds:
  - X_i is covariant or invariant and A_i is output-unsafe.
  - X_i is contravariant or invariant and A_i is input-unsafe.

A type T is ***input-unsafe*** if one of the following holds:

- `T` is a covariant type parameter
- `T` is an array type with an input-unsafe element type
- `T` is an interface or delegate type $S<A_i,... A_e>$ constructed from a generic type $S<X_i, ... X_e>$ where for at least one A_i one of the following holds:
  - X_i is covariant or invariant and A_i is input-unsafe.
  - X_i is contravariant or invariant and A_i is output-unsafe.

Intuitively, an output-unsafe type is prohibited in an output position, and an input-unsafe type is prohibited in an input position.

A type is ***output-safe*** if it is not output-unsafe, and ***input-safe*** if it is not input-unsafe.

18.2.3.3 Variance conversion

The purpose of variance annotations is to provide for more lenient (but still type safe) conversions to interface and delegate types. To this end the definitions of implicit (§10.2) and explicit conversions (§10.3) make use of the notion of variance-convertibility, which is defined as follows:

A type $T<A_i, ..., A_v>$ is variance-convertible to a type $T<B_i, ..., B_v>$ if `T` is either an interface or a delegate type declared with the variant type parameters $T<X_i, ..., X_v>$, and for each variant type parameter X_i one of the following holds:

- X_i is covariant and an implicit reference or identity conversion exists from A_i to B_i
- X_i is contravariant and an implicit reference or identity conversion exists from B_i to A_i
- X_i is invariant and an identity conversion exists from A_i to B_i

18.2.4 Base interfaces

An interface can inherit from zero or more interface types, which are called the ***explicit base interfaces*** of the interface. When an interface has one or more explicit base interfaces, then in the declaration of that interface, the interface identifier is followed by a colon and a comma-separated list of base interface types.

ANTLR

```
interface_base
    : ':' interface_type_list
    ;
```

The explicit base interfaces can be constructed interface types (§8.4, §18.2). A base interface cannot be a type parameter on its own, though it can involve the type parameters that are in scope.

For a constructed interface type, the explicit base interfaces are formed by taking the explicit base interface declarations on the generic type declaration, and substituting, for each *type_parameter* in the base interface declaration, the corresponding *type_argument* of the constructed type.

The explicit base interfaces of an interface shall be at least as accessible as the interface itself (§7.5.5).

> *Note*: For example, it is a compile-time error to specify a `private` or `internal` interface in the *interface_base* of a `public` interface. *end note*

It is a compile-time error for an interface to directly or indirectly inherit from itself.

The ***base interfaces*** of an interface are the explicit base interfaces and their base interfaces. In other words, the set of base interfaces is the complete transitive closure of the explicit base interfaces, their explicit base interfaces, and so on. An interface inherits all members of its base interfaces.

> *Example*: In the following code
>
> C#
>
> ```
> interface IControl
> {
>     void Paint();
> }
>
> interface ITextBox : IControl
> {
>     void SetText(string text);
> ```

```
}

interface IListBox : IControl
{
    void SetItems(string[] items);
}

interface IComboBox: ITextBox, IListBox {}
```

the base interfaces of `IComboBox` are `IControl`, `ITextBox`, and `IListBox`. In other words, the `IComboBox` interface above inherits members `SetText` and `SetItems` as well as `Paint`.

end example

Members inherited from a constructed generic type are inherited after type substitution. That is, any constituent types in the member have the base class declaration's type parameters replaced with the corresponding type arguments used in the *class_base* specification.

Example: In the following code

```C#
interface IBase<T>
{
    T[] Combine(T a, T b);
}

interface IDerived : IBase<string[,]>
{
    // Inherited: string[][,] Combine(string[,] a, string[,] b);
}
```

the interface `IDerived` inherits the `Combine` method after the type parameter `T` is replaced with `string[,]`.

end example

A class or struct that implements an interface also implicitly implements all of the interface's base interfaces.

The handling of interfaces on multiple parts of a partial interface declaration (§15.2.7) are discussed further in §15.2.4.3.

Every base interface of an interface shall be output-safe (§18.2.3.2).

18.3 Interface body

The *interface_body* of an interface defines the members of the interface.

ANTLR

```
interface_body
    : '{' interface_member_declaration* '}'
    ;
```

18.4 Interface members

18.4.1 General

The members of an interface are the members inherited from the base interfaces and the members declared by the interface itself.

ANTLR

```
interface_member_declaration
    : interface_method_declaration
    | interface_property_declaration
    | interface_event_declaration
    | interface_indexer_declaration
    ;
```

An interface declaration declares zero or more members. The members of an interface shall be methods, properties, events, or indexers. An interface cannot contain constants, fields, operators, instance constructors, finalizers, or types, nor can an interface contain static members of any kind.

All interface members implicitly have public access. It is a compile-time error for interface member declarations to include any modifiers.

An *interface_declaration* creates a new declaration space (§7.3), and the type parameters and *interface_member_declaration*s immediately contained by the *interface_declaration* introduce new members into this declaration space. The following rules apply to *interface_member_declaration*s:

- The name of a type parameter in the *variant_type_parameter_list* of an interface declaration shall differ from the names of all other type parameters in the same *variant_type_parameter_list* and shall differ from the names of all members of the interface.

- The name of a method shall differ from the names of all properties and events declared in the same interface. In addition, the signature (§7.6) of a method shall differ from the signatures of all other methods declared in the same interface, and two methods declared in the same interface shall not have signatures that differ solely by `in`, `out`, and `ref`.
- The name of a property or event shall differ from the names of all other members declared in the same interface.
- The signature of an indexer shall differ from the signatures of all other indexers declared in the same interface.

The inherited members of an interface are specifically not part of the declaration space of the interface. Thus, an interface is allowed to declare a member with the same name or signature as an inherited member. When this occurs, the derived interface member is said to *hide* the base interface member. Hiding an inherited member is not considered an error, but it does cause the compiler to issue a warning. To suppress the warning, the declaration of the derived interface member shall include a `new` modifier to indicate that the derived member is intended to hide the base member. This topic is discussed further in §7.7.2.3.

If a `new` modifier is included in a declaration that doesn't hide an inherited member, a warning is issued to that effect. This warning is suppressed by removing the `new` modifier.

> *Note*: The members in class `object` are not, strictly speaking, members of any interface (§18.4). However, the members in class `object` are available via member lookup in any interface type (§12.5). *end note*

The set of members of an interface declared in multiple parts (§15.2.7) is the union of the members declared in each part. The bodies of all parts of the interface declaration share the same declaration space (§7.3), and the scope of each member (§7.7) extends to the bodies of all the parts.

18.4.2 Interface methods

Interface methods are declared using *interface_method_declaration*s:

ANTLR

```
interface_method_declaration
    : attributes? 'new'? return_type interface_method_header
    | attributes? 'new'? ref_kind ref_return_type interface_method_header
    ;
```

```
interface_method_header
    : identifier '(' formal_parameter_list? ')' ';'
    | identifier type_parameter_list '(' formal_parameter_list? ')'
      type_parameter_constraints_clause* ';'
    ;
```

The *attributes*, *return_type*, *ref_return_type*, *identifier*, and *formal_parameter_list* of an interface method declaration have the same meaning as those of a method declaration in a class (§15.6). An interface method declaration is not permitted to specify a method body, and the declaration therefore always ends with a semicolon.

All formal parameter types of an interface method shall be input-safe (§18.2.3.2), and the return type shall be either `void` or output-safe. In addition, any output or reference formal parameter types shall also be output-safe.

> *Note*: Output parameters are required to be input-safe due to common implementation restrictions. *end note*

Furthermore, each class type constraint, interface type constraint and type parameter constraint on any type parameters of the method shall be input-safe.

Furthermore, each class type constraint, interface type constraint and type parameter constraint on any type parameter of the method shall be input-safe.

These rules ensure that any covariant or contravariant usage of the interface remains typesafe.

> *Example*:
>
> C#
>
> ```
> interface I<out T>
> {
>     void M<U>() where U : T;     // Error
> }
> ```
>
> is ill-formed because the usage of `T` as a type parameter constraint on `U` is not input-safe.
>
> Were this restriction not in place it would be possible to violate type safety in the following manner:
>
> C#

```
class B {}
class D : B {}
class E : B {}
class C : I<D>
{
    public void M<U>() {...}
}

...

I<B> b = new C();
b.M<E>();
```

This is actually a call to `C.M<E>`. But that call requires that `E` derive from `D`, so type safety would be violated here.

end example

18.4.3 Interface properties

Interface properties are declared using *interface_property_declarations*:

ANTLR

```
interface_property_declaration
    : attributes? 'new'? type identifier '{' interface_accessors '}'
    | attributes? 'new'? ref_kind type identifier '{' ref_interface_accessor
'}'
    ;

interface_accessors
    : attributes? 'get' ';'
    | attributes? 'set' ';'
    | attributes? 'get' ';' attributes? 'set' ';'
    | attributes? 'set' ';' attributes? 'get' ';'
    ;

ref_interface_accessor
    : attributes? 'get' ';'
    ;
```

The *attributes*, *type*, and *identifier* of an interface property declaration have the same meaning as those of a property declaration in a class (§15.7).

The accessors of an interface property declaration correspond to the accessors of a class property declaration (§15.7.3), except that the *accessor_body* shall always be a semicolon.

Thus, the accessors simply indicate whether the property is read-write, read-only, or write-only.

The type of an interface property shall be output-safe if there is a get accessor, and shall be input-safe if there is a set accessor.

18.4.4 Interface events

Interface events are declared using *interface_event_declarations*:

ANTLR

```
interface_event_declaration
    : attributes? 'new'? 'event' type identifier ';'
    ;
```

The *attributes*, *type*, and *identifier* of an interface event declaration have the same meaning as those of an event declaration in a class (§15.8).

The type of an interface event shall be input-safe.

18.4.5 Interface indexers

Interface indexers are declared using *interface_indexer_declarations*:

ANTLR

```
interface_indexer_declaration
    : attributes? 'new'? type 'this' '[' formal_parameter_list ']'
      '{' interface_accessors '}'
    | attributes? 'new'? ref_kind type 'this' '[' formal_parameter_list ']'
      '{' ref_interface_accessor '}'
    ;
```

The *attributes*, *type*, and *formal_parameter_list* of an interface indexer declaration have the same meaning as those of an indexer declaration in a class (§15.9).

The accessors of an interface indexer declaration correspond to the accessors of a class indexer declaration (§15.9), except that the *accessor_body* shall always be a semicolon. Thus, the accessors simply indicate whether the indexer is read-write, read-only, or write-only.

All the formal parameter types of an interface indexer shall be input-safe (§18.2.3.2). In addition, any output or reference formal parameter types shall also be output-safe.

> *Note*: Output parameters are required to be input-safe due to common implementation restrictions. *end note*

The type of an interface indexer shall be output-safe if there is a get accessor, and shall be input-safe if there is a set accessor.

18.4.6 Interface member access

Interface members are accessed through member access (§12.8.7) and indexer access (§12.8.11.3) expressions of the form `I.M` and `I[A]`, where `I` is an interface type, `M` is a method, property, or event of that interface type, and `A` is an indexer argument list.

For interfaces that are strictly single-inheritance (each interface in the inheritance chain has exactly zero or one direct base interface), the effects of the member lookup (§12.5), method invocation (§12.8.9.2), and indexer access (§12.8.11.3) rules are exactly the same as for classes and structs: More derived members hide less derived members with the same name or signature. However, for multiple-inheritance interfaces, ambiguities can occur when two or more unrelated base interfaces declare members with the same name or signature. This subclause shows several examples, some of which lead to ambiguities and others which don't. In all cases, explicit casts can be used to resolve the ambiguities.

> *Example*: In the following code

C#

```
interface IList
{
    int Count { get; set; }
}

interface ICounter
{
    void Count(int i);
}

interface IListCounter : IList, ICounter {}

class C
{
    void Test(IListCounter x)
    {
        x.Count(1);             // Error
        x.Count = 1;            // Error
        ((IList)x).Count = 1;   // Ok, invokes IList.Count.set
        ((ICounter)x).Count(1); // Ok, invokes ICounter.Count
```

```
    }
}
```

the first two statements cause compile-time errors because the member lookup (§12.5) of `Count` in `IListCounter` is ambiguous. As illustrated by the example, the ambiguity is resolved by casting `x` to the appropriate base interface type. Such casts have no run-time costs—they merely consist of viewing the instance as a less derived type at compile-time.

end example

Example: In the following code

C#

```
interface IInteger
{
    void Add(int i);
}

interface IDouble
{
    void Add(double d);
}

interface INumber : IInteger, IDouble {}

class C
{
    void Test(INumber n)
    {
        n.Add(1);             // Invokes IInteger.Add
        n.Add(1.0);           // Only IDouble.Add is applicable
        ((IInteger)n).Add(1); // Only IInteger.Add is a candidate
        ((IDouble)n).Add(1);  // Only IDouble.Add is a candidate
    }
}
```

the invocation `n.Add(1)` selects `IInteger.Add` by applying overload resolution rules of §12.6.4. Similarly, the invocation `n.Add(1.0)` selects `IDouble.Add`. When explicit casts are inserted, there is only one candidate method, and thus no ambiguity.

end example

Example: In the following code

C#

```
interface IBase
{
    void F(int i);
}

interface ILeft : IBase
{
    new void F(int i);
}

interface IRight : IBase
{
    void G();
}

interface IDerived : ILeft, IRight {}

class A
{
    void Test(IDerived d)
    {
        d.F(1);           // Invokes ILeft.F
        ((IBase)d).F(1);  // Invokes IBase.F
        ((ILeft)d).F(1);  // Invokes ILeft.F
        ((IRight)d).F(1); // Invokes IBase.F
    }
}
```

the `IBase.F` member is hidden by the `ILeft.F` member. The invocation `d.F(1)` thus selects `ILeft.F`, even though `IBase.F` appears to not be hidden in the access path that leads through `IRight`.

The intuitive rule for hiding in multiple-inheritance interfaces is simply this: If a member is hidden in any access path, it is hidden in all access paths. Because the access path from `IDerived` to `ILeft` to `IBase` hides `IBase.F`, the member is also hidden in the access path from `IDerived` to `IRight` to `IBase`.

end example

18.5 Qualified interface member names

An interface member is sometimes referred to by its ***qualified interface member name***. The qualified name of an interface member consists of the name of the interface in which the member is declared, followed by a dot, followed by the name of the member. The qualified name of a member references the interface in which the member is declared.

Example: Given the declarations

```
C#
```

```
interface IControl
{
    void Paint();
}

interface ITextBox : IControl
{
    void SetText(string text);
}
```

the qualified name of `Paint` is `IControl.Paint` and the qualified name of SetText is `ITextBox.SetText`. In the example above, it is not possible to refer to `Paint` as `ITextBox.Paint`.

end example

When an interface is part of a namespace, a qualified interface member name can include the namespace name.

Example:

```
C#
```

```
namespace System
{
    public interface ICloneable
    {
        object Clone();
    }
}
```

Within the `System` namespace, both `ICloneable.Clone` and `System.ICloneable.Clone` are qualified interface member names for the `Clone` method.

end example

18.6 Interface implementations

18.6.1 General

Interfaces may be implemented by classes and structs. To indicate that a class or struct directly implements an interface, the interface is included in the base class list of the class or struct.

> *Example*:
>
> C#
>
> ```
> interface ICloneable
> {
>     object Clone();
> }
>
> interface IComparable
> {
>     int CompareTo(object other);
> }
>
> class ListEntry : ICloneable, IComparable
> {
>     public object Clone() {...}
>     public int CompareTo(object other) {...}
> }
> ```
>
> *end example*

A class or struct that directly implements an interface also implicitly implements all of the interface's base interfaces. This is true even if the class or struct doesn't explicitly list all base interfaces in the base class list.

> *Example*:
>
> C#
>
> ```
> interface IControl
> {
>     void Paint();
> }
>
> interface ITextBox : IControl
> {
>     void SetText(string text);
> }
>
> class TextBox : ITextBox
> {
>     public void Paint() {...}
> ```

```
    public void SetText(string text) {...}
}
```

Here, class `TextBox` implements both `IControl` and `ITextBox`.

end example

When a class `C` directly implements an interface, all classes derived from `C` also implement the interface implicitly.

The base interfaces specified in a class declaration can be constructed interface types (§8.4, §18.2).

Example: The following code illustrates how a class can implement constructed interface types:

```
C#

class C<U, V> {}
interface I1<V> {}
class D : C<string, int>, I1<string> {}
class E<T> : C<int, T>, I1<T> {}
```

end example

The base interfaces of a generic class declaration shall satisfy the uniqueness rule described in §18.6.3.

18.6.2 Explicit interface member implementations

For purposes of implementing interfaces, a class or struct may declare ***explicit interface member implementations***. An explicit interface member implementation is a method, property, event, or indexer declaration that references a qualified interface member name.

Example:

```
C#

interface IList<T>
{
    T[] GetElements();
}
```

```
interface IDictionary<K, V>
{
    V this[K key] { get; }
    void Add(K key, V value);
}

class List<T> : IList<T>, IDictionary<int, T>
{
    public T[] GetElements() {...}
    T IDictionary<int, T>.this[int index] {...}
    void IDictionary<int, T>.Add(int index, T value) {...}
}
```

Here `IDictionary<int,T>.this` and `IDictionary<int,T>.Add` are explicit interface member implementations.

end example

Example: In some cases, the name of an interface member might not be appropriate for the implementing class, in which case, the interface member may be implemented using explicit interface member implementation. A class implementing a file abstraction, for example, would likely implement a `Close` member function that has the effect of releasing the file resource, and implement the `Dispose` method of the `IDisposable` interface using explicit interface member implementation:

C#

```
interface IDisposable
{
    void Dispose();
}

class MyFile : IDisposable
{
    void IDisposable.Dispose() => Close();

    public void Close()
    {
        // Do what's necessary to close the file
        System.GC.SuppressFinalize(this);
    }
}
```

end example

It is not possible to access an explicit interface member implementation through its qualified interface member name in a method invocation, property access, event access, or indexer access. An explicit interface member implementation can only be accessed through an interface instance, and is in that case referenced simply by its member name.

It is a compile-time error for an explicit interface member implementation to include any modifiers (§15.6) other than `extern` or `async`.

It is a compile-time error for an explicit interface method implementation to include *type_parameter_constraints_clause*s. The constraints for a generic explicit interface method implementation are inherited from the interface method.

> *Note*: Explicit interface member implementations have different accessibility characteristics than other members. Because explicit interface member implementations are never accessible through a qualified interface member name in a method invocation or a property access, they are in a sense private. However, since they can be accessed through the interface, they are in a sense also as public as the interface in which they are declared. Explicit interface member implementations serve two primary purposes:
>
> - Because explicit interface member implementations are not accessible through class or struct instances, they allow interface implementations to be excluded from the public interface of a class or struct. This is particularly useful when a class or struct implements an internal interface that is of no interest to a consumer of that class or struct.
> - Explicit interface member implementations allow disambiguation of interface members with the same signature. Without explicit interface member implementations it would be impossible for a class or struct to have different implementations of interface members with the same signature and return type, as would it be impossible for a class or struct to have any implementation at all of interface members with the same signature but with different return types.
>
> *end note*

For an explicit interface member implementation to be valid, the class or struct shall name an interface in its base class list that contains a member whose qualified interface member name, type, number of type parameters, and parameter types exactly match those of the explicit interface member implementation. If an interface function member has a parameter array, the corresponding parameter of an associated explicit interface member implementation is allowed, but not required, to have the `params` modifier. If the

interface function member does not have a parameter array then an associated explicit interface member implementation shall not have a parameter array.

> *Example*: Thus, in the following class
>
> ```C#
> class Shape : ICloneable
> {
>     object ICloneable.Clone() {...}
>     int IComparable.CompareTo(object other) {...} // invalid
> }
> ```
>
> the declaration of `IComparable.CompareTo` results in a compile-time error because `IComparable` is not listed in the base class list of `Shape` and is not a base interface of `ICloneable`. Likewise, in the declarations
>
> ```C#
> class Shape : ICloneable
> {
>     object ICloneable.Clone() {...}
> }
>
> class Ellipse : Shape
> {
>     object ICloneable.Clone() {...} // invalid
> }
> ```
>
> the declaration of `ICloneable.Clone` in `Ellipse` results in a compile-time error because `ICloneable` is not explicitly listed in the base class list of `Ellipse`.
>
> *end example*

The qualified interface member name of an explicit interface member implementation shall reference the interface in which the member was declared.

> *Example*: Thus, in the declarations
>
> ```C#
> interface IControl
> {
>     void Paint();
> }
>
> interface ITextBox : IControl
> ```

```
{
    void SetText(string text);
}

class TextBox : ITextBox
{
    void IControl.Paint() {...}
    void ITextBox.SetText(string text) {...}
}
```

the explicit interface member implementation of Paint must be written as `IControl.Paint`, not `ITextBox.Paint`.

end example

18.6.3 Uniqueness of implemented interfaces

The interfaces implemented by a generic type declaration shall remain unique for all possible constructed types. Without this rule, it would be impossible to determine the correct method to call for certain constructed types.

Example: Suppose a generic class declaration were permitted to be written as follows:

C#

```
interface I<T>
{
    void F();
}

class X<U ,V> : I<U>, I<V> // Error: I<U> and I<V> conflict
{
    void I<U>.F() {...}
    void I<V>.F() {...}
}
```

Were this permitted, it would be impossible to determine which code to execute in the following case:

C#

```
I<int> x = new X<int, int>();
x.F();
```

end example

To determine if the interface list of a generic type declaration is valid, the following steps are performed:

- Let `L` be the list of interfaces directly specified in a generic class, struct, or interface declaration `C`.
- Add to `L` any base interfaces of the interfaces already in `L`.
- Remove any duplicates from `L`.
- If any possible constructed type created from `C` would, after type arguments are substituted into `L`, cause two interfaces in `L` to be identical, then the declaration of `C` is invalid. Constraint declarations are not considered when determining all possible constructed types.

> *Note*: In the class declaration `X` above, the interface list `L` consists of `I<U>` and `I<V>`. The declaration is invalid because any constructed type with `U` and `V` being the same type would cause these two interfaces to be identical types. *end note*

It is possible for interfaces specified at different inheritance levels to unify:

C#

```
interface I<T>
{
    void F();
}

class Base<U> : I<U>
{
    void I<U>.F() {...}
}

class Derived<U, V> : Base<U>, I<V> // Ok
{
    void I<V>.F() {...}
}
```

This code is valid even though `Derived<U,V>` implements both `I<U>` and `I<V>`. The code

C#

```
I<int> x = new Derived<int, int>();
x.F();
```

invokes the method in `Derived`, since `Derived<int,int>'` effectively re-implements `I<int>` (§18.6.7).

18.6.4 Implementation of generic methods

When a generic method implicitly implements an interface method, the constraints given for each method type parameter shall be equivalent in both declarations (after any interface type parameters are replaced with the appropriate type arguments), where method type parameters are identified by ordinal positions, left to right.

> *Example*: In the following code:

```
C#

interface I<X, Y, Z>
{
    void F<T>(T t) where T : X;
    void G<T>(T t) where T : Y;
    void H<T>(T t) where T : Z;
}

class C : I<object, C, string>
{
    public void F<T>(T t) {...}                  // Ok
    public void G<T>(T t) where T : C {...}      // Ok
    public void H<T>(T t) where T : string {...} // Error
}
```

> the method `C.F<T>` implicitly implements `I<object,C,string>.F<T>`. In this case, `C.F<T>` is not required (nor permitted) to specify the constraint `T: object` since `object` is an implicit constraint on all type parameters. The method `C.G<T>` implicitly implements `I<object,C,string>.G<T>` because the constraints match those in the interface, after the interface type parameters are replaced with the corresponding type arguments. The constraint for method `C.H<T>` is an error because sealed types (`string` in this case) cannot be used as constraints. Omitting the constraint would also be an error since constraints of implicit interface method implementations are required to match. Thus, it is impossible to implicitly implement `I<object,C,string>.H<T>`. This interface method can only be implemented using an explicit interface member implementation:

```
C#

class C : I<object, C, string>
{
    ...
    public void H<U>(U u) where U : class {...}

    void I<object, C, string>.H<T>(T t)
    {
```

```
            string s = t; // Ok
            H<T>(t);
        }
    }
```

In this case, the explicit interface member implementation invokes a public method having strictly weaker constraints. The assignment from t to s is valid since `T` inherits a constraint of `T: string`, even though this constraint is not expressible in source code. *end example*

Note: When a generic method explicitly implements an interface method no constraints are allowed on the implementing method (§15.7.1, §18.6.2). *end note*

18.6.5 Interface mapping

A class or struct shall provide implementations of all members of the interfaces that are listed in the base class list of the class or struct. The process of locating implementations of interface members in an implementing class or struct is known as ***interface mapping***.

Interface mapping for a class or struct `C` locates an implementation for each member of each interface specified in the base class list of `C`. The implementation of a particular interface member `I.M`, where `I` is the interface in which the member `M` is declared, is determined by examining each class or struct `S`, starting with `C` and repeating for each successive base class of `C`, until a match is located:

- If `S` contains a declaration of an explicit interface member implementation that matches `I` and `M`, then this member is the implementation of `I.M`.
- Otherwise, if `S` contains a declaration of a non-static public member that matches `M`, then this member is the implementation of `I.M`. If more than one member matches, it is unspecified which member is the implementation of `I.M`. This situation can only occur if `S` is a constructed type where the two members as declared in the generic type have different signatures, but the type arguments make their signatures identical.

A compile-time error occurs if implementations cannot be located for all members of all interfaces specified in the base class list of `C`. The members of an interface include those members that are inherited from base interfaces.

Members of a constructed interface type are considered to have any type parameters replaced with the corresponding type arguments as specified in §15.3.3.

Example: For example, given the generic interface declaration:

```
C#

interface I<T>
{
    T F(int x, T[,] y);
    T this[int y] { get; }
}
```

the constructed interface `I<string[]>` has the members:

```
C#

string[] F(int x, string[,][] y);
string[] this[int y] { get; }
```

end example

For purposes of interface mapping, a class or struct member `A` matches an interface member `B` when:

- `A` and `B` are methods, and the name, type, and formal parameter lists of `A` and `B` are identical.
- `A` and `B` are properties, the name and type of `A` and `B` are identical, and `A` has the same accessors as `B` (`A` is permitted to have additional accessors if it is not an explicit interface member implementation).
- `A` and `B` are events, and the name and type of `A` and `B` are identical.
- `A` and `B` are indexers, the type and formal parameter lists of `A` and `B` are identical, and `A` has the same accessors as `B` (`A` is permitted to have additional accessors if it is not an explicit interface member implementation).

Notable implications of the interface-mapping algorithm are:

- Explicit interface member implementations take precedence over other members in the same class or struct when determining the class or struct member that implements an interface member.
- Neither non-public nor static members participate in interface mapping.

Example: In the following code

```
C#
```

```
interface ICloneable
{
    object Clone();
}

class C : ICloneable
{
    object ICloneable.Clone() {...}
    public object Clone() {...}
}
```

the `ICloneable.Clone` member of `C` becomes the implementation of `Clone` in 'ICloneable' because explicit interface member implementations take precedence over other members.

end example

If a class or struct implements two or more interfaces containing a member with the same name, type, and parameter types, it is possible to map each of those interface members onto a single class or struct member.

Example:

C#

```
interface IControl
{
    void Paint();
}

interface IForm
{
    void Paint();
}

class Page : IControl, IForm
{
    public void Paint() {...}
}
```

Here, the `Paint` methods of both `IControl` and `IForm` are mapped onto the `Paint` method in `Page`. It is of course also possible to have separate explicit interface member implementations for the two methods.

end example

If a class or struct implements an interface that contains hidden members, then some members may need to be implemented through explicit interface member implementations.

> *Example*:
>
> C#
>
> ```
> interface IBase
> {
>     int P { get; }
> }
>
> interface IDerived : IBase
> {
>     new int P();
> }
> ```
>
> An implementation of this interface would require at least one explicit interface member implementation, and would take one of the following forms
>
> C#
>
> ```
> class C1 : IDerived
> {
>     int IBase.P { get; }
>     int IDerived.P() {...}
> }
> class C2 : IDerived
> {
>     public int P { get; }
>     int IDerived.P() {...}
> }
> class C3 : IDerived
> {
>     int IBase.P { get; }
>     public int P() {...}
> }
> ```
>
> *end example*

When a class implements multiple interfaces that have the same base interface, there can be only one implementation of the base interface.

> *Example*: In the following code
>
> C#

```
interface IControl
{
    void Paint();
}

interface ITextBox : IControl
{
    void SetText(string text);
}

interface IListBox : IControl
{
    void SetItems(string[] items);
}

class ComboBox : IControl, ITextBox, IListBox
{
    void IControl.Paint() {...}
    void ITextBox.SetText(string text) {...}
    void IListBox.SetItems(string[] items) {...}
}
```

it is not possible to have separate implementations for the `IControl` named in the base class list, the `IControl` inherited by `ITextBox`, and the `IControl` inherited by `IListBox`. Indeed, there is no notion of a separate identity for these interfaces. Rather, the implementations of `ITextBox` and `IListBox` share the same implementation of `IControl`, and `ComboBox` is simply considered to implement three interfaces, `IControl`, `ITextBox`, and `IListBox`.

end example

The members of a base class participate in interface mapping.

Example: In the following code

C#

```
interface Interface1
{
    void F();
}

class Class1
{
    public void F() {}
    public void G() {}
}
```

```
class Class2 : Class1, Interface1
{
    public new void G() {}
}
```

the method `F` in `Class1` is used in `Class2's` implementation of `Interface1`.

end example

18.6.6 Interface implementation inheritance

A class inherits all interface implementations provided by its base classes.

Without explicitly re-implementing an interface, a derived class cannot in any way alter the interface mappings it inherits from its base classes.

Example: In the declarations

C#

```
interface IControl
{
    void Paint();
}

class Control : IControl
{
    public void Paint() {...}
}

class TextBox : Control
{
    public new void Paint() {...}
}
```

the `Paint` method in `TextBox` hides the `Paint` method in `Control`, but it does not alter the mapping of `Control.Paint` onto `IControl.Paint`, and calls to `Paint` through class instances and interface instances will have the following effects

C#

```
Control c = new Control();
TextBox t = new TextBox();
IControl ic = c;
IControl it = t;
c.Paint();  // invokes Control.Paint();
t.Paint();  // invokes TextBox.Paint();
```

```
ic.Paint(); // invokes Control.Paint();
it.Paint(); // invokes Control.Paint();
```

end example

However, when an interface method is mapped onto a virtual method in a class, it is possible for derived classes to override the virtual method and alter the implementation of the interface.

Example: Rewriting the declarations above to

C#

```
interface IControl
{
    void Paint();
}

class Control : IControl
{
    public virtual void Paint() {...}
}

class TextBox : Control
{
    public override void Paint() {...}
}
```

the following effects will now be observed

C#

```
Control c = new Control();
TextBox t = new TextBox();
IControl ic = c;
IControl it = t;
c.Paint();  // invokes Control.Paint();
t.Paint();  // invokes TextBox.Paint();
ic.Paint(); // invokes Control.Paint();
it.Paint(); // invokes TextBox.Paint();
```

end example

Since explicit interface member implementations cannot be declared virtual, it is not possible to override an explicit interface member implementation. However, it is

perfectly valid for an explicit interface member implementation to call another method, and that other method can be declared virtual to allow derived classes to override it.

> *Example*:
>
> C#
>
> ```
> interface IControl
> {
>     void Paint();
> }
>
> class Control : IControl
> {
>     void IControl.Paint() { PaintControl(); }
>     protected virtual void PaintControl() {...}
> }
>
> class TextBox : Control
> {
>     protected override void PaintControl() {...}
> }
> ```
>
> Here, classes derived from `Control` can specialize the implementation of `IControl.Paint` by overriding the `PaintControl` method.
>
> *end example*

18.6.7 Interface re-implementation

A class that inherits an interface implementation is permitted to ***re-implement*** the interface by including it in the base class list.

A re-implementation of an interface follows exactly the same interface mapping rules as an initial implementation of an interface. Thus, the inherited interface mapping has no effect whatsoever on the interface mapping established for the re-implementation of the interface.

> *Example*: In the declarations
>
> C#
>
> ```
> interface IControl
> {
>     void Paint();
> }
> ```

```
class Control : IControl
{
    void IControl.Paint() {...}
}

class MyControl : Control, IControl
{
    public void Paint() {}
}
```

the fact that `Control` maps `IControl.Paint` onto `Control.IControl.Paint` doesn't affect the re-implementation in `MyControl`, which maps `IControl.Paint` onto `MyControl.Paint`.

end example

Inherited public member declarations and inherited explicit interface member declarations participate in the interface mapping process for re-implemented interfaces.

Example:

C#

```
interface IMethods
{
    void F();
    void G();
    void H();
    void I();
}

class Base : IMethods
{
    void IMethods.F() {}
    void IMethods.G() {}
    public void H() {}
    public void I() {}
}

class Derived : Base, IMethods
{
    public void F() {}
    void IMethods.H() {}
}
```

Here, the implementation of `IMethods` in `Derived` maps the interface methods onto `Derived.F`, `Base.IMethods.G`, `Derived.IMethods.H`, and `Base.I`.

end example

When a class implements an interface, it implicitly also implements all that interface's base interfaces. Likewise, a re-implementation of an interface is also implicitly a re-implementation of all of the interface's base interfaces.

Example:

C#

```
interface IBase
{
    void F();
}

interface IDerived : IBase
{
    void G();
}

class C : IDerived
{
    void IBase.F() {...}
    void IDerived.G() {...}
}

class D : C, IDerived
{
    public void F() {...}
    public void G() {...}
}
```

Here, the re-implementation of `IDerived` also re-implements `IBase`, mapping `IBase.F` onto `D.F`.

end example

18.6.8 Abstract classes and interfaces

Like a non-abstract class, an abstract class shall provide implementations of all members of the interfaces that are listed in the base class list of the class. However, an abstract class is permitted to map interface methods onto abstract methods.

Example:

C#

```
interface IMethods
{
    void F();
    void G();
}

abstract class C : IMethods
{
    public abstract void F();
    public abstract void G();
    }
```

Here, the implementation of `IMethods` maps `F` and `G` onto abstract methods, which shall be overridden in non-abstract classes that derive from `C`.

end example

Explicit interface member implementations cannot be abstract, but explicit interface member implementations are of course permitted to call abstract methods.

Example:

C#

```
interface IMethods
{
    void F();
    void G();
}

abstract class C: IMethods
{
    void IMethods.F() { FF(); }
    void IMethods.G() { GG(); }
    protected abstract void FF();
    protected abstract void GG();
}
```

Here, non-abstract classes that derive from `C` would be required to override `FF` and `GG`, thus providing the actual implementation of `IMethods`.

end example

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

- Open a documentation issue
- Provide product feedback

19 Enums

Article • 04/07/2023

19.1 General

An ***enum type*** is a distinct value type (§8.3) that declares a set of named constants.

> *Example:* The example
>
> C#
>
> ```
> enum Color
> {
>     Red,
>     Green,
>     Blue
> }
> ```
>
> declares an enum type named `Color` with members `Red`, `Green`, and `Blue`.
>
> *end example*

19.2 Enum declarations

An enum declaration declares a new enum type. An enum declaration begins with the keyword `enum`, and defines the name, accessibility, underlying type, and members of the enum.

ANTLR

```
enum_declaration
    : attributes? enum_modifier* 'enum' identifier enum_base? enum_body ';'?
    ;

enum_base
    : ':' integral_type
    | ':' integral_type_name
    ;

integral_type_name
    : type_name // Shall resolve to an integral type other than char
    ;

enum_body
    : '{' enum_member_declarations? '}'
```

```
    | '{' enum_member_declarations ',' '}'
    ;
```

Each enum type has a corresponding integral type called the ***underlying type*** of the enum type. This underlying type shall be able to represent all the enumerator values defined in the enumeration. If the *enum_base* is present, it explicitly declares the underlying type. The underlying type shall be one of the *integral types* (§8.3.6) other than `char`. The underlying type may be specified either by an `integral_type` (§8.3.5), or an `integral_type_name`. The `integral_type_name` is resolved in the same way as `type_name` (§7.8.1), including taking any using directives (§14.5) into account.

> *Note*: The `char` type cannot be used as an underlying type, either by keyword or via an `integral_type_name`. *end note*

An enum declaration that does not explicitly declare an underlying type has an underlying type of `int`.

> *Example*: The example
>
> ```
> C#
>
> enum Color : long
> {
>     Red,
>     Green,
>     Blue
> }
> ```
>
> declares an enum with an underlying type of `long`.
>
> *end example*

> *Note*: A developer might choose to use an underlying type of `long`, as in the example, to enable the use of values that are in the range of `long` but not in the range of `int`, or to preserve this option for the future. *end note*

> *Note*: C# allows a trailing comma in an *enum_body*, just like it allows one in an *array_initializer* (§17.7). *end note*

An enum declaration cannot include a type parameter list, but any enum nested inside a generic class declaration or a generic struct declaration is a generic enum declaration,

since type arguments for the containing type shall be supplied to create a constructed type (§8.4).

19.3 Enum modifiers

An *enum_declaration* may optionally include a sequence of enum modifiers:

ANTLR

```
enum_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    ;
```

It is a compile-time error for the same modifier to appear multiple times in an enum declaration.

The modifiers of an enum declaration have the same meaning as those of a class declaration (§15.2.2). However, the `abstract`, and `sealed`, and `static` modifiers are not permitted in an enum declaration. Enums cannot be abstract and do not permit derivation.

19.4 Enum members

The body of an enum type declaration defines zero or more enum members, which are the named constants of the enum type. No two enum members can have the same name.

ANTLR

```
enum_member_declarations
    : enum_member_declaration (',' enum_member_declaration)*
    ;
```

ANTLR

```
enum_member_declaration
    : attributes? identifier ('=' constant_expression)?
    ;
```

Each enum member has an associated constant value. The type of this value is the underlying type for the containing enum. The constant value for each enum member shall be in the range of the underlying type for the enum.

> *Example*: The example
>
> ```
> C#
> ```
>
> ```
> enum Color: uint
> {
>     Red = -1,
>     Green = -2,
>     Blue = -3
> }
> ```
>
> results in a compile-time error because the constant values `-1`, `-2`, and `-3` are not in the range of the underlying integral type `uint`.
>
> *end example*

Multiple enum members may share the same associated value.

> *Example*: The example
>
> ```
> C#
> ```
>
> ```
> enum Color
> {
>     Red,
>     Green,
>     Blue,
>     Max = Blue
> }
> ```
>
> shows an enum in which two enum members—`Blue` and `Max`—have the same associated value.
>
> *end example*

The associated value of an enum member is assigned either implicitly or explicitly. If the declaration of the enum member has a *constant_expression* initializer, the value of that constant expression, implicitly converted to the underlying type of the enum, is the associated value of the enum member. If the declaration of the enum member has no initializer, its associated value is set implicitly, as follows:

- If the enum member is the first enum member declared in the enum type, its associated value is zero.
- Otherwise, the associated value of the enum member is obtained by increasing the associated value of the textually preceding enum member by one. This increased value shall be within the range of values that can be represented by the underlying type, otherwise a compile-time error occurs.

Example: The example

C#

```
enum Color
{
    Red,
    Green = 10,
    Blue
}

class Test
{
    static void Main()
    {
        Console.WriteLine(StringFromColor(Color.Red));
        Console.WriteLine(StringFromColor(Color.Green));
        Console.WriteLine(StringFromColor(Color.Blue));
    }

    static string StringFromColor(Color c)
    {
        switch (c)
        {
            case Color.Red:
                return $"Red = {(int) c}";
            case Color.Green:
                return $"Green = {(int) c}";
            case Color.Blue:
                return $"Blue = {(int) c}";
            default:
                return "Invalid color";
      }
   }
}
```

prints out the enum member names and their associated values. The output is:

Console

```
Red = 0
Green = 10
```

```
Blue = 11
```

for the following reasons:

- the enum member `Red` is automatically assigned the value zero (since it has no initializer and is the first enum member);
- the enum member `Green` is explicitly given the value `10`;
- and the enum member `Blue` is automatically assigned the value one greater than the member that textually precedes it.

end example

The associated value of an enum member shall not, directly or indirectly, use the value of its own associated enum member. Other than this circularity restriction, enum member initializers may freely refer to other enum member initializers, regardless of their textual position. Within an enum member initializer, values of other enum members are always treated as having the type of their underlying type, so that casts are not necessary when referring to other enum members.

Example: The example

```
C#

enum Circular
{
    A = B,
    B
}
```

results in a compile-time error because the declarations of `A` and `B` are circular. `A` depends on `B` explicitly, and `B` depends on `A` implicitly.

end example

Enum members are named and scoped in a manner exactly analogous to fields within classes. The scope of an enum member is the body of its containing enum type. Within that scope, enum members can be referred to by their simple name. From all other code, the name of an enum member shall be qualified with the name of its enum type. Enum members do not have any declared accessibility—an enum member is accessible if its containing enum type is accessible.

19.5 The System.Enum type

The type `System.Enum` is the abstract base class of all enum types (this is distinct and different from the underlying type of the enum type), and the members inherited from `System.Enum` are available in any enum type. A boxing conversion (§10.2.9) exists from any enum type to `System.Enum`, and an unboxing conversion (§10.3.7) exists from `System.Enum` to any enum type.

Note that `System.Enum` is not itself an *enum_type*. Rather, it is a *class_type* from which all *enum_type*s are derived. The type `System.Enum` inherits from the type `System.ValueType` (§8.3.2), which, in turn, inherits from type `object`. At run-time, a value of type `System.Enum` can be `null` or a reference to a boxed value of any enum type.

19.6 Enum values and operations

Each enum type defines a distinct type; an explicit enumeration conversion (§10.3.3) is required to convert between an enum type and an integral type, or between two enum types. The set of values of the enum type is the same as the set of values of the underlying type and is not restricted to the values of the named constants. Any value of the underlying type of an enum can be cast to the enum type, and is a distinct valid value of that enum type.

Enum members have the type of their containing enum type (except within other enum member initializers: see §19.4). The value of an enum member declared in enum type `E` with associated value `v` is `(E)v`.

The following operators can be used on values of enum types:

- `==`, `!=`, `<`, `>`, `<=`, `>=` (§12.12.6)
- binary `+` (§12.10.5)
- binary `-` (§12.10.6)
- `^`, `&`, `|` (§12.13.3)
- `~` (§12.9.5)
- `++`, `--` (§12.8.15 and §12.9.6)
- `sizeof` (§23.6.9)

Every enum type automatically derives from the class `System.Enum` (which, in turn, derives from `System.ValueType` and `object`). Thus, inherited methods and properties of this class can be used on values of an enum type.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

C# Standard documentation feedback

C# Standard documentation is an open source project. Select a link to provide feedback:

- Open a documentation issue
- Provide product feedback

20 Delegates

Article • 02/07/2024

20.1 General

A delegate declaration defines a class that is derived from the class `System.Delegate`. A delegate instance encapsulates an ***invocation list***, which is a list of one or more methods, each of which is referred to as a ***callable entity***. For instance methods, a callable entity consists of an instance and a method on that instance. For static methods, a callable entity consists of just a method. Invoking a delegate instance with an appropriate set of arguments causes each of the delegate's callable entities to be invoked with the given set of arguments.

> *Note*: An interesting and useful property of a delegate instance is that it does not know or care about the classes of the methods it encapsulates; all that matters is that those methods be compatible (§20.4) with the delegate's type. This makes delegates perfectly suited for "anonymous" invocation. *end note*

20.2 Delegate declarations

A *delegate_declaration* is a *type_declaration* (§14.7) that declares a new delegate type.

ANTLR

```
delegate_declaration
    : attributes? delegate_modifier* 'delegate' return_type delegate_header
    | attributes? delegate_modifier* 'delegate' ref_kind ref_return_type
      delegate_header
    ;

delegate_header
    : identifier '(' formal_parameter_list? ')' ';'
    | identifier variant_type_parameter_list '(' formal_parameter_list? ')'
      type_parameter_constraints_clause* ';'
    ;

delegate_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | unsafe_modifier   // unsafe code support
    ;
```

unsafe_modifier is defined in §23.2.

It is a compile-time error for the same modifier to appear multiple times in a delegate declaration.

A delegate declaration that supplies a *variant_type_parameter_list* is a generic delegate declaration. Additionally, any delegate nested inside a generic class declaration or a generic struct declaration is itself a generic delegate declaration, since type arguments for the containing type shall be supplied to create a constructed type (§8.4).

The `new` modifier is only permitted on delegates declared within another type, in which case it specifies that such a delegate hides an inherited member by the same name, as described in §15.3.5.

The `public`, `protected`, `internal`, and `private` modifiers control the accessibility of the delegate type. Depending on the context in which the delegate declaration occurs, some of these modifiers might not be permitted (§7.5.2).

The delegate's type name is *identifier*.

As with methods (§15.6.1), if `ref` is present, the delegate returns-by-ref; otherwise, if *return_type* is `void`, the delegate returns-no-value; otherwise, the delegate returns-by-value.

The optional *formal_parameter_list* specifies the parameters of the delegate.

The *return_type* of a returns-by-value or returns-no-value delegate declaration specifies the type of the result, if any, returned by the delegate.

The *ref_return_type* of a returns-by-ref delegate declaration specifies the type of the variable referenced by the *variable_reference* (§9.5) returned by the delegate.

The optional *variant_type_parameter_list* (§18.2.3) specifies the type parameters to the delegate itself.

The return type of a delegate type shall be either `void`, or output-safe (§18.2.3.2).

All the formal parameter types of a delegate type shall be input-safe (§18.2.3.2). In addition, any output or reference parameter types shall also be output-safe.

> *Note*: Output parameters are required to be input-safe due to common implementation restrictions. *end note*

Furthermore, each class type constraint, interface type constraint and type parameter constraint on any type parameters of the delegate shall be input-safe.

Delegate types in C# are name equivalent, not structurally equivalent.

> *Example*:
>
> ```C#
> delegate int D1(int i, double d);
> delegate int D2(int c, double d);
> ```
>
> The delegate types `D1` and `D2` are two different types, so they are not interchangeable, despite their identical signatures.
>
> *end example*

Like other generic type declarations, type arguments shall be given to create a constructed delegate type. The parameter types and return type of a constructed delegate type are created by substituting, for each type parameter in the delegate declaration, the corresponding type argument of the constructed delegate type.

The only way to declare a delegate type is via a *delegate_declaration*. Every delegate type is a reference type that is derived from `System.Delegate`. The members required for every delegate type are detailed in §20.3. Delegate types are implicitly `sealed`, so it is not permissible to derive any type from a delegate type. It is also not permissible to declare a non-delegate class type deriving from `System.Delegate`. `System.Delegate` is not itself a delegate type; it is a class type from which all delegate types are derived.

20.3 Delegate members

Every delegate type inherits members from the `Delegate` class as described in §15.3.4. In addition, every delegate type shall provide a non-generic `Invoke` method whose parameter list matches the *formal_parameter_list* in the delegate declaration, whose return type matches the *return_type* or *ref_return_type* in the delegate declaration, and for returns-by-ref delegates whose *ref_kind* matches that in the delegate declaration. The `Invoke` method shall be at least as accessible as the containing delegate type. Calling the `Invoke` method on a delegate type is semantically equivalent to using the delegate invocation syntax (§20.6) .

Implementations may define additional members in the delegate type.

Except for instantiation, any operation that can be applied to a class or class instance can also be applied to a delegate class or instance, respectively. In particular, it is

possible to access members of the `System.Delegate` type via the usual member access syntax.

20.4 Delegate compatibility

A method or delegate type `M` is ***compatible*** with a delegate type `D` if all of the following are true:

- `D` and `M` have the same number of parameters, and each parameter in `D` has the same `in`, `out`, or `ref` modifiers as the corresponding parameter in `M`.
- For each value parameter, an identity conversion (§10.2.2) or implicit reference conversion (§10.2.8) exists from the parameter type in `D` to the corresponding parameter type in `M`.
- For each `in`, `out`, or `ref` parameter, the parameter type in `D` is the same as the parameter type in `M`.
- One of the following is true:
  - `D` and `M` are both *returns-no-value*
  - `D` and `M` are returns-by-value (§15.6.1, §20.2), and an identity or implicit reference conversion exists from the return type of `M` to the return type of `D`.
  - `D` and `M` are both returns-by-ref, an identity conversion exists between the return type of `M` and the return type of `D`, and both have the same *ref_kind*.

This definition of compatibility allows covariance in return type and contravariance in parameter types.

> *Example*:
>
> C#
>
> ```
> delegate int D1(int i, double d);
> delegate int D2(int c, double d);
> delegate object D3(string s);
>
> class A
> {
>     public static int M1(int a, double b) {...}
> }
>
> class B
> {
>     public static int M1(int f, double g) {...}
>     public static void M2(int k, double l) {...}
>     public static int M3(int g) {...}
>     public static void M4(int g) {...}
>     public static object M5(string s) {...}
> ```

```
    public static int[] M6(object o) {...}
}
```

The methods `A.M1` and `B.M1` are compatible with both the delegate types `D1` and `D2`, since they have the same return type and parameter list. The methods `B.M2`, `B.M3`, and `B.M4` are incompatible with the delegate types `D1` and `D2`, since they have different return types or parameter lists. The methods `B.M5` and `B.M6` are both compatible with delegate type `D3`.

end example

Example:

C#

```
delegate bool Predicate<T>(T value);

class X
{
    static bool F(int i) {...}
    static bool G(string s) {...}
}
```

The method `X.F` is compatible with the delegate type `Predicate<int>` and the method `X.G` is compatible with the delegate type `Predicate<string>`.

end example

Note: The intuitive meaning of delegate compatibility is that a method is compatible with a delegate type if every invocation of the delegate could be replaced with an invocation of the method without violating type safety, treating optional parameters and parameter arrays as explicit parameters. For example, in the following code:

C#

```
delegate void Action<T>(T arg);

class Test
{
    static void Print(object value) => Console.WriteLine(value);

    static void Main()
    {
        Action<string> log = Print;
        log("text");
```

```
    }
}
```

The `Print` method is compatible with the `Action<string>` delegate type because any invocation of an `Action<string>` delegate would also be a valid invocation of the `Print` method.

If the signature of the `Print` method above were changed to `Print(object value, bool prependTimestamp = false)` for example, the `Print` method would no longer be compatible with `Action<string>` by the rules of this clause.

end note

20.5 Delegate instantiation

An instance of a delegate is created by a *delegate_creation_expression* (§12.8.16.6), a conversion to a delegate type, delegate combination or delegate removal. The newly created delegate instance then refers to one or more of:

- The static method referenced in the *delegate_creation_expression*, or
- The target object (which cannot be `null`) and instance method referenced in the *delegate_creation_expression*, or
- Another delegate (§12.8.16.6).

Example:

C#

```
delegate void D(int x);

class C
{
    public static void M1(int i) {...}
    public void M2(int i) {...}
}

class Test
{
    static void Main()
    {
        D cd1 = new D(C.M1); // Static method
        C t = new C();
        D cd2 = new D(t.M2); // Instance method
        D cd3 = new D(cd2);  // Another delegate
    }
}
```

end example

The set of methods encapsulated by a delegate instance is called an *invocation list*. When a delegate instance is created from a single method, it encapsulates that method, and its invocation list contains only one entry. However, when two non-`null` delegate instances are combined, their invocation lists are concatenated—in the order left operand then right operand—to form a new invocation list, which contains two or more entries.

When a new delegate is created from a single delegate the resultant invocation list has just one entry, which is the source delegate (§12.8.16.6).

Delegates are combined using the binary `+` (§12.10.5) and `+=` operators (§12.21.4). A delegate can be removed from a combination of delegates, using the binary `-` (§12.10.6) and `-=` operators (§12.21.4). Delegates can be compared for equality (§12.12.9).

Example: The following example shows the instantiation of a number of delegates, and their corresponding invocation lists:

```
C#

delegate void D(int x);

class C
{
    public static void M1(int i) {...}
    public static void M2(int i) {...}
}

class Test
{
    static void Main()
    {
        D cd1 = new D(C.M1); // M1 - one entry in invocation list
        D cd2 = new D(C.M2); // M2 - one entry
        D cd3 = cd1 + cd2;   // M1 + M2 - two entries
        D cd4 = cd3 + cd1;   // M1 + M2 + M1 - three entries
        D cd5 = cd4 + cd3;   // M1 + M2 + M1 + M1 + M2 - five entries
        D td3 = new D(cd3);  // [M1 + M2] - ONE entry in invocation
                             // list, which is itself a list of two
methods.
        D td4 = td3 + cd1;   // [M1 + M2] + M1 - two entries
        D cd6 = cd4 - cd2;   // M1 + M1 - two entries in invocation list
        D td6 = td4 - cd2;   // [M1 + M2] + M1 - two entries in
invocation list,
                             // but still three methods called, M2 not
removed.
```

```
    }
}
```

When `cd1` and `cd2` are instantiated, they each encapsulate one method. When `cd3` is instantiated, it has an invocation list of two methods, `M1` and `M2`, in that order. `cd4`'s invocation list contains `M1`, `M2`, and `M1`, in that order. For `cd5`, the invocation list contains `M1`, `M2`, `M1`, `M1`, and `M2`, in that order.

When creating a delegate from another delegate with a *delegate_creation_expression* the result has an invocation list with a different structure from the original, but which results in the same methods being invoked in the same order. When `td3` is created from `cd3` its invocation list has just one member, but that member is a list of the methods `M1` and `M2` and those methods are invoked by `td3` in the same order as they are invoked by `cd3`. Similarly when `td4` is instantiated its invocation list has just two entries but it invokes the three methods `M1`, `M2`, and `M1`, in that order just as `cd4` does.

The structure of the invocation list affects delegate subtraction. Delegate `cd6`, created by subtracting `cd2` (which invokes `M2`) from `cd4` (which invokes `M1`, `M2`, and `M1`) invokes `M1` and `M1`. However delegate `td6`, created by subtracting `cd2` (which invokes `M2`) from `td4` (which invokes `M1`, `M2`, and `M1`) still invokes `M1`, `M2` and `M1`, in that order, as `M2` is not a single entry in the list but a member of a nested list. For more examples of combining (as well as removing) delegates, see §20.6.

end example

Once instantiated, a delegate instance always refers to the same invocation list.

> *Note*: Remember, when two delegates are combined, or one is removed from another, a new delegate results with its own invocation list; the invocation lists of the delegates combined or removed remain unchanged. *end note*

20.6 Delegate invocation

C# provides special syntax for invoking a delegate. When a non-`null` delegate instance whose invocation list contains one entry, is invoked, it invokes the one method with the same arguments it was given, and returns the same value as the referred to method. (See §12.8.9.4 for detailed information on delegate invocation.) If an exception occurs during the invocation of such a delegate, and that exception is not caught within the method that was invoked, the search for an exception catch clause continues in the

method that called the delegate, as if that method had directly called the method to which that delegate referred.

Invocation of a delegate instance whose invocation list contains multiple entries, proceeds by invoking each of the methods in the invocation list, synchronously, in order. Each method so called is passed the same set of arguments as was given to the delegate instance. If such a delegate invocation includes reference parameters (§15.6.2.3.3), each method invocation will occur with a reference to the same variable; changes to that variable by one method in the invocation list will be visible to methods further down the invocation list. If the delegate invocation includes output parameters or a return value, their final value will come from the invocation of the last delegate in the list. If an exception occurs during processing of the invocation of such a delegate, and that exception is not caught within the method that was invoked, the search for an exception catch clause continues in the method that called the delegate, and any methods further down the invocation list are not invoked.

Attempting to invoke a delegate instance whose value is `null` results in an exception of type `System.NullReferenceException`.

> *Example*: The following example shows how to instantiate, combine, remove, and invoke delegates:

C#

```
delegate void D(int x);

class C
{
    public static void M1(int i) => Console.WriteLine("C.M1: " + i);

    public static void M2(int i) => Console.WriteLine("C.M2: " + i);

    public void M3(int i) => Console.WriteLine("C.M3: " + i);
}

class Test
{
    static void Main()
    {
        D cd1 = new D(C.M1);
        cd1(-1);             // call M1
        D cd2 = new D(C.M2);
        cd2(-2);             // call M2
        D cd3 = cd1 + cd2;
        cd3(10);             // call M1 then M2
        cd3 += cd1;
        cd3(20);             // call M1, M2, then M1
        C c = new C();
```

```
        D cd4 = new D(c.M3);
        cd3 += cd4;
        cd3(30);             // call M1, M2, M1, then M3
        cd3 -= cd1;          // remove last M1
        cd3(40);             // call M1, M2, then M3
        cd3 -= cd4;
        cd3(50);             // call M1 then M2
        cd3 -= cd2;
        cd3(60);             // call M1
        cd3 -= cd2;          // impossible removal is benign
        cd3(60);             // call M1
        cd3 -= cd1;          // invocation list is empty so cd3 is null
        // cd3(70);          // System.NullReferenceException thrown
        cd3 -= cd1;          // impossible removal is benign
    }
}
```

As shown in the statement `cd3 += cd1;`, a delegate can be present in an invocation list multiple times. In this case, it is simply invoked once per occurrence. In an invocation list such as this, when that delegate is removed, the last occurrence in the invocation list is the one actually removed.

Immediately prior to the execution of the final statement, `cd3 -= cd1;`, the delegate `cd3` refers to an empty invocation list. Attempting to remove a delegate from an empty list (or to remove a non-existent delegate from a non-empty list) is not an error.

The output produced is:

Console

```
C.M1: -1
C.M2: -2
C.M1: 10
C.M2: 10
C.M1: 20
C.M2: 20
C.M1: 20
C.M1: 30
C.M2: 30
C.M1: 30
C.M3: 30
C.M1: 40
C.M2: 40
C.M3: 40
C.M1: 50
C.M2: 50
C.M1: 60
C.M1: 60
```

end example

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

21 Exceptions

Article • 02/07/2024

21.1 General

Exceptions in C# provide a structured, uniform, and type-safe way of handling both system level and application-level error conditions.

21.2 Causes of exceptions

Exceptions can be thrown in two different ways.

- A `throw` statement (§13.10.6) throws an exception immediately and unconditionally. Control never reaches the statement immediately following the `throw`.
- Certain exceptional conditions that arise during the processing of C# statements and expressions cause an exception to be thrown in certain circumstances when the operation cannot be completed normally. See §21.5 for a list of the various exceptions that can be thrown in this way.

> *Example*: An integer division operation (§12.10.3) throws a `System.DivideByZeroException` if the denominator is zero. *end example*

21.3 The System.Exception class

The `System.Exception` class is the base type of all exceptions. This class has a few notable properties that all exceptions share:

- `Message` is a read-only property of type `string` that contains a human-readable description of the reason for the exception.
- `InnerException` is a read-only property of type `Exception`. If its value is non-`null`, it refers to the exception that caused the current exception. (That is, the current exception was raised in a catch block handling the `InnerException`.) Otherwise, its value is `null`, indicating that this exception was not caused by another exception. The number of exception objects chained together in this manner can be arbitrary.

The value of these properties can be specified in calls to the instance constructor for `System.Exception`.

21.4 How exceptions are handled

Exceptions are handled by a `try` statement (§13.11).

When an exception is thrown (§21.2), the system searches for the nearest catch clause that can handle the exception, as determined by the run-time type of the exception. First, the current method is searched for a lexically enclosing `try` statement, and the associated `catch` clauses of the `try` statement are considered in order. If that fails, the method that called the current method is searched for a lexically enclosing `try` statement that encloses the point of the call to the current method. This search continues until a `catch` clause is found that can handle the current exception, by naming an exception class that is of the same class, or a base class, of the run-time type of the exception being thrown. A `catch` clause that doesn't name an exception class can handle any exception.

Once a matching `catch` clause is found, the system prepares to transfer control to the first statement of the `catch` clause. Before execution of the `catch` clause begins, the system first executes, in order, any `finally` clauses that were associated with `try` statements more nested that than the one that caught the exception.

If no matching `catch` clause is found:

- If the search for a matching `catch` clause reaches a static constructor (§15.12) or static field initializer, then a `System.TypeInitializationException` is thrown at the point that triggered the invocation of the static constructor. The inner exception of the `System.TypeInitializationException` contains the exception that was originally thrown.
- Otherwise, if an exception occurs during finalizer execution, and that exception is not caught, then the behavior is unspecified.
- Otherwise, if the search for matching `catch` clauses reaches the code that initially started the thread, then execution of the thread is terminated. The impact of such termination is implementation-defined.

21.5 Common exception classes

The following exceptions are thrown by certain C# operations.

Expand table

Exception Type	Description
`System.ArithmeticException`	A base class for exceptions that occur during arithmetic operations, such as `System.DivideByZeroException` and `System.OverflowException`.
`System.ArrayTypeMismatchException`	Thrown when a store into an array fails because the type of the stored element is incompatible with the type of the array.
`System.DivideByZeroException`	Thrown when an attempt to divide an integral value by zero occurs.
`System.IndexOutOfRangeException`	Thrown when an attempt to index an array via an index that is less than zero or outside the bounds of the array.
`System.InvalidCastException`	Thrown when an explicit conversion from a base type or interface to a derived type fails at run-time.
`System.NullReferenceException`	Thrown when a `null` reference is used in a way that causes the referenced object to be required.
`System.OutOfMemoryException`	Thrown when an attempt to allocate memory (via `new`) fails.
`System.OverflowException`	Thrown when an arithmetic operation in a `checked` context overflows.
`System.StackOverflowException`	Thrown when the execution stack is exhausted by having too many pending calls; typically indicative of very deep or unbounded recursion.
`System.TypeInitializationException`	Thrown when a static constructor or static field initializer throws an exception, and no `catch` clause exists to catch it.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

22 Attributes

Article • 02/07/2024

22.1 General

Much of the C# language enables the programmer to specify declarative information about the entities defined in the program. For example, the accessibility of a method in a class is specified by decorating it with the *method_modifier*s `public`, `protected`, `internal`, and `private`.

C# enables programmers to invent new kinds of declarative information, called ***attributes***. Programmers can then attach attributes to various program entities, and retrieve attribute information in a run-time environment.

> *Note*: For instance, a framework might define a `HelpAttribute` attribute that can be placed on certain program elements (such as classes and methods) to provide a mapping from those program elements to their documentation. *end note*

Attributes are defined through the declaration of attribute classes (§22.2), which can have positional and named parameters (§22.2.3). Attributes are attached to entities in a C# program using attribute specifications (§22.3), and can be retrieved at run-time as attribute instances (§22.4).

22.2 Attribute classes

22.2.1 General

A class that derives from the abstract class `System.Attribute`, whether directly or indirectly, is an ***attribute class***. The declaration of an attribute class defines a new kind of attribute that can be placed on program entities. By convention, attribute classes are named with a suffix of `Attribute`. Uses of an attribute may either include or omit this suffix.

A generic class declaration shall not use `System.Attribute` as a direct or indirect base class.

> *Example*:
>
> ```
> C#
> ```

```
public class B : Attribute {}
public class C<T> : B {} // Error – generic cannot be an attribute
```

end example

22.2.2 Attribute usage

The attribute `AttributeUsage` (§22.5.2) is used to describe how an attribute class can be used.

`AttributeUsage` has a positional parameter (§22.2.3) that enables an attribute class to specify the kinds of program entities on which it can be used.

Example: The following example defines an attribute class named `SimpleAttribute` that can be placed on *class_declaration*s and *interface_declaration*s only, and shows several uses of the `Simple` attribute.

C#

```
[AttributeUsage(AttributeTargets.Class | AttributeTargets.Interface)]
public class SimpleAttribute : Attribute
{
    ...
}

[Simple] class Class1 {...}
[Simple] interface Interface1 {...}
```

Although this attribute is defined with the name `SimpleAttribute`, when this attribute is used, the `Attribute` suffix may be omitted, resulting in the short name `Simple`. Thus, the example above is semantically equivalent to the following

C#

```
[SimpleAttribute] class Class1 {...}
[SimpleAttribute] interface Interface1 {...}
```

end example

`AttributeUsage` has a named parameter (§22.2.3), called `AllowMultiple`, which indicates whether the attribute can be specified more than once for a given entity. If `AllowMultiple` for an attribute class is true, then that attribute class is a ***multi-use***

attribute class, and can be specified more than once on an entity. If `AllowMultiple` for an attribute class is false or it is unspecified, then that attribute class is a ***single-use attribute class***, and can be specified at most once on an entity.

> *Example*: The following example defines a multi-use attribute class named `AuthorAttribute` and shows a class declaration with two uses of the `Author` attribute:
>
> C#
>
> ```
> [AttributeUsage(AttributeTargets.Class, AllowMultiple = true)]
> public class AuthorAttribute : Attribute
> {
>     public string Name { get; }
>     public AuthorAttribute(string name) => Name = name;
> }
>
> [Author("Brian Kernighan"), Author("Dennis Ritchie")]
> class Class1
> {
>     ...
> }
> ```
>
> *end example*

`AttributeUsage` has another named parameter (§22.2.3), called `Inherited`, which indicates whether the attribute, when specified on a base class, is also inherited by classes that derive from that base class. If `Inherited` for an attribute class is true, then that attribute is inherited. If `Inherited` for an attribute class is false then that attribute is not inherited. If it is unspecified, its default value is true.

An attribute class `X` not having an `AttributeUsage` attribute attached to it, as in

C#

```
class X : Attribute { ... }
```

is equivalent to the following:

C#

```
[AttributeUsage(
   AttributeTargets.All,
   AllowMultiple = false,
   Inherited = true)
```

```
]
class X : Attribute { ... }
```

22.2.3 Positional and named parameters

Attribute classes can have ***positional parameters*** and ***named parameters***. Each public instance constructor for an attribute class defines a valid sequence of positional parameters for that attribute class. Each non-static public read-write field and property for an attribute class defines a named parameter for the attribute class. For a property to define a named parameter, that property shall have both a public get accessor and a public set accessor.

> *Example*: The following example defines an attribute class named `HelpAttribute` that has one positional parameter, `url`, and one named parameter, `Topic`. Although it is non-static and public, the property `Url` does not define a named parameter, since it is not read-write. Two uses of this attribute are also shown:

C#

```
[AttributeUsage(AttributeTargets.Class)]
public class HelpAttribute : Attribute
{
    public HelpAttribute(string url) // url is a positional parameter
    {
        ...
    }

    // Topic is a named parameter
    public string Topic
    {
        get;
        set;
    }

    public string Url { get; }
}

[Help("http://www.mycompany.com/xxx/Class1.htm")]
class Class1
{
}

[Help("http://www.mycompany.com/xxx/Misc.htm", Topic ="Class2")]
class Class2
{
}
```

> *end example*

22.2.4 Attribute parameter types

The types of positional and named parameters for an attribute class are limited to the ***attribute parameter types***, which are:

- One of the following types: `bool`, `byte`, `char`, `double`, `float`, `int`, `long`, `sbyte`, `short`, `string`, `uint`, `ulong`, `ushort`.
- The type `object`.
- The type `System.Type`.
- Enum types.
- Single-dimensional arrays of the above types.
- A constructor argument or public field that does not have one of these types, shall not be used as a positional or named parameter in an attribute specification.

22.3 Attribute specification

Attribute specification is the application of a previously defined attribute to a program entity. An attribute is a piece of additional declarative information that is specified for a program entity. Attributes can be specified at global scope (to specify attributes on the containing assembly or module) and for *type_declaration*s (§14.7), *class_member_declaration*s (§15.3), *interface_member_declaration*s (§18.4), *struct_member_declaration*s (§16.3), *enum_member_declaration*s (§19.2), *accessor_declaration*s (§15.7.3), *event_accessor_declaration*s (§15.8), elements of *formal_parameter_list*s (§15.6.2), and elements of *type_parameter_list*s (§15.2.3).

Attributes are specified in ***attribute sections***. An attribute section consists of a pair of square brackets, which surround a comma-separated list of one or more attributes. The order in which attributes are specified in such a list, and the order in which sections attached to the same program entity are arranged, is not significant. For instance, the attribute specifications `[A][B]`, `[B][A]`, `[A, B]`, and `[B, A]` are equivalent.

ANTLR

```
global_attributes
    : global_attribute_section+
    ;

global_attribute_section
    : '[' global_attribute_target_specifier attribute_list ']'
    | '[' global_attribute_target_specifier attribute_list ',' ']'
    ;
```

```
global_attribute_target_specifier
    : global_attribute_target ':'
    ;

global_attribute_target
    : identifier
    ;

attributes
    : attribute_section+
    ;

attribute_section
    : '[' attribute_target_specifier? attribute_list ']'
    | '[' attribute_target_specifier? attribute_list ',' ']'
    ;

attribute_target_specifier
    : attribute_target ':'
    ;

attribute_target
    : identifier
    | keyword
    ;

attribute_list
    : attribute (',' attribute)*
    ;

attribute
    : attribute_name attribute_arguments?
    ;

attribute_name
    : type_name
    ;

attribute_arguments
    : '(' positional_argument_list? ')'
    | '(' positional_argument_list ',' named_argument_list ')'
    | '(' named_argument_list ')'
    ;

positional_argument_list
    : positional_argument (',' positional_argument)*
    ;

positional_argument
    : argument_name? attribute_argument_expression
    ;

named_argument_list
    : named_argument (','  named_argument)*
```

```
    ;

named_argument
    : identifier '=' attribute_argument_expression
    ;

attribute_argument_expression
    : expression
    ;
```

For the production *global_attribute_target*, and in the text below, *identifier* shall have a spelling equal to `assembly` or `module`, where equality is that defined in §6.4.3. For the production *attribute_target*, and in the text below, *identifier* shall have a spelling that is not equal to `assembly` or `module`, using the same definition of equality as above.

An attribute consists of an *attribute_name* and an optional list of positional and named arguments. The positional arguments (if any) precede the named arguments. A positional argument consists of an *attribute_argument_expression*; a named argument consists of a name, followed by an equal sign, followed by an *attribute_argument_expression*, which, together, are constrained by the same rules as simple assignment. The order of named arguments is not significant.

> *Note*: For convenience, a trailing comma is allowed in a *global_attribute_section* and an *attribute_section*, just as one is allowed in an *array_initializer* (§17.7). *end note*

The *attribute_name* identifies an attribute class.

When an attribute is placed at the global level, a *global_attribute_target_specifier* is required. When the *global_attribute_target* is equal to:

- `assembly` — the target is the containing assembly
- `module` — the target is the containing module

No other values for *global_attribute_target* are allowed.

The standardized *attribute_target* names are `event`, `field`, `method`, `param`, `property`, `return`, `type`, and `typevar`. These target names shall only be used in the following contexts:

- `event` — an event.
- `field` — a field. A field-like event (i.e., one without accessors) (§15.8.2) and an automatically implemented property (§15.7.4) can also have an attribute with this target.

- `method` — a constructor, finalizer, method, operator, property get and set accessors, indexer get and set accessors, and event add and remove accessors. A field-like event (i.e., one without accessors) can also have an attribute with this target.
- `param` — a property set accessor, an indexer set accessor, event add and remove accessors, and a parameter in a constructor, method, and operator.
- `property` — a property and an indexer.
- `return` — a delegate, method, operator, property get accessor, and indexer get accessor.
- `type` — a delegate, class, struct, enum, and interface.
- `typevar` — a type parameter.

Certain contexts permit the specification of an attribute on more than one target. A program can explicitly specify the target by including an *attribute_target_specifier*. Without an *attribute_target_specifier* a default is applied, but an *attribute_target_specifier* can be used to affirm or override the default. The contexts are resolved as follows:

- For an attribute on a delegate declaration the default target is the delegate. Otherwise when the *attribute_target* is equal to:
  - `type` — the target is the delegate
  - `return` — the target is the return value
- For an attribute on a method declaration the default target is the method. Otherwise when the *attribute_target* is equal to:
  - `method` — the target is the method
  - `return` — the target is the return value
- For an attribute on an operator declaration the default target is the operator. Otherwise when the *attribute_target* is equal to:
  - `method` — the target is the operator
  - `return` — the target is the return value
- For an attribute on a get accessor declaration for a property or indexer declaration the default target is the associated method. Otherwise when the *attribute_target* is equal to:
  - `method` — the target is the associated method
  - `return` — the target is the return value
- For an attribute specified on a set accessor for a property or indexer declaration the default target is the associated method. Otherwise when the *attribute_target* is equal to:
  - `method` — the target is the associated method
  - `param` — the target is the lone implicit parameter

- For an attribute on an automatically implemented property declaration the default target is the property. Otherwise when the *attribute_target* is equal to:
  - `field` — the target is the compiler-generated backing field for the property
- For an attribute specified on an event declaration that omits *event_accessor_declarations* the default target is the event declaration. Otherwise when the *attribute_target* is equal to:
  - `event` — the target is the event declaration
  - `field` — the target is the field
  - `method` — the targets are the methods
- In the case of an event declaration that does not omit *event_accessor_declarations* the default target is the method.
  - `method` — the target is the associated method
  - `param` — the target is the lone parameter

In all other contexts, inclusion of an *attribute_target_specifier* is permitted but unnecessary.

> *Example*: a class declaration may either include or omit the specifier `type`:
>
> C#
>
> ```
> [type: Author("Brian Kernighan")]
> class Class1 {}
>
> [Author("Dennis Ritchie")]
> class Class2 {}
> ```
>
> *end example*.

An implementation can accept other *attribute_target*s, the purposes of which are implementation defined. An implementation that does not recognize such an *attribute_target* shall issue a warning and ignore the containing *attribute_section*.

By convention, attribute classes are named with a suffix of `Attribute`. An *attribute_name* can either include or omit this suffix. Specifically, an *attribute_name* is resolved as follows:

- If the right-most identifier of the *attribute_name* is a verbatim identifier (§6.4.3), then the *attribute_name* is resolved as a *type_name* (§7.8). If the result is not a type derived from `System.Attribute`, a compile-time error occurs.
- Otherwise,

- The *attribute_name* is resolved as a *type_name* (§7.8) except any errors are suppressed. If this resolution is successful and results in a type derived from `System.Attribute` then the type is the result of this step.
- The characters `Attribute` are appended to the right-most identifier in the *attribute_name* and the resulting string of tokens is resolved as a *type_name* (§7.8) except any errors are suppressed. If this resolution is successful and results in a type derived from `System.Attribute` then the type is the result of this step.

If exactly one of the two steps above results in a type derived from `System.Attribute`, then that type is the result of the *attribute_name*. Otherwise a compile-time error occurs.

> *Example*: If an attribute class is found both with and without this suffix, an ambiguity is present, and a compile-time error results. If the *attribute_name* is spelled such that its right-most *identifier* is a verbatim identifier (§6.4.3), then only an attribute without a suffix is matched, thus enabling such an ambiguity to be resolved. The example
>
> ```
> C#
> ```
>
> ```
> [AttributeUsage(AttributeTargets.All)]
> public class Example : Attribute
> {}
>
> [AttributeUsage(AttributeTargets.All)]
> public class ExampleAttribute : Attribute
> {}
>
> [Example]               // Error: ambiguity
> class Class1 {}
>
> [ExampleAttribute]      // Refers to ExampleAttribute
> class Class2 {}
>
> [@Example]              // Refers to Example
> class Class3 {}
>
> [@ExampleAttribute]     // Refers to ExampleAttribute
> class Class4 {}
> ```
>
> shows two attribute classes named `Example` and `ExampleAttribute`. The attribute `[Example]` is ambiguous, since it could refer to either `Example` or `ExampleAttribute`. Using a verbatim identifier allows the exact intent to be specified in such rare cases. The attribute `[ExampleAttribute]` is not ambiguous (although it would be if there was an attribute class named `ExampleAttributeAttribute`!). If the declaration for

class `Example` is removed, then both attributes refer to the attribute class named `ExampleAttribute`, as follows:

C#

```
[AttributeUsage(AttributeTargets.All)]
public class ExampleAttribute : Attribute
{}

[Example]            // Refers to ExampleAttribute
class Class1 {}

[ExampleAttribute]   // Refers to ExampleAttribute
class Class2 {}

[@Example]           // Error: no attribute named “Example”
class Class3 {}
```

end example

It is a compile-time error to use a single-use attribute class more than once on the same entity.

Example: The example

C#

```
[AttributeUsage(AttributeTargets.Class)]
public class HelpStringAttribute : Attribute
{
    public HelpStringAttribute(string value)
    {
        Value = value;
    }

    public string Value { get; }
}
[HelpString("Description of Class1")]
[HelpString("Another description of Class1")]   // multiple uses not
allowed
public class Class1 {}
```

results in a compile-time error because it attempts to use `HelpString`, which is a single-use attribute class, more than once on the declaration of `Class1`.

end example

An expression `E` is an *attribute_argument_expression* if all of the following statements are true:

- The type of `E` is an attribute parameter type (§22.2.4).
- At compile-time, the value of `E` can be resolved to one of the following:
  - A constant value.
  - A `System.Type` object obtained using a *typeof_expression* (§12.8.17) specifying a non-generic type, a closed constructed type (§8.4.3), or an unbound generic type (§8.4.4), but not an open type (§8.4.3).
  - A single-dimensional array of *attribute_argument_expressions*.

> *Example*:
>
> C#
>
> ```
> [AttributeUsage(AttributeTargets.Class | AttributeTargets.Field)]
> public class TestAttribute : Attribute
> {
>     public int P1 { get; set; }
>
>     public Type P2 { get; set; }
>
>     public object P3 { get; set; }
> }
>
> [Test(P1 = 1234, P3 = new int[]{1, 3, 5}, P2 = typeof(float))]
> class MyClass {}
>
> class C<T> {
>     [Test(P2 = typeof(T))] // Error - T not a closed type.
>     int x1;
>
>     [Test(P2 = typeof(C<T>))] // Error - C<;T>; not a closed type.
>     int x2;
>
>     [Test(P2 = typeof(C<int>))] // Ok
>     int x3;
>
>     [Test(P2 = typeof(C<>))] // Ok
>     int x4;
> }
> ```
>
> *end example*

The attributes of a type declared in multiple parts are determined by combining, in an unspecified order, the attributes of each of its parts. If the same attribute is placed on multiple parts, it is equivalent to specifying that attribute multiple times on the type.

Example: The two parts:

```C#
[Attr1, Attr2("hello")]
partial class A {}

[Attr3, Attr2("goodbye")]
partial class A {}
```

are equivalent to the following single declaration:

```C#
[Attr1, Attr2("hello"), Attr3, Attr2("goodbye")]
class A {}
```

end example

Attributes on type parameters combine in the same way.

22.4 Attribute instances

22.4.1 General

An ***attribute instance*** is an instance that represents an attribute at run-time. An attribute is defined with an attribute class, positional arguments, and named arguments. An attribute instance is an instance of the attribute class that is initialized with the positional and named arguments.

Retrieval of an attribute instance involves both compile-time and run-time processing, as described in the following subclauses.

22.4.2 Compilation of an attribute

The compilation of an *attribute* with attribute class `T`, *positional_argument_list* `P`, *named_argument_list* `N`, and specified on a program entity `E` is compiled into an assembly `A` via the following steps:

- Follow the compile-time processing steps for compiling an *object_creation_expression* of the form new `T(P)`. These steps either result in a

compile-time error, or determine an instance constructor `C` on `T` that can be invoked at run-time.

- If `C` does not have public accessibility, then a compile-time error occurs.
- For each *named_argument* `Arg` in `N`:
  - Let `Name` be the *identifier* of the *named_argument* `Arg`.
  - `Name` shall identify a non-static read-write public field or property on `T`. If `T` has no such field or property, then a compile-time error occurs.
- If any of the values within *positional_argument_list* `P` or one of the values within *named_argument_list* `N` is of type `System.String` and the value is not well-formed as defined by the Unicode Standard, it is implementation-defined whether the value compiled is equal to the run-time value retrieved (§22.4.3).

> *Note*: As an example, a string which contains a high surrogate UTF-16 code unit which isn't immediately followed by a low surrogate code unit is not well-formed. *end note*

- Store the following information (for run-time instantiation of the attribute) in the assembly output by the compiler as a result of compiling the program containing the attribute: the attribute class `T`, the instance constructor `C` on `T`, the *positional_argument_list* `P`, the *named_argument_list* `N`, and the associated program entity `E`, with the values resolved completely at compile-time.

22.4.3 Run-time retrieval of an attribute instance

Using the terms defined in §22.4.2, the attribute instance represented by `T`, `C`, `P`, and `N`, and associated with `E` can be retrieved at run-time from the assembly `A` using the following steps:

- Follow the run-time processing steps for executing an *object_creation_expression* of the form `new T(P)`, using the instance constructor `C` and values as determined at compile-time. These steps either result in an exception, or produce an instance `O` of `T`.
- For each *named_argument* `Arg` in `N`, in order:
  - Let `Name` be the *identifier* of the *named_argument* `Arg`. If `Name` does not identify a non-static public read-write field or property on `O`, then an exception is thrown.
  - Let `Value` be the result of evaluating the *attribute_argument_expression* of `Arg`.
  - If `Name` identifies a field on `O`, then set this field to `Value`.
  - Otherwise, Name identifies a property on `O`. Set this property to Value.

- The result is `O`, an instance of the attribute class `T` that has been initialized with the *positional_argument_list* `P` and the *named_argument_list* `N`.

> *Note*: The format for storing `T`, `C`, `P`, `N` (and associating it with `E`) in `A` and the mechanism to specify `E` and retrieve `T`, `C`, `P`, `N` from `A` (and hence how an attribute instance is obtained at runtime) is beyond the scope of this specification. *end note*

> *Example*: In an implementation of the CLI, the `Help` attribute instances in the assembly created by compiling the example program in §22.2.3 can be retrieved with the following program:

```C#
public sealed class InterrogateHelpUrls
{
    public static void Main(string[] args)
    {
        Type helpType = typeof(HelpAttribute);
        string assemblyName = args[0];
        foreach (Type t in Assembly.Load(assemblyName).GetTypes())
        {
            Console.WriteLine($"Type : {t}");
            var attributes = t.GetCustomAttributes(helpType, false);
            var helpers = (HelpAttribute[]) attributes;
            foreach (var helper in helpers)
            {
                Console.WriteLine($"\tUrl : {helper.Url}");
            }
        }
    }
}
```

end example

22.5 Reserved attributes

22.5.1 General

A small number of attributes affect the language in some way. These attributes include:

- `System.AttributeUsageAttribute` (§22.5.2), which is used to describe the ways in which an attribute class can be used.

- `System.Diagnostics.ConditionalAttribute` (§22.5.3), is a multi-use attribute class which is used to define conditional methods and conditional attribute classes. This attribute indicates a condition by testing a conditional compilation symbol.
- `System.ObsoleteAttribute` (§22.5.4), which is used to mark a member as obsolete.
- `System.Runtime.CompilerServices.AsyncMethodBuilderAttribute` (§22.5.5), which is used to establish a task builder for an async method.
- `System.Runtime.CompilerServices.CallerLineNumberAttribute` (§22.5.6.2), `System.Runtime.CompilerServices.CallerFilePathAttribute` (§22.5.6.3), and `System.Runtime.CompilerServices.CallerMemberNameAttribute` (§22.5.6.4), which are used to supply information about the calling context to optional parameters.

An execution environment may provide additional implementation-specific attributes that affect the execution of a C# program.

22.5.2 The AttributeUsage attribute

The attribute `AttributeUsage` is used to describe the manner in which the attribute class can be used.

A class that is decorated with the `AttributeUsage` attribute shall derive from `System.Attribute`, either directly or indirectly. Otherwise, a compile-time error occurs.

> *Note*: For an example of using this attribute, see §22.2.2. *end note*

22.5.3 The Conditional attribute

22.5.3.1 General

The attribute `Conditional` enables the definition of ***conditional methods*** and ***conditional attribute classes***.

22.5.3.2 Conditional methods

A method decorated with the `Conditional` attribute is a conditional method. Each conditional method is thus associated with the conditional compilation symbols declared in its `Conditional` attributes.

> *Example*:
>
> ```
> C#
> ```

```
class Eg
{
    [Conditional("ALPHA")]
    [Conditional("BETA")]
    public static void M()
    {
        // ...
    }
}
```

declares `Eg.M` as a conditional method associated with the two conditional compilation symbols `ALPHA` and `BETA`.

end example

A call to a conditional method is included if one or more of its associated conditional compilation symbols is defined at the point of call, otherwise the call is omitted.

A conditional method is subject to the following restrictions:

- The conditional method shall be a method in a *class_declaration* or *struct_declaration*. A compile-time error occurs if the `Conditional` attribute is specified on a method in an interface declaration.
- The conditional method shall have a return type of `void`.
- The conditional method shall not be marked with the `override` modifier. A conditional method can be marked with the `virtual` modifier, however. Overrides of such a method are implicitly conditional, and shall not be explicitly marked with a `Conditional` attribute.
- The conditional method shall not be an implementation of an interface method. Otherwise, a compile-time error occurs.
- The parameters of the conditional method shall not have the `out` modifier.

In addition, a compile-time error occurs if a delegate is created from a conditional method.

Example: The example

C#

```
#define DEBUG
using System;
using System.Diagnostics;

class Class1
{
```

```
    [Conditional("DEBUG")]
    public static void M()
    {
        Console.WriteLine("Executed Class1.M");
    }
}

class Class2
{
    public static void Test()
    {
        Class1.M();
    }
}
```

declares `Class1.M` as a conditional method. `Class2`'s `Test` method calls this method. Since the conditional compilation symbol `DEBUG` is defined, if `Class2.Test` is called, it will call `M`. If the symbol `DEBUG` had not been defined, then `Class2.Test` would not call `Class1.M`.

end example

It is important to understand that the inclusion or exclusion of a call to a conditional method is controlled by the conditional compilation symbols at the point of the call.

Example: In the following code

C#

```
// File Class1.cs:
using System.Diagnostics;
class Class1
{
    [Conditional("DEBUG")]
    public static void F()
    {
        Console.WriteLine("Executed Class1.F");
    }
}

// File Class2.cs:
#define DEBUG
class Class2
{
    public static void G()
    {
        Class1.F(); // F is called
    }
}
```

```
// File Class3.cs:
#undef DEBUG
class Class3
{
    public static void H()
    {
        Class1.F(); // F is not called
    }
}
```

the classes `Class2` and `Class3` each contain calls to the conditional method `Class1.F`, which is conditional based on whether or not `DEBUG` is defined. Since this symbol is defined in the context of `Class2` but not `Class3`, the call to `F` in `Class2` is included, while the call to `F` in `Class3` is omitted.

end example

The use of conditional methods in an inheritance chain can be confusing. Calls made to a conditional method through `base`, of the form `base.M`, are subject to the normal conditional method call rules.

Example: In the following code

C#

```
// File Class1.cs
using System.Diagnostics;
class Class1
{
    [Conditional("DEBUG")]
    public virtual void M() => Console.WriteLine("Class1.M executed");
}

// File Class2.cs
class Class2 : Class1
{
    public override void M()
    {
        Console.WriteLine("Class2.M executed");
        base.M(); // base.M is not called!
    }
}

// File Class3.cs
#define DEBUG
class Class3
{
    public static void Main()
    {
```

```
        Class2 c = new Class2();
        c.M(); // M is called
    }
}
```

`Class2` includes a call to the `M` defined in its base class. This call is omitted because the base method is conditional based on the presence of the symbol `DEBUG`, which is undefined. Thus, the method writes to the console "`Class2.M executed`" only. Judicious use of *pp_declaration*s can eliminate such problems.

end example

22.5.3.3 Conditional attribute classes

An attribute class (§22.2) decorated with one or more `Conditional` attributes is a conditional attribute class. A conditional attribute class is thus associated with the conditional compilation symbols declared in its `Conditional` attributes.

Example:

C#

```
[Conditional("ALPHA")]
[Conditional("BETA")]
public class TestAttribute : Attribute {}
```

declares `TestAttribute` as a conditional attribute class associated with the conditional compilations symbols `ALPHA` and `BETA`.

end example

Attribute specifications (§22.3) of a conditional attribute are included if one or more of its associated conditional compilation symbols is defined at the point of specification, otherwise the attribute specification is omitted.

It is important to note that the inclusion or exclusion of an attribute specification of a conditional attribute class is controlled by the conditional compilation symbols at the point of the specification.

Example: In the example

C#

```
// File Test.cs:
using System.Diagnostics;
[Conditional("DEBUG")]
public class TestAttribute : Attribute {}

// File Class1.cs:
#define DEBUG
[Test] // TestAttribute is specified
class Class1 {}

// File Class2.cs:
#undef DEBUG
[Test] // TestAttribute is not specified
class Class2 {}
```

the classes `Class1` and `Class2` are each decorated with attribute `Test`, which is conditional based on whether or not `DEBUG` is defined. Since this symbol is defined in the context of `Class1` but not `Class2`, the specification of the Test attribute on `Class1` is included, while the specification of the `Test` attribute on `Class2` is omitted.

end example

22.5.4 The Obsolete attribute

The attribute `Obsolete` is used to mark types and members of types that should no longer be used.

If a program uses a type or member that is decorated with the `Obsolete` attribute, the compiler shall issue a warning or an error. Specifically, the compiler shall issue a warning if no error parameter is provided, or if the error parameter is provided and has the value `false`. The compiler shall issue an error if the error parameter is specified and has the value `true`.

Example: In the following code

C#

```
[Obsolete("This class is obsolete; use class B instead")]
class A
{
    public void F() {}
}

class B
{
```

```
    public void F() {}
}

class Test
{
    static void Main()
    {
        A a = new A(); // Warning
        a.F();
    }
}
```

the class `A` is decorated with the `Obsolete` attribute. Each use of `A` in `Main` results in a warning that includes the specified message, "This class is obsolete; use class `B` instead".

end example

22.5.5 The AsyncMethodBuilder attribute

This attribute is described in §15.15.1.

22.5.6 Caller-info attributes

22.5.6.1 General

For purposes such as logging and reporting, it is sometimes useful for a function member to obtain certain compile-time information about the calling code. The caller-info attributes provide a way to pass such information transparently.

When an optional parameter is annotated with one of the caller-info attributes, omitting the corresponding argument in a call does not necessarily cause the default parameter value to be substituted. Instead, if the specified information about the calling context is available, that information will be passed as the argument value.

Example:

C#

```
public void Log(
    [CallerLineNumber] int line = -1,
    [CallerFilePath] string path = null,
    [CallerMemberName] string name = null
)
{
    Console.WriteLine((line < 0) ? "No line" : "Line "+ line);
```

```
    Console.WriteLine((path == null) ? "No file path" : path);
    Console.WriteLine((name == null) ? "No member name" : name);
}
```

A call to `Log()` with no arguments would print the line number and file path of the call, as well as the name of the member within which the call occurred.

end example

Caller-info attributes can occur on optional parameters anywhere, including in delegate declarations. However, the specific caller-info attributes have restrictions on the types of the parameters they can attribute, so that there will always be an implicit conversion from a substituted value to the parameter type.

It is an error to have the same caller-info attribute on a parameter of both the defining and implementing part of a partial method declaration. Only caller-info attributes in the defining part are applied, whereas caller-info attributes occurring only in the implementing part are ignored.

Caller information does not affect overload resolution. As the attributed optional parameters are still omitted from the source code of the caller, overload resolution ignores those parameters in the same way it ignores other omitted optional parameters (§12.6.4).

Caller information is only substituted when a function is explicitly invoked in source code. Implicit invocations such as implicit parent constructor calls do not have a source location and will not substitute caller information. Also, calls that are dynamically bound will not substitute caller information. When a caller-info attributed parameter is omitted in such cases, the specified default value of the parameter is used instead.

One exception is query expressions. These are considered syntactic expansions, and if the calls they expand to omit optional parameters with caller-info attributes, caller information will be substituted. The location used is the location of the query clause which the call was generated from.

If more than one caller-info attribute is specified on a given parameter, they are recognized in the following order: `CallerLineNumber`, `CallerFilePath`, `CallerMemberName`. Consider the following parameter declaration:

C#

```
[CallerMemberName, CallerFilePath, CallerLineNumber] object p = ...
```

`CallerLineNumber` takes precedence, and the other two attributes are ignored. If `CallerLineNumber` were omitted, `CallerFilePath` would take precedence, and `CallerMemberName` would be ignored. The lexical ordering of these attributes is irrelevant.

22.5.6.2 The CallerLineNumber attribute

The attribute `System.Runtime.CompilerServices.CallerLineNumberAttribute` is allowed on optional parameters when there is a standard implicit conversion (§10.4.2) from the constant value `int.MaxValue` to the parameter's type. This ensures that any non-negative line number up to that value can be passed without error.

If a function invocation from a location in source code omits an optional parameter with the `CallerLineNumberAttribute`, then a numeric literal representing that location's line number is used as an argument to the invocation instead of the default parameter value.

If the invocation spans multiple lines, the line chosen is implementation-dependent.

The line number may be affected by `#line` directives (§6.5.8).

22.5.6.3 The CallerFilePath attribute

The attribute `System.Runtime.CompilerServices.CallerFilePathAttribute` is allowed on optional parameters when there is a standard implicit conversion (§10.4.2) from `string` to the parameter's type.

If a function invocation from a location in source code omits an optional parameter with the `CallerFilePathAttribute`, then a string literal representing that location's file path is used as an argument to the invocation instead of the default parameter value.

The format of the file path is implementation-dependent.

The file path may be affected by `#line` directives (§6.5.8).

22.5.6.4 The CallerMemberName attribute

The attribute `System.Runtime.CompilerServices.CallerMemberNameAttribute` is allowed on optional parameters when there is a standard implicit conversion (§10.4.2) from `string` to the parameter's type.

If a function invocation from a location within the body of a function member or within an attribute applied to the function member itself or its return type, parameters or type parameters in source code omits an optional parameter with the

`CallerMemberNameAttribute`, then a string literal representing the name of that member is used as an argument to the invocation instead of the default parameter value.

For invocations that occur within generic methods, only the method name itself is used, without the type parameter list.

For invocations that occur within explicit interface member implementations, only the method name itself is used, without the preceding interface qualification.

For invocations that occur within property or event accessors, the member name used is that of the property or event itself.

For invocations that occur within indexer accessors, the member name used is that supplied by an `IndexerNameAttribute` (§22.6) on the indexer member, if present, or the default name `Item` otherwise.

For invocations that occur within field or event initializers, the member name used is the name of the field or event being initialized.

For invocations that occur within declarations of instance constructors, static constructors, finalizers and operators the member name used is implementation-dependent.

22.6 Attributes for interoperation

For interoperation with other languages, an indexer may be implemented using indexed properties. If no `IndexerName` attribute is present for an indexer, then the name `Item` is used by default. The `IndexerName` attribute enables a developer to override this default and specify a different name.

> *Example*: By default, an indexer's name is `Item`. This can be overridden, as follows:
>
> C#
>
> ```
> [System.Runtime.CompilerServices.IndexerName("TheItem")]
> public int this[int index]
> {
>     get { ... }
>     set { ... }
> }
> ```
>
> Now, the indexer's name is `TheItem`.
>
> *end example*

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

- Open a documentation issue
- Provide product feedback

23 Unsafe code

Article • 02/07/2024

23.1 General

An implementation that does not support unsafe code is required to diagnose any usage of the syntactic rules defined in this clause.

The remainder of this clause, including all of its subclauses, is conditionally normative.

> *Note*: The core C# language, as defined in the preceding clauses, differs notably from C and C++ in its omission of pointers as a data type. Instead, C# provides references and the ability to create objects that are managed by a garbage collector. This design, coupled with other features, makes C# a much safer language than C or C++. In the core C# language, it is simply not possible to have an uninitialized variable, a "dangling" pointer, or an expression that indexes an array beyond its bounds. Whole categories of bugs that routinely plague C and C++ programs are thus eliminated.
>
> While practically every pointer type construct in C or C++ has a reference type counterpart in C#, nonetheless, there are situations where access to pointer types becomes a necessity. For example, interfacing with the underlying operating system, accessing a memory-mapped device, or implementing a time-critical algorithm might not be possible or practical without access to pointers. To address this need, C# provides the ability to write ***unsafe code***.
>
> In unsafe code, it is possible to declare and operate on pointers, to perform conversions between pointers and integral types, to take the address of variables, and so forth. In a sense, writing unsafe code is much like writing C code within a C# program.
>
> Unsafe code is in fact a "safe" feature from the perspective of both developers and users. Unsafe code shall be clearly marked with the modifier `unsafe`, so developers can't possibly use unsafe features accidentally, and the execution engine works to ensure that unsafe code cannot be executed in an untrusted environment.
>
> *end note*

23.2 Unsafe contexts

The unsafe features of C# are available only in unsafe contexts. An unsafe context is introduced by including an `unsafe` modifier in the declaration of a type, member, or local function, or by employing an *unsafe_statement*:

- A declaration of a class, struct, interface, or delegate may include an `unsafe` modifier, in which case, the entire textual extent of that type declaration (including the body of the class, struct, or interface) is considered an unsafe context.

  > *Note*: If the *type_declaration* is partial, only that part is an unsafe context. *end note*

- A declaration of a field, method, property, event, indexer, operator, instance constructor, finalizer, static constructor, or local function may include an `unsafe` modifier, in which case, the entire textual extent of that member declaration is considered an unsafe context.
- An *unsafe_statement* enables the use of an unsafe context within a *block*. The entire textual extent of the associated *block* is considered an unsafe context. A local function declared within an unsafe context is itself unsafe.

The associated grammar extensions are shown below and in subsequent subclauses.

ANTLR

```
unsafe_modifier
    : 'unsafe'
    ;

unsafe_statement
    : 'unsafe' block
    ;
```

> *Example*: In the following code
>
> C#
>
> ```
> public unsafe struct Node
> {
>     public int Value;
>     public Node* Left;
>     public Node* Right;
> }
> ```
>
> the `unsafe` modifier specified in the struct declaration causes the entire textual extent of the struct declaration to become an unsafe context. Thus, it is possible to

declare the `Left` and `Right` fields to be of a pointer type. The example above could also be written

```
C#

public struct Node
{
    public int Value;
    public unsafe Node* Left;
    public unsafe Node* Right;
}
```

Here, the `unsafe` modifiers in the field declarations cause those declarations to be considered unsafe contexts.

end example

Other than establishing an unsafe context, thus permitting the use of pointer types, the `unsafe` modifier has no effect on a type or a member.

Example: In the following code

```
C#

public class A
{
    public unsafe virtual void F()
    {
        char* p;
        ...
    }
}

public class B : A
{
    public override void F()
    {
        base.F();
        ...
    }
}
```

the unsafe modifier on the `F` method in `A` simply causes the textual extent of `F` to become an unsafe context in which the unsafe features of the language can be used. In the override of `F` in `B`, there is no need to re-specify the `unsafe` modifier—unless, of course, the `F` method in `B` itself needs access to unsafe features.

> The situation is slightly different when a pointer type is part of the method's signature
>
> ```C#
> public unsafe class A
> {
>     public virtual void F(char* p) {...}
> }
>
> public class B: A
> {
>     public unsafe override void F(char* p) {...}
> }
> ```
>
> Here, because `F`'s signature includes a pointer type, it can only be written in an unsafe context. However, the unsafe context can be introduced by either making the entire class unsafe, as is the case in `A`, or by including an `unsafe` modifier in the method declaration, as is the case in `B`.
>
> *end example*

When the `unsafe` modifier is used on a partial type declaration (§15.2.7), only that particular part is considered an unsafe context.

23.3 Pointer types

In an unsafe context, a *type* (§8.1) can be a *pointer_type* as well as a *value_type*, a *reference_type*, or a *type_parameter*. In an unsafe context a *pointer_type* may also be the element type of an array (§17). A *pointer_type* may also be used in a typeof expression (§12.8.17) outside of an unsafe context (as such usage is not unsafe).

A *pointer_type* is written as an *unmanaged_type* (§8.8) or the keyword `void`, followed by a `*` token:

```ANTLR
pointer_type
    : value_type ('*')+
    | 'void' ('*')+
    ;
```

The type specified before the `*` in a pointer type is called the ***referent type*** of the pointer type. It represents the type of the variable to which a value of the pointer type

points.

A *pointer_type* may only be used in an *array_type* in an unsafe context (§23.2). A *non_array_type* is any type that is not itself an *array_type*.

Unlike references (values of reference types), pointers are not tracked by the garbage collector—the garbage collector has no knowledge of pointers and the data to which they point. For this reason a pointer is not permitted to point to a reference or to a struct that contains references, and the referent type of a pointer shall be an *unmanaged_type*. Pointer types themselves are unmanaged types, so a pointer type may be used as the referent type for another pointer type.

The intuitive rule for mixing of pointers and references is that referents of references (objects) are permitted to contain pointers, but referents of pointers are not permitted to contain references.

> *Example*: Some examples of pointer types are given in the table below:
>
> Expand table
>
> | Example | Description |
> | --- | --- |
> | `byte*` | Pointer to `byte` |
> | `char*` | Pointer to `char` |
> | `int**` | Pointer to pointer to `int` |
> | `int*[]` | Single-dimensional array of pointers to `int` |
> | `void*` | Pointer to unknown type |
>
> *end example*

For a given implementation, all pointer types shall have the same size and representation.

> *Note*: Unlike C and C++, when multiple pointers are declared in the same declaration, in C# the `*` is written along with the underlying type only, not as a prefix punctuator on each pointer name. For example:
>
> ```
> C#
> ```
>
> ```
> int* pi, pj; // NOT as int *pi, *pj;
> ```

> *end note*

The value of a pointer having type `T*` represents the address of a variable of type `T`. The pointer indirection operator `*` (§23.6.2) can be used to access this variable.

> *Example*: Given a variable `P` of type `int*`, the expression `*P` denotes the `int` variable found at the address contained in `P`. *end example*

Like an object reference, a pointer may be `null`. Applying the indirection operator to a `null`-valued pointer results in implementation-defined behavior (§23.6.2). A pointer with value `null` is represented by all-bits-zero.

The `void*` type represents a pointer to an unknown type. Because the referent type is unknown, the indirection operator cannot be applied to a pointer of type `void*`, nor can any arithmetic be performed on such a pointer. However, a pointer of type `void*` can be cast to any other pointer type (and vice versa) and compared to values of other pointer types (§23.6.8).

Pointer types are a separate category of types. Unlike reference types and value types, pointer types do not inherit from `object` and no conversions exist between pointer types and `object`. In particular, boxing and unboxing (§8.3.13) are not supported for pointers. However, conversions are permitted between different pointer types and between pointer types and the integral types. This is described in §23.5.

A *pointer_type* cannot be used as a type argument (§8.4), and type inference (§12.6.3) fails on generic method calls that would have inferred a type argument to be a pointer type.

A *pointer_type* cannot be used as a type of a subexpression of a dynamically bound operation (§12.3.3).

A *pointer_type* cannot be used as the type of the first parameter in an extension method (§15.6.10).

A *pointer_type* may be used as the type of a volatile field (§15.5.4).

The *dynamic erasure* of a type `E*` is the pointer type with referent type of the dynamic erasure of `E`.

An expression with a pointer type cannot be used to provide the value in a *member_declarator* within an *anonymous_object_creation_expression* (§12.8.16.7).

The default value (§9.3) for any pointer type is `null`.

> *Note*: Although pointers can be passed as `in`, `ref` or `out` parameters, doing so can cause undefined behavior, since the pointer might well be set to point to a local variable that no longer exists when the called method returns, or the fixed object to which it used to point, is no longer fixed. For example:
>
> C#
>
> ```
> class Test
> {
>     static int value = 20;
>
>     unsafe static void F(out int* pi1, ref int* pi2)
>     {
>         int i = 10;
>         pi1 = &i;       // return address of local variable
>         fixed (int* pj = &value)
>         {
>             // ...
>             pi2 = pj;   // return address that will soon not be fixed
>         }
>     }
>
>     static void Main()
>     {
>         int i = 15;
>         unsafe
>         {
>             int* px1;
>             int* px2 = &i;
>             F(out px1, ref px2);
>             int v1 = *px1; // undefined
>             int v2 = *px2; // undefined
>         }
>     }
> }
> ```
>
> *end note*

A method can return a value of some type, and that type can be a pointer.

> *Example*: When given a pointer to a contiguous sequence of `int`s, that sequence's element count, and some other `int` value, the following method returns the address of that value in that sequence, if a match occurs; otherwise it returns `null`:
>
> C#

```
unsafe static int* Find(int* pi, int size, int value)
{
    for (int i = 0; i < size; ++i)
    {
        if (*pi == value)
        {
            return pi;
        }
        ++pi;
    }
    return null;
}
```

end example

In an unsafe context, several constructs are available for operating on pointers:

- The unary `*` operator may be used to perform pointer indirection (§23.6.2).
- The `->` operator may be used to access a member of a struct through a pointer (§23.6.3).
- The `[]` operator may be used to index a pointer (§23.6.4).
- The unary `&` operator may be used to obtain the address of a variable (§23.6.5).
- The `++` and `--` operators may be used to increment and decrement pointers (§23.6.6).
- The binary `+` and `-` operators may be used to perform pointer arithmetic (§23.6.7).
- The `==`, `!=`, `<`, `>`, `<=`, and `>=` operators may be used to compare pointers (§23.6.8).
- The `stackalloc` operator may be used to allocate memory from the call stack (§23.9).
- The `fixed` statement may be used to temporarily fix a variable so its address can be obtained (§23.7).

23.4 Fixed and moveable variables

The address-of operator (§23.6.5) and the `fixed` statement (§23.7) divide variables into two categories: ***Fixed variables*** and ***moveable variables***.

Fixed variables reside in storage locations that are unaffected by operation of the garbage collector. (Examples of fixed variables include local variables, value parameters, and variables created by dereferencing pointers.) On the other hand, moveable variables reside in storage locations that are subject to relocation or disposal by the garbage

collector. (Examples of moveable variables include fields in objects and elements of arrays.)

The `&` operator (§23.6.5) permits the address of a fixed variable to be obtained without restrictions. However, because a moveable variable is subject to relocation or disposal by the garbage collector, the address of a moveable variable can only be obtained using a `fixed statement` (§23.7), and that address remains valid only for the duration of that `fixed` statement.

In precise terms, a fixed variable is one of the following:

- A variable resulting from a *simple_name* (§12.8.4) that refers to a local variable, value parameter, or parameter array, unless the variable is captured by an anonymous function (§12.19.6.2).
- A variable resulting from a *member_access* (§12.8.7) of the form `V.I`, where `V` is a fixed variable of a *struct_type*.
- A variable resulting from a *pointer_indirection_expression* (§23.6.2) of the form `*P`, a *pointer_member_access* (§23.6.3) of the form `P->I`, or a *pointer_element_access* (§23.6.4) of the form `P[E]`.

All other variables are classified as moveable variables.

A static field is classified as a moveable variable. Also, an `in`, `out`, or `ref` parameter is classified as a moveable variable, even if the argument given for the parameter is a fixed variable. Finally, a variable produced by dereferencing a pointer is always classified as a fixed variable.

23.5 Pointer conversions

23.5.1 General

In an unsafe context, the set of available implicit conversions (§10.2) is extended to include the following implicit pointer conversions:

- From any *pointer_type* to the type `void*`.
- From the `null` literal (§6.4.5.7) to any *pointer_type*.

Additionally, in an unsafe context, the set of available explicit conversions (§10.3) is extended to include the following explicit pointer conversions:

- From any *pointer_type* to any other *pointer_type*.
- From `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, or `ulong` to any *pointer_type*.

- From any *pointer_type* to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, or `ulong`.

Finally, in an unsafe context, the set of standard implicit conversions (§10.4.2) includes the following pointer conversions:

- From any *pointer_type* to the type `void*`.
- From the `null` literal to any *pointer_type*.

Conversions between two pointer types never change the actual pointer value. In other words, a conversion from one pointer type to another has no effect on the underlying address given by the pointer.

When one pointer type is converted to another, if the resulting pointer is not correctly aligned for the pointed-to type, the behavior is undefined if the result is dereferenced. In general, the concept "correctly aligned" is transitive: if a pointer to type `A` is correctly aligned for a pointer to type `B`, which, in turn, is correctly aligned for a pointer to type `C`, then a pointer to type `A` is correctly aligned for a pointer to type `C`.

> *Example*: Consider the following case in which a variable having one type is accessed via a pointer to a different type:
>
> C#
>
> ```
> unsafe static void M()
> {
>     char c = 'A';
>     char* pc = &c;
>     void* pv = pc;
>     int* pi = (int*)pv; // pretend a 16-bit char is a 32-bit int
>     int i = *pi;        // read 32-bit int; undefined
>     *pi = 123456;       // write 32-bit int; undefined
> }
> ```
>
> *end example*

When a pointer type is converted to a pointer to `byte`, the result points to the lowest addressed `byte` of the variable. Successive increments of the result, up to the size of the variable, yield pointers to the remaining bytes of that variable.

> *Example*: The following method displays each of the eight bytes in a `double` as a hexadecimal value:
>
> C#

```
class Test
{
    static void Main()
    {
        double d = 123.456e23;
        unsafe
        {
            byte* pb = (byte*)&d;
            for (int i = 0; i < sizeof(double); ++i)
            {
                Console.Write($" {*pb++:X2}");
            }
            Console.WriteLine();
        }
    }
}
```

Of course, the output produced depends on endianness. One possibility is `" BA FF 51 A2 90 6C 24 45"`.

end example

Mappings between pointers and integers are implementation-defined.

> *Note*: However, on 32- and 64-bit CPU architectures with a linear address space, conversions of pointers to or from integral types typically behave exactly like conversions of `uint` or `ulong` values, respectively, to or from those integral types. *end note*

23.5.2 Pointer arrays

Arrays of pointers can be constructed using *array_creation_expression* (§12.8.16.5) in an unsafe context. Only some of the conversions that apply to other array types are allowed on pointer arrays:

- The implicit reference conversion (§10.2.8) from any *array_type* to `System.Array` and the interfaces it implements also applies to pointer arrays. However, any attempt to access the array elements through `System.Array` or the interfaces it implements may result in an exception at run-time, as pointer types are not convertible to `object`.
- The implicit and explicit reference conversions (§10.2.8, §10.3.5) from a single-dimensional array type `S[]` to `System.Collections.Generic.IList<T>` and its generic base interfaces never apply to pointer arrays.

- The explicit reference conversion (§10.3.5) from `System.Array` and the interfaces it implements to any *array_type* applies to pointer arrays.
- The explicit reference conversions (§10.3.5) from `System.Collections.Generic.IList<S>` and its base interfaces to a single-dimensional array type `T[]` never applies to pointer arrays, since pointer types cannot be used as type arguments, and there are no conversions from pointer types to non-pointer types.

These restrictions mean that the expansion for the `foreach` statement over arrays described in §9.4.4.17 cannot be applied to pointer arrays. Instead, a `foreach` statement of the form

`foreach (V v in x)` *embedded_statement*

where the type of `x` is an array type of the form `T[,,...,]`, *n* is the number of dimensions minus 1 and `T` or `V` is a pointer type, is expanded using nested for-loops as follows:

C#

```
{
    T[,,...,] a = x;
    for (int i0 = a.GetLowerBound(0); i0 <= a.GetUpperBound(0); i0++)
    {
        for (int i1 = a.GetLowerBound(1); i1 <= a.GetUpperBound(1); i1++)
        {
            ...
            for (int in = a.GetLowerBound(n); in <= a.GetUpperBound(n);
in++)
            {
                V v = (V)a[i0,i1,...,in];
                *embedded_statement*
            }
        }
    }
}
```

The variables `a`, `i0`, `i1`, ... `in` are not visible to or accessible to `x` or the *embedded_statement* or any other source code of the program. The variable `v` is read-only in the embedded statement. If there is not an explicit conversion (§23.5) from `T` (the element type) to `V`, an error is produced and no further steps are taken. If `x` has the value `null`, a `System.NullReferenceException` is thrown at run-time.

> *Note*: Although pointer types are not permitted as type arguments, pointer arrays may be used as type arguments. *end note*

23.6 Pointers in expressions

23.6.1 General

In an unsafe context, an expression may yield a result of a pointer type, but outside an unsafe context, it is a compile-time error for an expression to be of a pointer type. In precise terms, outside an unsafe context a compile-time error occurs if any *simple_name* (§12.8.4), *member_access* (§12.8.7), *invocation_expression* (§12.8.9), or *element_access* (§12.8.11) is of a pointer type.

In an unsafe context, the *primary_no_array_creation_expression* (§12.8) and *unary_expression* (§12.9) productions permit additional constructs, which are described in the following subclauses.

> *Note*: The precedence and associativity of the unsafe operators is implied by the grammar. *end note*

23.6.2 Pointer indirection

A *pointer_indirection_expression* consists of an asterisk (`*`) followed by a *unary_expression*.

ANTLR

```
pointer_indirection_expression
    : '*' unary_expression
    ;
```

The unary `*` operator denotes pointer indirection and is used to obtain the variable to which a pointer points. The result of evaluating `*P`, where `P` is an expression of a pointer type `T*`, is a variable of type `T`. It is a compile-time error to apply the unary `*` operator to an expression of type `void*` or to an expression that isn't of a pointer type.

The effect of applying the unary `*` operator to a `null`-valued pointer is implementation-defined. In particular, there is no guarantee that this operation throws a `System.NullReferenceException`.

If an invalid value has been assigned to the pointer, the behavior of the unary `*` operator is undefined.

> *Note*: Among the invalid values for dereferencing a pointer by the unary `*` operator are an address inappropriately aligned for the type pointed to (see example in

§23.5), and the address of a variable after the end of its lifetime.

For purposes of definite assignment analysis, a variable produced by evaluating an expression of the form `*P` is considered initially assigned (§9.4.2).

23.6.3 Pointer member access

A *pointer_member_access* consists of a *primary_expression*, followed by a "`->`" token, followed by an *identifier* and an optional *type_argument_list*.

ANTLR

```
pointer_member_access
    : primary_expression '->' identifier type_argument_list?
    ;
```

In a pointer member access of the form `P->I`, `P` shall be an expression of a pointer type, and `I` shall denote an accessible member of the type to which `P` points.

A pointer member access of the form `P->I` is evaluated exactly as `(*P).I`. For a description of the pointer indirection operator (`*`), see §23.6.2. For a description of the member access operator (`.`), see §12.8.7.

Example: In the following code

C#

```
struct Point
{
    public int x;
    public int y;
    public override string ToString() => $"({x},{y})";
}

class Test
{
    static void Main()
    {
        Point point;
        unsafe
        {
            Point* p = &point;
            p->x = 10;
            p->y = 20;
            Console.WriteLine(p->ToString());
        }
```

```
    }
}
```

the `->` operator is used to access fields and invoke a method of a struct through a pointer. Because the operation `P->I` is precisely equivalent to `(*P).I`, the `Main` method could equally well have been written:

C#

```
class Test
{
    static void Main()
    {
        Point point;
        unsafe
        {
            Point* p = &point;
            (*p).x = 10;
            (*p).y = 20;
            Console.WriteLine((*p).ToString());
        }
    }
}
```

end example

23.6.4 Pointer element access

A *pointer_element_access* consists of a *primary_no_array_creation_expression* followed by an expression enclosed in "`[`" and "`]`".

ANTLR

```
pointer_element_access
    : primary_no_array_creation_expression '[' expression ']'
    ;
```

In a pointer element access of the form `P[E]`, `P` shall be an expression of a pointer type other than `void*`, and `E` shall be an expression that can be implicitly converted to `int`, `uint`, `long`, or `ulong`.

A pointer element access of the form `P[E]` is evaluated exactly as `*(P + E)`. For a description of the pointer indirection operator (`*`), see §23.6.2. For a description of the pointer addition operator (`+`), see §23.6.7.

Example: In the following code

```
C#
```

```
class Test
{
    static void Main()
    {
        unsafe
        {
            char* p = stackalloc char[256];
            for (int i = 0; i < 256; i++)
            {
                p[i] = (char)i;
            }
        }
    }
}
```

a pointer element access is used to initialize the character buffer in a `for` loop. Because the operation `P[E]` is precisely equivalent to `*(P + E)`, the example could equally well have been written:

```
C#
```

```
class Test
{
    static void Main()
    {
        unsafe
        {
            char* p = stackalloc char[256];
            for (int i = 0; i < 256; i++)
            {
                *(p + i) = (char)i;
            }
        }
    }
}
```

end example

The pointer element access operator does not check for out-of-bounds errors and the behavior when accessing an out-of-bounds element is undefined.

Note: This is the same as C and C++. *end note*

23.6.5 The address-of operator

An *addressof_expression* consists of an ampersand (`&`) followed by a *unary_expression*.

ANTLR

```
addressof_expression
    : '&' unary_expression
    ;
```

Given an expression `E` which is of a type `T` and is classified as a fixed variable (§23.4), the construct `&E` computes the address of the variable given by `E`. The type of the result is `T*` and is classified as a value. A compile-time error occurs if `E` is not classified as a variable, if `E` is classified as a read-only local variable, or if `E` denotes a moveable variable. In the last case, a fixed statement (§23.7) can be used to temporarily "fix" the variable before obtaining its address.

> *Note*: As stated in §12.8.7, outside an instance constructor or static constructor for a struct or class that defines a `readonly` field, that field is considered a value, not a variable. As such, its address cannot be taken. Similarly, the address of a constant cannot be taken.

The `&` operator does not require its argument to be definitely assigned, but following an `&` operation, the variable to which the operator is applied is considered definitely assigned in the execution path in which the operation occurs. It is the responsibility of the programmer to ensure that correct initialization of the variable actually does take place in this situation.

> *Example*: In the following code

C#

```
class Test
{
    static void Main()
    {
        int i;
        unsafe
        {
            int* p = &i;
            *p = 123;
        }
        Console.WriteLine(i);
```

```
    }
}
```

`i` is considered definitely assigned following the `&i` operation used to initialize `p`. The assignment to `*p` in effect initializes `i`, but the inclusion of this initialization is the responsibility of the programmer, and no compile-time error would occur if the assignment was removed.

end example

Note: The rules of definite assignment for the `&` operator exist such that redundant initialization of local variables can be avoided. For example, many external APIs take a pointer to a structure which is filled in by the API. Calls to such APIs typically pass the address of a local struct variable, and without the rule, redundant initialization of the struct variable would be required. *end note*

Note: When a local variable, value parameter, or parameter array is captured by an anonymous function (§12.8.23), that local variable, parameter, or parameter array is no longer considered to be a fixed variable (§23.7), but is instead considered to be a moveable variable. Thus it is an error for any unsafe code to take the address of a local variable, value parameter, or parameter array that has been captured by an anonymous function. *end note*

23.6.6 Pointer increment and decrement

In an unsafe context, the `++` and `--` operators (§12.8.15 and §12.9.6) can be applied to pointer variables of all types except `void*`. Thus, for every pointer type `T*`, the following operators are implicitly defined:

```
C#

T* operator ++(T* x);
T* operator --(T* x);
```

The operators produce the same results as `x+1` and `x-1`, respectively (§23.6.7). In other words, for a pointer variable of type `T*`, the `++` operator adds `sizeof(T)` to the address contained in the variable, and the `--` operator subtracts `sizeof(T)` from the address contained in the variable.

If a pointer increment or decrement operation overflows the domain of the pointer type, the result is implementation-defined, but no exceptions are produced.

23.6.7 Pointer arithmetic

In an unsafe context, the `+` operator (§12.10.5) and `-` operator (§12.10.6) can be applied to values of all pointer types except `void*`. Thus, for every pointer type `T*`, the following operators are implicitly defined:

```
C#
```

```
T* operator +(T* x, int y);
T* operator +(T* x, uint y);
T* operator +(T* x, long y);
T* operator +(T* x, ulong y);
T* operator +(int x, T* y);
T* operator +(uint x, T* y);
T* operator +(long x, T* y);
T* operator +(ulong x, T* y);
T* operator -(T* x, int y);
T* operator -(T* x, uint y);
T* operator -(T* x, long y);
T* operator -(T* x, ulong y);
long operator -(T* x, T* y);
```

Given an expression `P` of a pointer type `T*` and an expression `N` of type `int`, `uint`, `long`, or `ulong`, the expressions `P + N` and `N + P` compute the pointer value of type `T*` that results from adding `N * sizeof(T)` to the address given by `P`. Likewise, the expression `P - N` computes the pointer value of type `T*` that results from subtracting `N * sizeof(T)` from the address given by `P`.

Given two expressions, `P` and `Q`, of a pointer type `T*`, the expression `P - Q` computes the difference between the addresses given by `P` and `Q` and then divides that difference by `sizeof(T)`. The type of the result is always `long`. In effect, `P - Q` is computed as `((long)(P) - (long)(Q)) / sizeof(T)`.

> *Example*:
>
> ```
> C#
> ```
>
> ```
> class Test
> {
>     static void Main()
>     {
>         unsafe
>         {
>             int* values = stackalloc int[20];
>             int* p = &values[1];
>             int* q = &values[15];
>             Console.WriteLine($"p - q = {p - q}");
> ```

```
            Console.WriteLine($"q - p = {q - p}");
        }
    }
}
```

which produces the output:

Console

```
p - q = -14
q - p = 14
```

end example

If a pointer arithmetic operation overflows the domain of the pointer type, the result is truncated in an implementation-defined fashion, but no exceptions are produced.

23.6.8 Pointer comparison

In an unsafe context, the `==`, `!=`, `<`, `>`, `<=`, and `>=` operators (§12.12) can be applied to values of all pointer types. The pointer comparison operators are:

C#

```
bool operator ==(void* x, void* y);
bool operator !=(void* x, void* y);
bool operator <(void* x, void* y);
bool operator >(void* x, void* y);
bool operator <=(void* x, void* y);
bool operator >=(void* x, void* y);
```

Because an implicit conversion exists from any pointer type to the `void*` type, operands of any pointer type can be compared using these operators. The comparison operators compare the addresses given by the two operands as if they were unsigned integers.

23.6.9 The sizeof operator

For certain predefined types (§12.8.18), the `sizeof` operator yields a constant `int` value. For all other types, the result of the `sizeof` operator is implementation-defined and is classified as a value, not a constant.

The order in which members are packed into a struct is unspecified.

For alignment purposes, there may be unnamed padding at the beginning of a struct, within a struct, and at the end of the struct. The contents of the bits used as padding are indeterminate.

When applied to an operand that has struct type, the result is the total number of bytes in a variable of that type, including any padding.

23.7 The fixed statement

In an unsafe context, the *embedded_statement* (§13.1) production permits an additional construct, the fixed statement, which is used to "fix" a moveable variable such that its address remains constant for the duration of the statement.

ANTLR

```
fixed_statement
    : 'fixed' '(' pointer_type fixed_pointer_declarators ')'
embedded_statement
    ;

fixed_pointer_declarators
    : fixed_pointer_declarator (','  fixed_pointer_declarator)*
    ;

fixed_pointer_declarator
    : identifier '=' fixed_pointer_initializer
    ;

fixed_pointer_initializer
    : '&' variable_reference
    | expression
    ;
```

Each *fixed_pointer_declarator* declares a local variable of the given *pointer_type* and initializes that local variable with the address computed by the corresponding *fixed_pointer_initializer*. A local variable declared in a fixed statement is accessible in any *fixed_pointer_initializer*s occurring to the right of that variable's declaration, and in the *embedded_statement* of the fixed statement. A local variable declared by a fixed statement is considered read-only. A compile-time error occurs if the embedded statement attempts to modify this local variable (via assignment or the `++` and `--` operators) or pass it as a `ref` or `out` parameter.

It is an error to use a captured local variable (§12.19.6.2), value parameter, or parameter array in a *fixed_pointer_initializer*. A *fixed_pointer_initializer* can be one of the following:

- The token "`&`" followed by a *variable_reference* (§9.5) to a moveable variable (§23.4) of an unmanaged type `T`, provided the type `T*` is implicitly convertible to the pointer type given in the `fixed` statement. In this case, the initializer computes the address of the given variable, and the variable is guaranteed to remain at a fixed address for the duration of the fixed statement.
- An expression of an *array_type* with elements of an unmanaged type `T`, provided the type `T*` is implicitly convertible to the pointer type given in the fixed statement. In this case, the initializer computes the address of the first element in the array, and the entire array is guaranteed to remain at a fixed address for the duration of the `fixed` statement. If the array expression is `null` or if the array has zero elements, the initializer computes an address equal to zero.
- An expression of type `string`, provided the type `char*` is implicitly convertible to the pointer type given in the `fixed` statement. In this case, the initializer computes the address of the first character in the string, and the entire string is guaranteed to remain at a fixed address for the duration of the `fixed` statement. The behavior of the `fixed` statement is implementation-defined if the string expression is `null`.
- An expression of type other than *array_type* or `string`, provided there exists an accessible method or accessible extension method matching the signature `ref [readonly] T GetPinnableReference()`, where `T` is an *unmanaged_type*, and `T*` is implicitly convertible to the pointer type given in the `fixed` statement. In this case, the initializer computes the address of the returned variable, and that variable is guaranteed to remain at a fixed address for the duration of the `fixed` statement. A `GetPinnableReference()` method can be used by the `fixed` statement when overload resolution (§12.6.4) produces exactly one function member and that function member satisfies the preceding conditions. The `GetPinnableReference` method should return a reference to an address equal to zero, such as that returned from `System.Runtime.CompilerServices.Unsafe.NullRef<T>()` when there is no data to pin.
- A *simple_name* or *member_access* that references a fixed-size buffer member of a moveable variable, provided the type of the fixed-size buffer member is implicitly convertible to the pointer type given in the `fixed` statement. In this case, the initializer computes a pointer to the first element of the fixed-size buffer (§23.8.3), and the fixed-size buffer is guaranteed to remain at a fixed address for the duration of the `fixed` statement.

For each address computed by a *fixed_pointer_initializer* the `fixed` statement ensures that the variable referenced by the address is not subject to relocation or disposal by the garbage collector for the duration of the `fixed` statement.

> *Example*: If the address computed by a *fixed_pointer_initializer* references a field of an object or an element of an array instance, the fixed statement guarantees that the containing object instance is not relocated or disposed of during the lifetime of the statement. *end example*

It is the programmer's responsibility to ensure that pointers created by fixed statements do not survive beyond execution of those statements.

> *Example*: When pointers created by `fixed` statements are passed to external APIs, it is the programmer's responsibility to ensure that the APIs retain no memory of these pointers. *end example*

Fixed objects can cause fragmentation of the heap (because they can't be moved). For that reason, objects should be fixed only when absolutely necessary and then only for the shortest amount of time possible.

> *Example*: The example

C#

```
class Test
{
    static int x;
    int y;

    unsafe static void F(int* p)
    {
        *p = 1;
    }

    static void Main()
    {
        Test t = new Test();
        int[] a = new int[10];
        unsafe
        {
            fixed (int* p = &x) F(p);
            fixed (int* p = &t.y) F(p);
            fixed (int* p = &a[0]) F(p);
            fixed (int* p = a) F(p);
        }
    }
}
```

> demonstrates several uses of the `fixed` statement. The first statement fixes and obtains the address of a static field, the second statement fixes and obtains the

address of an instance field, and the third statement fixes and obtains the address of an array element. In each case, it would have been an error to use the regular `&` operator since the variables are all classified as moveable variables.

The third and fourth `fixed` statements in the example above produce identical results. In general, for an array instance `a`, specifying `a[0]` in a `fixed` statement is the same as simply specifying `a`.

end example

In an unsafe context, array elements of single-dimensional arrays are stored in increasing index order, starting with index `0` and ending with index `Length – 1`. For multi-dimensional arrays, array elements are stored such that the indices of the rightmost dimension are increased first, then the next left dimension, and so on to the left.

Within a `fixed` statement that obtains a pointer `p` to an array instance `a`, the pointer values ranging from `p` to `p + a.Length - 1` represent addresses of the elements in the array. Likewise, the variables ranging from `p[0]` to `p[a.Length - 1]` represent the actual array elements. Given the way in which arrays are stored, an array of any dimension can be treated as though it were linear.

Example:

C#

```
class Test
{
    static void Main()
    {
        int[,,] a = new int[2,3,4];
        unsafe
        {
            fixed (int* p = a)
            {
                for (int i = 0; i < a.Length; ++i) // treat as linear
                {
                    p[i] = i;
                }
            }
        }
        for (int i = 0; i < 2; ++i)
        {
            for (int j = 0; j < 3; ++j)
            {
                for (int k = 0; k < 4; ++k)
                {
```

```
                    Console.Write($"[{i},{j},{k}] = {a[i,j,k],2} ");
                }
                Console.WriteLine();
            }
        }
    }
}
```

which produces the output:

Console

```
[0,0,0] =  0 [0,0,1] =  1 [0,0,2] =  2 [0,0,3] =  3
[0,1,0] =  4 [0,1,1] =  5 [0,1,2] =  6 [0,1,3] =  7
[0,2,0] =  8 [0,2,1] =  9 [0,2,2] = 10 [0,2,3] = 11
[1,0,0] = 12 [1,0,1] = 13 [1,0,2] = 14 [1,0,3] = 15
[1,1,0] = 16 [1,1,1] = 17 [1,1,2] = 18 [1,1,3] = 19
[1,2,0] = 20 [1,2,1] = 21 [1,2,2] = 22 [1,2,3] = 23
```

end example

Example: In the following code

C#

```
class Test
{
    unsafe static void Fill(int* p, int count, int value)
    {
        for (; count != 0; count--)
        {
            *p++ = value;
        }
    }

    static void Main()
    {
        int[] a = new int[100];
        unsafe
        {
            fixed (int* p = a) Fill(p, 100, -1);
        }
    }
}
```

a `fixed` statement is used to fix an array so its address can be passed to a method that takes a pointer.

end example

A `char*` value produced by fixing a string instance always points to a null-terminated string. Within a fixed statement that obtains a pointer `p` to a string instance `s`, the pointer values ranging from `p` to `p + s.Length - 1` represent addresses of the characters in the string, and the pointer value `p + s.Length` always points to a null character (the character with value '\0').

> *Example*:
>
> C#
>
> ```
> class Test
> {
>     static string name = "xx";
>
>     unsafe static void F(char* p)
>     {
>         for (int i = 0; p[i] != '\0'; ++i)
>         {
>             System.Console.WriteLine(p[i]);
>         }
>     }
>
>     static void Main()
>     {
>         unsafe
>         {
>             fixed (char* p = name) F(p);
>             fixed (char* p = "xx") F(p);
>         }
>     }
> }
> ```
>
> *end example*

> *Example*: The following code shows a *fixed_pointer_initializer* with an expression of type other than *array_type* or `string`:
>
> C#
>
> ```
> public class C
> {
>     private int _value;
>     public C(int value) => _value = value;
>     public ref int GetPinnableReference() => ref _value;
> }
>
> public class Test
> {
>     unsafe private static void Main()
> ```

```
    {
        C c = new C(10);
        fixed (int* p = c)
        {
            // ...
        }
    }
}
```

Type `C` has an accessible `GetPinnableReference` method with the correct signature. In the `fixed` statement, the `ref int` returned from that method when it is called on `c` is used to initialize the `int*` pointer `p`. *end example*

Modifying objects of managed type through fixed pointers can result in undefined behavior.

> *Note*: For example, because strings are immutable, it is the programmer's responsibility to ensure that the characters referenced by a pointer to a fixed string are not modified. *end note*

> *Note*: The automatic null-termination of strings is particularly convenient when calling external APIs that expect "C-style" strings. Note, however, that a string instance is permitted to contain null characters. If such null characters are present, the string will appear truncated when treated as a null-terminated `char*`. *end note*

23.8 Fixed-size buffers

23.8.1 General

Fixed-size buffers are used to declare "C-style" in-line arrays as members of structs, and are primarily useful for interfacing with unmanaged APIs.

23.8.2 Fixed-size buffer declarations

A ***fixed-size buffer*** is a member that represents storage for a fixed-length buffer of variables of a given type. A fixed-size buffer declaration introduces one or more fixed-size buffers of a given element type.

> *Note*: Like an array, a fixed-size buffer can be thought of as containing elements. As such, the term *element type* as defined for an array is also used with a fixed-size buffer. *end note*

Fixed-size buffers are only permitted in struct declarations and may only occur in unsafe contexts (§23.2).

ANTLR

```
fixed_size_buffer_declaration
    : attributes? fixed_size_buffer_modifier* 'fixed' buffer_element_type
      fixed_size_buffer_declarators ';'
    ;

fixed_size_buffer_modifier
    : 'new'
    | 'public'
    | 'internal'
    | 'private'
    | 'unsafe'
    ;

buffer_element_type
    : type
    ;

fixed_size_buffer_declarators
    : fixed_size_buffer_declarator (',' fixed_size_buffer_declarator)*
    ;

fixed_size_buffer_declarator
    : identifier '[' constant_expression ']'
    ;
```

A fixed-size buffer declaration may include a set of attributes (§22), a `new` modifier (§15.3.5), accessibility modifiers corresponding to any of the declared accessibilities permitted for struct members (§16.4.3) and an `unsafe` modifier (§23.2). The attributes and modifiers apply to all of the members declared by the fixed-size buffer declaration. It is an error for the same modifier to appear multiple times in a fixed-size buffer declaration.

A fixed-size buffer declaration is not permitted to include the `static` modifier.

The buffer element type of a fixed-size buffer declaration specifies the element type of the buffer(s) introduced by the declaration. The buffer element type shall be one of the predefined types `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, `double`, or `bool`.

The buffer element type is followed by a list of fixed-size buffer declarators, each of which introduces a new member. A fixed-size buffer declarator consists of an identifier that names the member, followed by a constant expression enclosed in `[` and `]` tokens.

The constant expression denotes the number of elements in the member introduced by that fixed-size buffer declarator. The type of the constant expression shall be implicitly convertible to type `int`, and the value shall be a non-zero positive integer.

The elements of a fixed-size buffer shall be laid out sequentially in memory.

A fixed-size buffer declaration that declares multiple fixed-size buffers is equivalent to multiple declarations of a single fixed-size buffer declaration with the same attributes, and element types.

> *Example*:
>
> ```C#
> unsafe struct A
> {
>     public fixed int x[5], y[10], z[100];
> }
> ```
>
> is equivalent to
>
> ```C#
> unsafe struct A
> {
>     public fixed int x[5];
>     public fixed int y[10];
>     public fixed int z[100];
> }
> ```
>
> *end example*

23.8.3 Fixed-size buffers in expressions

Member lookup (§12.5) of a fixed-size buffer member proceeds exactly like member lookup of a field.

A fixed-size buffer can be referenced in an expression using a *simple_name* (§12.8.4), a *member_access* (§12.8.7), or an *element_access* (§12.8.11).

When a fixed-size buffer member is referenced as a simple name, the effect is the same as a member access of the form `this.I`, where `I` is the fixed-size buffer member.

In a member access of the form `E.I` where `E.` may be the implicit `this.`, if `E` is of a struct type and a member lookup of `I` in that struct type identifies a fixed-size member,

then `E.I` is evaluated and classified as follows:

- If the expression `E.I` does not occur in an unsafe context, a compile-time error occurs.
- If `E` is classified as a value, a compile-time error occurs.
- Otherwise, if `E` is a moveable variable (§23.4) then:
  - If the expression `E.I` is a *fixed_pointer_initializer* (§23.7), then the result of the expression is a pointer to the first element of the fixed size buffer member `I` in `E`.
  - Otherwise if the expression `E.I` is a *primary_no_array_creation_expression* (§12.8.11.1) within an *element_access* (§12.8.11) of the form `E.I[J]`, then the result of `E.I` is a pointer, `P`, to the first element of the fixed size buffer member `I` in `E`, and the enclosing *element_access* is then evaluated as the *pointer_element_access* (§23.6.4) `P[J]`.
  - Otherwise a compile-time error occurs.
- Otherwise, `E` references a fixed variable and the result of the expression is a pointer to the first element of the fixed-size buffer member `I` in `E`. The result is of type `S*`, where S is the element type of `I`, and is classified as a value.

The subsequent elements of the fixed-size buffer can be accessed using pointer operations from the first element. Unlike access to arrays, access to the elements of a fixed-size buffer is an unsafe operation and is not range checked.

> *Example*: The following declares and uses a struct with a fixed-size buffer member.

C#

```
unsafe struct Font
{
    public int size;
    public fixed char name[32];
}

class Test
{
    unsafe static void PutString(string s, char* buffer, int bufSize)
    {
        int len = s.Length;
        if (len > bufSize)
        {
            len = bufSize;
        }
        for (int i = 0; i < len; i++)
        {
            buffer[i] = s[i];
        }
```

```
            for (int i = len; i < bufSize; i++)
            {
                buffer[i] = (char)0;
            }
        }

        unsafe static void Main()
        {
            Font f;
            f.size = 10;
            PutString("Times New Roman", f.name, 32);
        }
    }
```

end example

23.8.4 Definite assignment checking

Fixed-size buffers are not subject to definite assignment-checking (§9.4), and fixed-size buffer members are ignored for purposes of definite-assignment checking of struct type variables.

When the outermost containing struct variable of a fixed-size buffer member is a static variable, an instance variable of a class instance, or an array element, the elements of the fixed-size buffer are automatically initialized to their default values (§9.3). In all other cases, the initial content of a fixed-size buffer is undefined.

23.9 Stack allocation

See §12.8.21 for general information about the operator `stackalloc`. Here, the ability of that operator to result in a pointer is discussed.

In an unsafe context if a *stackalloc_expression* (§12.8.21) occurs as the initializing expression of a *local_variable_declaration* (§13.6.2), where the *local_variable_type* is either a pointer type (§23.3) or inferred (`var`), then the result of the *stackalloc_expression* is a pointer of type `T *` to be beginning of the allocated block, where `T` is the *unmanaged_type* of the *stackalloc_expression*.

In all other respects the semantics of *local_variable_declaration*s (§13.6.2) and *stackalloc_expressions* (§12.8.21) in unsafe contexts follow those defined for safe contexts.

Example:

C#

```
unsafe
{
    // Memory uninitialized
    int* p1 = stackalloc int[3];
    // Memory initialized
    int* p2 = stackalloc int[3] { -10, -15, -30 };
    // Type int is inferred
    int* p3 = stackalloc[] { 11, 12, 13 };
    // Can't infer context, so pointer result assumed
    var p4 = stackalloc[] { 11, 12, 13 };
    // Error; no conversion exists
    long* p5 = stackalloc[] { 11, 12, 13 };
    // Converts 11 and 13, and returns long*
    long* p6 = stackalloc[] { 11, 12L, 13 };
    // Converts all and returns long*
    long* p7 = stackalloc long[] { 11, 12, 13 };
}
```

end example

Unlike access to arrays or `stackalloc`'ed blocks of `Span<T>` type, access to the elements of a `stackalloc`'ed block of pointer type is an unsafe operation and is not range checked.

Example: In the following code

C#

```
class Test
{
    static string IntToString(int value)
    {
        if (value == int.MinValue)
        {
            return "-2147483648";
        }
        int n = value >= 0 ? value : -value;
        unsafe
        {
            char* buffer = stackalloc char[16];
            char* p = buffer + 16;
            do
            {
                *--p = (char)(n % 10 + '0');
                n /= 10;
            } while (n != 0);
            if (value < 0)
            {
```

```
                *--p = '-';
            }
            return new string(p, 0, (int)(buffer + 16 - p));
        }
    }

    static void Main()
    {
        Console.WriteLine(IntToString(12345));
        Console.WriteLine(IntToString(-999));
    }
}
```

a `stackalloc` expression is used in the `IntToString` method to allocate a buffer of 16 characters on the stack. The buffer is automatically discarded when the method returns.

Note, however, that `IntToString` can be rewritten in safe mode; that is, without using pointers, as follows:

C#

```
class Test
{
    static string IntToString(int value)
    {
        if (value == int.MinValue)
        {
            return "-2147483648";
        }
        int n = value >= 0 ? value : -value;
        Span<char> buffer = stackalloc char[16];
        int idx = 16;
        do
        {
            buffer[--idx] = (char)(n % 10 + '0');
            n /= 10;
        } while (n != 0);
        if (value < 0)
        {
            buffer[--idx] = '-';
        }
        return buffer.Slice(idx).ToString();
    }
}
```

end example

End of conditionally normative text.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

- Open a documentation issue
- Provide product feedback

Annex A Grammar

Article • 02/07/2024

This clause is informative.

A.1 General

This annex contains the grammar productions found in the specification, including the optional ones for unsafe code. Productions appear here in the same order in which they appear in the specification.

A.2 Lexical grammar

ANTLR

```
// Source: §6.3.1 General
DEFAULT  : 'default' ;
NULL     : 'null' ;
TRUE     : 'true' ;
FALSE    : 'false' ;
ASTERISK : '*' ;
SLASH    : '/' ;

// Source: §6.3.1 General
input
    : input_section?
    ;

input_section
    : input_section_part+
    ;

input_section_part
    : input_element* New_Line
    | PP_Directive
    ;

input_element
    : Whitespace
    | Comment
    | token
    ;

// Source: §6.3.2 Line terminators
New_Line
    : New_Line_Character
```

```
    | '\u000D\u000A'    // carriage return, line feed
    ;

// Source: §6.3.3 Comments
Comment
    : Single_Line_Comment
    | Delimited_Comment
    ;

fragment Single_Line_Comment
    : '//' Input_Character*
    ;

fragment Input_Character
    // anything but New_Line_Character
    : ~('\u000D' | '\u000A'   | '\u0085' | '\u2028' | '\u2029')
    ;

fragment New_Line_Character
    : '\u000D'  // carriage return
    | '\u000A'  // line feed
    | '\u0085'  // next line
    | '\u2028'  // line separator
    | '\u2029'  // paragraph separator
    ;

fragment Delimited_Comment
    : '/*' Delimited_Comment_Section* ASTERISK+ '/'
    ;

fragment Delimited_Comment_Section
    : SLASH
    | ASTERISK* Not_Slash_Or_Asterisk
    ;

fragment Not_Slash_Or_Asterisk
    : ~('/' | '*')    // Any except SLASH or ASTERISK
    ;

// Source: §6.3.4 White space
Whitespace
    : [\p{Zs}]  // any character with Unicode class Zs
    | '\u0009'  // horizontal tab
    | '\u000B'  // vertical tab
    | '\u000C'  // form feed
    ;

// Source: §6.4.1 General
token
    : identifier
    | keyword
    | Integer_Literal
    | Real_Literal
    | Character_Literal
    | String_Literal
```

```
    | operator_or_punctuator
    ;

// Source: §6.4.2 Unicode character escape sequences
fragment Unicode_Escape_Sequence
    : '\\u' Hex_Digit Hex_Digit Hex_Digit Hex_Digit
    | '\\U' Hex_Digit Hex_Digit Hex_Digit Hex_Digit
            Hex_Digit Hex_Digit Hex_Digit Hex_Digit
    ;

// Source: §6.4.3 Identifiers
identifier
    : Simple_Identifier
    | contextual_keyword
    ;

Simple_Identifier
    : Available_Identifier
    | Escaped_Identifier
    ;

fragment Available_Identifier
    // excluding keywords or contextual keywords, see note below
    : Basic_Identifier
    ;

fragment Escaped_Identifier
    // Includes keywords and contextual keywords prefixed by '@'.
    // See note below.
    : '@' Basic_Identifier
    ;

fragment Basic_Identifier
    : Identifier_Start_Character Identifier_Part_Character*
    ;

fragment Identifier_Start_Character
    : Letter_Character
    | Underscore_Character
    ;

fragment Underscore_Character
    : '_'               // underscore
    | '\\u005' [fF]     // Unicode_Escape_Sequence for underscore
    | '\\U0000005' [fF] // Unicode_Escape_Sequence for underscore
    ;

fragment Identifier_Part_Character
    : Letter_Character
    | Decimal_Digit_Character
    | Connecting_Character
    | Combining_Character
    | Formatting_Character
    ;
```

```
fragment Letter_Character
    // Category Letter, all subcategories; category Number, subcategory
letter.
    : [\p{L}\p{Nl}]
    // Only escapes for categories L & Nl allowed. See note below.
    | Unicode_Escape_Sequence
    ;

fragment Combining_Character
    // Category Mark, subcategories non-spacing and spacing combining.
    : [\p{Mn}\p{Mc}]
    // Only escapes for categories Mn & Mc allowed. See note below.
    | Unicode_Escape_Sequence
    ;

fragment Decimal_Digit_Character
    // Category Number, subcategory decimal digit.
    : [\p{Nd}]
    // Only escapes for category Nd allowed. See note below.
    | Unicode_Escape_Sequence
    ;

fragment Connecting_Character
    // Category Punctuation, subcategory connector.
    : [\p{Pc}]
    // Only escapes for category Pc allowed. See note below.
    | Unicode_Escape_Sequence
    ;

fragment Formatting_Character
    // Category Other, subcategory format.
    : [\p{Cf}]
    // Only escapes for category Cf allowed, see note below.
    | Unicode_Escape_Sequence
    ;

// Source: §6.4.4 Keywords
keyword
    : 'abstract' | 'as'       | 'base'       | 'bool'      | 'break'
    | 'byte'     | 'case'     | 'catch'      | 'char'      | 'checked'
    | 'class'    | 'const'    | 'continue'   | 'decimal'   | DEFAULT
    | 'delegate' | 'do'       | 'double'     | 'else'      | 'enum'
    | 'event'    | 'explicit' | 'extern'     | FALSE       | 'finally'
    | 'fixed'    | 'float'    | 'for'        | 'foreach'   | 'goto'
    | 'if'       | 'implicit' | 'in'         | 'int'       | 'interface'
    | 'internal' | 'is'       | 'lock'       | 'long'      | 'namespace'
    | 'new'      | NULL       | 'object'     | 'operator'  | 'out'
    | 'override' | 'params'   | 'private'    | 'protected' | 'public'
    | 'readonly' | 'ref'      | 'return'     | 'sbyte'     | 'sealed'
    | 'short'    | 'sizeof'   | 'stackalloc' | 'static'    | 'string'
    | 'struct'   | 'switch'   | 'this'       | 'throw'     | TRUE
    | 'try'      | 'typeof'   | 'uint'       | 'ulong'     | 'unchecked'
    | 'unsafe'   | 'ushort'   | 'using'      | 'virtual'   | 'void'
    | 'volatile' | 'while'
    ;
```

```
// Source: §6.4.4 Keywords
contextual_keyword
    : 'add'    | 'alias'      | 'ascending' | 'async'     | 'await'
    | 'by'     | 'descending' | 'dynamic'   | 'equals'    | 'from'
    | 'get'    | 'global'     | 'group'     | 'into'      | 'join'
    | 'let'    | 'nameof'     | 'on'        | 'orderby'   | 'partial'
    | 'remove' | 'select'     | 'set'       | 'unmanaged' | 'value'
    | 'var'    | 'when'       | 'where'     | 'yield'
    ;

// Source: §6.4.5.1 General
literal
    : boolean_literal
    | Integer_Literal
    | Real_Literal
    | Character_Literal
    | String_Literal
    | null_literal
    ;

// Source: §6.4.5.2 Boolean literals
boolean_literal
    : TRUE
    | FALSE
    ;

// Source: §6.4.5.3 Integer literals
Integer_Literal
    : Decimal_Integer_Literal
    | Hexadecimal_Integer_Literal
    | Binary_Integer_Literal
    ;

fragment Decimal_Integer_Literal
    : Decimal_Digit Decorated_Decimal_Digit* Integer_Type_Suffix?
    ;

fragment Decorated_Decimal_Digit
    : '_'* Decimal_Digit
    ;

fragment Decimal_Digit
    : '0'..'9'
    ;

fragment Integer_Type_Suffix
    : 'U' | 'u' | 'L' | 'l' |
      'UL' | 'Ul' | 'uL' | 'ul' | 'LU' | 'Lu' | 'lU' | 'lu'
    ;

fragment Hexadecimal_Integer_Literal
    : ('0x' | '0X') Decorated_Hex_Digit+ Integer_Type_Suffix?
    ;
```

```
fragment Decorated_Hex_Digit
    : '_'* Hex_Digit
    ;

fragment Hex_Digit
    : '0'..'9' | 'A'..'F' | 'a'..'f'
    ;

fragment Binary_Integer_Literal
    : ('0b' | '0B') Decorated_Binary_Digit+ Integer_Type_Suffix?
    ;

fragment Decorated_Binary_Digit
    : '_'* Binary_Digit
    ;

fragment Binary_Digit
    : '0' | '1'
    ;

// Source: §6.4.5.4 Real literals
Real_Literal
    : Decimal_Digit Decorated_Decimal_Digit* '.'
      Decimal_Digit Decorated_Decimal_Digit* Exponent_Part?
Real_Type_Suffix?
    | '.' Decimal_Digit Decorated_Decimal_Digit* Exponent_Part?
Real_Type_Suffix?
    | Decimal_Digit Decorated_Decimal_Digit* Exponent_Part Real_Type_Suffix?
    | Decimal_Digit Decorated_Decimal_Digit* Real_Type_Suffix
    ;

fragment Exponent_Part
    : ('e' | 'E') Sign? Decimal_Digit Decorated_Decimal_Digit*
    ;

fragment Sign
    : '+' | '-'
    ;

fragment Real_Type_Suffix
    : 'F' | 'f' | 'D' | 'd' | 'M' | 'm'
    ;

// Source: §6.4.5.5 Character literals
Character_Literal
    : '\'' Character '\''
    ;

fragment Character
    : Single_Character
    | Simple_Escape_Sequence
    | Hexadecimal_Escape_Sequence
    | Unicode_Escape_Sequence
    ;
```

```
fragment Single_Character
    // anything but ', \, and New_Line_Character
    : ~['\\\u000D\u000A\u0085\u2028\u2029]
    ;

fragment Simple_Escape_Sequence
    : '\\'' | '\\"' | '\\\\' | '\\0' | '\\a' | '\\b' |
      '\\f' | '\\n' | '\\r' | '\\t' | '\\v'
    ;

fragment Hexadecimal_Escape_Sequence
    : '\\x' Hex_Digit Hex_Digit? Hex_Digit? Hex_Digit?
    ;

// Source: §6.4.5.6 String literals
String_Literal
    : Regular_String_Literal
    | Verbatim_String_Literal
    ;

fragment Regular_String_Literal
    : '"' Regular_String_Literal_Character* '"'
    ;

fragment Regular_String_Literal_Character
    : Single_Regular_String_Literal_Character
    | Simple_Escape_Sequence
    | Hexadecimal_Escape_Sequence
    | Unicode_Escape_Sequence
    ;

fragment Single_Regular_String_Literal_Character
    // anything but ", \, and New_Line_Character
    : ~["\\\u000D\u000A\u0085\u2028\u2029]
    ;

fragment Verbatim_String_Literal
    : '@"' Verbatim_String_Literal_Character* '"'
    ;

fragment Verbatim_String_Literal_Character
    : Single_Verbatim_String_Literal_Character
    | Quote_Escape_Sequence
    ;

fragment Single_Verbatim_String_Literal_Character
    : ~["]     // anything but quotation mark (U+0022)
    ;

fragment Quote_Escape_Sequence
    : '""'
    ;

// Source: §6.4.5.7 The null literal
null_literal
```

```
    : NULL
    ;

// Source: §6.4.6 Operators and punctuators
operator_or_punctuator
    : '{'  | '}'  | '['  | ']'  | '('  | ')'  | '.'  | ','  | ':'  | ';'
    | '+'  | '-'  | ASTERISK   | SLASH | '%'  | '&'  | '|'  | '^'  | '!' |
'~'
    | '='  | '<'  | '>'  | '?'  | '??'  | '::' | '++' | '--' | '&&' | '||'
    | '->' | '==' | '!=' | '<=' | '>='  | '+=' | '-=' | '*=' | '/=' | '%='
    | '&=' | '|=' | '^=' | '<<' | '<<=' | '=>'
    ;

right_shift
    : '>'  '>'
    ;

right_shift_assignment
    : '>' '>='
    ;

// Source: §6.5.1 General
PP_Directive
    : PP_Start PP_Kind PP_New_Line
    ;

fragment PP_Kind
    : PP_Declaration
    | PP_Conditional
    | PP_Line
    | PP_Diagnostic
    | PP_Region
    | PP_Pragma
    ;

// Only recognised at the beginning of a line
fragment PP_Start
    // See note below.
    : { getCharPositionInLine() == 0 }? PP_Whitespace? '#' PP_Whitespace?
    ;

fragment PP_Whitespace
    : ( [\p{Zs}]  // any character with Unicode class Zs
      | '\u0009'  // horizontal tab
      | '\u000B'  // vertical tab
      | '\u000C'  // form feed
      )+
    ;

fragment PP_New_Line
    : PP_Whitespace? Single_Line_Comment? New_Line
    ;

// Source: §6.5.2 Conditional compilation symbols
fragment PP_Conditional_Symbol
```

```
    // Must not be equal to tokens TRUE or FALSE. See note below.
    : Basic_Identifier
    ;

// Source: §6.5.3 Pre-processing expressions
fragment PP_Expression
    : PP_Whitespace? PP_Or_Expression PP_Whitespace?
    ;

fragment PP_Or_Expression
    : PP_And_Expression (PP_Whitespace? '||' PP_Whitespace?
PP_And_Expression)*
    ;

fragment PP_And_Expression
    : PP_Equality_Expression (PP_Whitespace? '&&' PP_Whitespace?
      PP_Equality_Expression)*
    ;

fragment PP_Equality_Expression
    : PP_Unary_Expression (PP_Whitespace? ('==' | '!=') PP_Whitespace?
      PP_Unary_Expression)*
    ;

fragment PP_Unary_Expression
    : PP_Primary_Expression
    | '!' PP_Whitespace? PP_Unary_Expression
    ;

fragment PP_Primary_Expression
    : TRUE
    | FALSE
    | PP_Conditional_Symbol
    | '(' PP_Whitespace? PP_Expression PP_Whitespace? ')'
    ;

// Source: §6.5.4 Definition directives
fragment PP_Declaration
    : 'define' PP_Whitespace PP_Conditional_Symbol
    | 'undef' PP_Whitespace PP_Conditional_Symbol
    ;

// Source: §6.5.5 Conditional compilation directives
fragment PP_Conditional
    : PP_If_Section
    | PP_Elif_Section
    | PP_Else_Section
    | PP_Endif
    ;

fragment PP_If_Section
    : 'if' PP_Whitespace PP_Expression
    ;

fragment PP_Elif_Section
```

```
    : 'elif' PP_Whitespace PP_Expression
    ;

fragment PP_Else_Section
    : 'else'
    ;

fragment PP_Endif
    : 'endif'
    ;

// Source: §6.5.6 Diagnostic directives
fragment PP_Diagnostic
    : 'error' PP_Message?
    | 'warning' PP_Message?
    ;

fragment PP_Message
    : PP_Whitespace Input_Character*
    ;

// Source: §6.5.7 Region directives
fragment PP_Region
    : PP_Start_Region
    | PP_End_Region
    ;

fragment PP_Start_Region
    : 'region' PP_Message?
    ;

fragment PP_End_Region
    : 'endregion' PP_Message?
    ;

// Source: §6.5.8 Line directives
fragment PP_Line
    : 'line' PP_Whitespace PP_Line_Indicator
    ;

fragment PP_Line_Indicator
    : Decimal_Digit+ PP_Whitespace PP_Compilation_Unit_Name
    | Decimal_Digit+
    | DEFAULT
    | 'hidden'
    ;

fragment PP_Compilation_Unit_Name
    : '"' PP_Compilation_Unit_Name_Character+ '"'
    ;

fragment PP_Compilation_Unit_Name_Character
    // Any Input_Character except "
    : ~('\u000D' | '\u000A'   | '\u0085' | '\u2028' | '\u2029' | '#')
    ;
```

```
// Source: §6.5.9 Pragma directives
fragment PP_Pragma
    : 'pragma' PP_Pragma_Text?
    ;

fragment PP_Pragma_Text
    : PP_Whitespace Input_Character*
    ;
```

A.3 Syntactic grammar

ANTLR

```
// Source: §7.8.1 General
namespace_name
    : namespace_or_type_name
    ;

type_name
    : namespace_or_type_name
    ;

namespace_or_type_name
    : identifier type_argument_list?
    | namespace_or_type_name '.' identifier type_argument_list?
    | qualified_alias_member
    ;

// Source: §8.1 General
type
    : reference_type
    | value_type
    | type_parameter
    | pointer_type     // unsafe code support
    ;

// Source: §8.2.1 General
reference_type
    : class_type
    | interface_type
    | array_type
    | delegate_type
    | 'dynamic'
    ;

class_type
    : type_name
    | 'object'
    | 'string'
    ;
```

```
interface_type
    : type_name
    ;

array_type
    : non_array_type rank_specifier+
    ;

non_array_type
    : value_type
    | class_type
    | interface_type
    | delegate_type
    | 'dynamic'
    | type_parameter
    | pointer_type      // unsafe code support
    ;

rank_specifier
    : '[' ','* ']'
    ;

delegate_type
    : type_name
    ;

// Source: §8.3.1 General
value_type
    : non_nullable_value_type
    | nullable_value_type
    ;

non_nullable_value_type
    : struct_type
    | enum_type
    ;

struct_type
    : type_name
    | simple_type
    | tuple_type
    ;

simple_type
    : numeric_type
    | 'bool'
    ;

numeric_type
    : integral_type
    | floating_point_type
    | 'decimal'
    ;
```

```
integral_type
    : 'sbyte'
    | 'byte'
    | 'short'
    | 'ushort'
    | 'int'
    | 'uint'
    | 'long'
    | 'ulong'
    | 'char'
    ;

floating_point_type
    : 'float'
    | 'double'
    ;

tuple_type
    : '(' tuple_type_element (',' tuple_type_element)+ ')'
    ;

tuple_type_element
    : type identifier?
    ;

enum_type
    : type_name
    ;

nullable_value_type
    : non_nullable_value_type '?'
    ;

// Source: §8.4.2 Type arguments
type_argument_list
    : '<' type_arguments '>'
    ;

type_arguments
    : type_argument (',' type_argument)*
    ;

type_argument
    : type
    ;

// Source: §8.5 Type parameters
type_parameter
    : identifier
    ;

// Source: §8.8 Unmanaged types
unmanaged_type
    : value_type
    | pointer_type     // unsafe code support
```

```
    ;

// Source: §9.5 Variable references
variable_reference
    : expression
    ;

// Source: §11.2.1 General
pattern
    : declaration_pattern
    | constant_pattern
    | var_pattern
    ;

// Source: §11.2.2 Declaration pattern
declaration_pattern
    : type simple_designation
    ;
simple_designation
    : single_variable_designation
    ;
single_variable_designation
    : identifier
    ;

// Source: §11.2.3 Constant pattern
constant_pattern
    : constant_expression
    ;

// Source: §11.2.4 Var pattern
var_pattern
    : 'var' designation
    ;
designation
    : simple_designation
    ;

// Source: §12.6.2.1 General
argument_list
    : argument (',' argument)*
    ;

argument
    : argument_name? argument_value
    ;

argument_name
    : identifier ':'
    ;

argument_value
    : expression
    | 'in' variable_reference
    | 'ref' variable_reference
```

```
    | 'out' variable_reference
    ;

// Source: §12.8.1 General
primary_expression
    : primary_no_array_creation_expression
    | array_creation_expression
    ;

primary_no_array_creation_expression
    : literal
    | interpolated_string_expression
    | simple_name
    | parenthesized_expression
    | tuple_expression
    | member_access
    | null_conditional_member_access
    | invocation_expression
    | element_access
    | null_conditional_element_access
    | this_access
    | base_access
    | post_increment_expression
    | post_decrement_expression
    | object_creation_expression
    | delegate_creation_expression
    | anonymous_object_creation_expression
    | typeof_expression
    | sizeof_expression
    | checked_expression
    | unchecked_expression
    | default_value_expression
    | nameof_expression
    | anonymous_method_expression
    | pointer_member_access     // unsafe code support
    | pointer_element_access    // unsafe code support
    | stackalloc_expression
    ;

// Source: §12.8.3 Interpolated string expressions
interpolated_string_expression
    : interpolated_regular_string_expression
    | interpolated_verbatim_string_expression
    ;

// interpolated regular string expressions

interpolated_regular_string_expression
    : Interpolated_Regular_String_Start Interpolated_Regular_String_Mid?
      ('{' regular_interpolation '}' Interpolated_Regular_String_Mid?)*
      Interpolated_Regular_String_End
    ;

regular_interpolation
    : expression (',' interpolation_minimum_width)?
```

```
      Regular_Interpolation_Format?
    ;

interpolation_minimum_width
    : constant_expression
    ;

Interpolated_Regular_String_Start
    : '$"'
    ;

// the following three lexical rules are context sensitive, see details
below

Interpolated_Regular_String_Mid
    : Interpolated_Regular_String_Element+
    ;

Regular_Interpolation_Format
    : ':' Interpolated_Regular_String_Element+
    ;

Interpolated_Regular_String_End
    : '"'
    ;

fragment Interpolated_Regular_String_Element
    : Interpolated_Regular_String_Character
    | Simple_Escape_Sequence
    | Hexadecimal_Escape_Sequence
    | Unicode_Escape_Sequence
    | Open_Brace_Escape_Sequence
    | Close_Brace_Escape_Sequence
    ;

fragment Interpolated_Regular_String_Character
    // Any character except " (U+0022), \\ (U+005C),
    // { (U+007B), } (U+007D), and New_Line_Character.
    : ~["\\{}\u000D\u000A\u0085\u2028\u2029]
    ;

// interpolated verbatim string expressions

interpolated_verbatim_string_expression
    : Interpolated_Verbatim_String_Start Interpolated_Verbatim_String_Mid?
      ('{' verbatim_interpolation '}' Interpolated_Verbatim_String_Mid?)*
      Interpolated_Verbatim_String_End
    ;

verbatim_interpolation
    : expression (',' interpolation_minimum_width)?
      Verbatim_Interpolation_Format?
    ;

Interpolated_Verbatim_String_Start
```

```
    : '$@"'
    | '@$"'
    ;

// the following three lexical rules are context sensitive, see details
below

Interpolated_Verbatim_String_Mid
    : Interpolated_Verbatim_String_Element+
    ;

Verbatim_Interpolation_Format
    : ':' Interpolated_Verbatim_String_Element+
    ;

Interpolated_Verbatim_String_End
    : '"'
    ;

fragment Interpolated_Verbatim_String_Element
    : Interpolated_Verbatim_String_Character
    | Quote_Escape_Sequence
    | Open_Brace_Escape_Sequence
    | Close_Brace_Escape_Sequence
    ;

fragment Interpolated_Verbatim_String_Character
    : ~["{}]    // Any character except " (U+0022), { (U+007B) and }
(U+007D)
    ;

// lexical fragments used by both regular and verbatim interpolated strings

fragment Open_Brace_Escape_Sequence
    : '{{'
    ;

fragment Close_Brace_Escape_Sequence
    : '}}'
    ;

// Source: §12.8.4 Simple names
simple_name
    : identifier type_argument_list?
    ;

// Source: §12.8.5 Parenthesized expressions
parenthesized_expression
    : '(' expression ')'
    ;

// Source: §12.8.6 Tuple expressions
tuple_expression
    : '(' tuple_element (',' tuple_element)+ ')'
    | deconstruction_expression
```

```
    ;

tuple_element
    : (identifier ':')? expression
    ;

deconstruction_expression
    : 'var' deconstruction_tuple
    ;

deconstruction_tuple
    : '(' deconstruction_element (',' deconstruction_element)+ ')'
    ;

deconstruction_element
    : deconstruction_tuple
    | identifier
    ;

// Source: §12.8.7.1 General
member_access
    : primary_expression '.' identifier type_argument_list?
    | predefined_type '.' identifier type_argument_list?
    | qualified_alias_member '.' identifier type_argument_list?
    ;

predefined_type
    : 'bool' | 'byte' | 'char' | 'decimal' | 'double' | 'float' | 'int'
    | 'long' | 'object' | 'sbyte' | 'short' | 'string' | 'uint' | 'ulong'
    | 'ushort'
    ;

// Source: §12.8.8 Null Conditional Member Access
null_conditional_member_access
    : primary_expression '?' '.' identifier type_argument_list?
      dependent_access*
    ;

dependent_access
    : '.' identifier type_argument_list?    // member access
    | '[' argument_list ']'                 // element access
    | '(' argument_list? ')'                // invocation
    ;

null_conditional_projection_initializer
    : primary_expression '?' '.' identifier type_argument_list?
    ;

// Source: §12.8.9.1 General
invocation_expression
    : primary_expression '(' argument_list? ')'
    ;

// Source: §12.8.10 Null Conditional Invocation Expression
null_conditional_invocation_expression
```

```
    : null_conditional_member_access '(' argument_list? ')'
    | null_conditional_element_access '(' argument_list? ')'
    ;

// Source: §12.8.11.1 General
element_access
    : primary_no_array_creation_expression '[' argument_list ']'
    ;

// Source: §12.8.12 Null Conditional Element Access
null_conditional_element_access
    : primary_no_array_creation_expression '?' '[' argument_list ']'
      dependent_access*
    ;

// Source: §12.8.13 This access
this_access
    : 'this'
    ;

// Source: §12.8.14 Base access
base_access
    : 'base' '.' identifier type_argument_list?
    | 'base' '[' argument_list ']'
    ;

// Source: §12.8.15 Postfix increment and decrement operators
post_increment_expression
    : primary_expression '++'
    ;

post_decrement_expression
    : primary_expression '--'
    ;

// Source: §12.8.16.2 Object creation expressions
object_creation_expression
    : 'new' type '(' argument_list? ')' object_or_collection_initializer?
    | 'new' type object_or_collection_initializer
    ;

object_or_collection_initializer
    : object_initializer
    | collection_initializer
    ;

// Source: §12.8.16.3 Object initializers
object_initializer
    : '{' member_initializer_list? '}'
    | '{' member_initializer_list ',' '}'
    ;

member_initializer_list
    : member_initializer (',' member_initializer)*
    ;
```

```
member_initializer
    : initializer_target '=' initializer_value
    ;

initializer_target
    : identifier
    | '[' argument_list ']'
    ;

initializer_value
    : expression
    | object_or_collection_initializer
    ;

// Source: §12.8.16.4 Collection initializers
collection_initializer
    : '{' element_initializer_list '}'
    | '{' element_initializer_list ',' '}'
    ;

element_initializer_list
    : element_initializer (',' element_initializer)*
    ;

element_initializer
    : non_assignment_expression
    | '{' expression_list '}'
    ;

expression_list
    : expression
    | expression_list ',' expression
    ;

// Source: §12.8.16.5 Array creation expressions
array_creation_expression
    : 'new' non_array_type '[' expression_list ']' rank_specifier*
      array_initializer?
    | 'new' array_type array_initializer
    | 'new' rank_specifier array_initializer
    ;

// Source: §12.8.16.6 Delegate creation expressions
delegate_creation_expression
    : 'new' delegate_type '(' expression ')'
    ;

// Source: §12.8.16.7 Anonymous object creation expressions
anonymous_object_creation_expression
    : 'new' anonymous_object_initializer
    ;

anonymous_object_initializer
    : '{' member_declarator_list? '}'
```

```
    | '{' member_declarator_list ',' '}'
    ;

member_declarator_list
    : member_declarator (',' member_declarator)*
    ;

member_declarator
    : simple_name
    | member_access
    | null_conditional_projection_initializer
    | base_access
    | identifier '=' expression
    ;

// Source: §12.8.17 The typeof operator
typeof_expression
    : 'typeof' '(' type ')'
    | 'typeof' '(' unbound_type_name ')'
    | 'typeof' '(' 'void' ')'
    ;

unbound_type_name
    : identifier generic_dimension_specifier?
    | identifier '::' identifier generic_dimension_specifier?
    | unbound_type_name '.' identifier generic_dimension_specifier?
    ;

generic_dimension_specifier
    : '<' comma* '>'
    ;

comma
    : ','
    ;


// Source: §12.8.18 The sizeof operator
sizeof_expression
    : 'sizeof' '(' unmanaged_type ')'
    ;

// Source: §12.8.19 The checked and unchecked operators
checked_expression
    : 'checked' '(' expression ')'
    ;

unchecked_expression
    : 'unchecked' '(' expression ')'
    ;

// Source: §12.8.20 Default value expressions
default_value_expression
    : explictly_typed_default
    | default_literal
```

```
    ;

explictly_typed_default
    : 'default' '(' type ')'
    ;

default_literal
    : 'default'
    ;

// Source: §12.8.21 Stack allocation
stackalloc_expression
    : 'stackalloc' unmanaged_type '[' expression ']'
    | 'stackalloc' unmanaged_type? '[' constant_expression? ']'
      stackalloc_initializer
    ;

stackalloc_initializer
     : '{' stackalloc_initializer_element_list '}'
     ;

stackalloc_initializer_element_list
     : stackalloc_element_initializer (',' stackalloc_element_initializer)*
','?
     ;

stackalloc_element_initializer
    : expression
    ;

// Source: §12.8.22 The nameof operator
nameof_expression
    : 'nameof' '(' named_entity ')'
    ;

named_entity
    : named_entity_target ('.' identifier type_argument_list?)*
    ;

named_entity_target
    : simple_name
    | 'this'
    | 'base'
    | predefined_type
    | qualified_alias_member
    ;

// Source: §12.9.1 General
unary_expression
    : primary_expression
    | '+' unary_expression
    | '-' unary_expression
    | '!' unary_expression
    | '~' unary_expression
    | pre_increment_expression
```

```
    | pre_decrement_expression
    | cast_expression
    | await_expression
    | pointer_indirection_expression    // unsafe code support
    | addressof_expression              // unsafe code support
    ;

// Source: §12.9.6 Prefix increment and decrement operators
pre_increment_expression
    : '++' unary_expression
    ;

pre_decrement_expression
    : '--' unary_expression
    ;

// Source: §12.9.7 Cast expressions
cast_expression
    : '(' type ')' unary_expression
    ;

// Source: §12.9.8.1 General
await_expression
    : 'await' unary_expression
    ;

// Source: §12.10.1 General
multiplicative_expression
    : unary_expression
    | multiplicative_expression '*' unary_expression
    | multiplicative_expression '/' unary_expression
    | multiplicative_expression '%' unary_expression
    ;

additive_expression
    : multiplicative_expression
    | additive_expression '+' multiplicative_expression
    | additive_expression '-' multiplicative_expression
    ;

// Source: §12.11 Shift operators
shift_expression
    : additive_expression
    | shift_expression '<<' additive_expression
    | shift_expression right_shift additive_expression
    ;

// Source: §12.12.1 General
relational_expression
    : shift_expression
    | relational_expression '<' shift_expression
    | relational_expression '>' shift_expression
    | relational_expression '<=' shift_expression
    | relational_expression '>=' shift_expression
    | relational_expression 'is' type
```

```
    | relational_expression 'is' pattern
    | relational_expression 'as' type
    ;

equality_expression
    : relational_expression
    | equality_expression '==' relational_expression
    | equality_expression '!=' relational_expression
    ;

// Source: §12.13.1 General
and_expression
    : equality_expression
    | and_expression '&' equality_expression
    ;

exclusive_or_expression
    : and_expression
    | exclusive_or_expression '^' and_expression
    ;

inclusive_or_expression
    : exclusive_or_expression
    | inclusive_or_expression '|' exclusive_or_expression
    ;

// Source: §12.14.1 General
conditional_and_expression
    : inclusive_or_expression
    | conditional_and_expression '&&' inclusive_or_expression
    ;

conditional_or_expression
    : conditional_and_expression
    | conditional_or_expression '||' conditional_and_expression
    ;

// Source: §12.15 The null coalescing operator
null_coalescing_expression
    : conditional_or_expression
    | conditional_or_expression '??' null_coalescing_expression
    | throw_expression
    ;

// Source: §12.16 The throw expression operator
throw_expression
    : 'throw' null_coalescing_expression
    ;

// Source: §12.17 Declaration expressions
declaration_expression
    : local_variable_type identifier
    ;

local_variable_type
```

```
    : type
    | 'var'
    ;

// Source: §12.18 Conditional operator
conditional_expression
    : null_coalescing_expression
    | null_coalescing_expression '?' expression ':' expression
    | null_coalescing_expression '?' 'ref' variable_reference ':'
      'ref' variable_reference
    ;

// Source: §12.19.1 General
lambda_expression
    : 'async'? anonymous_function_signature '=>' anonymous_function_body
    ;

anonymous_method_expression
    : 'async'? 'delegate' explicit_anonymous_function_signature? block
    ;

anonymous_function_signature
    : explicit_anonymous_function_signature
    | implicit_anonymous_function_signature
    ;

explicit_anonymous_function_signature
    : '(' explicit_anonymous_function_parameter_list? ')'
    ;

explicit_anonymous_function_parameter_list
    : explicit_anonymous_function_parameter
      (',' explicit_anonymous_function_parameter)*
    ;

explicit_anonymous_function_parameter
    : anonymous_function_parameter_modifier? type identifier
    ;

anonymous_function_parameter_modifier
    : 'ref'
    | 'out'
    | 'in'
    ;

implicit_anonymous_function_signature
    : '(' implicit_anonymous_function_parameter_list? ')'
    | implicit_anonymous_function_parameter
    ;

implicit_anonymous_function_parameter_list
    : implicit_anonymous_function_parameter
      (',' implicit_anonymous_function_parameter)*
    ;
```

```
implicit_anonymous_function_parameter
    : identifier
    ;

anonymous_function_body
    : null_conditional_invocation_expression
    | expression
    | 'ref' variable_reference
    | block
    ;

// Source: §12.20.1 General
query_expression
    : from_clause query_body
    ;

from_clause
    : 'from' type? identifier 'in' expression
    ;

query_body
    : query_body_clauses? select_or_group_clause query_continuation?
    ;

query_body_clauses
    : query_body_clause
    | query_body_clauses query_body_clause
    ;

query_body_clause
    : from_clause
    | let_clause
    | where_clause
    | join_clause
    | join_into_clause
    | orderby_clause
    ;

let_clause
    : 'let' identifier '=' expression
    ;

where_clause
    : 'where' boolean_expression
    ;

join_clause
    : 'join' type? identifier 'in' expression 'on' expression
      'equals' expression
    ;

join_into_clause
    : 'join' type? identifier 'in' expression 'on' expression
      'equals' expression 'into' identifier
    ;
```

```
orderby_clause
    : 'orderby' orderings
    ;

orderings
    : ordering (',' ordering)*
    ;

ordering
    : expression ordering_direction?
    ;

ordering_direction
    : 'ascending'
    | 'descending'
    ;

select_or_group_clause
    : select_clause
    | group_clause
    ;

select_clause
    : 'select' expression
    ;

group_clause
    : 'group' expression 'by' expression
    ;

query_continuation
    : 'into' identifier query_body
    ;

// Source: §12.21.1 General
assignment
    : unary_expression assignment_operator expression
    ;

assignment_operator
    : '=' 'ref'? | '+=' | '-=' | '*=' | '/=' | '%=' | '&=' | '|=' | '^=' |
'<<='
    | right_shift_assignment
    ;

// Source: §12.22 Expression
expression
    : non_assignment_expression
    | assignment
    ;

non_assignment_expression
    : declaration_expression
    | conditional_expression
```

```
    | lambda_expression
    | query_expression
    ;

// Source: §12.23 Constant expressions
constant_expression
    : expression
    ;

// Source: §12.24 Boolean expressions
boolean_expression
    : expression
    ;

// Source: §13.1 General
statement
    : labeled_statement
    | declaration_statement
    | embedded_statement
    ;

embedded_statement
    : block
    | empty_statement
    | expression_statement
    | selection_statement
    | iteration_statement
    | jump_statement
    | try_statement
    | checked_statement
    | unchecked_statement
    | lock_statement
    | using_statement
    | yield_statement
    | unsafe_statement   // unsafe code support
    | fixed_statement    // unsafe code support
    ;

// Source: §13.3.1 General
block
    : '{' statement_list? '}'
    ;

// Source: §13.3.2 Statement lists
statement_list
    : statement+
    ;

// Source: §13.4 The empty statement
empty_statement
    : ';'
    ;

// Source: §13.5 Labeled statements
labeled_statement
```

```
    : identifier ':' statement
    ;

// Source: §13.6.1 General
declaration_statement
    : local_variable_declaration ';'
    | local_constant_declaration ';'
    | local_function_declaration
    ;

// Source: §13.6.2.1 General
local_variable_declaration
    : implicitly_typed_local_variable_declaration
    | explicitly_typed_local_variable_declaration
    | ref_local_variable_declaration
    ;

// Source: §13.6.2.2 Implicitly typed local variable declarations
implicitly_typed_local_variable_declaration
    : 'var' implicitly_typed_local_variable_declarator
    | ref_kind 'var' ref_local_variable_declarator
    ;

implicitly_typed_local_variable_declarator
    : identifier '=' expression
    ;

// Source: §13.6.2.3 Explicitly typed local variable declarations
explicitly_typed_local_variable_declaration
    : type explicitly_typed_local_variable_declarators
    ;

explicitly_typed_local_variable_declarators
    : explicitly_typed_local_variable_declarator
      (',' explicitly_typed_local_variable_declarator)*
    ;

explicitly_typed_local_variable_declarator
    : identifier ('=' local_variable_initializer)?
    ;

local_variable_initializer
    : expression
    | array_initializer
    ;

// Source: §13.6.2.4 Ref local variable declarations
ref_local_variable_declaration
    : ref_kind type ref_local_variable_declarators
    ;

ref_local_variable_declarators
    : ref_local_variable_declarator (',' ref_local_variable_declarator)*
    ;
```

```
ref_local_variable_declarator
    : identifier '=' 'ref' variable_reference
    ;

// Source: §13.6.3 Local constant declarations
local_constant_declaration
    : 'const' type constant_declarators
    ;

constant_declarators
    : constant_declarator (',' constant_declarator)*
    ;

constant_declarator
    : identifier '=' constant_expression
    ;

// Source: §13.6.4 Local function declarations
local_function_declaration
    : local_function_modifier* return_type local_function_header
      local_function_body
    | ref_local_function_modifier* ref_kind ref_return_type
      local_function_header ref_local_function_body
    ;

local_function_header
    : identifier '(' formal_parameter_list? ')'
    | identifier type_parameter_list '(' formal_parameter_list? ')'
      type_parameter_constraints_clause*
    ;

local_function_modifier
    : ref_local_function_modifier
    | 'async'
    ;

ref_local_function_modifier
    : 'static'
    | unsafe_modifier   // unsafe code support
    ;

local_function_body
    : block
    | '=>' null_conditional_invocation_expression ';'
    | '=>' expression ';'
    ;

ref_local_function_body
    : block
    | '=>' 'ref' variable_reference ';'
    ;

// Source: §13.7 Expression statements
expression_statement
    : statement_expression ';'
```

```
    ;

statement_expression
    : null_conditional_invocation_expression
    | invocation_expression
    | object_creation_expression
    | assignment
    | post_increment_expression
    | post_decrement_expression
    | pre_increment_expression
    | pre_decrement_expression
    | await_expression
    ;

// Source: §13.8.1 General
selection_statement
    : if_statement
    | switch_statement
    ;

// Source: §13.8.2 The if statement
if_statement
    : 'if' '(' boolean_expression ')' embedded_statement
    | 'if' '(' boolean_expression ')' embedded_statement
      'else' embedded_statement
    ;

// Source: §13.8.3 The switch statement
switch_statement
    : 'switch' '(' expression ')' switch_block
    ;

switch_block
    : '{' switch_section* '}'
    ;

switch_section
    : switch_label+ statement_list
    ;

switch_label
    : 'case' pattern case_guard?  ':'
    | 'default' ':'
    ;

case_guard
    : 'when' expression
    ;

// Source: §13.9.1 General
iteration_statement
    : while_statement
    | do_statement
    | for_statement
    | foreach_statement
```

```
    ;

// Source: §13.9.2 The while statement
while_statement
    : 'while' '(' boolean_expression ')' embedded_statement
    ;

// Source: §13.9.3 The do statement
do_statement
    : 'do' embedded_statement 'while' '(' boolean_expression ')' ';'
    ;

// Source: §13.9.4 The for statement
for_statement
    : 'for' '(' for_initializer? ';' for_condition? ';' for_iterator? ')'
      embedded_statement
    ;

for_initializer
    : local_variable_declaration
    | statement_expression_list
    ;

for_condition
    : boolean_expression
    ;

for_iterator
    : statement_expression_list
    ;

statement_expression_list
    : statement_expression (',' statement_expression)*
    ;

// Source: §13.9.5 The foreach statement
foreach_statement
    : 'foreach' '(' ref_kind? local_variable_type identifier 'in'
      expression ')' embedded_statement
    ;

// Source: §13.10.1 General
jump_statement
    : break_statement
    | continue_statement
    | goto_statement
    | return_statement
    | throw_statement
    ;

// Source: §13.10.2 The break statement
break_statement
    : 'break' ';'
    ;
```

```
// Source: §13.10.3 The continue statement
continue_statement
    : 'continue' ';'
    ;

// Source: §13.10.4 The goto statement
goto_statement
    : 'goto' identifier ';'
    | 'goto' 'case' constant_expression ';'
    | 'goto' 'default' ';'
    ;

// Source: §13.10.5 The return statement
return_statement
    : 'return' ';'
    | 'return' expression ';'
    | 'return' 'ref' variable_reference ';'
    ;

// Source: §13.10.6 The throw statement
throw_statement
    : 'throw' expression? ';'
    ;

// Source: §13.11 The try statement
try_statement
    : 'try' block catch_clauses
    | 'try' block catch_clauses? finally_clause
    ;

catch_clauses
    : specific_catch_clause+
    | specific_catch_clause* general_catch_clause
    ;

specific_catch_clause
    : 'catch' exception_specifier exception_filter? block
    | 'catch' exception_filter block
    ;

exception_specifier
    : '(' type identifier? ')'
    ;

exception_filter
    : 'when' '(' boolean_expression ')'
    ;

general_catch_clause
    : 'catch' block
    ;

finally_clause
    : 'finally' block
    ;
```

```
// Source: §13.12 The checked and unchecked statements
checked_statement
    : 'checked' block
    ;

unchecked_statement
    : 'unchecked' block
    ;

// Source: §13.13 The lock statement
lock_statement
    : 'lock' '(' expression ')' embedded_statement
    ;

// Source: §13.14 The using statement
using_statement
    : 'using' '(' resource_acquisition ')' embedded_statement
    ;

resource_acquisition
    : local_variable_declaration
    | expression
    ;

// Source: §13.15 The yield statement
yield_statement
    : 'yield' 'return' expression ';'
    | 'yield' 'break' ';'
    ;

// Source: §14.2 Compilation units
compilation_unit
    : extern_alias_directive* using_directive* global_attributes?
      namespace_member_declaration*
    ;

// Source: §14.3 Namespace declarations
namespace_declaration
    : 'namespace' qualified_identifier namespace_body ';'?
    ;

qualified_identifier
    : identifier ('.' identifier)*
    ;

namespace_body
    : '{' extern_alias_directive* using_directive*
      namespace_member_declaration* '}'
    ;

// Source: §14.4 Extern alias directives
extern_alias_directive
    : 'extern' 'alias' identifier ';'
    ;
```

```
// Source: §14.5.1 General
using_directive
    : using_alias_directive
    | using_namespace_directive
    | using_static_directive
    ;

// Source: §14.5.2 Using alias directives
using_alias_directive
    : 'using' identifier '=' namespace_or_type_name ';'
    ;

// Source: §14.5.3 Using namespace directives
using_namespace_directive
    : 'using' namespace_name ';'
    ;

// Source: §14.5.4 Using static directives
using_static_directive
    : 'using' 'static' type_name ';'
    ;

// Source: §14.6 Namespace member declarations
namespace_member_declaration
    : namespace_declaration
    | type_declaration
    ;

// Source: §14.7 Type declarations
type_declaration
    : class_declaration
    | struct_declaration
    | interface_declaration
    | enum_declaration
    | delegate_declaration
    ;

// Source: §14.8.1 General
qualified_alias_member
    : identifier '::' identifier type_argument_list?
    ;

// Source: §15.2.1 General
class_declaration
    : attributes? class_modifier* 'partial'? 'class' identifier
        type_parameter_list? class_base? type_parameter_constraints_clause*
        class_body ';'?
    ;

// Source: §15.2.2.1 General
class_modifier
    : 'new'
    | 'public'
    | 'protected'
```

```
    | 'internal'
    | 'private'
    | 'abstract'
    | 'sealed'
    | 'static'
    | unsafe_modifier   // unsafe code support
    ;

// Source: §15.2.3 Type parameters
type_parameter_list
    : '<' type_parameters '>'
  ;

type_parameters
    : attributes? type_parameter
    | type_parameters ',' attributes? type_parameter
    ;

// Source: §15.2.4.1 General
class_base
    : ':' class_type
    | ':' interface_type_list
    | ':' class_type ',' interface_type_list
    ;

interface_type_list
    : interface_type (',' interface_type)*
    ;

// Source: §15.2.5 Type parameter constraints
type_parameter_constraints_clauses
    : type_parameter_constraints_clause
    | type_parameter_constraints_clauses type_parameter_constraints_clause
    ;

type_parameter_constraints_clause
    : 'where' type_parameter ':' type_parameter_constraints
    ;

type_parameter_constraints
    : primary_constraint
    | secondary_constraints
    | constructor_constraint
    | primary_constraint ',' secondary_constraints
    | primary_constraint ',' constructor_constraint
    | secondary_constraints ',' constructor_constraint
    | primary_constraint ',' secondary_constraints ','
constructor_constraint
    ;

primary_constraint
    : class_type
    | 'class'
    | 'struct'
    | 'unmanaged'
```

```
    ;

secondary_constraints
    : interface_type
    | type_parameter
    | secondary_constraints ',' interface_type
    | secondary_constraints ',' type_parameter
    ;

constructor_constraint
    : 'new' '(' ')'
    ;

// Source: §15.2.6 Class body
class_body
    : '{' class_member_declaration* '}'
    ;

// Source: §15.3.1 General
class_member_declaration
    : constant_declaration
    | field_declaration
    | method_declaration
    | property_declaration
    | event_declaration
    | indexer_declaration
    | operator_declaration
    | constructor_declaration
    | finalizer_declaration
    | static_constructor_declaration
    | type_declaration
    ;

// Source: §15.4 Constants
constant_declaration
    : attributes? constant_modifier* 'const' type constant_declarators ';'
    ;

constant_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    ;

// Source: §15.5.1 General
field_declaration
    : attributes? field_modifier* type variable_declarators ';'
    ;

field_modifier
    : 'new'
    | 'public'
    | 'protected'
```

```
    | 'internal'
    | 'private'
    | 'static'
    | 'readonly'
    | 'volatile'
    | unsafe_modifier   // unsafe code support
    ;

variable_declarators
    : variable_declarator (',' variable_declarator)*
    ;

variable_declarator
    : identifier ('=' variable_initializer)?
    ;

// Source: §15.6.1 General
method_declaration
    : attributes? method_modifiers return_type method_header method_body
    | attributes? ref_method_modifiers ref_kind ref_return_type
method_header
      ref_method_body
    ;

method_modifiers
    : method_modifier* 'partial'?
    ;

ref_kind
    : 'ref'
    | 'ref' 'readonly'
    ;

ref_method_modifiers
    : ref_method_modifier*
    ;

method_header
    : member_name '(' formal_parameter_list? ')'
    | member_name type_parameter_list '(' formal_parameter_list? ')'
      type_parameter_constraints_clause*
    ;

method_modifier
    : ref_method_modifier
    | 'async'
    ;

ref_method_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'static'
```

```
    | 'virtual'
    | 'sealed'
    | 'override'
    | 'abstract'
    | 'extern'
    | unsafe_modifier   // unsafe code support
    ;

return_type
    : ref_return_type
    | 'void'
    ;

ref_return_type
    : type
    ;

member_name
    : identifier
    | interface_type '.' identifier
    ;

method_body
    : block
    | '=>' null_conditional_invocation_expression ';'
    | '=>' expression ';'
    | ';'
    ;

ref_method_body
    : block
    | '=>' 'ref' variable_reference ';'
    | ';'
    ;

// Source: §15.6.2.1 General
formal_parameter_list
    : fixed_parameters
    | fixed_parameters ',' parameter_array
    | parameter_array
    ;

fixed_parameters
    : fixed_parameter (',' fixed_parameter)*
    ;

fixed_parameter
    : attributes? parameter_modifier? type identifier default_argument?
    ;

default_argument
    : '=' expression
    ;

parameter_modifier
```

```
    : parameter_mode_modifier
    | 'this'
    ;

parameter_mode_modifier
    : 'ref'
    | 'out'
    | 'in'
    ;

parameter_array
    : attributes? 'params' array_type identifier
    ;

// Source: §15.7.1 General
property_declaration
    : attributes? property_modifier* type member_name property_body
    | attributes? property_modifier* ref_kind type member_name
ref_property_body
    ;

property_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'static'
    | 'virtual'
    | 'sealed'
    | 'override'
    | 'abstract'
    | 'extern'
    | unsafe_modifier   // unsafe code support
    ;

property_body
    : '{' accessor_declarations '}' property_initializer?
    | '=>' expression ';'
    ;

property_initializer
    : '=' variable_initializer ';'
    ;

ref_property_body
    : '{' ref_get_accessor_declaration '}'
    | '=>' 'ref' variable_reference ';'
    ;

// Source: §15.7.3 Accessors
accessor_declarations
    : get_accessor_declaration set_accessor_declaration?
    | set_accessor_declaration get_accessor_declaration?
    ;
```

```
get_accessor_declaration
    : attributes? accessor_modifier? 'get' accessor_body
    ;

set_accessor_declaration
    : attributes? accessor_modifier? 'set' accessor_body
    ;

accessor_modifier
    : 'protected'
    | 'internal'
    | 'private'
    | 'protected' 'internal'
    | 'internal' 'protected'
    | 'protected' 'private'
    | 'private' 'protected'
    ;

accessor_body
    : block
    | '=>' expression ';'
    | ';'
    ;

ref_get_accessor_declaration
    : attributes? accessor_modifier? 'get' ref_accessor_body
    ;

ref_accessor_body
    : block
    | '=>' 'ref' variable_reference ';'
    | ';'
    ;

// Source: §15.8.1 General
event_declaration
    : attributes? event_modifier* 'event' type variable_declarators ';'
    | attributes? event_modifier* 'event' type member_name
        '{' event_accessor_declarations '}'
    ;

event_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'static'
    | 'virtual'
    | 'sealed'
    | 'override'
    | 'abstract'
    | 'extern'
    | unsafe_modifier   // unsafe code support
```

```
    ;

event_accessor_declarations
    : add_accessor_declaration remove_accessor_declaration
    | remove_accessor_declaration add_accessor_declaration
    ;

add_accessor_declaration
    : attributes? 'add' block
    ;

remove_accessor_declaration
    : attributes? 'remove' block
    ;

// Source: §15.9.1 General
indexer_declaration
    : attributes? indexer_modifier* indexer_declarator indexer_body
    | attributes? indexer_modifier* ref_kind indexer_declarator
ref_indexer_body
    ;

indexer_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'virtual'
    | 'sealed'
    | 'override'
    | 'abstract'
    | 'extern'
    | unsafe_modifier   // unsafe code support
    ;

indexer_declarator
    : type 'this' '[' formal_parameter_list ']'
    | type interface_type '.' 'this' '[' formal_parameter_list ']'
    ;

indexer_body
    : '{' accessor_declarations '}'
    | '=>' expression ';'
    ;

ref_indexer_body
    : '{' ref_get_accessor_declaration '}'
    | '=>' 'ref' variable_reference ';'
    ;

// Source: §15.10.1 General
operator_declaration
    : attributes? operator_modifier+ operator_declarator operator_body
    ;
```

```
operator_modifier
    : 'public'
    | 'static'
    | 'extern'
    | unsafe_modifier   // unsafe code support
    ;

operator_declarator
    : unary_operator_declarator
    | binary_operator_declarator
    | conversion_operator_declarator
    ;

unary_operator_declarator
    : type 'operator' overloadable_unary_operator '(' fixed_parameter ')'
    ;

overloadable_unary_operator
    : '+' | '-' | '!' | '~' | '++' | '--' | 'true' | 'false'
    ;

binary_operator_declarator
    : type 'operator' overloadable_binary_operator
        '(' fixed_parameter ',' fixed_parameter ')'
    ;

overloadable_binary_operator
    : '+'  | '-'  | '*'  | '/'  | '%'  | '&' | '|' | '^'  | '<<'
    | right_shift | '==' | '!=' | '>' | '<' | '>=' | '<='
    ;

conversion_operator_declarator
    : 'implicit' 'operator' type '(' fixed_parameter ')'
    | 'explicit' 'operator' type '(' fixed_parameter ')'
    ;

operator_body
    : block
    | '=>' expression ';'
    | ';'
    ;

// Source: §15.11.1 General
constructor_declaration
    : attributes? constructor_modifier* constructor_declarator
constructor_body
    ;

constructor_modifier
    : 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'extern'
```

```
    | unsafe_modifier   // unsafe code support
    ;

constructor_declarator
    : identifier '(' formal_parameter_list? ')' constructor_initializer?
    ;

constructor_initializer
    : ':' 'base' '(' argument_list? ')'
    | ':' 'this' '(' argument_list? ')'
    ;

constructor_body
    : block
    | '=>' expression ';'
    | ';'
    ;

// Source: §15.12 Static constructors
static_constructor_declaration
    : attributes? static_constructor_modifiers identifier '(' ')'
      static_constructor_body
    ;

static_constructor_modifiers
    : 'static'
    | 'static' 'extern' unsafe_modifier?
    | 'static' unsafe_modifier 'extern'?
    | 'extern' 'static' unsafe_modifier?
    | 'extern' unsafe_modifier 'static'
    | unsafe_modifier 'static' 'extern'?
    | unsafe_modifier 'extern' 'static'
    ;

static_constructor_body
    : block
    | '=>' expression ';'
    | ';'
    ;

// Source: §15.13 Finalizers
finalizer_declaration
    : attributes? '~' identifier '(' ')' finalizer_body
    | attributes? 'extern' unsafe_modifier? '~' identifier '(' ')'
      finalizer_body
    | attributes? unsafe_modifier 'extern'? '~' identifier '(' ')'
      finalizer_body
    ;

finalizer_body
    : block
    | '=>' expression ';'
    | ';'
    ;
```

```
// Source: §16.2.1 General
struct_declaration
    : attributes? struct_modifier* 'ref'? 'partial'? 'struct'
      identifier type_parameter_list? struct_interfaces?
      type_parameter_constraints_clause* struct_body ';'?
    ;

// Source: §16.2.2 Struct modifiers
struct_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'readonly'
    | unsafe_modifier   // unsafe code support
    ;

// Source: §16.2.5 Struct interfaces
struct_interfaces
    : ':' interface_type_list
    ;

// Source: §16.2.6 Struct body
struct_body
    : '{' struct_member_declaration* '}'
    ;

// Source: §16.3 Struct members
struct_member_declaration
    : constant_declaration
    | field_declaration
    | method_declaration
    | property_declaration
    | event_declaration
    | indexer_declaration
    | operator_declaration
    | constructor_declaration
    | static_constructor_declaration
    | type_declaration
    | fixed_size_buffer_declaration   // unsafe code support
    ;

// Source: §17.7 Array initializers
array_initializer
    : '{' variable_initializer_list? '}'
    | '{' variable_initializer_list ',' '}'
    ;

variable_initializer_list
    : variable_initializer (',' variable_initializer)*
    ;

variable_initializer
    : expression
```

```
    | array_initializer
    ;

// Source: §18.2.1 General
interface_declaration
    : attributes? interface_modifier* 'partial'? 'interface'
      identifier variant_type_parameter_list? interface_base?
      type_parameter_constraints_clause* interface_body ';'?
    ;

// Source: §18.2.2 Interface modifiers
interface_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | unsafe_modifier   // unsafe code support
    ;

// Source: §18.2.3.1 General
variant_type_parameter_list
    : '<' variant_type_parameters '>'
    ;

// Source: §18.2.3.1 General
variant_type_parameters
    : attributes? variance_annotation? type_parameter
    | variant_type_parameters ',' attributes? variance_annotation?
      type_parameter
    ;

// Source: §18.2.3.1 General
variance_annotation
    : 'in'
    | 'out'
    ;

// Source: §18.2.4 Base interfaces
interface_base
    : ':' interface_type_list
    ;

// Source: §18.3 Interface body
interface_body
    : '{' interface_member_declaration* '}'
    ;

// Source: §18.4.1 General
interface_member_declaration
    : interface_method_declaration
    | interface_property_declaration
    | interface_event_declaration
    | interface_indexer_declaration
    ;
```

```
// Source: §18.4.2 Interface methods
interface_method_declaration
    : attributes? 'new'? return_type interface_method_header
    | attributes? 'new'? ref_kind ref_return_type interface_method_header
    ;

interface_method_header
    : identifier '(' formal_parameter_list? ')' ';'
    | identifier type_parameter_list '(' formal_parameter_list? ')'
      type_parameter_constraints_clause* ';'
    ;

// Source: §18.4.3 Interface properties
interface_property_declaration
    : attributes? 'new'? type identifier '{' interface_accessors '}'
    | attributes? 'new'? ref_kind type identifier '{' ref_interface_accessor
'}'
    ;

interface_accessors
    : attributes? 'get' ';'
    | attributes? 'set' ';'
    | attributes? 'get' ';' attributes? 'set' ';'
    | attributes? 'set' ';' attributes? 'get' ';'
    ;

ref_interface_accessor
    : attributes? 'get' ';'
    ;

// Source: §18.4.4 Interface events
interface_event_declaration
    : attributes? 'new'? 'event' type identifier ';'
    ;

// Source: §18.4.5 Interface indexers
interface_indexer_declaration
    : attributes? 'new'? type 'this' '[' formal_parameter_list ']'
      '{' interface_accessors '}'
    | attributes? 'new'? ref_kind type 'this' '[' formal_parameter_list ']'
      '{' ref_interface_accessor '}'
    ;

// Source: §19.2 Enum declarations
enum_declaration
    : attributes? enum_modifier* 'enum' identifier enum_base? enum_body ';'?
    ;

enum_base
    : ':' integral_type
    | ':' integral_type_name
    ;

integral_type_name
```

```
    : type_name // Shall resolve to an integral type other than char
    ;

enum_body
    : '{' enum_member_declarations? '}'
    | '{' enum_member_declarations ',' '}'
    ;

// Source: §19.3 Enum modifiers
enum_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    ;

// Source: §19.4 Enum members
enum_member_declarations
    : enum_member_declaration (',' enum_member_declaration)*
    ;

// Source: §19.4 Enum members
enum_member_declaration
    : attributes? identifier ('=' constant_expression)?
    ;

// Source: §20.2 Delegate declarations
delegate_declaration
    : attributes? delegate_modifier* 'delegate' return_type delegate_header
    | attributes? delegate_modifier* 'delegate' ref_kind ref_return_type
      delegate_header
    ;

delegate_header
    : identifier '(' formal_parameter_list? ')' ';'
    | identifier variant_type_parameter_list '(' formal_parameter_list? ')'
      type_parameter_constraints_clause* ';'
    ;

delegate_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | unsafe_modifier   // unsafe code support
    ;

// Source: §22.3 Attribute specification
global_attributes
    : global_attribute_section+
    ;

global_attribute_section
```

```
    : '[' global_attribute_target_specifier attribute_list ']'
    | '[' global_attribute_target_specifier attribute_list ',' ']'
    ;

global_attribute_target_specifier
    : global_attribute_target ':'
    ;

global_attribute_target
    : identifier
    ;

attributes
    : attribute_section+
    ;

attribute_section
    : '[' attribute_target_specifier? attribute_list ']'
    | '[' attribute_target_specifier? attribute_list ',' ']'
    ;

attribute_target_specifier
    : attribute_target ':'
    ;

attribute_target
    : identifier
    | keyword
    ;

attribute_list
    : attribute (',' attribute)*
    ;

attribute
    : attribute_name attribute_arguments?
    ;

attribute_name
    : type_name
    ;

attribute_arguments
    : '(' positional_argument_list? ')'
    | '(' positional_argument_list ',' named_argument_list ')'
    | '(' named_argument_list ')'
    ;

positional_argument_list
    : positional_argument (',' positional_argument)*
    ;

positional_argument
    : argument_name? attribute_argument_expression
    ;
```

```
named_argument_list
    : named_argument (','  named_argument)*
    ;

named_argument
    : identifier '=' attribute_argument_expression
    ;

attribute_argument_expression
    : expression
    ;
```

A.4 Grammar extensions for unsafe code

ANTLR

```
// Source: §23.2 Unsafe contexts
unsafe_modifier
    : 'unsafe'
    ;

unsafe_statement
    : 'unsafe' block
    ;

// Source: §23.3 Pointer types
pointer_type
    : value_type ('*')+
    | 'void' ('*')+
    ;

// Source: §23.6.2 Pointer indirection
pointer_indirection_expression
    : '*' unary_expression
    ;

// Source: §23.6.3 Pointer member access
pointer_member_access
    : primary_expression '->' identifier type_argument_list?
    ;

// Source: §23.6.4 Pointer element access
pointer_element_access
    : primary_no_array_creation_expression '[' expression ']'
    ;

// Source: §23.6.5 The address-of operator
addressof_expression
    : '&' unary_expression
    ;
```

```
// Source: §23.7 The fixed statement
fixed_statement
    : 'fixed' '(' pointer_type fixed_pointer_declarators ')'
embedded_statement
    ;

fixed_pointer_declarators
    : fixed_pointer_declarator (','  fixed_pointer_declarator)*
    ;

fixed_pointer_declarator
    : identifier '=' fixed_pointer_initializer
    ;

fixed_pointer_initializer
    : '&' variable_reference
    | expression
    ;

// Source: §23.8.2 Fixed-size buffer declarations
fixed_size_buffer_declaration
    : attributes? fixed_size_buffer_modifier* 'fixed' buffer_element_type
      fixed_size_buffer_declarators ';'
    ;

fixed_size_buffer_modifier
    : 'new'
    | 'public'
    | 'internal'
    | 'private'
    | 'unsafe'
    ;

buffer_element_type
    : type
    ;

fixed_size_buffer_declarators
    : fixed_size_buffer_declarator (',' fixed_size_buffer_declarator)*
    ;

fixed_size_buffer_declarator
    : identifier '[' constant_expression ']'
    ;
```

End of informative text.

Collaborate with us on GitHub

.NET ECMA C# draft specification feedback

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Annex B Portability issues

Article • 02/07/2024

This clause is informative.

B.1 General

This annex collects some information about portability that appears in this specification.

B.2 Undefined behavior

The behavior is undefined in the following circumstances:

1. The behavior of the enclosing async function when an awaiter's implementation of the interface methods `INotifyCompletion.OnCompleted` and `ICriticalNotifyCompletion.UnsafeOnCompleted` does not cause the resumption delegate to be invoked at most once (§12.9.8.4).
2. Passing pointers as `ref` or `out` parameters (§23.3).
3. When dereferencing the result of converting one pointer type to another and the resulting pointer is not correctly aligned for the pointed-to type. (§23.5.1).
4. When the unary `*` operator is applied to a pointer containing an invalid value (§23.6.2).
5. When a pointer is subscripted to access an out-of-bounds element (§23.6.4).
6. Modifying objects of managed type through fixed pointers (§23.7).
7. The content of memory newly allocated by `stackalloc` (§12.8.21).
8. Attempting to allocate a negative number of items using `stackalloc` (§12.8.21).
9. Implicit dynamic conversions (§10.2.10) of `in` parameters with value arguments (§12.6.4.2).

B.3 Implementation-defined behavior

A conforming implementation is required to document its choice of behavior in each of the areas listed in this subclause. The following are implementation-defined:

1. The behavior when an identifier not in Normalization Form C is encountered (§6.4.3).
2. The maximum value allowed for `Decimal_Digit+` in `PP_Line_Indicator` (§6.5.8).
3. The interpretation of the *input_characters* in the *pp_pragma-text* of a #pragma directive (§6.5.9).

4. The values of any application parameters passed to `Main` by the host environment prior to application startup (§7.1).
5. The precise structure of the expression tree, as well as the exact process for creating it, when an anonymous function is converted to an expression-tree (§10.7.3).
6. The value returned when a stack allocation of size zero is made (§12.8.21).
7. Whether a `System.ArithmeticException` (or a subclass thereof) is thrown or the overflow goes unreported with the resulting value being that of the left operand, when in an `unchecked` context and the left operand of an integer division is the maximum negative `int` or `long` value and the right operand is `-1` (§12.10.3).
8. When a `System.ArithmeticException` (or a subclass thereof) is thrown when performing a decimal remainder operation (§12.10.4).
9. The impact of thread termination when a thread has no handler for an exception, and the thread is itself terminated (§13.10.6).
10. The mechanism by which linkage to an external method is achieved (§15.6.8).
11. The impact of thread termination when no matching `catch` clause is found for an exception and the code that initially started that thread is reached. (§21.4).
12. The mappings between pointers and integers (§23.5.1).
13. The effect of applying the unary `*` operator to a `null` pointer (§23.6.2).
14. The behavior when pointer arithmetic overflows the domain of the pointer type (§23.6.6, §23.6.7).
15. The result of the `sizeof` operator for non-pre-defined value types (§23.6.9).
16. The behavior of the `fixed` statement if the array expression is `null` or if the array has zero elements (§23.7).
17. The behavior of the `fixed` statement if the string expression is `null` (§23.7).

B.4 Unspecified behavior

1. The time at which the finalizer (if any) for an object is run, once that object has become eligible for finalization (§7.9).
2. The representation of `true` (§8.3.9).
3. The value of the result when converting out-of-range values from `float` or `double` values to an integral type in an `unchecked` context (§10.3.2).
4. The exact target object and target method of the delegate produced from an *anonymous_method_expression* contains (§10.7.2).
5. The layout of arrays, except in an unsafe context (§12.8.16.5).
6. Whether there is any way to execute the *block* of an anonymous function other than through evaluation and invocation of the *lambda_expression* or *anonymous_method-expression* (§12.19.3).

7. The exact timing of static field initialization (§15.5.6.2).
8. The result of invoking `MoveNext` when an enumerator object is running (§15.14.5.2).
9. The result of accessing `Current` when an enumerator object is in the before, running, or after states (§15.14.5.3).
10. The result of invoking `Dispose` when an enumerator object is in the running state (§15.14.5.4).
11. The attributes of a type declared in multiple parts are determined by combining, in an unspecified order, the attributes of each of its parts (§22.3).
12. The order in which members are packed into a struct (§23.6.9).
13. An exception occurs during finalizer execution, and that exception is not caught (§21.4).
14. If more than one member matches, which member is the implementation of `I.M` (§18.6.5).

B.5 Other issues

1. The exact results of floating-point expression evaluation can vary from one implementation to another, because an implementation is permitted to evaluate such expressions using a greater range and/or precision than is required (§8.3.7).
2. The CLI reserves certain signatures for compatibility with other programming languages (§15.3.10).

End of informative text.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Annex C Standard library

Article • 02/07/2024

C.1 General

A conforming C# implementation shall provide a minimum set of types having specific semantics. These types and their members are listed here, in alphabetical order by namespace and type. For a formal definition of these types and their members, refer to ISO/IEC 23271:2012 *Common Language Infrastructure (CLI), Partition IV; Base Class Library (BCL), Extended Numerics Library, and Extended Array Library*, which are included by reference in this specification.

This text is informative.

The standard library is intended to be the minimum set of types and members required by a conforming C# implementation. As such, it contains only those members that are explicitly required by the C# language specification.

It is expected that a conforming C# implementation will supply a significantly more extensive library that enables useful programs to be written. For example, a conforming implementation might extend this library by

- Adding namespaces.
- Adding types.
- Adding members to non-interface types.
- Adding intervening base classes or interfaces.
- Having struct and class types implement additional interfaces.
- Adding attributes (other than the `ConditionalAttribute`) to existing types and members.

End of informative text.

C.2 Standard Library Types defined in ISO/IEC 23271

> *Note:* Some `struct` types below have the `readonly` modifier. This modifier was not available when ISO/IEC 23271 was released, but is required for conforming implementations of this specification. *end note*

C#

```
namespace System
{
    public delegate void Action();

    public class ArgumentException : SystemException
    {
        public ArgumentException();
        public ArgumentException(string message);
        public ArgumentException(string message, Exception innerException);
    }

    public class ArithmeticException : Exception
    {
        public ArithmeticException();
        public ArithmeticException(string message);
        public ArithmeticException(string message, Exception
innerException);
    }

    public abstract class Array : IList, ICollection, IEnumerable
    {
        public int Length { get; }
        public int Rank { get; }
        public int GetLength(int dimension);
    }

    public class ArrayTypeMismatchException : Exception
    {
        public ArrayTypeMismatchException();
        public ArrayTypeMismatchException(string message);
        public ArrayTypeMismatchException(string message,
            Exception innerException);
    }

    [AttributeUsageAttribute(AttributeTargets.All, Inherited = true,
        AllowMultiple = false)]
    public abstract class Attribute
    {
        protected Attribute();
    }

    public enum AttributeTargets
    {
        Assembly = 0x1,
        Module = 0x2,
        Class = 0x4,
        Struct = 0x8,
        Enum = 0x10,
        Constructor = 0x20,
        Method = 0x40,
        Property = 0x80,
        Field = 0x100,
```

```
        Event = 0x200,
        Interface = 0x400,
        Parameter = 0x800,
        Delegate = 0x1000,
        ReturnValue = 0x2000,
        GenericParameter = 0x4000,
        All = 0x7FFF
    }

    [AttributeUsageAttribute(AttributeTargets.Class, Inherited = true)]
    public sealed class AttributeUsageAttribute : Attribute
    {
        public AttributeUsageAttribute(AttributeTargets validOn);
        public bool AllowMultiple { get; set; }
        public bool Inherited { get; set; }
        public AttributeTargets ValidOn { get; }
    }

    public readonly struct Boolean { }
    public readonly struct Byte { }
    public readonly struct Char { }
    public readonly struct Decimal { }
    public abstract class Delegate { }

    public class DivideByZeroException : ArithmeticException
    {
        public DivideByZeroException();
        public DivideByZeroException(string message);
        public DivideByZeroException(string message, Exception
innerException);
    }

    public readonly struct Double { }

    public abstract class Enum : ValueType
    {
        protected Enum();
    }

    public class Exception
    {
        public Exception();
        public Exception(string message);
        public Exception(string message, Exception innerException);
        public sealed Exception InnerException { get; }
        public virtual string Message { get; }
    }

    public class GC { }

    public interface IDisposable
    {
        void Dispose();
    }
```

```
    public interface IFormattable { }

    public sealed class IndexOutOfRangeException : Exception
    {
        public IndexOutOfRangeException();
        public IndexOutOfRangeException(string message);
        public IndexOutOfRangeException(string message,
            Exception innerException);
    }

    public readonly struct Int16 { }
    public readonly struct Int32 { }
    public readonly struct Int64 { }
    public readonly struct IntPtr { }

    public class InvalidCastException : Exception
    {
        public InvalidCastException();
        public InvalidCastException(string message);
        public InvalidCastException(string message, Exception
innerException);
    }

    public class InvalidOperationException : Exception
    {
        public InvalidOperationException();
        public InvalidOperationException(string message);
        public InvalidOperationException(string message,
            Exception innerException);
    }

    public class NotSupportedException : Exception
    {
        public NotSupportedException();
        public NotSupportedException(string message);
        public NotSupportedException(string message,
            Exception innerException);
    }

    public class OperationCanceledException : Exception
    {
        public OperationCanceledException();
        public OperationCanceledException(string message);
        public OperationCanceledException(string message, Exception
innerException);
    }

    public struct Nullable<T>
    {
        public bool HasValue { get; }
        public T Value { get; }
    }

    public class NullReferenceException : Exception
    {
```

```
        public NullReferenceException();
        public NullReferenceException(string message);
        public NullReferenceException(string message, Exception
innerException);
    }

    public class Object
    {
        public Object();
        ~Object();
        public virtual bool Equals(object obj);
        public virtual int GetHashCode();
        public Type GetType();
        public virtual string ToString();
    }

    [AttributeUsageAttribute(AttributeTargets.Class |
AttributeTargets.Struct |
        AttributeTargets.Enum | AttributeTargets.Interface |
        AttributeTargets.Constructor | AttributeTargets.Method |
        AttributeTargets.Property | AttributeTargets.Field |
        AttributeTargets.Event | AttributeTargets.Delegate, Inherited =
false)]
    public sealed class ObsoleteAttribute : Attribute
    {
        public ObsoleteAttribute();
        public ObsoleteAttribute(string message);
        public ObsoleteAttribute(string message, bool error);
        public bool IsError { get; }
        public string Message { get; }
    }

    public class OutOfMemoryException : Exception
    {
        public OutOfMemoryException();
        public OutOfMemoryException(string message);
        public OutOfMemoryException(string message, Exception
innerException);
    }

    public class OverflowException : ArithmeticException
    {
        public OverflowException();
        public OverflowException(string message);
        public OverflowException(string message, Exception innerException);
    }

    public readonly struct SByte { }
    public readonly struct Single { }

    public sealed class StackOverflowException : Exception
    {
        public StackOverflowException();
        public StackOverflowException(string message);
        public StackOverflowException(string message, Exception
```

```
innerException);
    }

    public sealed class String : IEnumerable<Char>, IEnumerable
    {
        public int Length { get; }
        public char this [int index] { get; }
        public static string Format(string format, params object[] args);
    }

    public abstract class Type : MemberInfo { }

    public sealed class TypeInitializationException : Exception
    {
        public TypeInitializationException(string fullTypeName,
            Exception innerException);
    }

    public readonly struct UInt16 { }
    public readonly struct UInt32 { }
    public readonly struct UInt64 { }
    public readonly struct UIntPtr { }

    public struct ValueTuple<T1>
    {
        public T1 Item1;
        public ValueTuple(T1 item1);
    }
    public struct ValueTuple<T1, T2>
    {
        public T1 Item1;
        public T2 Item2;
        public ValueTuple(T1 item1, T2 item2);
    }
    public struct ValueTuple<T1, T2, T3>
    {
        public T1 Item1;
        public T2 Item2;
        public T3 Item3;
        public ValueTuple(T1 item1, T2 item2, T3 item3);
    }
    public struct ValueTuple<T1, T2, T3, T4>
    {
        public T1 Item1;
        public T2 Item2;
        public T3 Item3;
        public T4 Item4;
        public ValueTuple(T1 item1, T2 item2, T3 item3, T4 item4);
    }
    public struct ValueTuple<T1, T2, T3, T4, T5>
    {
        public T1 Item1;
        public T2 Item2;
        public T3 Item3;
        public T4 Item4;
```

```
        public T5 Item5;
        public ValueTuple(T1 item1, T2 item2, T3 item3, T4 item4, T5 item5);
    }
    public struct ValueTuple<T1, T2, T3, T4, T5, T6>
    {
        public T1 Item1;
        public T2 Item2;
        public T3 Item3;
        public T4 Item4;
        public T5 Item5;
        public T6 Item6;
        public ValueTuple(T1 item1, T2 item2, T3 item3, T4 item4, T5 item5,
            T6 item6);
    }
    public struct ValueTuple<T1, T2, T3, T4, T5, T6, T7>
    {
        public T1 Item1;
        public T2 Item2;
        public T3 Item3;
        public T4 Item4;
        public T5 Item5;
        public T6 Item6;
        public T7 Item7;
        public ValueTuple(T1 item1, T2 item2, T3 item3, T4 item4, T5 item5,
            T6 item6, T7 item7);
    }
    public struct ValueTuple<T1, T2, T3, T4, T5, T6, T7, TRest>
    {
        public T1 Item1;
        public T2 Item2;
        public T3 Item3;
        public T4 Item4;
        public T5 Item5;
        public T6 Item6;
        public T7 Item7;
        public TRest Rest;
        public ValueTuple(T1 item1, T2 item2, T3 item3, T4 item4, T5 item5,
            T6 item6, T7 item7, TRest rest);
    }

    public abstract class ValueType
    {
        protected ValueType();
    }
}

namespace System.Collections
{
    public interface ICollection : IEnumerable
    {
        int Count { get; }
        bool IsSynchronized { get; }
        object SyncRoot { get; }
        void CopyTo(Array array, int index);
    }
```

```
    public interface IEnumerable
    {
        IEnumerator GetEnumerator();
    }

    public interface IEnumerator
    {
        object Current { get; }
        bool MoveNext();
        void Reset();
    }

    public interface IList : ICollection, IEnumerable
    {
        bool IsFixedSize { get; }
        bool IsReadOnly { get; }
        object this [int index] { get; set; }
        int Add(object value);
        void Clear();
        bool Contains(object value);
        int IndexOf(object value);
        void Insert(int index, object value);
        void Remove(object value);
        void RemoveAt(int index);
    }
}

namespace System.Collections.Generic
{
    public interface ICollection<T> : IEnumerable<T>
    {
        int Count { get; }
        bool IsReadOnly { get; }
        void Add(T item);
        void Clear();
        bool Contains(T item);
        void CopyTo(T[] array, int arrayIndex);
        bool Remove(T item);
    }

    public interface IEnumerable<T> : IEnumerable
    {
        IEnumerator<T> GetEnumerator();
    }

    public interface IEnumerator<T> : IDisposable, IEnumerator
    {
        T Current { get; }
    }

    public interface IList<T> : ICollection<T>
    {
        T this [int index] { get; set; }
        int IndexOf(T item);
```

```
        void Insert(int index, T item);
        void RemoveAt(int index);
    }

    public interface IReadOnlyCollection<out T> : IEnumerable<T>
    {
        int Count { get; }
    }

    public interface IReadOnlyList<out T> : IReadOnlyCollection<T>
    {
        T this [int index] { get; }
    }
}

namespace System.Diagnostics
{
    [AttributeUsageAttribute(AttributeTargets.Method |
AttributeTargets.Class,
                             AllowMultiple = true)]
    public sealed class ConditionalAttribute : Attribute
    {
        public ConditionalAttribute(string conditionString);
        public string ConditionString { get; }
    }
}

namespace System.Reflection
{
    public abstract class MemberInfo
    {
        protected MemberInfo();
    }
}

namespace System.Runtime.CompilerServices
{
    public sealed class IndexerNameAttribute : Attribute
    {
        public IndexerNameAttribute(String indexerName);
    }

    public static class Unsafe
    {
        public static ref T NullRef<T>();
    }
}

namespace System.Threading
{
    public static class Monitor
    {
        public static void Enter(object obj);
        public static void Exit(object obj);
```

```
    }
}
```

C.3 Standard Library Types not defined in ISO/IEC 23271

The following types, including the members listed, shall be defined in a conforming standard library. (These types might be defined in a future edition of ISO/IEC 23271.) It is expected that many of these types will have more members available than are listed.

A conforming implementation may provide `Task.GetAwaiter()` and `Task<TResult>.GetAwaiter()` as extension methods.

C#

```
namespace System
{
    public class FormattableString : IFormattable { }
    public readonly ref struct ReadOnlySpan<T>
    {
        public int Length { get; }
        public ref readonly T this[int index] { get; }
    }
    public readonly ref struct Span<T>
    {
        public int Length { get; }
        public ref T this[int index] { get; }
        public static implicit operator ReadOnlySpan<T>(Span<T> span);
    }
}
namespace System.Linq.Expressions
{
    public sealed class Expression<TDelegate>
    {
        public TDelegate Compile();
    }
}
namespace System.Runtime.CompilerServices
{
    [AttributeUsage(AttributeTargets.Class | AttributeTargets.Struct |
        AttributeTargets.Interface,
        Inherited = false, AllowMultiple = false)]
    public sealed class AsyncMethodBuilderAttribute : Attribute
    {
        public AsyncMethodBuilderAttribute(Type builderType) {}

        public Type BuilderType { get; }
    }
    [AttributeUsage(AttributeTargets.Parameter, Inherited = false)]
```

```
    public sealed class CallerFilePathAttribute : Attribute
    {
        public CallerFilePathAttribute() { }
    }
    [AttributeUsage(AttributeTargets.Parameter, Inherited = false)]
    public sealed class CallerLineNumberAttribute : Attribute
    {
        public CallerLineNumberAttribute() { }
    }
    [AttributeUsage(AttributeTargets.Parameter, Inherited = false)]
    public sealed class CallerMemberNameAttribute : Attribute
    {
        public CallerMemberNameAttribute() { }
    }
    public static class FormattableStringFactory
    {
        public static FormattableString Create(string format,
            params object[] arguments);
    }
    public interface ICriticalNotifyCompletion : INotifyCompletion
    {
        void UnsafeOnCompleted(Action continuation);
    }
    public interface INotifyCompletion
    {
        void OnCompleted(Action continuation);
    }
    public readonly struct TaskAwaiter : ICriticalNotifyCompletion,
        INotifyCompletion
    {
        public bool IsCompleted { get; }
        public void GetResult();
    }
    public readonly struct TaskAwaiter<TResult> : ICriticalNotifyCompletion,
        INotifyCompletion
    {
        public bool IsCompleted { get; }
        public TResult GetResult();
    }
    public readonly struct ValueTaskAwaiter : ICriticalNotifyCompletion,
        INotifyCompletion
    {
        public bool IsCompleted { get; }
        public void GetResult();
    }
    public readonly struct ValueTaskAwaiter<TResult>
        : ICriticalNotifyCompletion, INotifyCompletion
    {
        public bool IsCompleted { get; }
        public TResult GetResult();
    }
}
namespace System.Threading.Tasks
{
    public class Task
```

```
    {
        public System.Runtime.CompilerServices.TaskAwaiter GetAwaiter();
    }
    public class Task<TResult> : Task
    {
        public new System.Runtime.CompilerServices.TaskAwaiter<T>
GetAwaiter();
    }
    public readonly struct ValueTask : System.IEquatable<ValueTask>
    {
        public System.Runtime.CompilerServices.ValueTaskAwaiter
GetAwaiter();
    }
    public readonly struct ValueTask<TResult>
        : System.IEquatable<ValueTask<TResult>>
    {
        public new System.Runtime.CompilerServices.ValueTaskAwaiter<TResult>
            GetAwaiter();
    }
}
```

C.4 Format Specifications

The meaning of the formats, as used in interpolated string expressions (§12.8.3), are defined in ISO/IEC 23271:2012. For convenience the following text is copied from the description of `System.IFormattable`.

This text is informative.

A *format* is a string that describes the appearance of an object when it is converted to a string. Either standard or custom formats can be used. A standard format takes the form *Axx*, where *A* is a single alphabetic character called the *format specifier*, and *xx* is an integer between zero and 99 inclusive, called the *precision specifier*. The format specifier controls the type of formatting applied to the value being represented as a string. The *precision specifier* controls the number of significant digits or decimal places in the string, if applicable.

> *Note*: For the list of standard format specifiers, see the table below. Note that a given data type, such as `System.Int32`, might not support one or more of the standard format specifiers. *end note*

> *Note*: When a format includes symbols that vary by culture, such as the currencysymbol included by the 'C' and 'c' formats, a formatting object supplies the actual characters used in the string representation. A method might include a parameter to pass a `System.IFormatProvider` object that supplies a formatting

> object, or the method might use the default formatting object, which contains the symbol definitions for the current culture. The current culture typically uses the same set of symbols used system-wide by default. In the Base Class Library, the formatting object for system-supplied numeric types is a `System.Globalization.NumberFormatInfo` instance. For `System.DateTime` instances, a `System.Globalization.DateTimeFormatInfo` is used. *end note*

The following table describes the standard format specifiers and associated formatting object members that are used with numeric data types in the Base Class Library.

Expand table

Format Specifier	Description

`C`
 	
`c`
	Currency Format: Used for strings containing a monetary value. The `System.Globalization.NumberFormatInfo.CurrencySymbol`, `System.Globalization.NumberFormatInfo.CurrencyGroupSizes`, `System.Globalization.NumberFormatInfo.CurrencyGroupSeparator`, and `System.Globalization.NumberFormatInfo.CurrencyDecimalSeparator` members of a `System.Globalization.NumberFormatInfo` supply the currency symbol, size and separator for digit groupings, and decimal separator, respectively.
	
	`System.Globalization.NumberFormatInfo.CurrencyNegativePattern` and `System.Globalization.NumberFormatInfo.CurrencyPositivePattern` determine the symbols used to represent negative and positive values. For example, a negative value can be prefixed with a minus sign, or enclosed in parentheses.
	
	If the precision specifier is omitted, `System.Globalization.NumberFormatInfo.CurrencyDecimalDigits` determines the number of decimal places in the string. Results are rounded to the nearest representable value when necessary.

`D`
 	
`d`
	Decimal Format: (This format is valid only when specified with integral data types.) Used for strings containing integer values. Negative numbers are prefixed with the negative number symbol specified by the `System.Globalization.NumberFormatInfo.NegativeSign` property.
	
	The precision specifier determines the minimum number of digits that appear in the string. If the specified precision requires more digits than the value contains, the string is left-padded with zeros. If the precision specifier specifies fewer digits than are in the value, the precision specifier is ignored.

`E`
 	
`e`
	Scientific (Engineering) Format: Used for strings in one of the following forms:
	
	[-]*m.ddddddE+xxx*
	
	[-]*m.ddddddE-xxx*

[-]*m.dddddd*e+*xxx*

[-]*m.dddddd*e-*xxx*

The negative number symbol ('-') appears only if the value is negative, and is supplied by the `System.Globalization.NumberFormatInfo.NegativeSign` property.

Exactly one non-zero decimal digit (*m*) precedes the decimal separator ('.'), which is supplied by the `System.Globalization.NumberFormatInfo.NumberDecimalSeparator` property.

The precision specifier determines the number of decimal places (*dddddd*) in the string. If the precision specifier is omitted, six decimal places are included in the string.

The exponent (*+/-xxx*) consists of either a positive or negative number symbol followed by a minimum of three digits (*xxx*). The exponent is left-padded with zeros, if necessary. The case of the format specifier ('E' or 'e') determines the case used for the exponent prefix (E or e) in the string. Results are rounded to the nearest representable value when necessary. The positive number symbol is supplied by the `System.Globalization.NumberFormatInfo.PositiveSign` property.

`F`

`f`

Fixed-Point Format: Used for strings in the following form:

[-]*m.dd...d*

At least one non-zero decimal digit (*m*) precedes the decimal separator ('.'), which is supplied by the `System.Globalization.NumberFormatInfo.NumberDecimalSeparator` property.

A negative number symbol sign ('-') precedes *m* only if the value is negative. This symbol is supplied by the `System.Globalization.NumberFormatInfo.NegativeSign` property.

The precision specifier determines the number of decimal places (*dd...d*) in the string. If the precision specifier is omitted, `System.Globalization.NumberFormatInfo.NumberDecimalDigits` determines the number of decimal places in the string. Results are rounded to the nearest representable value when necessary.

`G`

`g`

General Format: The string is formatted in either fixed-point format ('F' or 'f') or scientific format ('E' or 'e').

For integral types:

Values are formatted using fixed-point format if *exponent* < precision specifier, where *exponent* is the exponent of the value in scientific format. For all other values, scientific format is used.

If the precision specifier is omitted, a default precision equal to the field width required to display the maximum value for the data type is used, which results in the

value being formatted in fixed-point format. The default precisions for integral types are as follows:

`System.Int16`, `System.UInt16` : 5

`System.Int32`, `System.UInt32` : 10

`System.Int64`, `System.UInt64` : 19

For Single, Decimal and Double types:

Values are formatted using fixed-point format if *exponent* ≥ -4 and *exponent* < precision specifier, where *exponent* is the exponent of the value in scientific format. For all other values, scientific format is used. Results are rounded to the nearest representable value when necessary.

If the precision specifier is omitted, the following default precisions are used:

`System.Single` : 7

`System.Double` : 15

`System.Decimal` : 29

For all types:

- The number of digits that appear in the result (not including the exponent) will not exceed the value of the precision specifier; values are rounded as necessary.
- The decimal point and any trailing zeros after the decimal point are removed whenever possible.
- The case of the format specifier ('G' or 'g') determines whether 'E' or 'e' prefixes the scientific format exponent.

`N`

`n`

Number Format: Used for strings in the following form:

[-]*d,ddd,ddd.dd...d*

The representation of negative values is determined by the `System.Globalization.NumberFormatInfo.NumberNegativePattern` property. If the pattern includes a negative number symbol ('-'), this symbol is supplied by the `System.Globalization.NumberFormatInfo.NegativeSign` property.

At least one non-zero decimal digit (*d*) precedes the decimal separator ('.'), which is supplied by the `System.Globalization.NumberFormatInfo.NumberDecimalSeparator` property. Digits between the decimal point and the most significant digit in the value are grouped using the group size specified by the `System.Globalization.NumberFormatInfo.NumberGroupSizes` property. The group separator (',') is inserted between each digit group, and is supplied by the `System.Globalization.NumberFormatInfo.NumberGroupSeparator` property.

The precision specifier determines the number of decimal places (*dd...d*). If the precision specifier is omitted,

	`System.Globalization.NumberFormatInfo.NumberDecimalDigits` determines the number of decimal places in the string. Results are rounded to the nearest representable value when necessary.

`P`
 	
`p`
	Percent Format: Used for strings containing a percentage. The `System.Globalization.NumberFormatInfo.PercentSymbol`, `System.Globalization.NumberFormatInfo.PercentGroupSizes`, `System.Globalization.NumberFormatInfo.PercentGroupSeparator`, and `System.Globalization.NumberFormatInfo.PercentDecimalSeparator` members of a `System.Globalization.NumberFormatInfo` supply the percent symbol, size and separator for digit groupings, and decimal separator, respectively.
	
	`System.Globalization.NumberFormatInfo.PercentNegativePattern` and `System.Globalization.NumberFormatInfo.PercentPositivePattern` determine the symbols used to represent negative and positive values. For example, a negative value can be prefixed with a minus sign, or enclosed in parentheses.
	
	If no precision is specified, the number of decimal places in the result is determined by `System.Globalization.NumberFormatInfo.PercentDecimalDigits`. Results are rounded to the nearest representable value when necessary.
	
	The result is scaled by 100 (.99 becomes 99%).

`R`
 	
`r`
	Round trip Format: (This format is valid only when specified with `System.Double` or `System.Single`.) Used to ensure that the precision of the string representation of a floating-point value is such that parsing the string does not result in a loss of precision when compared to the original value. If the maximum precision of the data type (7 for `System.Single`, and 15 for `System.Double`) would result in a loss of precision, the precision is increased by two decimal places. If a precision specifier is supplied with this format specifier, it is ignored. This format is otherwise identical to the fixed-point format.

`X`
 	
`x`
	Hexadecimal Format: (This format is valid only when specified with integral data types.) Used for string representations of numbers in Base 16. The precision determines the minimum number of digits in the string. If the precision specifies more digits than the number contains, the number is left-padded with zeros. The case of the format specifier ('X' or 'x') determines whether upper case or lower case letters are used in the hexadecimal representation.

If the numerical value is a `System.Single` or `System.Double` with a value of `NaN`, `PositiveInfinity`, or `NegativeInfinity`, the format specifier is ignored, and one of the following is returned: `System.Globalization.NumberFormatInfo.NaNSymbol`, `System.Globalization.NumberFormatInfo.PositiveInfinitySymbol`, or `System.Globalization.NumberFormatInfo.NegativeInfinitySymbol`.

A custom format is any string specified as a format that is not in the form of a standard format string (Axx) described above. The following table describes the characters that are used in constructing custom formats.

Expand table

Format Specifier	Description
`0` (zero)	**Zero placeholder:** If the value being formatted has a digit in the position where a '0' appears in the custom format, then that digit is copied to the output string; otherwise a zero is stored in that position in the output string. The position of the leftmost '0' before the decimal separator and the rightmost '0' after the decimal separator determine the range of digits that are always present in the output string. The number of Zero and/or Digit placeholders after the decimal separator determines the number of digits that appear after the decimal separator. Values are rounded as necessary.
`#`	**Digit placeholder:** If the value being formatted has a digit in the position where a '#' appears in the custom format, then that digit is copied to the output string; otherwise, nothing is stored in that position in the output string. Note that this specifier never stores the '0' character if it is not a significant digit, even if '0' is the only digit in the string. (It does display the '0' character in the output string if it is a significant digit.) The number of Zero and/or Digit placeholders after the decimal separator determines the number of digits that appear after the decimal separator. Values are rounded as necessary.
`.` (period)	**Decimal separator:** The left most '.' character in the format string determines the location of the decimal separator in the formatted value; any additional '.' characters are ignored. The `System.Globalization.NumberFormatInfo.NumberDecimalSeparator` property determines the symbol used as the decimal separator.
`,` (comma)	**Group separator and number scaling:** The ',' character serves two purposes. First, if the custom format contains this character between two Zero or Digit placeholders (0 or #) and to the left of the decimal separator if one is present, then the output will have group separators inserted between each group of digits to the left of the decimal separator. The `System.Globalization.NumberFormatInfo.NumberGroupSeparator` and `System.Globalization.NumberFormatInfo.NumberGroupSizes` properties determine the symbol used as the group separator and the number of digits in each group, respectively. If the format string contains one or more ',' characters immediately to the left of the decimal separator, then the number will be scaled. The scale factor is determined by the number of group separator characters immediately to the left of the decimal separator. If there are x characters, then the value is divided by 1000^X before it is formatted. For example, the format string '0,,' will divide a value by one million. Note that the presence of the ',' character to indicate scaling does not insert group separators in the output string. Thus, to scale a number by 1 million and insert group separators, use a custom format similar to '#,##0,,'.

`%` (percent)	**Percentage placeholder:** The presence of a '%' character in a custom format causes a number to be multiplied by 100 before it is formatted. The percent symbol is inserted in the output string at the location where the '%' appears in the format string. The `System.Globalization.NumberFormatInfo.PercentSymbol` property determines the percent symbol.

`E0`
`E+0`
`E-0`
`e0`
`e+0`
`e-0`
	Engineering format: If any of the strings 'E', 'E+', 'E-', 'e', 'e+', or 'e-' are present in a custom format and is followed immediately by at least one '0' character, then the value is formatted using scientific notation. The number of '0' characters following the exponent prefix (E or e) determines the minimum number of digits in the exponent. The 'E+' and 'e+' formats indicate that a positive or negative number symbol always precedes the exponent. The 'E', 'E-', 'e', or 'e-' formats indicate that a negative number symbol precedes negative exponents; no symbol is precedes positive exponents. The positive number symbol is supplied by the `System.Globalization.NumberFormatInfo.PositiveSign` property. The negative number symbol is supplied by the `System.Globalization.NumberFormatInfo.NegativeSign` property.

`\` (backslash)	**Escape character:** In some languages, such as C#, the backslash character causes the next character in the custom format to be interpreted as an escape sequence. It is used with C language formatting sequences, such as '\n' (newline). In some languages, the escape character itself is required to be preceded by an escape character when used as a literal. Otherwise, the compiler interprets the character as an escape sequence. This escape character is not required to be supported in all programming languages.
`'ABC'` `"ABC"`	**Literal string:** Characters enclosed in single or double quotes are copied to the output string literally, and do not affect formatting.
`;` (semicolon)	**Section separator:** The ';' character is used to separate sections for positive, negative, and zero numbers in the format string. (This feature is described in detail below.)
Other	**All other characters:** All other characters are stored in the output string as literals in the position in which they appear.

Note that for fixed-point format strings (strings not containing an 'E0', 'E+0', 'E-0', 'e0', 'e+0', or 'e-0'), numbers are rounded to as many decimal places as there are Zero or Digit placeholders to the right of the decimal separator. If the custom format does not contain a decimal separator, the number is rounded to the nearest integer. If the number has more digits than there are Zero or Digit placeholders to the left of the decimal separator, the extra digits are copied to the output string immediately before the first Zero or Digit placeholder.

A custom format can contain up to three sections separated by section separator characters, to specify different formatting for positive, negative, and zero values. The sections are interpreted as follows:

- **One section**: The custom format applies to all values (positive, negative and zero). Negative values include a negative sign.
- **Two sections**: The first section applies to positive values and zeros, and the second section applies to negative values. If the value to be formatted is negative, but becomes zero after rounding according to the format in the second section, then the resulting zero is formatted according to the first section. Negative values do not include a negative sign to allow full control over representations of negative values. For example, a negative can be represented in parenthesis using a custom format similar to '####.####;(####.####)'.
- **Three sections**: The first section applies to positive values, the second section applies to negative values, and the third section applies to zeros. The second section can be empty (nothing appears between the semicolons), in which case the first section applies to all nonzero values, and negative values include a negative sign. If the number to be formatted is nonzero, but becomes zero after rounding according to the format in the first or second section, then the resulting zero is formatted according to the third section.

The `System.Enum` and `System.DateTime` types also support using format specifiers to format string representations of values. The meaning of a specific format specifier varies according to the kind of data (numeric, date/time, enumeration) being formatted. See `System.Enum` and `System.Globalization.DateTimeFormatInfo` for a comprehensive list of the format specifiers supported by each type.

C.5 Library Type Abbreviations

The following library types are referenced in this specification. The full names of those types, including the global namespace qualifier are listed below. Throughout this specification, these types appear as either the fully qualified name; with the global namespace qualifier omitted; or as a simple unqualified type name, with the namespace omitted as well. For example, the type `ICollection<T>`, when used in this specification, always means the type `global::System.Collections.Generic.ICollection<T>`.

- `global::System.Action`
- `global::System.ArgumentException`
- `global::System.ArithmeticException`
- `global::System.Array`
- `global::System.ArrayTypeMisMatchException`
- `global::System.Attribute`
- `global::System.AttributeTargets`

- `global::System.AttributeUsageAttribute`
- `global::System.Boolean`
- `global::System.Byte`
- `global::System.Char`
- `global::System.Collections.Generic.ICollection<T>`
- `global::System.Collections.Generic.IEnumerable<T>`
- `global::System.Collections.Generic.IEnumerator<T>`
- `global::System.Collections.Generic.IList<T>`
- `global::System.Collections.Generic.IReadonlyCollection<out T>`
- `global::System.Collections.Generic.IReadOnlyList<out T>`
- `global::System.Collections.ICollection`
- `global::System.Collections.IEnumerable`
- `global::System.Collections.IList`
- `global::System.Collections.IEnumerator`
- `global::System.Decimal`
- `global::System.Delegate`
- `global::System.Diagnostics.ConditionalAttribute`
- `global::System.DivideByZeroException`
- `global::System.Double`
- `global::System.Enum`
- `global::System.Exception`
- `global::System.GC`
- `global::System.ICollection`
- `global::System.IDisposable`
- `global::System.IEnumerable`
- `global::System.IEnumerable<out T>`
- `global::System.IList`
- `global::System.IndexOutOfRangeException`
- `global::System.Int16`
- `global::System.Int32`
- `global::System.Int64`
- `global::System.IntPtr`
- `global::System.InvalidCastException`
- `global::System.InvalidOperationException`
- `global::System.Linq.Expressions.Expression<TDelegate>`
- `global::System.MemberInfo`
- `global::System.NotSupportedException`
- `global::System.Nullable<T>`

- `global::System.NullReferenceException`
- `global::System.Object`
- `global::System.ObsoleteAttribute`
- `global::System.OutOfMemoryException`
- `global::System.OverflowException`
- `global::System.Runtime.CompilerServices.CallerFileAttribute`
- `global::System.Runtime.CompilerServices.CallerLineNumberAttribute`
- `global::System.Runtime.CompilerServices.CallerMemberNameAttribute`
- `global::System.Runtime.CompilerServices.ICriticalNotifyCompletion`
- `global::System.Runtime.CompilerServices.IndexerNameAttribute`
- `global::System.Runtime.CompilerServices.INotifyCompletion`
- `global::System.Runtime.CompilerServices.TaskAwaiter`
- `global::System.Runtime.CompilerServices.TaskAwaiter<T>`
- `global::System.SByte`
- `global::System.Single`
- `global::System.StackOverflowException`
- `global::System.String`
- `global::System.SystemException`
- `global::System.Threading.Monitor`
- `global::System.Threading.Tasks.Task`
- `global::System.Threading.Tasks.Task<TResult>`
- `global::System.Type`
- `global::System.TypeInitializationException`
- `global::System.UInt16`
- `global::System.UInt32`
- `global::System.UInt64`
- `global::System.UIntPtr`
- `global::System.ValueType`

End of informative text.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

more information, see our contributor guide.

Open a documentation issue

Provide product feedback

Annex D Documentation comments

Article • 02/07/2024

This annex is informative.

D.1 General

C# provides a mechanism for programmers to document their code using a comment syntax that contains XML text. In source code files, comments having a certain form can be used to direct a tool to produce XML from those comments and the source code elements, which they precede. Comments using such syntax are called ***documentation comments***. They must immediately precede a user-defined type (such as a class, delegate, or interface) or a member (such as a field, event, property, or method). The XML generation tool is called the ***documentation generator***. (This generator could be, but need not be, the C# compiler itself.) The output produced by the documentation generator is called the ***documentation file***. A documentation file is used as input to a ***documentation viewer***; a tool intended to produce some sort of visual display of type information and its associated documentation.

A conforming C# compiler is not required to check the syntax of documentation comments; such comments are simply ordinary comments. A conforming compiler is permitted to do such checking, however.

This specification suggests a set of standard tags to be used in documentation comments, but use of these tags is not required, and other tags may be used if desired, as long as the rules of well-formed XML are followed. For C# implementations targeting the CLI, it also provides information about the documentation generator and the format of the documentation file. No information is provided about the documentation viewer.

D.2 Introduction

Comments having a certain form can be used to direct a tool to produce XML from those comments and the source code elements that they precede. Such comments are *Single_Line_Comment*s (§6.3.3) that start with three slashes (`///`), or *Delimited_Comment*s (§6.3.3) that start with a slash and two asterisks (`/**`). They must immediately precede a user-defined type or a member that they annotate. Attribute sections (§22.3) are considered part of declarations, so documentation comments must precede attributes applied to a type or member.

For expository purposes, the format of document comments is shown below as two grammar rules: *Single_Line_Doc_Comment* and *Delimited_Doc_Comment*. However, these rules are *not* part of the C# grammar, but rather, they represent particular formats of *Single_Line_Comment* and *Delimited_Comment* lexer rules, respectively.

Syntax:

ANTLR

```
Single_Line_Doc_Comment
    : '///' Input_Character*
    ;

Delimited_Doc_Comment
    : '/**' Delimited_Comment_Section* ASTERISK+ '/'
    ;
```

In a *Single_Line_Doc_Comment*, if there is a *Whitespace* character following the `///` characters on each of the *Single_Line_Doc_Comments* adjacent to the current *Single_Line_Doc_Comment*, then that *Whitespace* character is not included in the XML output.

In a *Delimited_Doc_Comment*, if the first non-*Whitespace* character on the second line is an *ASTERISK* and the same pattern of optional *Whitespace* characters and an *ASTERISK* character is repeated at the beginning of each of the lines within the *Delimited_Doc_Comment*, then the characters of the repeated pattern are not included in the XML output. The pattern can include *Whitespace* characters after, as well as before, the *ASTERISK* character.

Example:

C#

```
/// <summary>
/// Class <c>Point</c> models a point in a two-dimensional plane.
/// </summary>
public class Point
{
    /// <summary>
    /// Method <c>Draw</c> renders the point.
    /// </summary>
    void Draw() {...}
}
```

The text within documentation comments must be well formed according to the rules of XML (http://www.w3.org/TR/REC-xml). If the XML is ill formed, a warning is generated

and the documentation file will contain a comment saying that an error was encountered.

Although developers are free to create their own set of tags, a recommended set is defined in §D.3. Some of the recommended tags have special meanings:

- The `<param>` tag is used to describe parameters. If such a tag is used, the documentation generator must verify that the specified parameter exists and that all parameters are described in documentation comments. If such verification fails, the documentation generator issues a warning.
- The `cref` attribute can be attached to any tag to provide a reference to a code element. The documentation generator must verify that this code element exists. If the verification fails, the documentation generator issues a warning. When looking for a name described in a `cref` attribute, the documentation generator must respect namespace visibility according to using statements appearing within the source code. For code elements that are generic, the normal generic syntax (e.g., "`List<T>`") cannot be used because it produces invalid XML. Braces can be used instead of brackets (e.g.; "`List{T}`"), or the XML escape syntax can be used (e.g., "`List<T>`").
- The `<summary>` tag is intended to be used by a documentation viewer to display additional information about a type or member.
- The `<include>` tag includes information from an external XML file.

Note carefully that the documentation file does not provide full information about the type and members (for example, it does not contain any type information). To get such information about a type or member, the documentation file must be used in conjunction with reflection on the type or member.

D.3 Recommended tags

D.3.1 General

The documentation generator must accept and process any tag that is valid according to the rules of XML. The following tags provide commonly used functionality in user documentation. (Of course, other tags are possible.)

Expand table

Tag	Reference	Purpose
`<c>`	§D.3.2	Set text in a code-like font
`<code>`	§D.3.3	Set one or more lines of source code or program output
`<example>`	§D.3.4	Indicate an example
`<exception>`	§D.3.5	Identifies the exceptions a method can throw
`<include>`	§D.3.6	Includes XML from an external file
`<list>`	§D.3.7	Create a list or table
`<para>`	§D.3.8	Permit structure to be added to text
`<param>`	§D.3.9	Describe a parameter for a method or constructor
`<paramref>`	§D.3.10	Identify that a word is a parameter name
`<permission>`	§D.3.11	Document the security accessibility of a member
`<remarks>`	§D.3.12	Describe additional information about a type
`<returns>`	§D.3.13	Describe the return value of a method
`<see>`	§D.3.14	Specify a link
`<seealso>`	§D.3.15	Generate a *See Also* entry
`<summary>`	§D.3.16	Describe a type or a member of a type
`<typeparam>`	§D.3.17	Describe a type parameter for a generic type or method
`<typeparamref>`	§D.3.18	Identify that a word is a type parameter name
`<value>`	§D.3.19	Describe a property

D.3.2 <c>

This tag provides a mechanism to indicate that a fragment of text within a description should be set in a special font such as that used for a block of code. For lines of actual code, use `<code>` (§D.3.3).

Syntax:

`<c>` *text* `</c>`

Example:

C#

```
/// <summary>
/// Class <c>Point</c> models a point in a two-dimensional plane.
/// </summary>
public class Point
{
}
```

D.3.3 <code>

This tag is used to set one or more lines of source code or program output in some special font. For small code fragments in narrative, use `<c>` (§D.3.2).

Syntax:

`<code>`*source code or program output*`</code>`

Example:

C#

```
public class Point
{
    /// <summary>
    /// This method changes the point's location by the given x- and y-
offsets.
    /// <example>
    /// For example:
    /// <code>
    /// Point p = new Point(3,5);
    /// p.Translate(-1,3);
    /// </code>
    /// results in <c>p</c>'s having the value (2,8).
    /// </example>
    /// </summary>
    public void Translate(int dx, int dy)
    {
        ...
    }
}
```

D.3.4 <example>

This tag allows example code within a comment, to specify how a method or other library member might be used. Ordinarily, this would also involve use of the tag `<code>` (§D.3.3) as well.

Syntax:

`<example>` *description* `</example>`

Example:

See `<code>` (§D.3.3) for an example.

D.3.5 <exception>

This tag provides a way to document the exceptions a method can throw.

Syntax:

`<exception cref="` *member* `">` *description* `</exception>`

where

- `cref="` *member* `"` is the name of a member. The documentation generator checks that the given member exists and translates *member* to the canonical element name in the documentation file.
- *description* is a description of the circumstances in which the exception is thrown.

Example:

C#

```
class PrimaryFileFormatCorruptException : System.Exception { ... }
class PrimaryFileLockedOpenException : System.Exception { ... }

public class DataBaseOperations
{
    /// <exception cref="PrimaryFileFormatCorruptException">
    /// Thrown when the primary file is corrupted.
    /// </exception>
    /// <exception cref="PrimaryFileLockedOpenException">
    /// Thrown when the primary file is already open.
    /// </exception>
    public static void ReadRecord(int flag)
    {
        if (flag == 1)
        {
            throw new PrimaryFileFormatCorruptException();
        }
        else if (flag == 2)
        {
            throw new PrimaryFileLockedOpenException();
        }
        ...
```

```
    }
}
```

D.3.6 <include>

This tag allows including information from an XML document that is external to the source code file. The external file must be a well-formed XML document, and an XPath expression is applied to that document to specify what XML from that document to include. The `<include>` tag is then replaced with the selected XML from the external document.

Syntax:

`<include file="`*filename*`" path="`*xpath*`" />`

where

- `file="`*filename*`"` is the file name of an external XML file. The file name is interpreted relative to the file that contains the include tag.
- `path="`*xpath*`"` is an XPath expression that selects some of the XML in the external XML file.

Example:

If the source code contained a declaration like:

C#

```
/// <include file="docs.xml" path='extradoc/class[@name="IntList"]/*' />
public class IntList { ... }
```

and the external file “docs.xml” had the following contents:

XML

```
<?xml version="1.0"?>
<extradoc>
    <class name="IntList">
        <summary>
            Contains a list of integers.
        </summary>
    </class>
    <class name="StringList">
        <summary>
            Contains a list of strings.
        </summary>
```

```
    </class>
</extradoc>
```

then the same documentation is output as if the source code contained:

C#

```
/// <summary>
/// Contains a list of integers.
/// </summary>
public class IntList { ... }
```

D.3.7 <list>

This tag is used to create a list or table of items. It can contain a `<listheader>` block to define the heading row of either a table or definition list. (When defining a table, only an entry for *term* in the heading need be supplied.)

Each item in the list is specified with an `<item>` block. When creating a definition list, both *term* and *description* must be specified. However, for a table, bulleted list, or numbered list, only *description* need be specified.

Syntax:

XML

```
<list type="bullet" | "number" | "table">
    <listheader>
        <term>term</term>
        <description>description</description>
    </listheader>
    <item>
        <term>term</term>
        <description>description</description>
    </item>
    ...
    <item>
        <term>term</term>
        <description>description</description>
    </item>
</list>
```

where

- *term* is the term to define, whose definition is in *description*.
- *description* is either an item in a bullet or numbered list, or the definition of a *term*.

Example:

```
C#
```

```
public class MyClass
{
    /// <summary>Here is an example of a bulleted list:
    /// <list type="bullet">
    /// <item>
    /// <description>Item 1.</description>
    /// </item>
    /// <item>
    /// <description>Item 2.</description>
    /// </item>
    /// </list>
    /// </summary>
    public static void Main()
    {
        ...
    }
}
```

D.3.8 <para>

This tag is for use inside other tags, such as `<summary>` (§D.3.16) or `<returns>` (§D.3.13), and permits structure to be added to text.

Syntax:

`<para>`*content*`</para>`

where

- *content* is the text of the paragraph.

Example:

```
C#
```

```
public class Point
{
    /// <summary>This is the entry point of the Point class testing program.
    /// <para>
    /// This program tests each method and operator, and
    /// is intended to be run after any non-trivial maintenance has
    /// been performed on the Point class.
    /// </para>
    /// </summary>
    public static void Main()
    {
```

```
        ...
    }
}
```

D.3.9 <param>

This tag is used to describe a parameter for a method, constructor, or indexer.

Syntax:

`<param name="`*name*`">`*description*`</param>`

where

- *name* is the name of the parameter.
- *description* is a description of the parameter.

Example:

C#

```
public class Point
{
    /// <summary>
    /// This method changes the point's location to
    /// the given coordinates.
    /// </summary>
    /// <param name="xPosition">the new x-coordinate.</param>
    /// <param name="yPosition">the new y-coordinate.</param>
    public void Move(int xPosition, int yPosition)
    {
        ...
    }
}
```

D.3.10 <paramref>

This tag is used to indicate that a word is a parameter. The documentation file can be processed to format this parameter in some distinct way.

Syntax:

`<paramref name="`*name*`"/>`

where

- *name* is the name of the parameter.

Example:

C#

```
public class Point
{
    /// <summary>This constructor initializes the new Point to
    /// (<paramref name="xPosition"/>,<paramref name="yPosition"/>).
    /// </summary>
    /// <param name="xPosition">the new Point's x-coordinate.</param>
    /// <param name="yPosition">the new Point's y-coordinate.</param>
    public Point(int xPosition, int yPosition)
    {
        ...
    }
}
```

D.3.11 <permission>

This tag allows the security accessibility of a member to be documented.

Syntax:

`<permission cref="`*member*`">`*description*`</permission>`

where

- *member* is the name of a member. The documentation generator checks that the given code element exists and translates *member* to the canonical element name in the documentation file.
- *description* is a description of the access to the member.

Example:

C#

```
public class MyClass
{
    /// <permission cref="System.Security.PermissionSet">
    /// Everyone can access this method.
    /// </permission>
    public static void Test()
    {
        ...
    }
}
```

D.3.12 <remarks>

This tag is used to specify extra information about a type. Use `<summary>` (§D.3.16) to describe the type itself and the members of a type.

Syntax:

`<remarks>` *description* `</remarks>`

where

- *description* is the text of the remark.

Example:

C#

```
/// <summary>
/// Class <c>Point</c> models a point in a two-dimensional plane.
/// </summary>
/// <remarks>
/// Uses polar coordinates
/// </remarks>
public class Point
{
    ...
}
```

D.3.13 <returns>

This tag is used to describe the return value of a method.

Syntax:

`<returns>` *description* `</returns>`

where

- *description* is a description of the return value.

Example:

C#

```
public class Point
{
    /// <summary>
    /// Report a point's location as a string.
```

```
    /// </summary>
    /// <returns>
    /// A string representing a point's location, in the form (x,y),
    /// without any leading, trailing, or embedded whitespace.
    /// </returns>
    public override string ToString() => $"({X},{Y})";
    public int X { get; set; }
    public int Y { get; set; }
}
```

D.3.14 <see>

This tag allows a link to be specified within text. Use `<seealso>` (§D.3.15) to indicate text that is to appear in a *See Also* subclause.

Syntax:

`<see cref="`*member*`" href="`*url*`" langword="`*keyword*`" />`

where

- *member* is the name of a member. The documentation generator checks that the given code element exists and changes *member* to the element name in the generated documentation file.
- *url* is a reference to an external source.
- *langword* is a word to be highlighted somehow.

Example:

C#

```
public class Point
{
    /// <summary>
    /// This method changes the point's location to
    /// the given coordinates. <see cref="Translate"/>
    /// </summary>
    public void Move(int xPosition, int yPosition)
    {
        ...
    }
    /// <summary>This method changes the point's location by
    /// the given x- and y-offsets. <see cref="Move"/>
    /// </summary>
    public void Translate(int dx, int dy)
    {
        ...
```

```
    }
}
```

D.3.15 <seealso>

This tag allows an entry to be generated for the *See Also* subclause. Use `<see>` (§D.3.14) to specify a link from within text.

Syntax:

`<seealso cref="`*member*`" href="`*url*`" />`

where

- *member* is the name of a member. The documentation generator checks that the given code element exists and changes *member* to the element name in the generated documentation file.
- *url* is a reference to an external source.

Example:

C#

```
public class Point
{
    /// <summary>
    /// This method determines whether two Points have the same location.
    /// </summary>
    /// <seealso cref="operator=="/>
    /// <seealso cref="operator!="/>
    public override bool Equals(object o)
    {
        ...
    }
}
```

D.3.16 <summary>

This tag can be used to describe a type or a member of a type. Use `<remarks>` (§D.3.12) to specify extra information about the type or member.

Syntax:

`<summary>`*description*`</summary>`

where

- *description* is a summary of the type or member.

Example:

C#

```
public class Point
{

    /// <summary>
    /// This constructor initializes the new Point to
    /// (<paramref name="xPosition"/>,<paramref name="yPosition"/>).
    /// </summary>
    public Point(int xPosition, int yPosition)
    {
        ...
    }

    /// <summary>This constructor initializes the new Point to (0,0).
</summary>
    public Point() : this(0, 0)
    {
    }
}
```

D.3.17 <typeparam>

This tag is used to describe a type parameter for a generic type or method.

Syntax:

`<typeparam name="`*name*`">`*description*`</typeparam>`

where

- *name* is the name of the type parameter.
- *description* is a description of the type parameter.

Example:

C#

```
/// <summary>A generic list class.</summary>
/// <typeparam name="T">The type stored by the list.</typeparam>
public class MyList<T>
{
   ...
}
```

D.3.18 <typeparamref>

This tag is used to indicate that a word is a type parameter. The documentation file can be processed to format this type parameter in some distinct way.

Syntax:

`<typeparamref name="`*name*`"/>`

where

- *name* is the name of the type parameter.

Example:

```
C#

public class MyClass
{
    /// <summary>
    /// This method fetches data and returns a list of
    /// <typeparamref name="T"/>.
    /// </summary>
    /// <param name="query">query to execute</param>
    public List<T> FetchData<T>(string query)
    {
        ...
    }
}
```

D.3.19 <value>

This tag allows a property to be described.

Syntax:

`<value>`*property description*`</value>`

where

- *property description* is a description for the property.

Example:

```
C#

public class Point
{
```

```
    /// <value>Property <c>X</c> represents the point's x-coordinate.
</value>
    public int X { get; set; }
}
```

D.4 Processing the documentation file

D.4.1 General

The following information is intended for C# implementations targeting the CLI.

The documentation generator generates an ID string for each element in the source code that is tagged with a documentation comment. This ID string uniquely identifies a source element. A documentation viewer can use an ID string to identify the corresponding item to which the documentation applies.

The documentation file is not a hierarchical representation of the source code; rather, it is a flat list with a generated ID string for each element.

D.4.2 ID string format

The documentation generator observes the following rules when it generates the ID strings:

- No white space is placed in the string.
- The first part of the string identifies the kind of member being documented, via a single character followed by a colon. The following kinds of members are defined:

Expand table

Character	Description
E	Event
F	Field
M	Method (including constructors, finalizers, and operators)
N	Namespace
P	Property (including indexers)
T	Type (such as class, delegate, enum, interface, and struct)

Character	Description
!	Error string; the rest of the string provides information about the error. For example, the documentation generator generates error information for links that cannot be resolved.

- The second part of the string is the fully qualified name of the element, starting at the root of the namespace. The name of the element, its enclosing type(s), and namespace are separated by periods. If the name of the item itself has periods, they are replaced by # (U+0023) characters. (It is assumed that no element has this character in its name.) Type arguments in the fully qualified name, for when a member explicitly implements a member of a generic interface, are encoded by replacing the "`<`" and "`>`" surrounding them with the "`{`" and "`}`" characters.

- For methods and properties with arguments, the argument list follows, enclosed in parentheses. For those without arguments, the parentheses are omitted. The arguments are separated by commas. The encoding of each argument is the same as a CLI signature, as follows:
  - Arguments are represented by their documentation name, which is based on their fully qualified name, modified as follows:
    - Arguments that represent generic types have an appended "`'`" character followed by the number of type parameters
    - Arguments having the `in`, `out` or `ref` modifier have an `@` following their type name. Arguments passed by value or via `params` have no special notation.
    - Arguments that are arrays are represented as `[` *lowerbound* `:` *size* `,` ... `,` *lowerbound* `:` *size* `]` where the number of commas is the rank less one, and the lower bounds and size of each dimension, if known, are represented in decimal. If a lower bound or size is not specified, it is omitted. If the lower bound and size for a particular dimension are omitted, the "`:`" is omitted as well. Jagged arrays are represented by one "`[]`" per level. Single dimensional arrays omit the lower bound when the lower bound is 0 (the default) (§17.1).
    - Arguments that have pointer types other than `void` are represented using a `*` following the type name. A `void` pointer is represented using a type name of `System.Void`.
    - Arguments that refer to generic type parameters defined on types are encoded using the "`` ` ``" character followed by the zero-based index of the type parameter.
    - Arguments that use generic type parameters defined in methods use a double-backtick "`` `` ``" instead of the "`` ` ``" used for types.

- Arguments that refer to constructed generic types are encoded using the generic type, followed by "`{`", followed by a comma-separated list of type arguments, followed by "`}`".

D.4.3 ID string examples

The following examples each show a fragment of C# code, along with the ID string produced from each source element capable of having a documentation comment:

Types are represented using their fully qualified name, augmented with generic information:

C#

```
enum Color { Red, Blue, Green }

namespace Acme
{
    interface IProcess { ... }

    struct ValueType { ... }

    class Widget : IProcess
    {
        public class NestedClass { ... }
        public interface IMenuItem { ... }
        public delegate void Del(int i);
        public enum Direction { North, South, East, West }
    }

    class MyList<T>
    {
        class Helper<U,V> { ... }
    }
}
```

IDs:

Console

```
"T:Color"
"T:Acme.IProcess"
"T:Acme.ValueType"
"T:Acme.Widget"
"T:Acme.Widget.NestedClass"
"T:Acme.Widget.IMenuItem"
"T:Acme.Widget.Del"
"T:Acme.Widget.Direction"
```

```
"T:Acme.MyList`1"
"T:Acme.MyList`1.Helper`2"
```

Fields are represented by their fully qualified name.

C#

```
namespace Acme
{
    struct ValueType
    {
        private int total;
    }

    class Widget : IProcess
    {
        public class NestedClass
        {
            private int value;
        }

        private string message;
        private static Color defaultColor;
        private const double PI = 3.14159;
        protected readonly double monthlyAverage;
        private long[] array1;
        private Widget[,] array2;
        private unsafe int *pCount;
        private unsafe float **ppValues;
    }
}
```

IDs:

Console

```
"F:Acme.ValueType.total"
"F:Acme.Widget.NestedClass.value"
"F:Acme.Widget.message"
"F:Acme.Widget.defaultColor"
"F:Acme.Widget.PI"
"F:Acme.Widget.monthlyAverage"
"F:Acme.Widget.array1"
"F:Acme.Widget.array2"
"F:Acme.Widget.pCount"
"F:Acme.Widget.ppValues"
```

Constructors

C#

```
namespace Acme
{
    class Widget : IProcess
    {
        static Widget() { ... }
        public Widget() { ... }
        public Widget(string s) { ... }
    }
}
```

IDs:

Console

```
"M:Acme.Widget.#cctor"
"M:Acme.Widget.#ctor"
"M:Acme.Widget.#ctor(System.String)"
```

Finalizers

C#

```
namespace Acme
{
    class Widget : IProcess
    {
        ~Widget() { ... }
    }
}
```

IDs:

Console

```
"M:Acme.Widget.Finalize"
```

Methods

C#

```
namespace Acme
{
    struct ValueType
    {
        public void M(int i) { ... }
    }

    class Widget : IProcess
```

```
    {
        public class NestedClass
        {
            public void M(int i) { ... }
        }

        public static void M0() { ... }
        public void M1(char c, out float f, ref ValueType v, in int i) { ...
}
        public void M2(short[] x1, int[,] x2, long[][] x3) { ... }
        public void M3(long[][] x3, Widget[][,,] x4) { ... }
        public unsafe void M4(char *pc, Color **pf) { ... }
        public unsafe void M5(void *pv, double *[][,] pd) { ... }
        public void M6(int i, params object[] args) { ... }
    }

    class MyList<T>
    {
        public void Test(T t) { ... }
    }

    class UseList
    {
        public void Process(MyList<int> list) { ... }
        public MyList<T> GetValues<T>(T value) { ... }
    }
}
```

IDs:

Console

```
"M:Acme.ValueType.M(System.Int32)"
"M:Acme.Widget.NestedClass.M(System.Int32)"
"M:Acme.Widget.M0"
"M:Acme.Widget.M1(System.Char,System.Single@,Acme.ValueType@,System.Int32@)"
"M:Acme.Widget.M2(System.Int16[],System.Int32[0:,0:],System.Int64[][])"
"M:Acme.Widget.M3(System.Int64[][],Acme.Widget[0:,0:,0:][])"
"M:Acme.Widget.M4(System.Char*,Color**)"
"M:Acme.Widget.M5(System.Void*,System.Double*[0:,0:][])"
"M:Acme.Widget.M6(System.Int32,System.Object[])"
"M:Acme.MyList`1.Test(`0)"
"M:Acme.UseList.Process(Acme.MyList{System.Int32})"
"M:Acme.UseList.GetValues``1(``0)"
```

Properties and indexers

C#

```
namespace Acme
{
    class Widget : IProcess
```

```
    {
        public int Width { get { ... } set { ... } }
        public int this[int i] { get { ... } set { ... } }
        public int this[string s, int i] { get { ... } set { ... } }
    }
}
```

IDs:

Console

```
"P:Acme.Widget.Width"
"P:Acme.Widget.Item(System.Int32)"
"P:Acme.Widget.Item(System.String,System.Int32)"
```

Events

C#

```
namespace Acme
{
    class Widget : IProcess
    {
        public event Del AnEvent;
    }
}
```

IDs:

Console

```
"E:Acme.Widget.AnEvent"
```

Unary operators

C#

```
namespace Acme
{
    class Widget : IProcess
    {
        public static Widget operator+(Widget x) { ... }
    }
}
```

IDs:

Console

```
"M:Acme.Widget.op_UnaryPlus(Acme.Widget)"
```

The complete set of unary operator function names used is as follows: `op_UnaryPlus`, `op_UnaryNegation`, `op_LogicalNot`, `op_OnesComplement`, `op_Increment`, `op_Decrement`, `op_True`, and `op_False`.

Binary operators

C#

```
namespace Acme
{
    class Widget : IProcess
    {
        public static Widget operator+(Widget x1, Widget x2) { ... }
    }
}
```

IDs:

Console

```
"M:Acme.Widget.op_Addition(Acme.Widget,Acme.Widget)"
```

The complete set of binary operator function names used is as follows: `op_Addition`, `op_Subtraction`, `op_Multiply`, `op_Division`, `op_Modulus`, `op_BitwiseAnd`, `op_BitwiseOr`, `op_ExclusiveOr`, `op_LeftShift`, `op_RightShift`, `op_Equality`, `op_Inequality`, `op_LessThan`, `op_LessThanOrEqual`, `op_GreaterThan`, and `op_GreaterThanOrEqual`.

Conversion operators have a trailing "`~`" followed by the return type. When either the source or destination of a conversion operator is a generic type, the "`<`" and "`">`" characters are replaced by the "`{`" and "`}`" characters, respectively.

C#

```
namespace Acme
{
    class Widget : IProcess
    {
        public static explicit operator int(Widget x) { ... }
        public static implicit operator long(Widget x) { ... }
    }
}
```

IDs:

Console

```
"M:Acme.Widget.op_Explicit(Acme.Widget)~System.Int32"
"M:Acme.Widget.op_Implicit(Acme.Widget)~System.Int64"
```

D.5 An example

D.5.1 C# source code

The following example shows the source code of a Point class:

C#

```
namespace Graphics
{
    /// <summary>
    /// Class <c>Point</c> models a point in a two-dimensional plane.
    /// </summary>
    public class Point
    {
        /// <value>
        /// Property <c>X</c> represents the point's x-coordinate.
        /// </value>
        public int X { get; set; }

        /// <value>
        /// Property <c>Y</c> represents the point's y-coordinate.
        /// </value>
        public int Y { get; set; }

        /// <summary>
        /// This constructor initializes the new Point to (0,0).
        /// </summary>
        public Point() : this(0, 0) {}

        /// <summary>
        /// This constructor initializes the new Point to
        /// (<paramref name="xPosition"/>,<paramref name="yPosition"/>).
        /// </summary>
        /// <param name="xPosition">The new Point's x-coordinate.</param>
        /// <param name="yPosition">The new Point's y-coordinate.</param>
        public Point(int xPosition, int yPosition)
        {
            X = xPosition;
            Y = yPosition;
        }
```

```
        /// <summary>
        /// This method changes the point's location to
        /// the given coordinates. <see cref="Translate"/>
        /// </summary>
        /// <param name="xPosition">The new x-coordinate.</param>
        /// <param name="yPosition">The new y-coordinate.</param>
        public void Move(int xPosition, int yPosition)
        {
            X = xPosition;
            Y = yPosition;
        }

        /// <summary>
        /// This method changes the point's location by
        /// the given x- and y-offsets.
        /// <example>For example:
        /// <code>
        /// Point p = new Point(3, 5);
        /// p.Translate(-1, 3);
        /// </code>
        /// results in <c>p</c>'s having the value (2, 8).
        /// <see cref="Move"/>
        /// </example>
        /// </summary>
        /// <param name="dx">The relative x-offset.</param>
        /// <param name="dy">The relative y-offset.</param>
        public void Translate(int dx, int dy)
        {
            X += dx;
            Y += dy;
        }

        /// <summary>
        /// This method determines whether two Points have the same
location.
        /// </summary>
        /// <param name="o">
        /// The object to be compared to the current object.
        /// </param>
        /// <returns>
        /// True if the Points have the same location and they have
        /// the exact same type; otherwise, false.
        /// </returns>
        /// <seealso cref="operator=="/>
        /// <seealso cref="operator!="/>
        public override bool Equals(object o)
        {
            if (o == null)
            {
                return false;
            }
            if ((object)this == o)
            {
                return true;
            }
```

```
            if (GetType() == o.GetType())
            {
                Point p = (Point)o;
                return (X == p.X) && (Y == p.Y);
            }
            return false;
        }

        /// <summary>
        /// This method returns a Point's hashcode.
        /// </summary>
        /// <returns>
        /// The int hashcode.
        /// </returns>
        public override int GetHashCode()
        {
            return X + (Y >> 4);    // a crude version
        }

        /// <summary>Report a point's location as a string.</summary>
        /// <returns>
        /// A string representing a point's location, in the form (x,y),
        /// without any leading, training, or embedded whitespace.
        /// </returns>
        public override string ToString() => $"({X},{Y})";

        /// <summary>
        /// This operator determines whether two Points have the same
location.
        /// </summary>
        /// <param name="p1">The first Point to be compared.</param>
        /// <param name="p2">The second Point to be compared.</param>
        /// <returns>
        /// True if the Points have the same location and they have
        /// the exact same type; otherwise, false.
        /// </returns>
        /// <seealso cref="Equals"/>
        /// <seealso cref="operator!="/>
        public static bool operator==(Point p1, Point p2)
        {
            if ((object)p1 == null || (object)p2 == null)
            {
                return false;
            }
            if (p1.GetType() == p2.GetType())
            {
                return (p1.X == p2.X) && (p1.Y == p2.Y);
            }
            return false;
        }

        /// <summary>
        /// This operator determines whether two Points have the same
location.
        /// </summary>
```

```
        /// <param name="p1">The first Point to be compared.</param>
        /// <param name="p2">The second Point to be compared.</param>
        /// <returns>
        /// True if the Points do not have the same location and the
        /// exact same type; otherwise, false.
        /// </returns>
        /// <seealso cref="Equals"/>
        /// <seealso cref="operator=="/>
        public static bool operator!=(Point p1, Point p2) => !(p1 == p2);
    }
}
```

D.5.2 Resulting XML

Here is the output produced by one documentation generator when given the source code for class `Point`, shown above:

XML

```
<?xml version="1.0"?>
<doc>
  <assembly>
    <name>Point</name>
  </assembly>
  <members>
    <member name="T:Graphics.Point">
    <summary>Class <c>Point</c> models a point in a two-dimensional
    plane.
    </summary>
    </member>
    <member name="M:Graphics.Point.#ctor">
      <summary>This constructor initializes the new Point to (0, 0).
</summary>
    </member>
    <member name="M:Graphics.Point.#ctor(System.Int32,System.Int32)">
      <summary>
        This constructor initializes the new Point to
        (<paramref name="xPosition"/>,<paramref name="yPosition"/>).
      </summary>
      <param name="xPosition">The new Point's x-coordinate.</param>
      <param name="yPosition">The new Point's y-coordinate.</param>
    </member>
    <member name="M:Graphics.Point.Move(System.Int32,System.Int32)">
      <summary>
        This method changes the point's location to
        the given coordinates.
        <see cref="M:Graphics.Point.Translate(System.Int32,System.Int32)"/>
      </summary>
      <param name="xPosition">The new x-coordinate.</param>
      <param name="yPosition">The new y-coordinate.</param>
      </member>
```

```
    <member name="M:Graphics.Point.Translate(System.Int32,System.Int32)">
      <summary>
        This method changes the point's location by
        the given x- and y-offsets.
        <example>For example:
        <code>
        Point p = new Point(3,5);
        p.Translate(-1,3);
        </code>
        results in <c>p</c>'s having the value (2,8).
        </example>
        <see cref="M:Graphics.Point.Move(System.Int32,System.Int32)"/>
      </summary>
      <param name="dx">The relative x-offset.</param>
      <param name="dy">The relative y-offset.</param>
    </member>
    <member name="M:Graphics.Point.Equals(System.Object)">
      <summary>
        This method determines whether two Points have the same location.
      </summary>
      <param name="o">
        The object to be compared to the current object.
      </param>
      <returns>
        True if the Points have the same location and they have
        the exact same type; otherwise, false.
      </returns>
      <seealso
        cref="M:Graphics.Point.op_Equality(Graphics.Point,Graphics.Point)"
/>
      <seealso

cref="M:Graphics.Point.op_Inequality(Graphics.Point,Graphics.Point)"/>
    </member>
     <member name="M:Graphics.Point.ToString">
      <summary>
        Report a point's location as a string.
      </summary>
      <returns>
        A string representing a point's location, in the form (x,y),
        without any leading, training, or embedded whitespace.
      </returns>
     </member>
    <member
name="M:Graphics.Point.op_Equality(Graphics.Point,Graphics.Point)">
      <summary>
        This operator determines whether two Points have the same location.
      </summary>
      <param name="p1">The first Point to be compared.</param>
      <param name="p2">The second Point to be compared.</param>
      <returns>
        True if the Points have the same location and they have
        the exact same type; otherwise, false.
      </returns>
      <seealso cref="M:Graphics.Point.Equals(System.Object)"/>
```

```
      <seealso

cref="M:Graphics.Point.op_Inequality(Graphics.Point,Graphics.Point)"/>
    </member>
    <member

name="M:Graphics.Point.op_Inequality(Graphics.Point,Graphics.Point)">
      <summary>
        This operator determines whether two Points have the same location.
      </summary>
      <param name="p1">The first Point to be compared.</param>
      <param name="p2">The second Point to be compared.</param>
      <returns>
        True if the Points do not have the same location and the
        exact same type; otherwise, false.
      </returns>
      <seealso cref="M:Graphics.Point.Equals(System.Object)"/>
      <seealso
        cref="M:Graphics.Point.op_Equality(Graphics.Point,Graphics.Point)"/>
      </member>
      <member name="M:Graphics.Point.Main">
        <summary>
          This is the entry point of the Point class testing program.
          <para>
            This program tests each method and operator, and
            is intended to be run after any non-trivial maintenance has
            been performed on the Point class.
          </para>
        </summary>
      </member>
      <member name="P:Graphics.Point.X">
        <value>
          Property <c>X</c> represents the point's x-coordinate.
        </value>
      </member>
      <member name="P:Graphics.Point.Y">
        <value>
          Property <c>Y</c> represents the point's y-coordinate.
        </value>
    </member>
  </members>
</doc>
```

End of informative text.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to

can also create and review issues and pull requests. For more information, see our contributor guide.

provide feedback:

- Open a documentation issue
- Provide product feedback

Annex E Bibliography

Article • 02/07/2024

This annex is informative.

ANSI X3.274-1996, *Programming Language REXX*. (This document is useful in understanding floating-point decimal arithmetic rules.)

ISO/IEC 9075-1, *Information technology — Database languages — SQL — Part 1: Framework (SQL/Framework)*

ISO/IEC 9899, *Programming languages — C*.

ISO/IEC 14882 *Programming languages — C++*

ISO 80000-1, Quantities and units — Part 1: General. (This document defines "banker's rounding.")

End of informative text.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

ECMA C# draft specification feedback

ECMA C# draft specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Feature specifications

Article • 05/08/2024

Because the C# standard specification has fallen behind the latest C# implementation, this section contains the Microsoft specifications for those newer features that haven't yet been incorporated into the standard. You can read these specifications to get information on newer features.

The feature specifications began as proposals for the design. They include proposed changes to the standard. The C# language design team and compiler team produce these feature specifications. The purpose of the proposals is to guide the design and implementation of the feature. They might include proposed features that haven't yet been implemented. The actual implementation might have modified behavior. Those changes are captured in the language design meeting (LDM) notes. The LDM notes are the minutes of the language design meetings. In most cases, the pertinent LDM notes are linked from the feature specifications.

As the committee works on newer versions, the feature specifications are removed from this site, and those links are redirected to the updated sections of the standard. In the meantime, the feature specifications represent the best information on those features.

See also

- C# standard specification

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

.NET feedback

.NET is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Enhanced #line directives

Article • 01/24/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

The compiler applies the mapping defined by `#line` directives to diagnostic locations and sequence points emitted to the PDB.

Currently only the line number and file path can be mapped while the starting character is inferred from the source code. The proposal is to allow specifying full span mapping.

Motivation

DSLs that generate C# source code (such as ASP.NET Razor) can't currently produce precise source mapping using `#line` directives. This results in degraded debugging experience in some cases as the sequence points emitted to the PDB can't map to the precise location in the original source code.

For example, the following Razor code

```
@page "/"
Time: @DateTime.Now
```

generates code like so (simplified):

```
C#
```

```
#line hidden
void Render()
{
   _builder.Add("Time:");
#line 2 "page.razor"
   _builder.Add(DateTime.Now);
#line hidden
}
```

The above directive would map the sequence point emitted by the compiler for the `_builder.Add(DateTime.Now);` statement to the line 2, but the column would be off (16 instead of 7).

The Razor source generator actually incorrectly generates the following code:

C#

```
#line hidden
void Render()
{
   _builder.Add("Time:");
   _builder.Add(
#line 2 "page.razor"
      DateTime.Now
#line hidden
);
}
```

The intent was to preserve the starting character and it works for diagnostic location mapping. However, this does not work for sequence points since `#line` directive only applies to the sequence points that follow it. There is no sequence point in the middle of the `_builder.Add(DateTime.Now);` statement (sequence points can only be emitted at IL instructions with empty evaluation stack). The `#line 2` directive in above code thus has no effect on the generated PDB and the debugger won't place a breakpoint or stop on the `@DateTime.Now` snippet in the Razor page.

Issues addressed by this proposal: https://github.com/dotnet/roslyn/issues/43432 https://github.com/dotnet/roslyn/issues/46526

Detailed design

We amend the syntax of `line_indicator` used in `pp_line` directive like so:

Current:

```
line_indicator
    : decimal_digit+ whitespace file_name
    | decimal_digit+
    | 'default'
    | 'hidden'
    ;
```

Proposed:

```
line_indicator
    : '(' decimal_digit+ ',' decimal_digit+ ')' '-' '(' decimal_digit+ ','
decimal_digit+ ')' whitespace decimal_digit+ whitespace file_name
    | '(' decimal_digit+ ',' decimal_digit+ ')' '-' '(' decimal_digit+ ','
decimal_digit+ ')' whitespace file_name
    | decimal_digit+ whitespace file_name
    | decimal_digit+
    | 'default'
    | 'hidden'
    ;
```

That is, the `#line` directive would accept either 5 decimal numbers (*start line*, *start character*, *end line*, *end character*, *character offset*), 4 decimal numbers (*start line*, *start character*, *end line*, *end character*), or a single one (*line*).

If *character offset* is not specified its default value is 0, otherwise it specifies the number of UTF-16 characters. The number must be non-negative and less then length of the line following the #line directive in the unmapped file.

(*start line*, *start character*)-(*end line*, *end character*) specifies a span in the mapped file. *start line* and *end line* are positive integers that specify line numbers. *start character*, *end character* are positive integers that specify UTF-16 character numbers. *start line*, *start character*, *end line*, *end character* are 1-based, meaning that the first line of the file and the first UTF-16 character on each line is assigned number 1.

> The implementation would constraint these numbers so that they specify a valid sequence point source span:
>
> - *start line* - 1 is within range [0, 0x20000000) and not equal to 0xfeefee.
> - *end line* - 1 is within range [0, 0x20000000) and not equal to 0xfeefee.
> - *start character* - 1 is within range [0, 0x10000)
> - *end character* - 1 is within range [0, 0x10000)
> - *end line* is greater or equal to *start line*.

> - *start line* is equal to *end line* then *end character* is greater than *start character*.

> Note that the numbers specified in the directive syntax are 1-based numbers but the actual spans in the PDB are zero-based. Hence the -1 adjustments above.

The mapped spans of sequence points and the locations of diagnostics that `#line` directive applies to are calculated as follows.

Let *d* be the zero-based number of the unmapped line containing the `#line` directive. Let span L = (start: (*start line* - 1, *start character* - 1), end: (*end line* - 1, *end character* - 1)) be zero-based span specified by the directive.

Function M that maps a position (line, character) within the scope of the `#line` directive in the source file containing the #line directive to a mapped position (mapped line, mapped character) is defined as follows:

M(*l*, *c*) =

l == *d* + 1 => (*L.start.line* + *l* – *d* – 1, *L.start.character* + max(*c* – *character offset*, 0))

l > *d* + 1 => (*L.start.line* + *l* – *d* – 1, *c*)

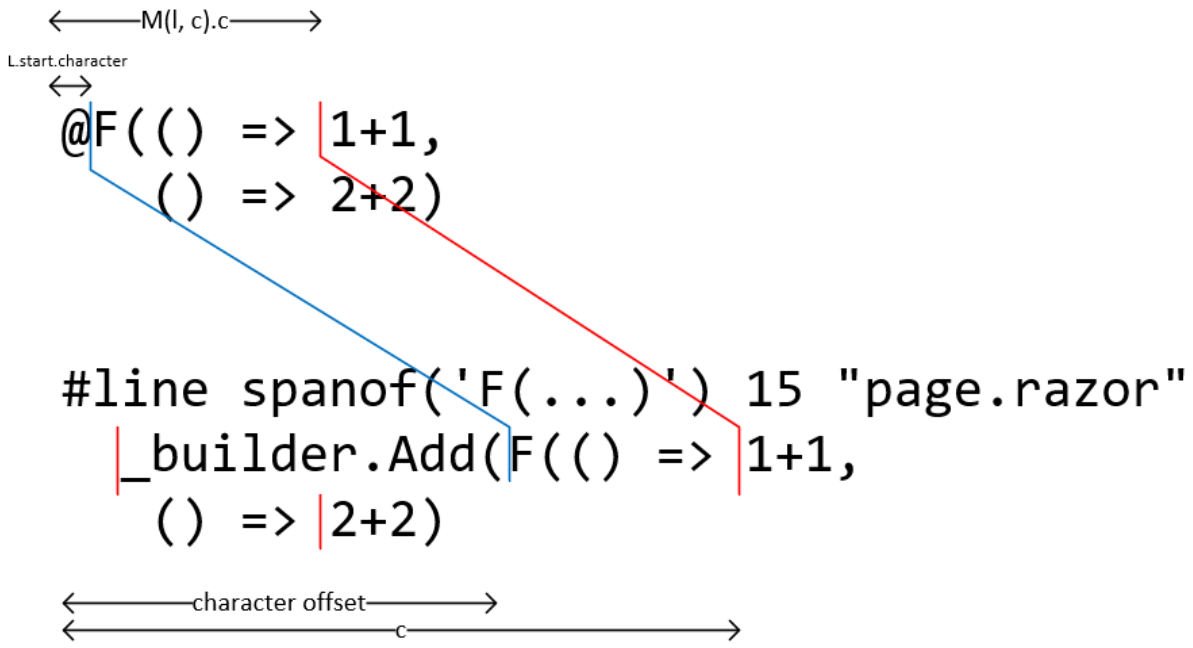
The syntax constructs that sequence points are associated with are determined by the compiler implementation and not covered by this specification. The compiler also decides for each sequence point its unmapped span. This span may partially or fully cover the associated syntax construct.

Once the unmapped spans are determined by the compiler the function *M* defined above is applied to their starting and ending positions, with the exception of the ending position of all sequence points within the scope of the #line directive whose unmapped location is at line *d* + 1 and character less than character offset. The end position of all these sequence points is *L.end*.

> Example [5.i] demonstrates why it is necessary to provide the ability to specify the end position of the first sequence point span.

> The above definition allows the generator of the unmapped source code to avoid intimate knowledge of which exact source constructs of the C# language produce sequence points. The mapped spans of the sequence points in the scope of the `#line` directive are derived from the relative position of the corresponding unmapped spans to the first unmapped span.

> Specifying the *character offset* allows the generator to insert any single-line prefix on the first line. This prefix is generated code that is not present in the mapped file. Such inserted prefix affects the value of the first unmapped sequence point span. Therefore the starting character of subsequent sequence point spans need to be offset by the length of the prefix (*character offset*). See example [2].



Examples

For clarity the examples use `spanof('...')` and `lineof('...')` pseudo-syntax to express the mapped span start position and line number, respectively, of the specified code snippet.

1. First and subsequent spans

Consider the following code with unmapped zero-based line numbers listed on the right:

```
#line (1,10)-(1,15) "a" // 3
  A();B(                // 4
);C();                  // 5
    D();                // 6
```

d = 3 L = (0, 9)..(0, 14)

There are 4 sequence point spans the directive applies to with following unmapped and mapped spans: (4, 2)..(4, 5) => (0, 9)..(0, 14) (4, 6)..(5, 1) => (0, 15)..(1, 1) (5, 2)..(5, 5) => (1, 2)..(1, 5) (6, 4)..(6, 7) => (2, 4)..(2, 7)

2. Character offset

Razor generates `_builder.Add(` prefix of length 15 (including two leading spaces).

Razor:

```
@page "/"
@F(() => 1+1,
   () => 2+2
)
```

Generated C#:

```C#
#line hidden
void Render()
{
#line spanof('F(...)') 15 "page.razor"  // 4
  _builder.Add(F(() => 1+1,            // 5
  () => 2+2                            // 6
));                                    // 7
#line hidden
}
);
}
```

d = 4 L = (1, 1)..(3,0) *character offset* = 15

Spans:

- `_builder.Add(F(…));` => `F(…)`: (5, 2)..(7, 2) => (1, 1)..(3, 0)
- `1+1` => `1+1`: (5, 23)..(5, 25) => (1, 9)..(1, 11)
- `2+2` => `2+2`: (6, 7)..(6, 9) => (2, 7)..(2, 9)

3. Razor: Single-line span

Razor:

```
@page "/"
Time: @DateTime.Now
```

Generated C#:

```C#
#line hidden
void Render()
{
  _builder.Add("Time:");
#line spanof('DateTime.Now') 15 "page.razor"
  _builder.Add(DateTime.Now);
#line hidden
);
}
```

4. Razor: Multi-line span

Razor:

```
@page "/"
@JsonToHtml(@"
{
  ""key1"": "value1",
  ""key2"": "value2"
}")
```

Generated C#:

```C#
#line hidden
void Render()
{
  _builder.Add("Time:");
#line spanof('JsonToHtml(@"...")') 15 "page.razor"
  _builder.Add(JsonToHtml(@"
{
  ""key1"": "value1",
  ""key2"": "value2"
}"));
#line hidden
}
);
}
```

5. Razor: block constructs

i. block containing expressions

In this example, the mapped span of the first sequence point that is associated with the IL instruction that is emitted for the `_builder.Add(Html.Helper(() =>` statement needs to cover the whole expression of `Html.Helper(...)` in the generated file `a.razor`. This is achieved by application of rule [1] to the end position of the sequence point.

```

```

```
@Html.Helper(() =>
{
    <p>Hello World</p>
    @DateTime.Now
})
```

C#

```
#line spanof('Html.Helper(() => { ... })') 13 "a.razor"
_builder.Add(Html.Helper(() =>
#line lineof('{') "a.razor"
{
#line spanof('DateTime.Now') 13 "a.razor"
_builder.Add(DateTime.Now);
#line lineof('}') "a.razor"
}
#line hidden
)
```

ii. block containing statements

Uses existing `#line line file` form since

a) Razor does not add any prefix, b) `{` is not present in the generated file and there can't be a sequence point placed on it, therefore the span of the first unmapped sequence point is unknown to Razor.

The starting character of `Console` in the generated file must be aligned with the Razor file.

```

```

```
@{Console.WriteLine(1);Console.WriteLine(2);}
```

C#

```
#line lineof('@{') "a.razor"
  Console.WriteLine(1);Console.WriteLine(2);
#line hidden
```

iii. block containing top-level code (@code, @functions)

Uses existing `#line line file` form since

a) Razor does not add any prefix, b) `{` is not present in the generated file and there can't be a sequence point placed on it, therefore the span of the first unmapped sequence point is unknown to Razor.

The starting character of `[Parameter]` in the generated file must be aligned with the Razor file.

```
@code {
    [Parameter]
    public int IncrementAmount { get; set; }
}
```

C#

```
#line lineof('[') "a.razor"
    [Parameter]
    public int IncrementAmount { get; set; }
#line hidden
```

6. Razor: `@for`, `@foreach`, `@while`, `@do`, `@if`, `@switch`, `@using`, `@try`, `@lock`

Uses existing `#line line file` form since a) Razor does not add any prefix. b) the span of the first unmapped sequence point may not be known to Razor (or shouldn't need to know).

The starting character of the keyword in the generated file must be aligned with the Razor file.

```
@for (var i = 0; i < 10; i++)
{
}
@if (condition)
{
}
else
{
}
```

C#

```
#line lineof('for') "a.razor"
 for (var i = 0; i < 10; i++)
{
}
#line lineof('if') "a.razor"
 if (condition)
{
}
else
{
}
#line hidden
```

Top-level statements

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow a sequence of *statements* to occur right before the *namespace_member_declaration*s of a *compilation_unit* (i.e. source file).

The semantics are that if such a sequence of *statements* is present, the following type declaration, modulo the actual method name, would be emitted:

```c#
partial class Program
{
    static async Task Main(string[] args)
    {
        // statements
    }
}
```

See also https://github.com/dotnet/csharplang/issues/3117.

Motivation

There's a certain amount of boilerplate surrounding even the simplest of programs, because of the need for an explicit `Main` method. This seems to get in the way of language learning and program clarity. The primary goal of the feature therefore is to allow C# programs without unnecessary boilerplate around them, for the sake of learners and the clarity of code.

Detailed design

Syntax

The only additional syntax is allowing a sequence of *statements* in a compilation unit, just before the *namespace_member_declarations*:

antlr

```
compilation_unit
    : extern_alias_directive* using_directive* global_attributes? statement*
namespace_member_declaration*
    ;
```

Only one *compilation_unit* is allowed to have *statements*.

Example:

c#

```
if (args.Length == 0
    || !int.TryParse(args[0], out int n)
    || n < 0) return;
Console.WriteLine(Fib(n).curr);

(int curr, int prev) Fib(int i)
{
    if (i == 0) return (1, 0);
    var (curr, prev) = Fib(i - 1);
    return (curr + prev, curr);
}
```

Semantics

If any top-level statements are present in any compilation unit of the program, the meaning is as if they were combined in the block body of a `Main` method of a `Program` class in the global namespace, as follows:

c#

```
partial class Program
{
    static async Task Main(string[] args)
    {
        // statements
```

```
    }
}
```

The type is named "Program", so can be referenced by name from source code. It is a partial type, so a type named "Program" in source code must also be declared as partial. But the method name "Main" is used only for illustration purposes, the actual name used by the compiler is implementation dependent and the method cannot be referenced by name from source code.

The method is designated as the entry point of the program. Explicitly declared methods that by convention could be considered as an entry point candidates are ignored. A warning is reported when that happens. It is an error to specify `-main:<type>` compiler switch when there are top-level statements.

The entry point method always has one formal parameter, `string[] args`. The execution environment creates and passes a `string[]` argument containing the command-line arguments that were specified when the application was started. The `string[]` argument is never null, but it may have a length of zero if no command-line arguments were specified. The 'args' parameter is in scope within top-level statements and is not in scope outside of them. Regular name conflict/shadowing rules apply.

Async operations are allowed in top-level statements to the degree they are allowed in statements within a regular async entry point method. However, they are not required, if `await` expressions and other async operations are omitted, no warning is produced.

The signature of the generated entry point method is determined based on operations used by the top level statements as follows:

Async-operations\Return-with-expression	Present	Absent
Present	`static Task<int> Main(string[] args)`	`static Task Main(string[] args)`
Absent	`static int Main(string[] args)`	`static void Main(string[] args)`

The example above would yield the following `$Main` method declaration:

```c#
partial class Program
{
    static void $Main(string[] args)
    {
```

```
        if (args.Length == 0
            || !int.TryParse(args[0], out int n)
            || n < 0) return;
        Console.WriteLine(Fib(n).curr);

        (int curr, int prev) Fib(int i)
        {
            if (i == 0) return (1, 0);
            var (curr, prev) = Fib(i - 1);
            return (curr + prev, curr);
        }
    }
}
```

At the same time an example like this:

```c#
await System.Threading.Tasks.Task.Delay(1000);
System.Console.WriteLine("Hi!");
```

would yield:

```c#
partial class Program
{
    static async Task $Main(string[] args)
    {
        await System.Threading.Tasks.Task.Delay(1000);
        System.Console.WriteLine("Hi!");
    }
}
```

An example like this:

```c#
await System.Threading.Tasks.Task.Delay(1000);
System.Console.WriteLine("Hi!");
return 0;
```

would yield:

```c#
partial class Program
{
    static async Task<int> $Main(string[] args)
```

```
    {
        await System.Threading.Tasks.Task.Delay(1000);
        System.Console.WriteLine("Hi!");
        return 0;
    }
}
```

And an example like this:

```c#
System.Console.WriteLine("Hi!");
return 2;
```

would yield:

```c#
partial class Program
{
    static int $Main(string[] args)
    {
        System.Console.WriteLine("Hi!");
        return 2;
    }
}
```

Scope of top-level local variables and local functions

Even though top-level local variables and functions are "wrapped" into the generated entry point method, they should still be in scope throughout the program in every compilation unit. For the purpose of simple-name evaluation, once the global namespace is reached:

- First, an attempt is made to evaluate the name within the generated entry point method and only if this attempt fails
- The "regular" evaluation within the global namespace declaration is performed.

This could lead to name shadowing of namespaces and types declared within the global namespace as well as to shadowing of imported names.

If the simple name evaluation occurs outside of the top-level statements and the evaluation yields a top-level local variable or function, that should lead to an error.

In this way we protect our future ability to better address "Top-level functions" (scenario 2 in https://github.com/dotnet/csharplang/issues/3117), and are able to give useful

diagnostics to users who mistakenly believe them to be supported.

Module Initializers

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Although the .NET platform has a feature that directly supports writing initialization code for the assembly (technically, the module), it is not exposed in C#. This is a rather niche scenario, but once you run into it the solutions appear to be pretty painful. There are reports of a number of customers (inside and outside Microsoft) struggling with the problem, and there are no doubt more undocumented cases.

Motivation

- Enable libraries to do eager, one-time initialization when loaded, with minimal overhead and without the user needing to explicitly call anything
- One particular pain point of current `static` constructor approaches is that the runtime must do additional checks on usage of a type with a static constructor, in order to decide whether the static constructor needs to be run or not. This adds measurable overhead.
- Enable source generators to run some global initialization logic without the user needing to explicitly call anything

Detailed design

A method can be designated as a module initializer by decorating it with a `[ModuleInitializer]` attribute.

```
C#
```

```
using System;
namespace System.Runtime.CompilerServices
{
    [AttributeUsage(AttributeTargets.Method, AllowMultiple = false)]
    public sealed class ModuleInitializerAttribute : Attribute { }
}
```

The attribute can be used like this:

C#

```
using System.Runtime.CompilerServices;
class C
{
    [ModuleInitializer]
    internal static void M1()
    {
        // ...
    }
}
```

Some requirements are imposed on the method targeted with this attribute:

1. The method must be `static`.
2. The method must be parameterless.
3. The method must return `void`.
4. The method must not be generic or be contained in a generic type.
5. The method must be accessible from the containing module.
    - This means the method's effective accessibility must be `internal` or `public`.
    - This also means the method cannot be a local function.

When one or more valid methods with this attribute are found in a compilation, the compiler will emit a module initializer which calls each of the attributed methods. The calls will be emitted in a reserved, but deterministic order.

Drawbacks

Why should we *not* do this?

- Perhaps the existing third-party tooling for "injecting" module initializers is sufficient for users who have been asking for this feature.

Design meetings

April 8th, 2020

Suppress emitting of `localsinit` flag.

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow suppressing emit of `localsinit` flag via `SkipLocalsInitAttribute` attribute.

Motivation

Background

Per CLR spec local variables that do not contain references are not initialized to a particular value by the VM/JIT. Reading from such variables without initialization is type-safe, but otherwise the behavior is undefined and implementation specific. Typically uninitialized locals contain whatever values were left in the memory that is now occupied by the stack frame. That could lead to nondeterministic behavior and hard to reproduce bugs.

There are two ways to "assign" a local variable:

- by storing a value or
- by specifying `localsinit` flag which forces everything that is allocated form the local memory pool to be zero-initialized NOTE: this includes both local variables and `stackalloc` data.

Use of uninitialized data is discouraged and is not allowed in verifiable code. While it might be possible to prove that by the means of flow analysis, it is permitted for the verification algorithm to be conservative and simply require that `localsinit` is set.

Historically C# compiler emits `localsinit` flag on all methods that declare locals.

While C# employs definite-assignment analysis which is more strict than what CLR spec would require (C# also needs to consider scoping of locals), it is not strictly guaranteed that the resulting code would be formally verifiable:

- CLR and C# rules may not agree on whether passing a local as `out` argument is a `use`.
- CLR and C# rules may not agree on treatment of conditional branches when conditions are known (constant propagation).
- CLR could as well simply require `localinits`, since that is permitted.

Problem

In high-performance application the cost of forced zero-initialization could be noticeable. It is particularly noticeable when `stackalloc` is used.

In some cases JIT can elide initial zero-initialization of individual locals when such initialization is "killed" by subsequent assignments. Not all JITs do this and such optimization has limits. It does not help with `stackalloc`.

To illustrate that the problem is real - there is a known bug where a method not containing any `IL` locals would not have `localsinit` flag. The bug is already being exploited by users by putting `stackalloc` into such methods - intentionally to avoid initialization costs. That is despite the fact that absence of `IL` locals is an unstable metric and may vary depending on changes in codegen strategy. The bug should be fixed and users should get a more documented and reliable way of suppressing the flag.

Detailed design

Allow specifying `System.Runtime.CompilerServices.SkipLocalsInitAttribute` as a way to tell the compiler to not emit `localsinit` flag.

The end result of this will be that the locals may not be zero-initialized by the JIT, which is in most cases unobservable in C#.
In addition to that `stackalloc` data will not be zero-initialized. That is definitely observable, but also is the most motivating scenario.

Permitted and recognized attribute targets are: `Method`, `Property`, `Module`, `Class`, `Struct`, `Interface`, `Constructor`. However compiler will not require that attribute is defined with the listed targets nor it will care in which assembly the attribute is defined.

When attribute is specified on a container (`class`, `module`, containing method for a nested method, ...), the flag affects all methods contained within the container.

Synthesized methods "inherit" the flag from the logical container/owner.

The flag affects only codegen strategy for actual method bodies. I.E. the flag has no effect on abstract methods and is not propagated to overriding/implementing methods.

This is explicitly a ***compiler feature*** and ***not a language feature***.
Similarly to compiler command line switches the feature controls implementation details of a particular codegen strategy and does not need to be required by the C# spec.

Drawbacks

- Old/other compilers may not honor the attribute. Ignoring the attribute is compatible behavior. Only may result in a slight perf hit.
- The code without `localinits` flag may trigger verification failures. Users that ask for this feature are generally unconcerned with verifiability.
- Applying the attribute at higher levels than an individual method has nonlocal effect, which is observable when `stackalloc` is used. Yet, this is the most requested scenario.

Alternatives

- omit `localinits` flag when method is declared in `unsafe` context. That could cause silent and dangerous behavior change from deterministic to nondeterministic in a case of `stackalloc`.
- omit `localinits` flag always. Even worse than above.
- omit `localinits` flag unless `stackalloc` is used in the method body. Does not address the most requested scenario and may turn code unverifiable with no option to revert that back.

Unresolved questions

- Should the attribute be actually emitted to metadata?

Design meetings

None yet.

Nullable reference types in C#

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

The goal of this feature is to:

- Allow developers to express whether a variable, parameter or result of a reference type is intended to be null or not.
- Provide warnings when such variables, parameters and results are not used according to that intent.

Expression of intent

The language already contains the `T?` syntax for value types. It is straightforward to extend this syntax to reference types.

It is assumed that the intent of an unadorned reference type `T` is for it to be non-null.

Checking of nullable references

A flow analysis tracks nullable reference variables. Where the analysis deems that they would not be null (e.g. after a check or an assignment), their value will be considered a non-null reference.

A nullable reference can also explicitly be treated as non-null with the postfix `x!` operator (the "damnit" operator), for when flow analysis cannot establish a non-null situation that the developer knows is there.

Otherwise, a warning is given if a nullable reference is dereferenced, or is converted to a non-null type.

A warning is given when converting from `S[]` to `T?[]` and from `S?[]` to `T[]`.

A warning is given when converting from `C<S>` to `C<T?>` except when the type parameter is covariant (`out`), and when converting from `C<S?>` to `C<T>` except when the type parameter is contravariant (`in`).

A warning is given on `C<T?>` if the type parameter has non-null constraints.

Checking of non-null references

A warning is given if a null literal is assigned to a non-null variable or passed as a non-null parameter.

A warning is also given if a constructor does not explicitly initialize non-null reference fields.

We cannot adequately track that all elements of an array of non-null references are initialized. However, we could issue a warning if no element of a newly created array is assigned to before the array is read from or passed on. That might handle the common case without being too noisy.

We need to decide whether `default(T)` generates a warning, or is simply treated as being of the type `T?`.

Metadata representation

Nullability adornments should be represented in metadata as attributes. This means that downlevel compilers will ignore them.

We need to decide if only nullable annotations are included, or there's also some indication of whether non-null was "on" in the assembly.

Generics

If a type parameter `T` has non-nullable constraints, it is treated as non-nullable within its scope.

If a type parameter is unconstrained or has only nullable constraints, the situation is a little more complex: this means that the corresponding type argument could be *either* nullable or non-nullable. The safe thing to do in that situation is to treat the type parameter as *both* nullable and non-nullable, giving warnings when either is violated.

It is worth considering whether explicit nullable reference constraints should be allowed. Note, however, that we cannot avoid having nullable reference types *implicitly* be constraints in certain cases (inherited constraints).

The `class` constraint is non-null. We can consider whether `class?` should be a valid nullable constraint denoting "nullable reference type".

Type inference

In type inference, if a contributing type is a nullable reference type, the resulting type should be nullable. In other words, nullness is propagated.

We should consider whether the `null` literal as a participating expression should contribute nullness. It doesn't today: for value types it leads to an error, whereas for reference types the null successfully converts to the plain type.

C#

```
string? n = "world";
var x = b ? "Hello" : n; // string?
var y = b ? "Hello" : null; // string? or error
var z = b ? 7 : null; // Error today, could be int?
```

Null guard guidance

As a feature, nullable reference types allow developers to express their intent, and provide warnings through flow analysis if that intent is contradicted. There is a common question as to whether or not null guards are necessary.

Example of null guard

C#

```
public void DoWork(Worker worker)
{
    // Guard against worker being null
    if (worker is null)
    {
        throw new ArgumentNullException(nameof(worker));
    }

    // Otherwise use worker argument
}
```

In the previous example, the `DoWork` function accepts a `Worker` and guards against it potentially being `null`. If the `worker` argument is `null`, the `DoWork` function will `throw`. With nullable reference types, the code in the previous example makes the intent that the `Worker` parameter would *not* be `null`. If the `DoWork` function was a public API, such as a NuGet package or a shared library - as guidance you should leave null guards in place. As a public API, the only guarantee that a caller isn't passing `null` is to guard against it.

Express intent

A more compelling use of the previous example is to express that the `Worker` parameter could be `null`, thus making the null guard more appropriate. If you remove the null guard in the following example, the compiler warns that you may be dereferencing null. Regardless, both null guards are still valid.

C#

```
public void DoWork(Worker? worker)
{
    // Guard against worker being null
    if (worker is null)
    {
        throw new ArgumentNullException(nameof(worker));
    }

    // Otherwise use worker argument
}
```

For non-public APIs, such as source code entirely in control by a developer or dev team - the nullable reference types could allow for the safe removal of null guards where the developers can guarantee it is not necessary. The feature can help with warnings, but it cannot guarantee that at runtime code execution could result in a `NullReferenceException`.

Breaking changes

Non-null warnings are an obvious breaking change on existing code, and should be accompanied with an opt-in mechanism.

Less obviously, warnings from nullable types (as described above) are a breaking change on existing code in certain scenarios where the nullability is implicit:

- Unconstrained type parameters will be treated as implicitly nullable, so assigning them to `object` or accessing e.g. `ToString` will yield warnings.
- if type inference infers nullness from `null` expressions, then existing code will sometimes yield nullable rather than non-nullable types, which can lead to new warnings.

So nullable warnings also need to be optional

Finally, adding annotations to an existing API will be a breaking change to users who have opted in to warnings, when they upgrade the library. This, too, merits the ability to opt in or out. "I want the bug fixes, but I am not ready to deal with their new annotations"

In summary, you need to be able to opt in/out of:

- Nullable warnings
- Non-null warnings
- Warnings from annotations in other files

The granularity of the opt-in suggests an analyzer-like model, where swaths of code can opt in and out with pragmas and severity levels can be chosen by the user. Additionally, per-library options ("ignore the annotations from JSON.NET until I'm ready to deal with the fall out") may be expressible in code as attributes.

The design of the opt-in/transition experience is crucial to the success and usefulness of this feature. We need to make sure that:

- Users can adopt nullability checking gradually as they want to
- Library authors can add nullability annotations without fear of breaking customers
- Despite these, there is not a sense of "configuration nightmare"

Tweaks

We could consider not using the `?` annotations on locals, but just observing whether they are used in accordance with what gets assigned to them. I don't favor this; I think we should uniformly let people express their intent.

We could consider a shorthand `T! x` on parameters, that auto-generates a runtime null check.

Certain patterns on generic types, such as `FirstOrDefault` or `TryGet`, have slightly weird behavior with non-nullable type arguments, because they explicitly yield default values in certain situations. We could try to nuance the type system to accommodate these

better. For instance, we could allow `?` on unconstrained type parameters, even though the type argument could already be nullable. I doubt that it is worth it, and it leads to weirdness related to interaction with nullable *value* types.

Nullable value types

We could consider adopting some of the above semantics for nullable value types as well.

We already mentioned type inference, where we could infer `int?` from `(7, null)`, instead of just giving an error.

Another opportunity is to apply the flow analysis to nullable value types. When they are deemed non-null, we could actually allow using as the non-nullable type in certain ways (e.g. member access). We just have to be careful that the things that you can *already* do on a nullable value type will be preferred, for back compat reasons.

Nullable Reference Types Specification

Article • 10/20/2021

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

This is a work in progress - several parts are missing or incomplete.

This feature adds two new kinds of nullable types (nullable reference types and nullable generic types) to the existing nullable value types, and introduces a static flow analysis for purpose of null-safety.

Syntax

Nullable reference types and nullable type parameters

Nullable reference types and nullable type parameters have the same syntax `T?` as the short form of nullable value types, but do not have a corresponding long form.

For the purposes of the specification, the current `nullable_type` production is renamed to `nullable_value_type`, and `nullable_reference_type` and `nullable_type_parameter` productions are added:

antlr

```
type
    : value_type
    | reference_type
    | nullable_type_parameter
    | type_parameter
    | type_unsafe
    ;

reference_type
    : ...
```

```
    | nullable_reference_type
    ;

nullable_reference_type
    : non_nullable_reference_type '?'
    ;

non_nullable_reference_type
    : reference_type
    ;

nullable_type_parameter
    : non_nullable_non_value_type_parameter '?'
    ;

non_nullable_non_value_type_parameter
    : type_parameter
    ;
```

The `non_nullable_reference_type` in a `nullable_reference_type` must be a nonnullable reference type (class, interface, delegate or array).

The `non_nullable_non_value_type_parameter` in `nullable_type_parameter` must be a type parameter that isn't constrained to be a value type.

Nullable reference types and nullable type parameters cannot occur in the following positions:

- as a base class or interface
- as the receiver of a `member_access`
- as the `type` in an `object_creation_expression`
- as the `delegate_type` in a `delegate_creation_expression`
- as the `type` in an `is_expression`, a `catch_clause` or a `type_pattern`
- as the `interface` in a fully qualified interface member name

A warning is given on a `nullable_reference_type` and `nullable_type_parameter` in a *disabled* nullable annotation context.

`class` and `class?` constraint

The `class` constraint has a nullable counterpart `class?`:

antlr

```
primary_constraint
    : ...
```

```
    | 'class' '?'
    ;
```

A type parameter constrained with `class` (in an *enabled* annotation context) must be instantiated with a nonnullable reference type.

A type parameter constrained with `class?` (or `class` in a *disabled* annotation context) may either be instantiated with a nullable or nonnullable reference type.

A warning is given on a `class?` constraint in a *disabled* annotation context.

The behavior of `?` annotations on type parameters that aren't constrained to be either a `struct` or `class` is covered in Unconstrained type parameter annotations.

`notnull` constraint

A type parameter constrained with `notnull` may not be a nullable type (nullable value type, nullable reference type or nullable type parameter).

antlr

```
primary_constraint
    : ...
    | 'notnull'
    ;
```

`default` constraint

The `default` constraint can be used on a method override or explicit implementation to disambiguate `T?` meaning "nullable type parameter" from "nullable value type" (`Nullable<T>`). Lacking the `default` constraint a `T?` syntax in an override or explicit implementation will be interpreted as `Nullable<T>`

See Unconstrained type parameter annotations.

The null-forgiving operator

The post-fix `!` operator is called the null-forgiving operator. It can be applied on a *primary_expression* or within a *null_conditional_expression*:

antlr

```
primary_expression
    : ...
    | null_forgiving_expression
    ;

null_forgiving_expression
    : primary_expression '!'
    ;

null_conditional_expression
    : primary_expression null_conditional_operations_no_suppression
suppression?
    ;

null_conditional_operations_no_suppression
    : null_conditional_operations? '?' '.' identifier type_argument_list?
    | null_conditional_operations? '?' '[' argument_list ']'
    | null_conditional_operations '.' identifier type_argument_list?
    | null_conditional_operations '[' argument_list ']'
    | null_conditional_operations '(' argument_list? ')'
    ;

null_conditional_operations
    : null_conditional_operations_no_suppression suppression?
    ;

suppression
    : '!'
    ;
```

For example:

C#

```
var v = expr!;
expr!.M();
_ = a?.b!.c;
```

The postfix `!` operator has no runtime effect - it evaluates to the result of the underlying expression. Its only role is to change the null state of the expression to "not null", and to limit warnings given on its use.

Nullable compiler directives

`#nullable` directives control the nullable annotation and warning contexts.

antlr

```
pp_directive
    : ...
    | pp_nullable
    ;

pp_nullable
    : whitespace? '#' whitespace? 'nullable' whitespace nullable_action
(whitespace nullable_target)? pp_new_line
    ;

nullable_action
    : 'disable'
    | 'enable'
    | 'restore'
    ;

nullable_target
    : 'warnings'
    | 'annotations'
    ;
```

`#pragma warning` directives are expanded to allow changing the nullable warning context:

antlr

```
pragma_warning_body
    : ...
    | 'warning' whitespace warning_action whitespace 'nullable'
    ;
```

For example:

C#

```
#pragma warning disable nullable
```

Nullable contexts

Every line of source code has a *nullable annotation context* and a *nullable warning context*. These control whether nullable annotations have effect, and whether nullability warnings are given. The annotation context of a given line is either *disabled* or *enabled*. The warning context of a given line is either *disabled* or *enabled*.

Both contexts can be specified at the project level (outside of C# source code), or anywhere within a source file via `#nullable` pre-processor directives. If no project level

settings are provided the default is for both contexts to be *disabled*.

The `#nullable` directive controls the annotation and warning contexts within the source text, and take precedence over the project-level settings.

A directive sets the context(s) it controls for subsequent lines of code, until another directive overrides it, or until the end of the source file.

The effect of the directives is as follows:

- `#nullable disable`: Sets the nullable annotation and warning contexts to *disabled*
- `#nullable enable`: Sets the nullable annotation and warning contexts to *enabled*
- `#nullable restore`: Restores the nullable annotation and warning contexts to project settings
- `#nullable disable annotations`: Sets the nullable annotation context to *disabled*
- `#nullable enable annotations`: Sets the nullable annotation context to *enabled*
- `#nullable restore annotations`: Restores the nullable annotation context to project settings
- `#nullable disable warnings`: Sets the nullable warning context to *disabled*
- `#nullable enable warnings`: Sets the nullable warning context to *enabled*
- `#nullable restore warnings`: Restores the nullable warning context to project settings

Nullability of types

A given type can have one of three nullabilities: *oblivious*, *nonnullable*, and *nullable*.

Nonnullable types may cause warnings if a potential `null` value is assigned to them. *Oblivious* and *nullable* types, however, are "*null-assignable*" and can have `null` values assigned to them without warnings.

Values of *oblivious* and *nonnullable* types can be dereferenced or assigned without warnings. Values of *nullable* types, however, are "*null-yielding*" and may cause warnings when dereferenced or assigned without proper null checking.

The *default null state* of a null-yielding type is "maybe null" or "maybe default". The default null state of a non-null-yielding type is "not null".

The kind of type and the nullable annotation context it occurs in determine its nullability:

- A nonnullable value type `S` is always *nonnullable*

- A nullable value type `S?` is always *nullable*
- An unannotated reference type `C` in a *disabled* annotation context is *oblivious*
- An unannotated reference type `C` in an *enabled* annotation context is *nonnullable*
- A nullable reference type `C?` is *nullable* (but a warning may be yielded in a *disabled* annotation context)

Type parameters additionally take their constraints into account:

- A type parameter `T` where all constraints (if any) are either nullable types or the `class?` constraint is *nullable*
- A type parameter `T` where at least one constraint is either *oblivious* or *nonnullable* or one of the `struct` or `class` or `notnull` constraints is
  - *oblivious* in a *disabled* annotation context
  - *nonnullable* in an *enabled* annotation context
- A nullable type parameter `T?` is *nullable*, but a warning is yielded in a *disabled* annotation context if `T` isn't a value type

Oblivious vs nonnullable

A `type` is deemed to occur in a given annotation context when the last token of the type is within that context.

Whether a given reference type `C` in source code is interpreted as oblivious or nonnullable depends on the annotation context of that source code. But once established, it is considered part of that type, and "travels with it" e.g. during substitution of generic type arguments. It is as if there is an annotation like `?` on the type, but invisible.

Constraints

Nullable reference types can be used as generic constraints.

`class?` is a new constraint denoting "possibly nullable reference type", whereas `class` in an *enabled* annotation context denotes "nonnullable reference type".

`default` is a new constraint denoting a type parameter that isn't known to be a reference or value type. It can only be used on overridden and explicitly implemented methods. With this constraint, `T?` means a nullable type parameter, as opposed to being a shorthand for `Nullable<T>`.

`notnull` is a new constraint denoting a type parameter that is nonnullable.

The nullability of a type argument or of a constraint does not impact whether the type satisfies the constraint, except where that is already the case today (nullable value types do not satisfy the `struct` constraint). However, if the type argument does not satisfy the nullability requirements of the constraint, a warning may be given.

Null state and null tracking

Every expression in a given source location has a *null state*, which indicated whether it is believed to potentially evaluate to null. The null state is either "not null", "maybe null", or "maybe default". The null state is used to determine whether a warning should be given about null-unsafe conversions and dereferences.

The distinction between "maybe null" and "maybe default" is subtle and applies to type parameters. The distinction is that a type parameter `T` which has the state "maybe null" means the value is in the domain of legal values for `T` however that legal value may include `null`. Where as a "maybe default" means that the value may be outside the legal domain of values for `T`.

Example:

```c#
// The value `t` here has the state "maybe null". It's possible for `T` to be instantiated
// with `string?` in which case `null` would be within the domain of legal values here. The
// assumption though is the value provided here is within the legal values of `T`. Hence
// if `T` is `string` then `null` will not be a value, just as we assume that `null` is not
// provided for a normal `string` parameter
void M<T>(T t)
{
    // There is no guarantee that default(T) is within the legal values for T hence the
    // state *must* be "maybe-default" and hence `local` must be `T?`
    T? local = default(T);
}
```

Null tracking for variables

For certain expressions denoting variables, fields or properties, the null state is tracked between occurrences, based on assignments to them, tests performed on them and the

control flow between them. This is similar to how definite assignment is tracked for variables. The tracked expressions are the ones of the following form:

```antlr
tracked_expression
    : simple_name
    | this
    | base
    | tracked_expression '.' identifier
    ;
```

Where the identifiers denote fields or properties.

The null state for tracked variables is "not null" in unreachable code. This follows other decisions around unreachable code like considering all locals to be definitely assigned.

Describe null state transitions similar to definite assignment

Null state for expressions

The null state of an expression is derived from its form and type, and from the null state of variables involved in it.

Literals

The null state of a `null` literal depends on the target type of the expression. If the target type is a type parameter constrained to a reference type then it's "maybe default". Otherwise it is "maybe null".

The null state of a `default` literal depends on the target type of the `default` literal. A `default` literal with target type `T` has the same null state as the `default(T)` expression.

The null state of any other literal is "not null".

Simple names

If a `simple_name` is not classified as a value, its null state is "not null". Otherwise it is a tracked expression, and its null state is its tracked null state at this source location.

Member access

If a `member_access` is not classified as a value, its null state is "not null". Otherwise, if it is a tracked expression, its null state is its tracked null state at this source location. Otherwise its null state is the default null state of its type.

c#

```
var person = new Person();

// The receiver is a tracked expression hence the member_access of the
property
// is tracked as well
if (person.FirstName is not null)
{
    Use(person.FirstName);
}

// The return of an invocation is not a tracked expression hence the
member_access
// of the return is also not tracked
if (GetAnonymous().FirstName is not null)
{
    // Warning: Cannot convert null literal to non-nullable reference type.
    Use(GetAnonymous().FirstName);
}

void Use(string s)
{
    // ...
}

public class Person
{
    public string? FirstName { get; set; }
    public string? LastName { get; set; }

    private static Person s_anonymous = new Person();
    public static Person GetAnonymous() => s_anonymous;
}
```

Invocation expressions

If an `invocation_expression` invokes a member that is declared with one or more attributes for special null behavior, the null state is determined by those attributes. Otherwise the null state of the expression is the default null state of its type.

The null state of an `invocation_expression` is not tracked by the compiler.

c#

```
// The result of an invocation_expression is not tracked
if (GetText() is not null)
{
    // Warning: Converting null literal or possible null value to non-
nullable type.
    string s = GetText();
    // Warning: Dereference of a possibly null reference.
    Use(s);
}

// Nullable friendly pattern
if (GetText() is string s)
{
    Use(s);
}

string? GetText() => ...
Use(string s) {  }
```

Element access

If an `element_access` invokes an indexer that is declared with one or more attributes for special null behavior, the null state is determined by those attributes. Otherwise the null state of the expression is the default null state of its type.

c#

```
object?[] array = ...;
if (array[0] != null)
{
    // Warning: Converting null literal or possible null value to non-
nullable type.
    object o = array[0];
    // Warning: Dereference of a possibly null reference.
    Console.WriteLine(o.ToString());
}

// Nullable friendly pattern
if (array[0] is {} o)
{
    Console.WriteLine(o.ToString());
}
```

Base access

If `B` denotes the base type of the enclosing type, `base.I` has the same null state as `((B)this).I` and `base[E]` has the same null state as `((B)this)[E]`.

Default expressions

`default(T)` has the null state based on the properties of the type `T`:

- If the type is a *nonnullable* type then it has the null state "not null"
- Else if the type is a type parameter then it has the null state "maybe default"
- Else it has the null state "maybe null"

Null-conditional expressions ?.

A `null_conditional_expression` has the null state based on the expression type. Note that this refers to the type of the `null_conditional_expression`, not the original type of the member being invoked:

- If the type is a *nullable* value type then it has the null state "maybe null"
- Else if the type is a *nullable* type parameter then it has the null state "maybe default"
- Else it has the null state "maybe null"

Cast expressions

If a cast expression `(T)E` invokes a user-defined conversion, then the null state of the expression is the default null state for the type of the user-defined conversion. Otherwise:

- If `T` is a *nonnullable* value type then `T` has the null state "not null"
- Else if `T` is a *nullable* value type then `T` has the null state "maybe null"
- Else if `T` is a *nullable* type in the form `U?` where `U` is a type parameter then `T` has the null state "maybe default"
- Else if `T` is a *nullable* type, and `E` has null state "maybe null" or "maybe default", then `T` has the null state "maybe null"
- Else if `T` is a type parameter, and `E` has null state "maybe null" or "maybe default", then `T` has the null state "maybe default"
- Else `T` has the same null state as `E`

Unary and binary operators

If a unary or binary operator invokes an user-defined operator then the null state of the expression is the default null state for the type of the user-defined operator. Otherwise it is the null state of the expression.

Something special to do for binary `+` over strings and delegates?

Await expressions

The null state of `await E` is the default null state of its type.

The `as` operator

The null state of an `E as T` expression depends first on properties of the type `T`. If the type of `T` is *nonnullable* then the null state is "not null". Otherwise the null state depends on the conversion from the type of `E` to type `T`:

- If the conversion is an identity, boxing, implicit reference, or implicit nullable conversion, then the null state is the null state of `E`
- Else if `T` is a type parameter then it has the null state "maybe default"
- Else it has the null state "maybe null"

The null-coalescing operator

The null state of `E1 ?? E2` is the null state of `E2`

The conditional operator

The null state of `E1 ? E2 : E3` is based on the null state of `E2` and `E3`:

- If both are "not null" then the null state is "not null"
- Else if either is "maybe default" then the null state is "maybe default"
- Else the null state is "not null"

Query expressions

The null state of a query expression is the default null state of its type.

Additional work needed here

Assignment operators

`E1 = E2` and `E1 op= E2` have the same null state as `E2` after any implicit conversions have been applied.

Expressions that propagate null state

`(E)`, `checked(E)` and `unchecked(E)` all have the same null state as `E`.

Expressions that are never null

The null state of the following expression forms is always "not null":

- `this` access
- interpolated strings
- `new` expressions (object, delegate, anonymous object and array creation expressions)
- `typeof` expressions
- `nameof` expressions
- anonymous functions (anonymous methods and lambda expressions)
- null-forgiving expressions
- `is` expressions

Nested functions

Nested functions (lambdas and local functions) are treated like methods, except in regards to their captured variables. The initial state of a captured variable inside a lambda or local function is the intersection of the nullable state of the variable at all the "uses" of that nested function or lambda. A use of a local function is either a call to that function, or where it is converted to a delegate. A use of a lambda is the point at which it is defined in source.

Type inference

nullable implicitly typed local variables

`var` infers an annotated type for reference types, and type parameters that aren't constrained to be a value type. For instance:

- in `var s = "";` the `var` is inferred as `string?`.
- in `var t = new T();` with an unconstrained `T` the `var` is inferred as `T?`.

Generic type inference

Generic type inference is enhanced to help decide whether inferred reference types should be nullable or not. This is a best effort. It may yield warnings regarding nullability constraints, and may lead to nullable warnings when the inferred types of the selected overload are applied to the arguments.

The first phase

Nullable reference types flow into the bounds from the initial expressions, as described below. In addition, two new kinds of bounds, namely `null` and `default` are introduced. Their purpose is to carry through occurrences of `null` or `default` in the input expressions, which may cause an inferred type to be nullable, even when it otherwise wouldn't.

The determination of what bounds to add in the first phase are enhanced as follows:

If an argument `Ei` has a reference type, the type `U` used for inference depends on the null state of `Ei` as well as its declared type:

- If the declared type is a nonnullable reference type `U0` or a nullable reference type `U0?` then
  - if the null state of `Ei` is "not null" then `U` is `U0`
  - if the null state of `Ei` is "maybe null" then `U` is `U0?`
- Otherwise if `Ei` has a declared type, `U` is that type
- Otherwise if `Ei` is `null` then `U` is the special bound `null`
- Otherwise if `Ei` is `default` then `U` is the special bound `default`
- Otherwise no inference is made.

Exact, upper-bound and lower-bound inferences

In inferences *from* the type `U` *to* the type `V`, if `V` is a nullable reference type `V0?`, then `V0` is used instead of `V` in the following clauses.

- If `V` is one of the unfixed type variables, `U` is added as an exact, upper or lower bound as before
- Otherwise, if `U` is `null` or `default`, no inference is made
- Otherwise, if `U` is a nullable reference type `U0?`, then `U0` is used instead of `U` in the subsequent clauses.

The essence is that nullability that pertains directly to one of the unfixed type variables is preserved into its bounds. For the inferences that recurse further into the source and target types, on the other hand, nullability is ignored. It may or may not match, but if it doesn't, a warning will be issued later if the overload is chosen and applied.

Fixing

The spec currently does not do a good job of describing what happens when multiple bounds are identity convertible to each other, but are different. This may happen between `object` and `dynamic`, between tuple types that differ only in element names, between types constructed thereof and now also between `C` and `C?` for reference types.

In addition we need to propagate "nullness" from the input expressions to the result type.

To handle these we add more phases to fixing, which is now:

1. Gather all the types in all the bounds as candidates, removing `?` from all that are nullable reference types
2. Eliminate candidates based on requirements of exact, lower and upper bounds (keeping `null` and `default` bounds)
3. Eliminate candidates that do not have an implicit conversion to all the other candidates
4. If the remaining candidates do not all have identity conversions to one another, then type inference fails
5. *Merge* the remaining candidates as described below
6. If the resulting candidate is a reference type and *all* of the exact bounds or *any* of the lower bounds are nullable reference types, `null` or `default`, then `?` is added to the resulting candidate, making it a nullable reference type.

Merging is described between two candidate types. It is transitive and commutative, so the candidates can be merged in any order with the same ultimate result. It is undefined if the two candidate types are not identity convertible to each other.

The *Merge* function takes two candidate types and a direction (+ or -):

- *Merge*(`T`, `T`, *d*) = T
- *Merge*(`S`, `T?`, +) = *Merge*(`S?`, `T`, +) = *Merge*(`S`, `T`, +)`?`
- *Merge*(`S`, `T?`, -) = *Merge*(`S?`, `T`, -) = *Merge*(`S`, `T`, -)
- *Merge*(`C<S1,...,Sn>`, `C<T1,...,Tn>`, +) = `C<`*Merge*(`S1`, `T1`, *d1*)`,...,`*Merge*(`Sn`, `Tn`, *dn*)`>`, *where*
  - `di` = + if the `i`'th type parameter of `C<...>` is covariant

    - `di` = - if the `i`'th type parameter of `C<...>` is contra- or invariant
- *Merge*(`C<S1,...,Sn>`, `C<T1,...,Tn>`, -) = `C<`*Merge*(`S1`, `T1`, *d1*)`,...,`*Merge*(`Sn`, `Tn`, *dn*)`>`, *where*
    - `di` = - if the `i`'th type parameter of `C<...>` is covariant
    - `di` = + if the `i`'th type parameter of `C<...>` is contra- or invariant
- *Merge*(`(S1 s1,..., Sn sn)`, `(T1 t1,..., Tn tn)`, *d*) = `(`*Merge*(`S1`, `T1`, *d*)`n1,...,`*Merge*(`Sn`, `Tn`, *d*) `nn)`, *where*
    - `ni` is absent if `si` and `ti` differ, or if both are absent
    - `ni` is `si` if `si` and `ti` are the same
- *Merge*(`object`, `dynamic`) = *Merge*(`dynamic`, `object`) = `dynamic`

Warnings

Potential null assignment

Potential null dereference

Constraint nullability mismatch

Nullable types in disabled annotation context

Override and implementation nullability mismatch

Attributes for special null behavior

Records

Article • 02/08/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

This proposal tracks the specification for the C# 9 records feature, as agreed to by the C# language design team.

The syntax for a record is as follows:

antlr

```
record_declaration
    : attributes? class_modifier* 'partial'? 'record' identifier
type_parameter_list?
      parameter_list? record_base? type_parameter_constraints_clause*
record_body
    ;

record_base
    : ':' class_type argument_list?
    | ':' interface_type_list
    | ':' class_type argument_list? ',' interface_type_list
    ;

record_body
    : '{' class_member_declaration* '}' ';'?
    | ';'
    ;
```

Record types are reference types, similar to a class declaration. It is an error for a record to provide a `record_base` `argument_list` if the `record_declaration` does not contain a `parameter_list`. At most one partial type declaration of a partial record may provide a `parameter_list`.

Record parameters cannot use `ref`, `out` or `this` modifiers (but `in` and `params` are allowed).

Inheritance

Records cannot inherit from classes, unless the class is `object`, and classes cannot inherit from records. Records can inherit from other records.

Members of a record type

In addition to the members declared in the record body, a record type has additional synthesized members. Members are synthesized unless a member with a "matching" signature is declared in the record body or an accessible concrete non-virtual member with a "matching" signature is inherited. A matching member prevents the compiler from generating that member, not any other synthesized members. Two members are considered matching if they have the same signature or would be considered "hiding" in an inheritance scenario. It is an error for a member of a record to be named "Clone". It is an error for an instance field of a record to have an unsafe type.

The synthesized members are as follows:

Equality members

If the record is derived from `object`, the record type includes a synthesized readonly property equivalent to a property declared as follows:

```
C#
```

```
Type EqualityContract { get; }
```

The property is `private` if the record type is `sealed`. Otherwise, the property is `virtual` and `protected`. The property can be declared explicitly. It is an error if the explicit declaration does not match the expected signature or accessibility, or if the explicit declaration doesn't allow overriding it in a derived type and the record type is not `sealed`.

If the record type is derived from a base record type `Base`, the record type includes a synthesized readonly property equivalent to a property declared as follows:

```
C#
```

```
protected override Type EqualityContract { get; }
```

The property can be declared explicitly. It is an error if the explicit declaration does not match the expected signature or accessibility, or if the explicit declaration doesn't allow overriding it in a derived type and the record type is not `sealed`. It is an error if either synthesized, or explicitly declared property doesn't override a property with this signature in the record type `Base` (for example, if the property is missing in the `Base`, or sealed, or not virtual, etc.). The synthesized property returns `typeof(R)` where `R` is the record type.

The record type implements `System.IEquatable<R>` and includes a synthesized strongly-typed overload of `Equals(R? other)` where `R` is the record type. The method is `public`, and the method is `virtual` unless the record type is `sealed`. The method can be declared explicitly. It is an error if the explicit declaration does not match the expected signature or accessibility, or the explicit declaration doesn't allow overriding it in a derived type and the record type is not `sealed`.

If `Equals(R? other)` is user-defined (not synthesized) but `GetHashCode` is not, a warning is produced.

C#

```
public virtual bool Equals(R? other);
```

The synthesized `Equals(R?)` returns `true` if and only if each of the following are `true`:

- `other` is not `null`, and
- For each instance field `fieldN` in the record type that is not inherited, the value of `System.Collections.Generic.EqualityComparer<TN>.Default.Equals(fieldN, other.fieldN)` where `TN` is the field type, and
- If there is a base record type, the value of `base.Equals(other)` (a non-virtual call to `public virtual bool Equals(Base? other)`); otherwise the value of `EqualityContract == other.EqualityContract`.

The record type includes synthesized `==` and `!=` operators equivalent to operators declared as follows:

C#

```
public static bool operator==(R? left, R? right)
    => (object)left == right || (left?.Equals(right) ?? false);
```

```
public static bool operator!=(R? left, R? right)
    => !(left == right);
```

The `Equals` method called by the `==` operator is the `Equals(R? other)` method specified above. The `!=` operator delegates to the `==` operator. It is an error if the operators are declared explicitly.

If the record type is derived from a base record type `Base`, the record type includes a synthesized override equivalent to a method declared as follows:

C#

```
public sealed override bool Equals(Base? other);
```

It is an error if the override is declared explicitly. It is an error if the method doesn't override a method with same signature in record type `Base` (for example, if the method is missing in the `Base`, or sealed, or not virtual, etc.). The synthesized override returns `Equals((object?)other)`.

The record type includes a synthesized override equivalent to a method declared as follows:

C#

```
public override bool Equals(object? obj);
```

It is an error if the override is declared explicitly. It is an error if the method doesn't override `object.Equals(object? obj)` (for example, due to shadowing in intermediate base types, etc.). The synthesized override returns `Equals(other as R)` where `R` is the record type.

The record type includes a synthesized override equivalent to a method declared as follows:

C#

```
public override int GetHashCode();
```

The method can be declared explicitly. It is an error if the explicit declaration doesn't allow overriding it in a derived type and the record type is not `sealed`. It is an error if either synthesized, or explicitly declared method doesn't override `object.GetHashCode()` (for example, due to shadowing in intermediate base types, etc.).

A warning is reported if one of `Equals(R?)` and `GetHashCode()` is explicitly declared but the other method is not explicit.

The synthesized override of `GetHashCode()` returns an `int` result of combining the following values:

- For each instance field `fieldN` in the record type that is not inherited, the value of `System.Collections.Generic.EqualityComparer<TN>.Default.GetHashCode(fieldN)` where `TN` is the field type, and
- If there is a base record type, the value of `base.GetHashCode()`; otherwise the value of `System.Collections.Generic.EqualityComparer<System.Type>.Default.GetHashCode(EqualityContract)`.

For example, consider the following record types:

C#

```
record R1(T1 P1);
record R2(T1 P1, T2 P2) : R1(P1);
record R3(T1 P1, T2 P2, T3 P3) : R2(P1, P2);
```

For those record types, the synthesized equality members would be something like:

C#

```
class R1 : IEquatable<R1>
{
    public T1 P1 { get; init; }
    protected virtual Type EqualityContract => typeof(R1);
    public override bool Equals(object? obj) => Equals(obj as R1);
    public virtual bool Equals(R1? other)
    {
        return !(other is null) &&
            EqualityContract == other.EqualityContract &&
            EqualityComparer<T1>.Default.Equals(P1, other.P1);
    }
    public static bool operator==(R1? left, R1? right)
        => (object)left == right || (left?.Equals(right) ?? false);
    public static bool operator!=(R1? left, R1? right)
        => !(left == right);
    public override int GetHashCode()
    {
        return
HashCode.Combine(EqualityComparer<Type>.Default.GetHashCode(EqualityContract
),
            EqualityComparer<T1>.Default.GetHashCode(P1));
    }
```

```
}

class R2 : R1, IEquatable<R2>
{
    public T2 P2 { get; init; }
    protected override Type EqualityContract => typeof(R2);
    public override bool Equals(object? obj) => Equals(obj as R2);
    public sealed override bool Equals(R1? other) => Equals((object?)other);
    public virtual bool Equals(R2? other)
    {
        return base.Equals((R1?)other) &&
            EqualityComparer<T2>.Default.Equals(P2, other.P2);
    }
    public static bool operator==(R2? left, R2? right)
        => (object)left == right || (left?.Equals(right) ?? false);
    public static bool operator!=(R2? left, R2? right)
        => !(left == right);
    public override int GetHashCode()
    {
        return HashCode.Combine(base.GetHashCode(),
            EqualityComparer<T2>.Default.GetHashCode(P2));
    }
}

class R3 : R2, IEquatable<R3>
{
    public T3 P3 { get; init; }
    protected override Type EqualityContract => typeof(R3);
    public override bool Equals(object? obj) => Equals(obj as R3);
    public sealed override bool Equals(R2? other) => Equals((object?)other);
    public virtual bool Equals(R3? other)
    {
        return base.Equals((R2?)other) &&
            EqualityComparer<T3>.Default.Equals(P3, other.P3);
    }
    public static bool operator==(R3? left, R3? right)
        => (object)left == right || (left?.Equals(right) ?? false);
    public static bool operator!=(R3? left, R3? right)
        => !(left == right);
    public override int GetHashCode()
    {
        return HashCode.Combine(base.GetHashCode(),
            EqualityComparer<T3>.Default.GetHashCode(P3));
    }
}
```

Copy and Clone members

A record type contains two copying members:

- A constructor taking a single argument of the record type. It is referred to as a "copy constructor".

- A synthesized public parameterless instance "clone" method with a compiler-reserved name

The purpose of the copy constructor is to copy the state from the parameter to the new instance being created. This constructor doesn't run any instance field/property initializers present in the record declaration. If the constructor is not explicitly declared, a constructor will be synthesized by the compiler. If the record is sealed, the constructor will be private, otherwise it will be protected. An explicitly declared copy constructor must be either public or protected, unless the record is sealed. The first thing the constructor must do, is to call a copy constructor of the base, or a parameter-less object constructor if the record inherits from object. An error is reported if a user-defined copy constructor uses an implicit or explicit constructor initializer that doesn't fulfill this requirement. After a base copy constructor is invoked, a synthesized copy constructor copies values for all instance fields implicitly or explicitly declared within the record type. The sole presence of a copy constructor, whether explicit or implicit, doesn't prevent an automatic addition of a default instance constructor.

If a virtual "clone" method is present in the base record, the synthesized "clone" method overrides it and the return type of the method is the current containing type if the "covariant returns" feature is supported and the override return type otherwise. An error is produced if the base record clone method is sealed. If a virtual "clone" method is not present in the base record, the return type of the clone method is the containing type and the method is virtual, unless the record is sealed or abstract. If the containing record is abstract, the synthesized clone method is also abstract. If the "clone" method is not abstract, it returns the result of a call to a copy constructor.

Printing members: PrintMembers and ToString methods

If the record is derived from `object`, the record includes a synthesized method equivalent to a method declared as follows:

```
C#

bool PrintMembers(System.Text.StringBuilder builder);
```

The method is `private` if the record type is `sealed`. Otherwise, the method is `virtual` and `protected`.

The method:

1. calls the method `System.Runtime.CompilerServices.RuntimeHelpers.EnsureSufficientExecutionStack`

`()` if the method is present and the record has printable members.
2. for each of the record's printable members (non-static public field and readable property members), appends that member's name followed by " = " followed by the member's value separated with ", ",
3. return true if the record has printable members.

For a member that has a value type, we will convert its value to a string representation using the most efficient method available to the target platform. At present that means calling `ToString` before passing to `StringBuilder.Append`.

If the record type is derived from a base record `Base`, the record includes a synthesized override equivalent to a method declared as follows:

```
C#

protected override bool PrintMembers(StringBuilder builder);
```

If the record has no printable members, the method calls the base `PrintMembers` method with one argument (its `builder` parameter) and returns the result.

Otherwise, the method:

1. calls the base `PrintMembers` method with one argument (its `builder` parameter),
2. if the `PrintMembers` method returned true, append ", " to the builder,
3. for each of the record's printable members, appends that member's name followed by " = " followed by the member's value: `this.member` (or `this.member.ToString()` for value types), separated with ", ",
4. return true.

The `PrintMembers` method can be declared explicitly. It is an error if the explicit declaration does not match the expected signature or accessibility, or if the explicit declaration doesn't allow overriding it in a derived type and the record type is not `sealed`.

The record includes a synthesized method equivalent to a method declared as follows:

```
C#

public override string ToString();
```

The method can be declared explicitly. It is an error if the explicit declaration does not match the expected signature or accessibility, or if the explicit declaration doesn't allow overriding it in a derived type and the record type is not `sealed`. It is an error if either

synthesized, or explicitly declared method doesn't override `object.ToString()` (for example, due to shadowing in intermediate base types, etc.).

The synthesized method:

1. creates a `StringBuilder` instance,
2. appends the record name to the builder, followed by " { ",
3. invokes the record's `PrintMembers` method giving it the builder, followed by " " if it returned true,
4. appends "}",
5. returns the builder's contents with `builder.ToString()`.

For example, consider the following record types:

C#

```
record R1(T1 P1);
record R2(T1 P1, T2 P2, T3 P3) : R1(P1);
```

For those record types, the synthesized printing members would be something like:

C#

```
class R1 : IEquatable<R1>
{
    public T1 P1 { get; init; }

    protected virtual bool PrintMembers(StringBuilder builder)
    {
        builder.Append(nameof(P1));
        builder.Append(" = ");
        builder.Append(this.P1); // or builder.Append(this.P1.ToString());
if P1 has a value type

        return true;
    }

    public override string ToString()
    {
        var builder = new StringBuilder();
        builder.Append(nameof(R1));
        builder.Append(" { ");

        if (PrintMembers(builder))
            builder.Append(" ");

        builder.Append("}");
        return builder.ToString();
    }
```

```
}

class R2 : R1, IEquatable<R2>
{
    public T2 P2 { get; init; }
    public T3 P3 { get; init; }

    protected override bool PrintMembers(StringBuilder builder)
    {
        if (base.PrintMembers(builder))
            builder.Append(", ");

        builder.Append(nameof(P2));
        builder.Append(" = ");
        builder.Append(this.P2); // or builder.Append(this.P2); if P2 has a
value type

        builder.Append(", ");

        builder.Append(nameof(P3));
        builder.Append(" = ");
        builder.Append(this.P3); // or builder.Append(this.P3); if P3 has a
value type

        return true;
    }

    public override string ToString()
    {
        var builder = new StringBuilder();
        builder.Append(nameof(R2));
        builder.Append(" { ");

        if (PrintMembers(builder))
            builder.Append(" ");

        builder.Append("}");
        return builder.ToString();
    }
}
```

Positional record members

In addition to the above members, records with a parameter list ("positional records") synthesize additional members with the same conditions as the members above.

Primary Constructor

A record type has a public constructor whose signature corresponds to the value parameters of the type declaration. This is called the primary constructor for the type,

and causes the implicitly declared default class constructor, if present, to be suppressed. It is an error to have a primary constructor and a constructor with the same signature already present in the class.

At runtime the primary constructor

1. executes the instance initializers appearing in the class-body
2. invokes the base class constructor with the arguments provided in the `record_base` clause, if present

If a record has a primary constructor, any user-defined constructor, except "copy constructor" must have an explicit `this` constructor initializer.

Parameters of the primary constructor as well as members of the record are in scope within the `argument_list` of the `record_base` clause and within initializers of instance fields or properties. Instance members would be an error in these locations (similar to how instance members are in scope in regular constructor initializers today, but an error to use), but the parameters of the primary constructor would be in scope and useable and would shadow members. Static members would also be useable, similar to how base calls and initializers work in ordinary constructors today.

A warning is produced if a parameter of the primary constructor is not read.

Expression variables declared in the `argument_list` are in scope within the `argument_list`. The same shadowing rules as within an argument list of a regular constructor initializer apply.

Properties

For each record parameter of a record type declaration there is a corresponding public property member whose name and type are taken from the value parameter declaration.

For a record:

- A public `get` and `init` auto-property is created (see separate `init` accessor specification). An inherited `abstract` property with matching type is overridden. It is an error if the inherited property does not have `public` overridable `get` and `init` accessors. It is an error if the inherited property is hidden.
  The auto-property is initialized to the value of the corresponding primary constructor parameter. Attributes can be applied to the synthesized auto-property

and its backing field by using `property:` or `field:` targets for attributes syntactically applied to the corresponding record parameter.

Deconstruct

A positional record with at least one parameter synthesizes a public void-returning instance method called Deconstruct with an out parameter declaration for each parameter of the primary constructor declaration. Each parameter of the Deconstruct method has the same type as the corresponding parameter of the primary constructor declaration. The body of the method assigns to each parameter of the Deconstruct method, the value of the instance property of the same name. The method can be declared explicitly. It is an error if the explicit declaration does not match the expected signature or accessibility, or is static.

The following example shows a positional record `R` with its compiler synthesized `Deconstruct` method, along with its usage:

C#

```
public record R(int P1, string P2 = "xyz")
{
    public void Deconstruct(out int P1, out string P2)
    {
        P1 = this.P1;
        P2 = this.P2;
    }
}

class Program
{
    static void Main()
    {
        R r = new R(12);
        (int p1, string p2) = r;
        Console.WriteLine($"p1: {p1}, p2: {p2}");
    }
}
```

`with` expression

A `with` expression is a new expression using the following syntax.

antlr

```
with_expression
    : switch_expression
    | switch_expression 'with' '{' member_initializer_list? '}'
    ;

member_initializer_list
    : member_initializer (',' member_initializer)*
    ;

member_initializer
    : identifier '=' expression
    ;
```

A `with` expression is not permitted as a statement.

A `with` expression allows for "non-destructive mutation", designed to produce a copy of the receiver expression with modifications in assignments in the `member_initializer_list`.

A valid `with` expression has a receiver with a non-void type. The receiver type must be a record.

On the right hand side of the `with` expression is a `member_initializer_list` with a sequence of assignments to *identifier*, which must be an accessible instance field or property of the receiver's type.

First, receiver's "clone" method (specified above) is invoked and its result is converted to the receiver's type. Then, each `member_initializer` is processed the same way as an assignment to a field or property access of the result of the conversion. Assignments are processed in lexical order.

Record structs

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

The syntax for a record struct is as follows:

antlr

```
record_struct_declaration
    : attributes? struct_modifier* 'partial'? 'record' 'struct' identifier
type_parameter_list?
      parameter_list? struct_interfaces? type_parameter_constraints_clause*
record_struct_body
    ;

record_struct_body
    : struct_body
    | ';'
    ;
```

Record struct types are value types, like other struct types. They implicitly inherit from the class `System.ValueType`. The modifiers and members of a record struct are subject to the same restrictions as those of structs (accessibility on type, modifiers on members, `base(...)` instance constructor initializers, definite assignment for `this` in constructor, destructors, ...). Record structs will also follow the same rules as structs for parameterless instance constructors and field initializers, but this document assumes that we will lift those restrictions for structs generally.

See §15.4.9 See parameterless struct constructors spec.

Record structs cannot use `ref` modifier.

At most one partial type declaration of a partial record struct may provide a `parameter_list`. The `parameter_list` may be empty.

Record struct parameters cannot use `ref`, `out` or `this` modifiers (but `in` and `params` are allowed).

Members of a record struct

In addition to the members declared in the record struct body, a record struct type has additional synthesized members. Members are synthesized unless a member with a "matching" signature is declared in the record struct body or an accessible concrete non-virtual member with a "matching" signature is inherited. Two members are considered matching if they have the same signature or would be considered "hiding" in an inheritance scenario. See Signatures and overloading §7.6. It is an error for a member of a record struct to be named "Clone".

It is an error for an instance field of a record struct to have an unsafe type.

A record struct is not permitted to declare a destructor.

The synthesized members are as follows:

Equality members

The synthesized equality members are similar as in a record class (`Equals` for this type, `Equals` for `object` type, `==` and `!=` operators for this type),
except for the lack of `EqualityContract`, null checks or inheritance.

The record struct implements `System.IEquatable<R>` and includes a synthesized strongly-typed overload of `Equals(R other)` where `R` is the record struct. The method is `public`. The method can be declared explicitly. It is an error if the explicit declaration does not match the expected signature or accessibility.

If `Equals(R other)` is user-defined (not synthesized) but `GetHashCode` is not, a warning is produced.

```
C#
```

```
public readonly bool Equals(R other);
```

The synthesized `Equals(R)` returns `true` if and only if for each instance field `fieldN` in the record struct the value of
`System.Collections.Generic.EqualityComparer<TN>.Default.Equals(fieldN, other.fieldN)` where `TN` is the field type is `true`.

The record struct includes synthesized `==` and `!=` operators equivalent to operators declared as follows:

```
C#
```

```
public static bool operator==(R r1, R r2)
    => r1.Equals(r2);
public static bool operator!=(R r1, R r2)
    => !(r1 == r2);
```

The `Equals` method called by the `==` operator is the `Equals(R other)` method specified above. The `!=` operator delegates to the `==` operator. It is an error if the operators are declared explicitly.

The record struct includes a synthesized override equivalent to a method declared as follows:

```
C#
```

```
public override readonly bool Equals(object? obj);
```

It is an error if the override is declared explicitly. The synthesized override returns `other is R temp && Equals(temp)` where `R` is the record struct.

The record struct includes a synthesized override equivalent to a method declared as follows:

```
C#
```

```
public override readonly int GetHashCode();
```

The method can be declared explicitly.

A warning is reported if one of `Equals(R)` and `GetHashCode()` is explicitly declared but the other method is not explicit.

The synthesized override of `GetHashCode()` returns an `int` result of combining the values of `System.Collections.Generic.EqualityComparer<TN>.Default.GetHashCode(fieldN)` for each instance field `fieldN` with `TN` being the type of `fieldN`.

For example, consider the following record struct:

```
C#
```

```
record struct R1(T1 P1, T2 P2);
```

For this record struct, the synthesized equality members would be something like:

C#

```
struct R1 : IEquatable<R1>
{
    public T1 P1 { get; set; }
    public T2 P2 { get; set; }
    public override bool Equals(object? obj) => obj is R1 temp &&
Equals(temp);
    public bool Equals(R1 other)
    {
        return
            EqualityComparer<T1>.Default.Equals(P1, other.P1) &&
            EqualityComparer<T2>.Default.Equals(P2, other.P2);
    }
    public static bool operator==(R1 r1, R1 r2)
        => r1.Equals(r2);
    public static bool operator!=(R1 r1, R1 r2)
        => !(r1 == r2);
    public override int GetHashCode()
    {
        return Combine(
            EqualityComparer<T1>.Default.GetHashCode(P1),
            EqualityComparer<T2>.Default.GetHashCode(P2));
    }
}
```

Printing members: PrintMembers and ToString methods

The record struct includes a synthesized method equivalent to a method declared as follows:

C#

```
private bool PrintMembers(System.Text.StringBuilder builder);
```

The method does the following:

1. for each of the record struct's printable members (non-static public field and readable property members), appends that member's name followed by " = " followed by the member's value separated with ", ",
2. return true if the record struct has printable members.

For a member that has a value type, we will convert its value to a string representation using the most efficient method available to the target platform. At present that means calling `ToString` before passing to `StringBuilder.Append`.

If the record's printable members do not include a readable property with a non-`readonly` `get` accessor, then the synthesized `PrintMembers` is `readonly`. There is no requirement for the record's fields to be `readonly` for the `PrintMembers` method to be `readonly`.

The `PrintMembers` method can be declared explicitly. It is an error if the explicit declaration does not match the expected signature or accessibility.

The record struct includes a synthesized method equivalent to a method declared as follows:

```
C#

public override string ToString();
```

If the record struct's `PrintMembers` method is `readonly`, then the synthesized `ToString()` method is `readonly`.

The method can be declared explicitly. It is an error if the explicit declaration does not match the expected signature or accessibility.

The synthesized method:

1. creates a `StringBuilder` instance,
2. appends the record struct name to the builder, followed by " { ",
3. invokes the record struct's `PrintMembers` method giving it the builder, followed by " " if it returned true,
4. appends "}",
5. returns the builder's contents with `builder.ToString()`.

For example, consider the following record struct:

```
C#

record struct R1(T1 P1, T2 P2);
```

For this record struct, the synthesized printing members would be something like:

```
C#
```

```
struct R1 : IEquatable<R1>
{
    public T1 P1 { get; set; }
    public T2 P2 { get; set; }

    private bool PrintMembers(StringBuilder builder)
    {
        builder.Append(nameof(P1));
        builder.Append(" = ");
        builder.Append(this.P1); // or builder.Append(this.P1.ToString());
if P1 has a value type
        builder.Append(", ");

        builder.Append(nameof(P2));
        builder.Append(" = ");
        builder.Append(this.P2); // or builder.Append(this.P2.ToString());
if P2 has a value type

        return true;
    }

    public override string ToString()
    {
        var builder = new StringBuilder();
        builder.Append(nameof(R1));
        builder.Append(" { ");

        if (PrintMembers(builder))
            builder.Append(" ");

        builder.Append("}");
        return builder.ToString();
    }
}
```

Positional record struct members

In addition to the above members, record structs with a parameter list ("positional records") synthesize additional members with the same conditions as the members above.

Primary Constructor

A record struct has a public constructor whose signature corresponds to the value parameters of the type declaration. This is called the primary constructor for the type. It is an error to have a primary constructor and a constructor with the same signature

already present in the struct. If the type declaration does not include a parameter list, no primary constructor is generated.

```C#
record struct R1
{
    public R1() { } // ok
}

record struct R2()
{
    public R2() { } // error: 'R2' already defines constructor with same
parameter types
}
```

Instance field declarations for a record struct are permitted to include variable initializers. If there is no primary constructor, the instance initializers execute as part of the parameterless constructor. Otherwise, at runtime the primary constructor executes the instance initializers appearing in the record-struct-body.

If a record struct has a primary constructor, any user-defined constructor must have an explicit `this` constructor initializer that calls the primary constructor or an explicitly declared constructor.

Parameters of the primary constructor as well as members of the record struct are in scope within initializers of instance fields or properties. Instance members would be an error in these locations (similar to how instance members are in scope in regular constructor initializers today, but an error to use), but the parameters of the primary constructor would be in scope and useable and would shadow members. Static members would also be useable.

A warning is produced if a parameter of the primary constructor is not read.

The definite assignment rules for struct instance constructors apply to the primary constructor of record structs. For instance, the following is an error:

```C#
record struct Pos(int X) // definite assignment error in primary constructor
{
    private int x;
    public int X { get { return x; } set { x = value; } } = X;
}
```

Properties

For each record struct parameter of a record struct declaration there is a corresponding public property member whose name and type are taken from the value parameter declaration.

For a record struct:

- A public `get` and `init` auto-property is created if the record struct has `readonly` modifier, `get` and `set` otherwise. Both kinds of set accessors (`set` and `init`) are considered "matching". So the user may declare an init-only property in place of a synthesized mutable one. An inherited `abstract` property with matching type is overridden. No auto-property is created if the record struct has an instance field with expected name and type. It is an error if the inherited property does not have `public` `get` and `set`/`init` accessors. It is an error if the inherited property or field is hidden.
  The auto-property is initialized to the value of the corresponding primary constructor parameter. Attributes can be applied to the synthesized auto-property and its backing field by using `property:` or `field:` targets for attributes syntactically applied to the corresponding record struct parameter.

Deconstruct

A positional record struct with at least one parameter synthesizes a public void-returning instance method called `Deconstruct` with an out parameter declaration for each parameter of the primary constructor declaration. Each parameter of the Deconstruct method has the same type as the corresponding parameter of the primary constructor declaration. The body of the method assigns each parameter of the Deconstruct method to the value from an instance member access to a member of the same name. If the instance members accessed in the body do not include a property with a non-`readonly` `get` accessor, then the synthesized `Deconstruct` method is `readonly`. The method can be declared explicitly. It is an error if the explicit declaration does not match the expected signature or accessibility, or is static.

Allow `with` expression on structs

It is now valid for the receiver in a `with` expression to have a struct type.

On the right hand side of the `with` expression is a `member_initializer_list` with a sequence of assignments to *identifier*, which must be an accessible instance field or

property of the receiver's type.

For a receiver with struct type, the receiver is first copied, then each `member_initializer` is processed the same way as an assignment to a field or property access of the result of the conversion. Assignments are processed in lexical order.

Improvements on records

Allow `record class`

The existing syntax for record types allows `record class` with the same meaning as `record`:

antlr

```
record_declaration
    : attributes? class_modifier* 'partial'? 'record' 'class'? identifier
type_parameter_list?
      parameter_list? record_base? type_parameter_constraints_clause*
record_body
    ;
```

Allow user-defined positional members to be fields

See https://github.com/dotnet/csharplang/blob/master/meetings/2020/LDM-2020-10-05.md#changing-the-member-type-of-a-primary-constructor-parameter

No auto-property is created if the record has or inherits an instance field with expected name and type.

Allow parameterless constructors and member initializers in structs

See parameterless struct constructors spec.

Open questions

- how to recognize record structs in metadata? (we don't have an unspeakable clone method to leverage...)

Answered

- confirm that we want to keep PrintMembers design (separate method returning `bool`) (answer: yes)
- confirm we won't allow `record ref struct` (issue with `IEquatable<RefStruct>` and ref fields) (answer: yes)
- confirm implementation of equality members. Alternative is that synthesized `bool Equals(R other)`, `bool Equals(object? other)` and operators all just delegate to `ValueType.Equals`. (answer: yes)
- confirm that we want to allow field initializers when there is a primary constructor. Do we also want to allow parameterless struct constructors while we're at it (the Activator issue was apparently fixed)? (answer: yes, updated spec should be reviewed in LDM)
- how much do we want to say about `Combine` method? (answer: as little as possible)
- should we disallow a user-defined constructor with a copy constructor signature? (answer: no, there is no notion of copy constructor in the record structs spec)
- confirm that we want to disallow members named "Clone". (answer: correct)
- double-check that synthesized `Equals` logic is functionally equivalent to runtime implementation (e.g. float.NaN) (answer: confirmed in LDM)
- could field- or property-targeting attributes be placed in the positional parameter list? (answer: yes, same as for record class)
- `with` on generics? (answer: out of scope for C# 10)
- should `GetHashCode` include a hash of the type itself, to get different values between `record struct S1;` and `record struct S2;`? (answer: no)

File-local types

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Permit a `file` modifier on top-level type declarations. The type only exists in the file where it is declared.

C#

```
// File1.cs
namespace NS;

file class Widget
{
}

// File2.cs
namespace NS;

file class Widget // different symbol than the Widget in File1
{
}

// File3.cs
using NS;

var widget = new Widget(); // error: The type or namespace name 'Widget'
could not be found.
```

Motivation

Our primary motivation is from source generators. Source generators work by adding files to the user's compilation.

1. Those files should be able to contain implementation details which are hidden from the rest of the compilation, yet are usable throughout the file they are declared in.
2. We want to reduce the need for generators to "search" for type names which won't collide with declarations in user code or code from other generators.

Detailed design

- We add the `file` modifier to the following modifier sets:
  - class
  - struct
  - interface
  - enum
  - delegate
  - record
  - record struct.
- The `file` modifier can only be used on a top-level type.

When a type has the `file` modifier, it is said to be a *file-local* type.

Accessibility

No accessibility modifiers can be used in combination with `file` on a type. `file` is treated as an independent concept from accessibility. Since file-local types can't be nested, only the default accessibility `internal` is usable with `file` types.

C#

```
public file class C1 { } // error
internal file class C2 { } // error
file class C3 { } // ok
```

Naming

The implementation guarantees that file-local types in different files with the same name will be distinct to the runtime. The type's accessibility and name in metadata is implementation-defined. The intention is to permit the compiler to adopt any future access-limitation features in the runtime which are suited to the feature. It's expected

that in the initial implementation, an `internal` accessibility would be used and an unspeakable generated name will be used which depends on the file the type is declared in.

Lookup

We amend the member lookup section as follows (new text in **bold**):

> - Next, if `K` is zero, all nested types whose declarations include type parameters are removed. If `K` is not zero, all members with a different number of type parameters are removed. When `K` is zero, methods having type parameters are not removed, since the type inference process (§11.6.3) might be able to infer the type arguments.
> - **Next, let *F* be the compilation unit which contains the expression where member lookup is occurring. All members which are file-local types and are not declared in *F* are removed from the set.**
> - **Next, if the set of accessible members contains file-local types, all members which are not file-local types are removed from the set.**

Remarks

These rules disallow usage of file-local types outside the file in which they are declared.

These rules also permit a file-local type to *shadow* a namespace or a non-file-local type:

```C#
// File1.cs
class C
{
    public static void M() { }
}
```

```C#
// File2.cs
file class C
{
    public static void M() { }
}

class Program
{
    static void Main()
    {
```

```
        C.M(); // refers to the 'C' in File2.cs
    }
}
```

Note that we don't update the scopes section of the spec. This is because, as the spec states:

> The ***scope*** of a name is the region of program text within which it is possible to refer to the entity declared by the name without qualification of the name.

In effect, scope only impacts the lookup of non-qualified names. This isn't quite the right concept for us to leverage because we need to also impact the lookup of qualified names:

C#

```
// File1.cs
namespace NS1
{
    file class C
    {
        public static void M() { }
    }
}

namespace NS2
{
    class Program
    {
        public static void M()
        {
            C.M(); // error: C is not in scope
            NS1.C.M(); // ok: C can be accessed through NS1.
        }
    }
}
```

C#

```
// File2.cs
namespace NS1
{
    class Program
    {
        C.M(); // error
        NS1.C.M(); // error
    }
}
```

Therefore, we don't specify the feature in terms of which scope the type is contained in, but rather as additional "filtering rules" in member lookup.

Attributes

File-local classes are permitted to be attribute types, and can be used as attributes within both file-local types and non-file-local types, just as if the attribute type were a non-file-local type. The metadata name of the file-local attribute type still goes through the same name generation strategy as other file-local types. This means detecting the presence of a file-local type by a hard-coded string name is likely to be impractical, because it requires depending on the internal name generation strategy of the compiler, which may change over time. However, detecting via `typeof(MyFileLocalAttribute)` works.

C#

```
using System;
using System.Linq;

file class MyFileLocalAttribute : Attribute { }

[MyFileLocalAttribute]
public class C
{
    public static void Main()
    {
        var attribute = typeof(C).CustomAttributes.Where(attr =>
attr.AttributeType == typeof(MyFileLocalAttribute)).First();
        Console.Write(attribute); // outputs the generated name of the file-
local attribute type
    }
}
```

Usage in signatures

There is a general need to prevent file-local types from appearing in member parameters, returns, and type parameter constraints where the file-local type might not be in scope at the point of usage of the member.

Note that non-file-local types are permitted to implement file-local interfaces, similar to how types can implement less-accessible interfaces. Depending on the types present in the interface members, it could result in a violation of the rules in the following section.

Only allow signature usage in members of file-local types

Perhaps the simplest way to ensure this is to enforce that file-local types can only appear in signatures or as base types of other file-local types:

```
C#

file class FileBase
{
}

public class Derived : FileBase // error
{
    private FileBase M2() => new FileBase() // error
}

file class FileDerived : FileBase // ok
{
    private FileBase M2() => new FileBase() // ok
}
```

Note that this does restrict usage in explicit implementations, even though such usages are safe. We do this in order to simplify the rules for the initial iteration of the feature.

```
C#

file interface I
{
    void M(I i);
}

class C : I
{
    void I.M(I i) { } // error
}
```

`global using static`

It is a compile-time error to use a file-local type in a `global using static` directive, i.e.

```
C#

global using static C; // error

file class C
{
    public static void M() { }
}
```

Implementation/overrides

file-local type declarations can implement interfaces, override virtual methods, etc. just like regular type declarations.

C#

```
file struct Widget : IEquatable<Widget>
{
    public bool Equals(Widget other) => true;
}
```

Unconstrained type parameter annotations

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow nullable annotations for type parameters that are not constrained to value types or reference types: `T?`.

```
C#

static T? FirstOrDefault<T>(this IEnumerable<T> collection) { ... }
```

`?` annotation

In C#8, `?` annotations could only be applied to type parameters that were explicitly constrained to value types or reference types. In C#9, `?` annotations can be applied to any type parameter, regardless of constraints.

Unless a type parameter is explicitly constrained to value types, annotations can only be applied within a `#nullable enable` context.

If a type parameter `T` is substituted with a reference type, then `T?` represents a nullable instance of that reference type.

```
C#
```

```
var s1 = new string[0].FirstOrDefault();  // string? s1
var s2 = new string?[0].FirstOrDefault(); // string? s2
```

If `T` is substituted with a value type, then `T?` represents an instance of `T`.

C#

```
var i1 = new int[0].FirstOrDefault();  // int i1
var i2 = new int?[0].FirstOrDefault(); // int? i2
```

If `T` is substituted with an annotated type `U?`, then `T?` represents the annotated type `U?` rather than `U??`.

C#

```
var u1 = new U[0].FirstOrDefault();  // U? u1
var u2 = new U?[0].FirstOrDefault(); // U? u2
```

If `T` is substituted with a type `U`, then `T?` represents `U?`, even within a `#nullable disable` context.

C#

```
#nullable disable
var u3 = new U[0].FirstOrDefault();  // U? u3
```

For return values, `T?` is equivalent to `[MaybeNull]T`; for argument values, `T?` is equivalent to `[AllowNull]T`. The equivalence is important when overriding or implementing interfaces from an assembly compiled with C#8.

C#

```
public abstract class A
{
    [return: MaybeNull] public abstract T F1<T>();
    public abstract void F2<T>([AllowNull] T t);
}

public class B : A
{
    public override T? F1<T>() where T : default { ... }       // matches
A.F1<T>()
    public override void F2<T>(T? t) where T : default { ... } // matches
A.F2<T>()
}
```

`default` constraint

For compatibility with existing code where overridden and explicitly implemented generic methods could not include explicit constraint clauses, `T?` in an overridden or explicitly implemented method is treated as `Nullable<T>` where `T` is a value type.

To allow annotations for type parameters constrained to reference types, C#8 allowed explicit `where T : class` and `where T : struct` constraints on the overridden or explicitly implemented method.

C#

```
class A1
{
    public virtual void F1<T>(T? t) where T : struct { }
    public virtual void F1<T>(T? t) where T : class { }
}

class B1 : A1
{
    public override void F1<T>(T? t) /*where T : struct*/ { }
    public override void F1<T>(T? t) where T : class { }
}
```

To allow annotations for type parameters that are not constrained to reference types or value types, C#9 allows a new `where T : default` constraint.

C#

```
class A2
{
    public virtual void F2<T>(T? t) where T : struct { }
    public virtual void F2<T>(T? t) { }
}

class B2 : A2
{
    public override void F2<T>(T? t) /*where T : struct*/ { }
    public override void F2<T>(T? t) where T : default { }
}
```

It is an error to use a `default` constraint other than on a method override or explicit implementation. It is an error to use a `default` constraint when the corresponding type parameter in the overridden or interface method is constrained to a reference type or value type.

Design meetings

- https://github.com/dotnet/csharplang/blob/master/meetings/2019/LDM-2019-11-25.md
- https://github.com/dotnet/csharplang/blob/master/meetings/2020/LDM-2020-06-17.md#t

Improved Interpolated Strings

Article • 03/21/2022

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

We introduce a new pattern for creating and using interpolated string expressions to allow for efficient formatting and use in both general `string` scenarios and more specialized scenarios such as logging frameworks, without incurring unnecessary allocations from formatting the string in the framework.

Motivation

Today, string interpolation mainly lowers down to a call to `string.Format`. This, while general purpose, can be inefficient for a number of reasons:

1. It boxes any struct arguments, unless the runtime has happened to introduce an overload of `string.Format` that takes exactly the correct types of arguments in exactly the correct order.
    - This ordering is why the runtime is hesitant to introduce generic versions of the method, as it would lead to combinatoric explosion of generic instantiations of a very common method.
2. It has to allocate an array for the arguments in most cases.
3. There is no opportunity to avoid instantiating the instance if it's not needed. Logging frameworks, for example, will recommend avoiding string interpolation because it will cause a string to be realized that may not be needed, depending on the current log-level of the application.

4. It can never use `Span` or other ref struct types today, because ref structs are not allowed as generic type parameters, meaning that if a user wants to avoid copying to intermediate locations they have to manually format strings.

Internally, the runtime has a type called `ValueStringBuilder` to help deal with the first 2 of these scenarios. They pass a stackalloc'd buffer to the builder, repeatedly call `AppendFormat` with every part, and then get a final string out. If the resulting string goes past the bounds of the stack buffer, they can then move to an array on the heap. However, this type is dangerous to expose directly, as incorrect usage could lead to a rented array to be double-disposed, which then will cause all sorts of undefined behavior in the program as two locations think they have sole access to the rented array. This proposal creates a way to use this type safely from native C# code by just writing an interpolated string literal, leaving written code unchanged while improving every interpolated string that a user writes. It also extends this pattern to allow for interpolated strings passed as arguments to other methods to use a handler pattern, defined by receiver of the method, that will allow things like logging frameworks to avoid allocating strings that will never be needed, and giving C# users familiar, convenient interpolation syntax.

Detailed Design

The handler pattern

We introduce a new handler pattern that can represent an interpolated string passed as an argument to a method. The simple English of the pattern is as follows:

When an *interpolated_string_expression* is passed as an argument to a method, we look at the type of the parameter. If the parameter type has a constructor that can be invoked with 2 int parameters, `literalLength` and `formattedCount`, optionally takes additional parameters specified by an attribute on the original parameter, optionally has an out boolean trailing parameter, and the type of the original parameter has instance `AppendLiteral` and `AppendFormatted` methods that can be invoked for every part of the interpolated string, then we lower the interpolation using that, instead of into a traditional call to `string.Format(formatStr, args)`. A more concrete example is helpful for picturing this:

C#

```
// The handler that will actually "build" the interpolated string"
[InterpolatedStringHandler]
public ref struct TraceLoggerParamsInterpolatedStringHandler
```

```
{
    // Storage for the built-up string

    private bool _logLevelEnabled;

    public TraceLoggerParamsInterpolatedStringHandler(int literalLength, int
formattedCount, Logger logger, out bool handlerIsValid)
    {
        if (!logger._logLevelEnabled)
        {
            handlerIsValid = false;
            return;
        }

        handlerIsValid = true;
        _logLevelEnabled = logger.EnabledLevel;
    }

    public void AppendLiteral(string s)
    {
        // Store and format part as required
    }

    public void AppendFormatted<T>(T t)
    {
        // Store and format part as required
    }
}

// The logger class. The user has an instance of this, accesses it via
static state, or some other access
// mechanism
public class Logger
{
    // Initialization code omitted
    public LogLevel EnabledLevel;

    public void
LogTrace([InterpolatedStringHandlerArguments("")]TraceLoggerParamsInterpolat
edStringHandler handler)
    {
        // Impl of logging
    }
}

Logger logger = GetLogger(LogLevel.Info);

// Given the above definitions, usage looks like this:
var name = "Fred Silberberg";
logger.LogTrace($"{name} will never be printed because info is < trace!");

// This is converted to:
var name = "Fred Silberberg";
var receiverTemp = logger;
var handler = new TraceLoggerParamsInterpolatedStringHandler(literalLength:
```

```
47, formattedCount: 1, receiverTemp, out var handlerIsValid);
if (handlerIsValid)
{
    handler.AppendFormatted(name);
    handler.AppendLiteral(" will never be printed because info is <
trace!");
}
receiverTemp.LogTrace(handler);
```

Here, because `TraceLoggerParamsInterpolatedStringHandler` has a constructor with the correct parameters, we say that the interpolated string has an implicit handler conversion to that parameter, and it lowers to the pattern shown above. The specese needed for this is a bit complicated, and is expanded below.

The rest of this proposal will use `Append...` to refer to either of `AppendLiteral` or `AppendFormatted` in cases when both are applicable.

New attributes

The compiler recognizes the `System.Runtime.CompilerServices.InterpolatedStringHandlerAttribute`:

C#

```
using System;
namespace System.Runtime.CompilerServices
{
    [AttributeUsage(AttributeTargets.Class | AttributeTargets.Struct,
AllowMultiple = false, Inherited = false)]
    public sealed class InterpolatedStringHandlerAttribute : Attribute
    {
        public InterpolatedStringHandlerAttribute()
        {
        }
    }
}
```

This attribute is used by the compiler to determine if a type is a valid interpolated string handler type.

The compiler also recognizes the `System.Runtime.CompilerServices.InterpolatedStringHandlerArgumentAttribute`:

C#

```
namespace System.Runtime.CompilerServices
{
```

```
    [AttributeUsage(AttributeTargets.Parameter, AllowMultiple = false,
Inherited = false)]
    public sealed class InterpolatedStringHandlerArgumentAttribute :
Attribute
    {
        public InterpolatedHandlerArgumentAttribute(string argument);
        public InterpolatedHandlerArgumentAttribute(params string[]
arguments);

        public string[] Arguments { get; }
    }
}
```

This attribute is used on parameters, to inform the compiler how to lower an interpolated string handler pattern used in a parameter position.

Interpolated string handler conversion

Type `T` is said to be an *applicable_interpolated_string_handler_type* if it is attributed with `System.Runtime.CompilerServices.InterpolatedStringHandlerAttribute`. There exists an implicit *interpolated_string_handler_conversion* to `T` from an *interpolated_string_expression*, or an *additive_expression* composed entirely of _interpolated_string_expression_s and using only `+` operators.

For simplicity in the rest of this speclet, *interpolated_string_expression* refers to both a simple *interpolated_string_expression*, and to an *additive_expression* composed entirely of _interpolated_string_expression_s and using only `+` operators.

Note that this conversion always exists, regardless of whether there will be later errors when actually attempting to lower the interpolation using the handler pattern. This is done to help ensure that there are predictable and useful errors and that runtime behavior doesn't change based on the content of an interpolated string.

Applicable function member adjustments

We adjust the wording of the applicable function member algorithm (§11.6.4.2) as follows (a new sub-bullet is added to each section, in bold):

A function member is said to be an ***applicable function member*** with respect to an argument list `A` when all of the following are true:

- Each argument in `A` corresponds to a parameter in the function member declaration as described in Corresponding parameters (§11.6.2.2), and any parameter to which no argument corresponds is an optional parameter.

- For each argument in `A`, the parameter passing mode of the argument (i.e., value, `ref`, or `out`) is identical to the parameter passing mode of the corresponding parameter, and
  - for a value parameter or a parameter array, an implicit conversion (§10.2) exists from the argument to the type of the corresponding parameter, or
  - **for a `ref` parameter whose type is a struct type, an implicit *interpolated_string_handler_conversion* exists from the argument to the type of the corresponding parameter, or**
  - for a `ref` or `out` parameter, the type of the argument is identical to the type of the corresponding parameter. After all, a `ref` or `out` parameter is an alias for the argument passed.

For a function member that includes a parameter array, if the function member is applicable by the above rules, it is said to be applicable in its ***normal form***. If a function member that includes a parameter array is not applicable in its normal form, the function member may instead be applicable in its ***expanded form***:

- The expanded form is constructed by replacing the parameter array in the function member declaration with zero or more value parameters of the element type of the parameter array such that the number of arguments in the argument list `A` matches the total number of parameters. If `A` has fewer arguments than the number of fixed parameters in the function member declaration, the expanded form of the function member cannot be constructed and is thus not applicable.
- Otherwise, the expanded form is applicable if for each argument in `A` the parameter passing mode of the argument is identical to the parameter passing mode of the corresponding parameter, and
  - for a fixed value parameter or a value parameter created by the expansion, an implicit conversion (§10.2) exists from the type of the argument to the type of the corresponding parameter, or
  - **for a `ref` parameter whose type is a struct type, an implicit *interpolated_string_handler_conversion* exists from the argument to the type of the corresponding parameter, or**
  - for a `ref` or `out` parameter, the type of the argument is identical to the type of the corresponding parameter.

Important note: this means that if there are 2 otherwise equivalent overloads, that only differ by the type of the *applicable_interpolated_string_handler_type*, these overloads will be considered ambiguous. Further, because we do not see through explicit casts, it is possible that there could arise an unresolvable scenario where both applicable overloads use `InterpolatedStringHandlerArguments` and are totally uncallable without manually performing the handler lowering pattern. We could potentially make changes

to the better function member algorithm to resolve this if we so choose, but this scenario unlikely to occur and isn't a priority to address.

Better conversion from expression adjustments

We change the better conversion from expression (§11.6.4.4) section to the following:

Given an implicit conversion `C1` that converts from an expression `E` to a type `T1`, and an implicit conversion `C2` that converts from an expression `E` to a type `T2`, `C1` is a ***better conversion*** than `C2` if:

1. `E` is a non-constant *interpolated_string_expression*, `C1` is an *implicit_string_handler_conversion*, `T1` is an *applicable_interpolated_string_handler_type*, and `C2` is not an *implicit_string_handler_conversion*, or
2. `E` does not exactly match `T2` and at least one of the following holds:
    - `E` exactly matches `T1` (§11.6.4.4)
    - `T1` is a better conversion target than `T2` (§11.6.4.6)

This does mean that there are some potentially non-obvious overload resolution rules, depending on whether the interpolated string in question is a constant-expression or not. For example:

C#

```
void Log(string s) { ... }
void Log(TraceLoggerParamsInterpolatedStringHandler p) { ... }

Log($""); // Calls Log(string s), because $"" is a constant expression
Log($"{"test"}"); // Calls Log(string s), because $"{"test"}" is a constant
expression
Log($"{1}"); // Calls Log(TraceLoggerParamsInterpolatedStringHandler p),
because $"{1}" is not a constant expression
```

This is introduced so that things that can simply be emitted as constants do so, and don't incur any overhead, while things that cannot be constant use the handler pattern.

InterpolatedStringHandler and Usage

We introduce a new type in `System.Runtime.CompilerServices`: `DefaultInterpolatedStringHandler`. This is a ref struct with many of the same semantics

as `ValueStringBuilder`, intended for direct use by the C# compiler. This struct would look approximately like this:

C#

```
// API Proposal issue: https://github.com/dotnet/runtime/issues/50601
namespace System.Runtime.CompilerServices
{
    [InterpolatedStringHandler]
    public ref struct DefaultInterpolatedStringHandler
    {
        public DefaultInterpolatedStringHandler(int literalLength, int
formattedCount);
        public string ToStringAndClear();

        public void AppendLiteral(string value);

        public void AppendFormatted<T>(T value);
        public void AppendFormatted<T>(T value, string? format);
        public void AppendFormatted<T>(T value, int alignment);
        public void AppendFormatted<T>(T value, int alignment, string?
format);

        public void AppendFormatted(ReadOnlySpan<char> value);
        public void AppendFormatted(ReadOnlySpan<char> value, int alignment
= 0, string? format = null);

        public void AppendFormatted(string? value);
        public void AppendFormatted(string? value, int alignment = 0,
string? format = null);

        public void AppendFormatted(object? value, int alignment = 0,
string? format = null);
    }
}
```

We make a slight change to the rules for the meaning of an *interpolated_string_expression* (§11.7.3 ↗):

If the type of an interpolated string is `string` and the type `System.Runtime.CompilerServices.DefaultInterpolatedStringHandler` exists, and the current context supports using that type, the string is lowered using the handler pattern. The final `string` value is then obtained by calling `ToStringAndClear()` on the handler type. Otherwise, if the type of an interpolated string is `System.IFormattable` or `System.FormattableString` [the rest is unchanged]

The "and the current context supports using that type" rule is intentionally vague to give the compiler leeway in optimizing usage of this pattern. The handler type is likely to be a ref struct type, and ref struct types are normally not permitted in async methods. For

this particular case, the compiler would be allowed to make use the handler if none of the interpolation holes contain an `await` expression, as we can statically determine that the handler type is safely used without additional complicated analysis because the handler will be dropped after the interpolated string expression is evaluated.

~~Open~~ **Question**:

Do we want to instead just make the compiler know about `DefaultInterpolatedStringHandler` and skip the `string.Format` call entirely? It would allow us to hide a method that we don't necessarily want to put in people's faces when they manually call `string.Format`.

Answer: Yes.

~~Open~~ **Question**:

Do we want to have handlers for `System.IFormattable` and `System.FormattableString` as well?

Answer: No.

Handler pattern codegen

In this section, method invocation resolution refers to the steps listed in §11.7.8.2.

Constructor resolution

Given an *applicable_interpolated_string_handler_type* `T` and an *interpolated_string_expression* `i`, method invocation resolution and validation for a valid constructor on `T` is performed as follows:

1. Member lookup for instance constructors is performed on `T`. The resulting method group is called `M`.
2. The argument list `A` is constructed as follows:
   a. The first two arguments are integer constants, representing the literal length of `i`, and the number of *interpolation* components in `i`, respectively.
   b. If `i` is used as an argument to some parameter `pi` in method `M1`, and parameter `pi` is attributed with `System.Runtime.CompilerServices.InterpolatedStringHandlerArgumentAttribute`, then for every name `Argx` in the `Arguments` array of that attribute the compiler matches it to a parameter `px` that has the same name. The empty string is matched to the receiver of `M1`.

- If any `Argx` is not able to be matched to a parameter of `M1`, or an `Argx` requests the receiver of `M1` and `M1` is a static method, an error is produced and no further steps are taken.
- Otherwise, the type of every resolved `px` is added to the argument list, in the order specified by the `Arguments` array. Each `px` is passed with the same `ref` semantics as is specified in `M1`.

c. The final argument is a `bool`, passed as an `out` parameter.

3. Traditional method invocation resolution is performed with method group `M` and argument list `A`. For the purposes of method invocation final validation, the context of `M` is treated as a *member_access* through type `T`.

   - If a single-best constructor `F` was found, the result of overload resolution is `F`.
   - If no applicable constructors were found, step 3 is retried, removing the final `bool` parameter from `A`. If this retry also finds no applicable members, an error is produced and no further steps are taken.
   - If no single-best method was found, the result of overload resolution is ambiguous, an error is produced, and no further steps are taken.

4. Final validation on `F` is performed.

   - If any element of `A` occurred lexically after `i`, an error is produced and no further steps are taken.
   - If any `A` requests the receiver of `F`, and `F` is an indexer being used as an *initializer_target* in a *member_initializer*, then an error is reported and no further steps are taken.

Note: the resolution here intentionally do *not* use the actual expressions passed as other arguments for `Argx` elements. We only consider the types post-conversion. This makes sure that we don't have double-conversion issues, or unexpected cases where a lambda is bound to one delegate type when passed to `M1` and bound to a different delegate type when passed to `M`.

Note: We report an error for indexers uses as member initializers because of the order of evaluation for nested member initializers. Consider this code snippet:

```
C#
```

```
var x1 = new C1 { C2 = { [GetString()] = { A = 2, B = 4 } } };

/* Lowering:
```

```
__c1 = new C1();
string argTemp = GetString();
__c1.C2[argTemp][1] = 2;
__c1.C2[argTemp][3] = 4;

Prints:
GetString
get_C2
get_C2
*/

string GetString()
{
    Console.WriteLine("GetString");
    return "";
}

class C1
{
    private C2 c2 = new C2();
    public C2 C2 { get { Console.WriteLine("get_C2"); return c2; } set { } }
}

class C2
{
    public C3 this[string s]
    {
        get => new C3();
        set { }
    }
}

class C3
{
    public int A
    {
        get => 0;
        set { }
    }
    public int B
    {
        get => 0;
        set { }
    }
}
```

The arguments to `__c1.C2[]` are evaluated *before* the receiver of the indexer. While we could come up with a lowering that works for this scenario (either by creating a temp for `__c1.C2` and sharing it across both indexer invocations, or only using it for the first indexer invocation and sharing the argument across both invocations) we think that any lowering would be confusing for what we believe is a pathological scenario. Therefore, we forbid the scenario entirely.

~~Open Question~~:

If we use a constructor instead of `Create`, we'd improve runtime codegen, at the expense of narrowing the pattern a bit.

Answer: We will restrict to constructors for now. We can revisit adding a general `Create` method later if the scenario arises.

`Append...` method overload resolution

Given an *applicable_interpolated_string_handler_type* `T` and an *interpolated_string_expression* `i`, overload resolution for a set of valid `Append...` methods on `T` is performed as follows:

1. If there are any *interpolated_regular_string_character* components in `i`:
   a. Member lookup on `T` with the name `AppendLiteral` is performed. The resulting method group is called `Ml`.
   b. The argument list `Al` is constructed with one value parameter of type `string`.
   c. Traditional method invocation resolution is performed with method group `Ml` and argument list `Al`. For the purposes of method invocation final validation, the context of `Ml` is treated as a *member_access* through an instance of `T`.
      - If a single-best method `Fi` is found and no errors were produced, the result of method invocation resolution is `Fi`.
      - Otherwise, an error is reported.
2. For every *interpolation* `ix` component of `i`:
   a. Member lookup on `T` with the name `AppendFormatted` is performed. The resulting method group is called `Mf`.
   b. The argument list `Af` is constructed:
      i. The first parameter is the `expression` of `ix`, passed by value.
      ii. If `ix` directly contains a *constant_expression* component, then an integer value parameter is added, with the name `alignment` specified.
      iii. If `ix` is directly followed by an *interpolation_format*, then a string value parameter is added, with the name `format` specified.
   c. Traditional method invocation resolution is performed with method group `Mf` and argument list `Af`. For the purposes of method invocation final validation, the context of `Mf` is treated as a *member_access* through an instance of `T`.

    - If a single-best method `Fi` is found, the result of method invocation resolution is `Fi`.
    - Otherwise, an error is reported.

3. Finally, for every `Fi` discovered in steps 1 and 2, final validation is performed:

    - If any `Fi` does not return `bool` by value or `void`, an error is reported.
    - If all `Fi` do not return the same type, an error is reported.

Note that these rules do not permit extension methods for the `Append...` calls. We could consider enabling that if we choose, but this is analogous to the enumerator pattern, where we allow `GetEnumerator` to be an extension method, but not `Current` or `MoveNext()`.

These rules *do* permit default parameters for the `Append...` calls, which will work with things like `CallerLineNumber` or `CallerArgumentExpression` (when supported by the language).

We have separate overload lookup rules for base elements vs interpolation holes because some handlers will want to be able to understand the difference between the components that were interpolated and the components that were part of the base string.

~~Open~~ Question

Some scenarios, like structured logging, want to be able to provide names for interpolation elements. For example, today a logging call might look like `Log("{name} bought {itemCount} items", name, items.Count);`. The names inside the `{}` provide important structure information for loggers that help with ensuring output is consistent and uniform. Some cases might be able to reuse the `:format` component of an interpolation hole for this, but many loggers already understand format specifiers and have existing behavior for output formatting based on this info. Is there some syntax we can use to enable putting these named specifiers in?

Some cases may be able to get away with `CallerArgumentExpression`, provided that support does land in C# 10. But for cases that invoke a method/property, that may not be sufficient.

Answer:

While there are some interesting parts to templated strings we could explore in an orthogonal language feature, we don't think a specific syntax here has much benefit over solutions such as using a tuple: `$"{("StructuredCategory", myExpression)}"`.

Performing the conversion

Given an *applicable_interpolated_string_handler_type* `T` and an *interpolated_string_expression* `i` that had a valid constructor `Fc` and `Append...` methods `Fa` resolved, lowering for `i` is performed as follows:

1. Any arguments to `Fc` that occur lexically before `i` are evaluated and stored into temporary variables in lexical order. In order to preserve lexical ordering, if `i` occurred as part of a larger expression `e`, any components of `e` that occurred before `i` will be evaluated as well, again in lexical order.
2. `Fc` is called with the length of the interpolated string literal components, the number of *interpolation* holes, any previously evaluated arguments, and a `bool` out argument (if `Fc` was resolved with one as the last parameter). The result is stored into a temporary value `ib`.
   a. The length of the literal components is calculated after replacing any *open_brace_escape_sequence* with a single `{`, and any *close_brace_escape_sequence* with a single `}`.
3. If `Fc` ended with a `bool` out argument, a check on that `bool` value is generated. If true, the methods in `Fa` will be called. Otherwise, they will not be called.
4. For every `Fax` in `Fa`, `Fax` is called on `ib` with either the current literal component or *interpolation* expression, as appropriate. If `Fax` returns a `bool`, the result is logically anded with all preceding `Fax` calls.
   a. If `Fax` is a call to `AppendLiteral`, the literal component is unescaped by replacing any *open_brace_escape_sequence* with a single `{`, and any *close_brace_escape_sequence* with a single `}`.
5. The result of the conversion is `ib`.

Again, note that arguments passed to `Fc` and arguments passed to `e` are the same temp. Conversions may occur on top of the temp to convert to a form that `Fc` requires, but for example lambdas cannot be bound to a different delegate type between `Fc` and `e`.

~~Open~~ Question

This lowering means that subsequent parts of the interpolated string after a false-returning `Append...` call don't get evaluated. This could potentially be very confusing, particularly if the format hole is side-effecting. We could instead evaluate all format holes first, then repeatedly call `Append...` with the results, stopping if it returns false. This would ensure that all expressions get evaluated as one might expect, but we call as

few methods as we need to. While the partial evaluation might be desirable for some more advanced cases, it is perhaps non-intuitive for the general case.

Another alternative, if we want to always evaluate all format holes, is to remove the `Append...` version of the API and just do repeated `Format` calls. The handler can track whether it should just be dropping the argument and immediately returning for this version.

Answer: We will have conditional evaluation of the holes.

~~Open~~ Question

Do we need to dispose of disposable handler types, and wrap calls with try/finally to ensure that Dispose is called? For example, the interpolated string handler in the bcl might have a rented array inside it, and if one of the interpolation holes throws an exception during evaluation, that rented array could be leaked if it wasn't disposed.

Answer: No. handlers can be assigned to locals (such as `MyHandler handler = $"{MyCode()};`), and the lifetime of such handlers is unclear. Unlike foreach enumerators, where the lifetime is obvious and no user-defined local is created for the enumerator.

Impact on nullable reference types

To minimize complexity of the implementation, we have a few limitations on how we perform nullable analysis on interpolated string handler constructors used as arguments to a method or indexer. In particular, we do not flow information from the constructor back through to the original slots of parameters or arguments from the original context, and we do not use constructor parameter types to inform generic type inference for type parameters in the containing method. An example of where this can have an impact is:

C#

```
string s = "";
C c = new C();
c.M(s, $"", c.ToString(), s.ToString()); // No warnings on c.ToString() or
s.ToString(), as the `MaybeNull` does not flow back.

public class C
{
    public void M(string s1, [InterpolatedStringHandlerArgument("", "s1")]
CustomHandler c1, string s2, string s3) { }
}

[InterpolatedStringHandler]
public partial struct CustomHandler
```

```
{
    public CustomHandler(int literalLength, int formattedCount, [MaybeNull]
C c, [MaybeNull] string s) : this()
    {
    }
}
```

C#

```
string? s = null;
M(s, $""); // Infers `string` for `T` because of the `T?` parameter, not
`string?`, as flow analysis does not consider the unannotated `T` parameter
of the constructor

void M<T>(T? t, [InterpolatedStringHandlerArgument("s1")] CustomHandler<T>
c) { }

[InterpolatedStringHandler]
public partial struct CustomHandler<T>
{
    public CustomHandler(int literalLength, int formattedCount, T t) :
this()
    {
    }
}
```

Other considerations

Allow `string` types to be convertible to handlers as well

For type author simplicity, we could consider allowing expressions of type `string` to be implicitly-convertible to *applicable_interpolated_string_handler_types*. As proposed today, authors will likely need to overload on both that handler type and regular `string` types, so their users don't have to understand the difference. This may be an annoying and non-obvious overhead, as a `string` expression can be viewed as an interpolation with `expression.Length` prefilled length and 0 holes to be filled.

This would allow new APIs to only expose a handler, without also having to expose a `string`-accepting overload. However, it won't get around the need for changes to better conversion from expression, so while it would work it may be unnecessary overhead.

Answer:

We think that this could end up being confusing, and there's an easy workaround for custom handler types: add a user-defined conversion from string.

Incorporating spans for heap-less strings

`ValueStringBuilder` as it exists today has 2 constructors: one that takes a count, and allocates on the heap eagerly, and one that takes a `Span<char>`. That `Span<char>` is usually a fixed size in the runtime codebase, around 250 elements on average. To truly replace that type, we should consider an extension to this where we also recognize `GetInterpolatedString` methods that take a `Span<char>`, instead of just the count version. However, we see a few potential thorny cases to resolve here:

- We don't want to stackalloc repeatedly in a hot loop. If we were to do this extension to the feature, we'd likely want to share the stackalloc'd span between loop iterations. We know this is safe, as `Span<T>` is a ref struct that can't be stored on the heap, and users would have to be pretty devious to manage to extract a reference to that `Span` (such as creating a method that accepts such a handler then deliberately retrieving the `Span` from the handler and returning it to the caller). However, allocating ahead of time produces other questions:
  - Should we eagerly stackalloc? What if the loop is never entered, or exits before it needs the space?
  - If we don't eagerly stackalloc, does that mean we introduce a hidden branch on every loop? Most loops likely won't care about this, but it could affect some tight loops that don't want to pay the cost.
- Some strings can be quite big, and the appropriate amount to `stackalloc` is dependent on a number of factors, including runtime factors. We don't really want the C# compiler and specification to have to determine this ahead of time, so we'd want to resolve https://github.com/dotnet/runtime/issues/25423 and add an API for the compiler to call in these cases. It also adds more pros and cons to the points from the previous loop, where we don't want to potentially allocate large arrays on the heap many times or before one is needed.

Answer:

This is out of scope for C# 10. We can look at this in general when we look at the more general `params Span<T>` feature.

Non-try version of the API

For simplicity, this spec currently just proposes recognizing a `Append...` method, and things that always succeed (like `InterpolatedStringHandler`) would always return true

from the method. This was done to support partial formatting scenarios where the user wants to stop formatting if an error occurs or if it's unnecessary, such as the logging case, but could potentially introduce a bunch of unnecessary branches in standard interpolated string usage. We could consider an addendum where we use just `FormatX` methods if no `Append...` method is present, but it does present questions about what we do if there's a mix of both `Append...` and `FormatX` calls.

Answer:

We want the non-try version of the API. The proposal has been updated to reflect this.

Passing previous arguments to the handler

There is unfortunate lack of symmetry in the proposal at it currently exists: invoking an extension method in reduced form produces different semantics than invoking the extension method in normal form. This is different from most other locations in the language, where reduced form is just a sugar. We propose adding an attribute to the framework that we will recognize when binding a method, that informs the compiler that certain parameters should be passed to the constructor on the handler. Usage looks like this:

C#

```
namespace System.Runtime.CompilerServices
{
    [AttributeUsage(AttributeTargets.Parameter, AllowMultiple = false,
Inherited = false)]
    public sealed class InterpolatedStringHandlerArgumentAttribute :
Attribute
    {
        public InterpolatedStringHandlerArgumentAttribute(string argument);
        public InterpolatedStringHandlerArgumentAttribute(params string[]
arguments);

        public string[] Arguments { get; }
    }
}
```

Usage of this is then:

C#

```
namespace System
{
    public sealed class String
    {
```

```
        public static string Format(IFormatProvider? provider,
[InterpolatedStringHandlerArgument("provider")] ref
DefaultInterpolatedStringHandler handler);
        …
    }
}

namespace System.Runtime.CompilerServices
{
    public ref struct DefaultInterpolatedStringHandler
    {
        public DefaultInterpolatedStringHandler(int baseLength, int
holeCount, IFormatProvider? provider); // additional factory
        …
    }
}

var formatted = string.Format(CultureInfo.InvariantCulture, $"{X} = {Y}");

// Is lowered to

var tmp1 = CultureInfo.InvariantCulture;
var handler = new DefaultInterpolatedStringHandler(3, 2, tmp1);
handler.AppendFormatted(X);
handler.AppendLiteral(" = ");
handler.AppendFormatted(Y);
var formatted = string.Format(tmp1, handler);
```

The questions we need to answer:

1. Do we like this pattern in general?
2. Do we want to allow these arguments to come from after the handler parameter? Some existing patterns in the BCL, such as `Utf8Formatter`, put the value to be formatted *before* the thing needed to format into. To fit in best with these patterns, we'd likely want to allow this, but we need to decide if this out-of-order evaluate is ok.

Answer:

We want to support this. The spec has been updated to reflect this. Arguments will be required to be specified in lexical order at the call site, and if a needed argument to the create method is specified after the interpolated string literal, an error is produced.

`await` usage in interpolation holes

Because `$"{await A()}"` is a valid expression today, we need to rationalize how interpolation holes with await. We could solve this with a few rules:

1. If an interpolated string used as a `string`, `IFormattable`, or `FormattableString` has an `await` in an interpolation hole, fall back to old-style formatter.
2. If an interpolated string is subject to an *implicit_string_handler_conversion* and *applicable_interpolated_string_handler_type* is a `ref struct`, `await` is not allowed to be used in the format holes.

Fundamentally, this desugaring could use a ref struct in an async method as long as we guarantee that the `ref struct` will not need to be saved to the heap, which should be possible if we forbid `await`s in the interpolation holes.

Alternatively, we could simply make all handler types non-ref structs, including the framework handler for interpolated strings. This would, however, preclude us from someday recognizing a `Span` version that does not need to allocate any scratch space at all.

Answer:

We will treat interpolated string handlers the same as any other type: this means that if the handler type is a ref struct and the current context doesn't allow the usage of ref structs, it is illegal to use handler here. The spec around lowering of string literals used as strings is intentionally vague to allow the compiler to decide on what rules it deems appropriate, but for custom handler types they will have to follow the same rules as the rest of the language.

Handlers as ref parameters

Some handlers might want to be passed as ref parameters (either `in` or `ref`). Should we allow either? And if so, what will a `ref` handler look like? `ref $""` is confusing, as you're not actually passing the string by ref, you're passing the handler that is created from the ref by ref, and has similar potential issues with async methods.

Answer:

We want to support this. The spec has been updated to reflect this. The rules should reflect the same rules that apply to extension methods on value types.

Interpolated strings through binary expressions and conversions

Because this proposal makes interpolated strings context sensitive, we would like to allow the compiler to treat a binary expression composed entirely of interpolated

strings, or an interpolated string subjected to a cast, as an interpolated string literal for the purposes of overload resolution. For example, take the following scenario:

```
C#
```

```
struct Handler1
{
    public Handler1(int literalLength, int formattedCount, C c) => ...;
    // AppendX... methods as necessary
}
struct Handler2
{
    public Handler2(int literalLength, int formattedCount, C c) => ...;
    // AppendX... methods as necessary
}

class C
{
    void M(Handler1 handler) => ...;
    void M(Handler2 handler) => ...;
}

c.M($"{X}"); // Ambiguous between the M overloads
```

This would be ambiguous, necessitating a cast to either `Handler1` or `Handler2` in order to resolve. However, in making that cast, we would potentially throw away the information that there is context from the method receiver, meaning that the cast would fail because there is nothing to fill in the information of `c`. A similar issue arises with binary concatenation of strings: the user could want to format the literal across several lines to avoid line wrapping, but would not be able to because that would no longer be an interpolated string literal convertible to the handler type.

To resolve these cases, we make the following changes:

- An *additive_expression* composed entirely of *interpolated_string_expressions* and using only `+` operators is considered to be an *interpolated_string_literal* for the purposes of conversions and overload resolution. The final interpolated string is created by logically concatinating all individual *interpolated_string_expression* components, from left to right.
- A *cast_expression* or a *relational_expression* with operator `as` whose operand is an *interpolated_string_expressions* is considered an *interpolated_string_expressions* for the purposes of conversions and overload resolution.

Open Questions:

Do we want to do this? We don't do this for `System.FormattableString`, for example, but that can be broken out onto a different line, whereas this can be context-dependent and therefore not able to be broken out into a different line. There are also no overload resolution concerns with `FormattableString` and `IFormattable`.

Answer:

We think that this is a valid use case for additive expressions, but that the cast version is not compelling enough at this time. We can add it later if necessary. The spec has been updated to reflect this decision.

Other use cases

See https://github.com/dotnet/runtime/issues/50635 for examples of proposed handler APIs using this pattern.

Constant Interpolated Strings

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Enables constants to be generated from interpolated strings of type string constant.

Motivation

The following code is already legal:

```
public class C
{
    const string S1 = "Hello world";
    const string S2 = "Hello" + " " + "World";
    const string S3 = S1 + " Kevin, welcome to the team!";
}
```

However, there have been many community requests to make the following also legal:

```
public class C
{
    const string S1 = $"Hello world";
    const string S2 = $"Hello{" "}World";
    const string S3 = $"{S1} Kevin, welcome to the team!";
}
```

This proposal represents the next logical step for constant string generation, where existing string syntax that works in other situations is made to work for constants.

Detailed design

The following represent the updated specifications for constant expressions under this new proposal. Current specifications from which this was directly based on can be found in §11.20.

Constant Expressions

A *constant_expression* is an expression that can be fully evaluated at compile-time.

antlr

```
constant_expression
    : expression
    ;
```

A constant expression must be the `null` literal or a value with one of the following types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, `double`, `decimal`, `bool`, `object`, `string`, or any enumeration type. Only the following constructs are permitted in constant expressions:

- Literals (including the `null` literal).
- References to `const` members of class and struct types.
- References to members of enumeration types.
- References to `const` parameters or local variables
- Parenthesized sub-expressions, which are themselves constant expressions.
- Cast expressions, provided the target type is one of the types listed above.
- `checked` and `unchecked` expressions
- Default value expressions
- Nameof expressions
- The predefined `+`, `-`, `!`, and `~` unary operators.
- The predefined `+`, `-`, `*`, `/`, `%`, `<<`, `>>`, `&`, `|`, `^`, `&&`, `||`, `==`, `!=`, `<`, `>`, `<=`, and `>=` binary operators, provided each operand is of a type listed above.
- The `?:` conditional operator.
- *Interpolated strings `$"{}"`, provided that all components are constant expressions of type `string` and all interpolated components lack alignment and format specifiers.*

The following conversions are permitted in constant expressions:

- Identity conversions
- Numeric conversions
- Enumeration conversions
- Constant expression conversions
- Implicit and explicit reference conversions, provided that the source of the conversions is a constant expression that evaluates to the null value.

Other conversions including boxing, unboxing and implicit reference conversions of non-null values are not permitted in constant expressions. For example:

```
C#
```

```
class C
{
    const object i = 5;          // error: boxing conversion not permitted
    const object str = "hello"; // error: implicit reference conversion
}
```

the initialization of i is an error because a boxing conversion is required. The initialization of str is an error because an implicit reference conversion from a non-null value is required.

Whenever an expression fulfills the requirements listed above, the expression is evaluated at compile-time. This is true even if the expression is a sub-expression of a larger expression that contains non-constant constructs.

The compile-time evaluation of constant expressions uses the same rules as run-time evaluation of non-constant expressions, except that where run-time evaluation would have thrown an exception, compile-time evaluation causes a compile-time error to occur.

Unless a constant expression is explicitly placed in an `unchecked` context, overflows that occur in integral-type arithmetic operations and conversions during the compile-time evaluation of the expression always cause compile-time errors (§11.20).

Constant expressions occur in the contexts listed below. In these contexts, a compile-time error occurs if an expression cannot be fully evaluated at compile-time.

- Constant declarations (§14.4).
- Enumeration member declarations (§18.4).
- Default arguments of formal parameter lists (§14.6.2)
- `case` labels of a `switch` statement (§12.8.3).
- `goto case` statements (§12.10.4).

- Dimension lengths in an array creation expression (§11.7.15.5) that includes an initializer.
- Attributes (§21).

An implicit constant expression conversion (§10.2.11) permits a constant expression of type `int` to be converted to `sbyte`, `byte`, `short`, `ushort`, `uint`, or `ulong`, provided the value of the constant expression is within the range of the destination type.

Drawbacks

This proposal adds additional complexity to the compiler in exchange for broader applicability of interpolated strings. As these strings are fully evaluated at compile time, the valuable automatic formatting features of interpolated strings are less neccesary. Most use cases can be largely replicated through the alternatives below.

Alternatives

The current `+` operator for string concatenation can combine strings in a similar manner to the current proposal.

Unresolved questions

What parts of the design are still undecided?

Design meetings

Link to design notes that affect this proposal, and describe in one sentence for each what changes they led to.

Raw string literal

Article • 01/10/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow a new form of string literal that starts with a minimum of three `"""` characters (but no maximum), optionally followed by a `new_line`, the content of the string, and then ends with the same number of quotes that the literal started with. For example:

```
var xml = """
          <element attr="content"/>
          """;
```

Because the nested contents might itself want to use `"""` then the starting/ending delimiters can be longer like so:

```
var xml = """"
          Ok to use """ here
          """";
```

To make the text easy to read and allow for indentation that developers like in code, these string literals will naturally remove the indentation specified on the last line when producing the final literal value. For example, a literal of the form:

```
var xml = """
          <element attr="content">
            <body>
            </body>
          </element>
          """;
```

Will have the contents:

```
<element attr="content">
  <body>
  </body>
</element>
```

This allows code to look natural, while still producing literals that are desired, and avoiding runtime costs if this required the use of specialized string manipulation routines.

If the indentation behavior is not desired, it is also trivial to disable like so:

```
var xml = """
          <element attr="content">
            <body>
            </body>
          </element>
""";
```

A single line form is also supported. It starts with a minimum of three `"""` characters (but no maximum), the content of the string (which cannot contain any `new_line` characters), and then ends with the same number of quotes that the literal started with. For example:

```
var xml = """<summary><element attr="content"/></summary>""";
```

Interpolated raw strings are also supported. In this case, the string specifies the number of braces needed to start an interpolation (determined by the number of dollar signs present at the start of the literal). Any brace sequence with fewer braces than that is just treated as content. For example:

```
var json = $$"""
             {
                "summary": "text",
                "length" : {{value.Length}},
             };
             """
```

Motivation

C# lacks a general way to create simple string literals that can contain effectively any arbitrary text. All C# string literal forms today need some form of escaping in case the contents use some special character (always if a delimiter is used). This prevents easily having literals containing other languages in them (for example, an XML, HTML or JSON literal).

All current approaches to form these literals in C# today always force the user to manually escape the contents. Editing at that point can be highly annoying as the escaping cannot be avoided and must be dealt with whenever it arises in the contents. This is particularly painful for regexes, especially when they contain quotes or backslashes. Even with a `@""` string, quotes themselves must be escaped leading to a mix of C# and regex interspersed. `{` and `}` are similarly frustrating in `$""` strings.

The crux of the problem is that all our strings have a fixed start/end delimiter. As long as that is the case, we will always have to have an escaping mechanism as the string contents may need to specify that end delimiter in their contents. This is particularly problematic as that delimiter `"` is exceedingly common in many languages.

To address this, this proposal allows for flexible start and end delimiters so that they can always be made in a way that will not conflict with the content of the string.

Goals

1. Provide a mechanism that will allow *all* string values to be provided by the user without the need for *any* escape-sequences whatsoever. Because all strings must be representable without escape-sequences, it must always be possible for the user to specify delimiters that will be guaranteed to not collide with any text contents.
2. Support interpolations in the same fashion. As above, because *all* strings must be representable without escapes, it must always be possible for the user to specify an `interpolation` delimiter that will be guaranteed to not collide with any text

contents. Importantly, languages that use our `interpolation` delimiter characters (`{` and `}`) should feel first-class and not painful to use.

3. Multiline string literals should look pleasant in code and should not make indentation within the compilation unit look strange. Importantly, literal values that themselves have no indentation should not be forced to occupy the first column of the file as that can break up the flow of code and will look unaligned with the rest of the code that surrounds it.

    - This behavior should be easy to override while keeping literals clear and easy to read.

4. For all strings that do not themselves contain a `new_line` or start or end with a quote (`"`) character, it should be possible to represent the string literal itself on a single line.

    - Optionally, with extra complexity, we could refine this to state that: For all strings that do not themselves contain a `new_line` (but can start or end with a quote `"` character), it should be possible to represent the string literal itself on a single line. For more details see the expanded proposal in the `Drawbacks` section.

Detailed design (non-interpolation case)

We will add a new `string_literal` production with the following form:

```
string_literal
    : regular_string_literal
    | verbatim_string_literal
    | raw_string_literal
    ;

raw_string_literal
    : single_line_raw_string_literal
    | multi_line_raw_string_literal
    ;

raw_string_literal_delimiter
    : """
    | """"
    | """""
    | etc.
    ;

raw_content
```

```
    : not_new_line+
    ;

single_line_raw_string_literal
    : raw_string_literal_delimiter raw_content raw_string_literal_delimiter
    ;

multi_line_raw_string_literal
    : raw_string_literal_delimiter whitespace* new_line (raw_content |
new_line)* new_line whitespace* raw_string_literal_delimiter
    ;

not_new_line
    : <any unicode character that is not new_line>
    ;
```

The ending delimiter to a `raw_string_literal` must match the starting delimiter. So if the starting delimiter is `"""""` the ending delimiter must be that as well.

The above grammar for a `raw_string_literal` should be interpreted as:

1. It starts with at least three quotes (but no upper bound on quotes).
2. It then continues with contents on the same line as the starting quotes. These contents on the same line can be blank, or non-blank. 'blank' is synonymous with 'entirely whitespace'.
3. If the contents on that same line is non-blank no further content can follow. In other words the literal is required to end with the same number of quotes on that same line.
4. If the contents on the same line is blank, then the literal can continue with a `new_line` and some number of subsequent content lines and `new_line`s.
    - A content line is any text except a `new_line`.
    - It then ends with a `new_line` some number (possibly zero) of `whitespace` and the same number of quotes that the literal started with.

Raw string literal value

The portions between the starting and ending `raw_string_literal_delimiter` are used to form the value of the `raw_string_literal` in the following fashion:

- In the case of `single_line_raw_string_literal` the value of the literal will exactly be the contents between the starting and ending `raw_string_literal_delimiter`.
- In the case of `multi_line_raw_string_literal` the initial `whitespace* new_line` and the final `new_line whitespace*` is not part of the value of the string. However, the

final `whitespace*` portion preceding the `raw_string_literal_delimiter` terminal is considered the 'indentation whitespace' and will affect how the other lines are interpreted.

- To get the final value the sequence of `(raw_content | new_line)*` is walked and the following is performed:
  - If it a `new_line` the content of the `new_line` is added to the final string value.
  - If it is not a 'blank' `raw_content` (i.e. `not_new_line+` contains a non-`whitespace` character):
    - the 'indentation whitespace' must be a prefix of the `raw_content`. It is an error otherwise.
    - the 'indentation whitespace' is stripped from the start of `raw_content` and the remainder is added to the final string value.
  - If it is a 'blank' `raw_content` (i.e. `not_new_line+` is entirely `whitespace`):
    - the 'indentation whitespace' must be a prefix of the `raw_content` or the `raw_content` must be a prefix of of the 'indentation whitespace'. It is an error otherwise.
    - as much of the 'indentation whitespace' is stripped from the start of `raw_content` and any remainder is added to the final string value.

Clarifications:

1. A `single_line_raw_string_literal` is not capable of representing a string with a `new_line` value in it. A `single_line_raw_string_literal` does not participate in the 'indentation whitespace' trimming. Its value is always the exact characters between the starting and ending delimiters.

2. Because a `multi_line_raw_string_literal` ignores the final `new_line` of the last content line, the following represents a string with no starting `new_line` and no terminating `new_line`

```
var v1 = """
         This is the entire content of the string.
         """
```

This maintains symmetry with how the starting `new_line` is ignored, and it also provides a uniform way to ensure the 'indentation whitespace' can always be adjusted. To represent a string with a terminal `new_line` an extra line must be provided like so:

```
```

```
var v1 = """
         This string ends with a new line.

         """
```

3. A `single_line_raw_string_literal` cannot represent a string value that starts or ends with a quote (`"`) though an augmentation to this proposal is provided in the `Drawbacks` section that shows how that could be supported.

4. A `multi_line_raw_string_literal` starts with `whitespace* new_line` following the initial `raw_string_literal_delimiter`. This content after the delimiter is entirely ignored and is not used in any way when determining the value of the string. This allows for a mechanism to specify a `raw_string_literal` whose content starts with a `"` character itself. For example:

```
var v1 = """
         "The content of this string starts with a quote
         """
```

5. A `raw_string_literal` can also represent content that end with a quote (`"`). This is supported as the terminating delimiter must be on its own line. For example:

```
var v1 = """
         "The content of this string starts and ends with a quote"
         """
```

```
var v1 = """
         ""The content of this string starts and ends with two quotes""
         """
```

5. The requirement that a 'blank' `raw_content` be either a prefix of the 'indentation whitespace' or the 'indentation whitespace' must be a prefix of it helps ensure confusing scenarios with mixed whitespace do not occur, especially as it would be unclear what should happen with that line. For example, the following case is illegal:

```
var v1 = """
         Start
<tab>
         End
         """
```

6. Here the 'indentation whitespace' is nine space characters, but the 'blank' `raw_content` does not start with a prefix of that. There is no clear answer as to how that `<tab>` line should be treated at all. Should it be ignored? Should it be the same as `.........<tab>`? As such, making it illegal seems the clearest for avoiding confusion.

7. The following cases are legal though and represent the same string:

```
var v1 = """
         Start
<four spaces>
         End
         """
```

```
var v1 = """
         Start
<nine spaces>
         End
         """
```

In both these cases, the 'indentation whitespace' will be nine spaces. And in both cases, we will remove as much of that prefix as possible, leading the 'blank' `raw_content` in each case to be empty (not counting every `new_line`). This allows users to not have to see and potentially fret about whitespace on these lines when they copy/paste or edit these lines.

8. In the case though of:

```
var v1 = """
         Start
<ten spaces>
         End
         """
```

The 'indentation whitespace' will still be nine spaces. Here though, we will remove as much of the 'indentation whitespace' as possible, and the 'blank' `raw_content` will contribute a single space to the final content. This allows for cases where the content does need whitespace on these lines that should be preserved.

9. The following is technically not legal:

```
var v1 = """
         """
```

This is because the start of the raw string must have a `new_line` (which it does) but the end must have a `new_line` as well (which it does not). The minimal legal `raw_string_literal` is:

```
var v1 = """

         """
```

However, this string is decidedly uninteresting as it is equivalent to `""`.

Indentation examples

The 'indentation whitespace' algorithm can be visualized on several inputs like so:

Example 1 - Standard case

```
var xml = """
          <element attr="content">
            <body>
            </body>
          </element>
          """;
```

is interpreted as

```
var xml = """
          |<element attr="content">
          |  <body>
          |  </body>
          |</element>
           """;
```

Example 2 - End delimiter on same line as content.

```
var xml = """
          <element attr="content">
            <body>
            </body>
          </element>""";
```

This is illegal. The last content line must end with a `new_line`.

Example 3 - End delimiter before start delimiter

```
var xml = """
          <element attr="content">
            <body>
            </body>
          </element>
""";
```

is interpreted as

```
var xml = """
|          <element attr="content">
|            <body>
|            </body>
|          </element>
""";
```

Example 4 - End delimiter after start delimiter

```
```

```
var xml = """
          <element attr="content">
            <body>
            </body>
          </element>
              """;
```

This is illegal. The lines of content must start with the 'indentation whitespace'

Example 5 - Empty blank line

```
var xml = """
          <element attr="content">
            <body>
            </body>

          </element>
          """;
```

is interpreted as

```
var xml = """
          |<element attr="content">
          |  <body>
          |  </body>
          |
          |</element>
           """;
```

Example 5 - Blank line with less whitespace than prefix (dots represent spaces)

```
var xml = """
          <element attr="content">
            <body>
            </body>
....
          </element>
          """;
```

is interpreted as

```
var xml = """
          |<element attr="content">
          |  <body>
          |  </body>
          |
          |</element>
           """;
```

Example 5 - Blank line with more whitespace than prefix (dots represent spaces)

```
var xml = """
          <element attr="content">
            <body>
            </body>
..............
          </element>
          """;
```

is interpreted as

```
var xml = """
          |<element attr="content">
          |  <body>
          |  </body>
          |....
          |</element>
           """;
```

Detailed design (interpolation case)

Interpolations in normal interpolated strings (e.g. `$"..."`) are supported today through the use of the `{` character to start an `interpolation` and the use of an `{{` escape-sequence to insert an actual open brace character. Using this same mechanism would violate goals '1' and '2' of this proposal. Languages that have `{` as a core character

(examples being JavaScript, JSON, Regex, and even embedded C#) would now need escaping, undoing the purpose of raw string literals.

To support interpolations we introduce them in a different fashion than normal `$"` interpolated strings. Specifically, an `interpolated_raw_string_literal` will start with some number of `$` characters. The count of these indicates how many `{` (and `}`) characters are needed in the content of the literal to delimit the `interpolation`. Importantly, there continues to be no escaping mechanism for curly braces. Rather, just as with quotes (`"`) the literal itself can always ensure it specifies delimiters for the interpolations that are certain to not collide with any of the rest of the content of the string. For example a JSON literal containing interpolation holes can be written like so:

```
c#

var v1 = $$"""
         {
            "orders":
            [
                { "number": {{order_number}} }
            ]
         }
         """
```

Here, the `{{...}}` matches the requisite count of two braces specified by the `$$` delimiter prefix. In the case of a single `$` that means the interpolation is specified just as `{...}` as in normal interpolated string literals. Importantly, this means that an interpolated literal with `N` `$` characters can have a sequence of `2*N-1` braces (of the same type in a row). The last `N` braces will start (or end) an interpolation, and the remaining `N-1` braces will just be content. For example:

```
c#

var v1 = $$"""X{{{1+1}}}Z""";
```

In this case the inner two `{{` and `}}` braces belong to the interpolation, and the outer singular braces are just content. So the above string is equivalent to the content `X{2}Z`. Having `2*N` (or more) braces is always an error. To have longer sequences of braces as content, the number of `$` characters must be increased accordingly.

Interpolated raw string literals are defined as:

```
```

```
interpolated_raw_string_literal
    : single_line_interpolated_raw_string_literal
    | multi_line_interpolated_raw_string_literal
    ;

interpolated_raw_string_start
    : $
    | $$
    | $$$
    | etc.
    ;

interpolated_raw_string_literal_delimiter
    : interpolated_raw_string_start raw_string_literal_delimiter
    ;

single_line_interpolated_raw_string_literal
    : interpolated_raw_string_literal_delimiter interpolated_raw_content
raw_string_literal_delimiter
    ;

multi_line_interpolated_raw_string_literal
    : interpolated_raw_string_literal_delimiter whitespace* new_line
(interpolated_raw_content | new_line)* new_line whitespace*
raw_string_literal_delimiter
    ;

interpolated_raw_content
    : (not_new_line | raw_interpolation)+
    ;

raw_interpolation
    : raw_interpolation_start interpolation raw_interpolation_end
    ;

raw_interpolation_start
    : {
    | {{
    | {{{
    | etc.
    ;

raw_interpolation_end
    : }
    | }}
    | }}}
    | etc.
    ;
```

The above is similar to the definition of `raw_string_literal` but with some important differences. A `interpolated_raw_string_literal` should be interpreted as:

1. It starts with at least one dollar sign (but no upper bound) and then three quotes (also with no upper bound).
2. It then continues with content on the same line as the starting quotes. This content on the same line can be blank, or non-blank. 'blank' is synonymous with 'entirely whitespace'.
3. If the content on that same line is non-blank no further content can follow. In other words the literal is required to end with the same number of quotes on that same line.
4. If the contents on the same line is blank, then the literal can continue with a `new_line` and some number of subsequent content lines and `new_line`s.

    - A content line is any text except a `new_line`.
    - A content line can contain multiple `raw_interpolation` occurrences at any position. The `raw_interpolation` must start with an equal number of open braces (`{`) as the number of dollar signs at the start of the literal.
    - If 'indentation whitespace' is not-empty, a `raw_interpolation` cannot immediately follow a `new_line`.
    - The `raw_interpolation` will following the normal rules specified at §11.7.3. Any `raw_interpolation` must end with the same number of close braces (`}`) as dollar signs and open braces.
    - Any `interpolation` can itself contain new-lines within in the same manner as an `interpolation` in a normal `verbatim_string_literal` (`@""`).
    - It then ends with a `new_line` some number (possibly zero) of `whitespace` and the same number of quotes that the literal started with.

Computation of the interpolated string value follows the same rules as a normal `raw_string_literal` except updated to handle lines containing `raw_interpolation`s. Building the string value happens in the same fashion, just with the interpolation holes replaced with whatever values those expressions produce at runtime. If the `interpolated_raw_string_literal` is converted to a `FormattableString` then the values of the interpolations are passed in their respective order to the `arguments` array to `FormattableString.Create`. The rest of the content of the `interpolated_raw_string_literal` *after* the 'indentation whitespace' has been stripped from all lines will be used to generate `format` string passed to `FormattableString.Create`, except with appropriately numbered `{N}` contents in each location where a `raw_interpolation` occurred (or `{N,constant}` in the case if its `interpolation` is of the form `expression ',' constant_expression`).

There is an ambiguity in the above specification. Specifically when a section of `{` in text and `{` of an interpolation abut. For example:

```
var v1 = $$"""
         {{{order_number}}}
         """
```

This could be interpreted as: `{{ {order_number } }}` or `{ {{order_number}} }`. However, as the former is illegal (no C# expression could start with `{`) it would be pointless to interpret that way. So we interpret in the latter fashion, where the innermost `{` and `}` braces form the interpolation, and any outermost ones form the text. In the future this might be an issue if the language ever supports any expressions that are surrounded by braces. However, in that case, the recommendation would be to write such a case like so: `{{{({some_new_expression_form})}}}`. Here, parentheses would help designate the expression portion from the rest of the literal/interpolation. This has precedence already with how ternary conditional expressions need to be wrapped to not conflict with the formatting/alignment specifier of an interpolation (e.g. `{(x ? y : z)}`).

Examples: (upcoming)

Drawbacks

Raw string literals add more complexity to the language. We already have many string literal forms already for numerous purposes. `""` strings, `@""` strings, and `$""` strings already have a lot of power and flexibility. But they all lack a way to provide raw contents that never need escaping.

The above rules do not support the case of 4.a:

4. ...

    - Optionally, with extra complexity, we could refine this to state that: For all strings that do not themselves contain a `new_line` (but can start or end with a quote `"` character), it should be possible to represent the string literal itself on a single line.

That's because we have no means to know that a starting or ending quote (`"`) should belong to the contents and not the delimiter itself. If this is an important scenario we want to support though, we can add a parallel `'''` construct to go along with the `"""` form. With that parallel construct, a single line string that start and ends with `"` can be written easily as `''''This string starts and ends with quotes"'''` along with the parallel construct `"""'This string starts and ends with apostrophes'"""`. This may also be desirable to support to help visually separate out quote characters, which may help

when embedding languages that primarily use one quote character much more than then other.

Alternatives

https://github.com/dotnet/csharplang/discussions/89 covers many options here. Alternatives are numerous, but i feel stray too far into complexity and poor ergonomics. This approach opts for simplicity where you just keep increasing the start/end quote length until there is no concern about a conflict with the string contents. It also allows the code you write to look well indented, while still producing a dedented literal that is what most code wants.

One of the most interesting potential variations though is the use of `` ` `` (or ```` ``` ````) fences for these raw string literals. This would have several benefits:

1. It would avoid all the issues with strings starting er ending with quotes.
2. It would look familiar to markdown. Though that in itself is potentially not a good thing as users might expect markdown interpretation.
3. A raw string literal would only have to start and end with a single character in most cases, and would only need multiple in the much rarer case of contents that contain back-ticks themselves.
4. It would feel natural to extend this in the future with ```` ```xml ````, again akin to markdown. Though, of course, that is also true of the `"""` form.

Overall though, the net benefit here seems small. In keeping with C# history, i think `"` should continue to be the `string literal` delimiter, just as it is for `@""` and `$""`.

Design meetings

~~Open issues to discuss~~ Resolved issues:

- [x] should we have a single line form? We technically could do without it. But it would mean simple strings not containing a newline would always take at least three lines. I think we should It's very heavyweight to force single line constructs to be three lines just to avoid escaping.

Design decision: Yes, we will have a single line form.

- [x] should we require that multiline *must* start with a newline? I think we should. It also gives us the ability to support things like `"""xml` in the future.

Design decision: Yes, we will require that multiline must start with a newline

- [x] should the automatic dedenting be done at all? I think we should. It makes code look so much more pleasant.

Design decision: Yes, automatic dedenting will be done.

- [x] should we restrict common-whitespace from mixing whitespace types? I don't think we should. Indeed, there is a common indentation strategy called "tab for indentation, space for alignment". It would be very natural to use this to align the end delimiter with the start delimiter in a case where the start delimiter doesn't start on a tab stop.

Design decision: We will not have any restrictions on mixing whitespace.

- [x] should we use something else for the fences? `` ` `` would match markdown syntax, and would mean we didn't need to always start these strings with three quotes. Just one would suffice for the common case.

Design decision: We will use `"""`

- [x] should we have a requirement that the delimiter have more quotes than the longest sequence of quotes in the string value? Technically it's not required. for example:

```
var v = """
        contents""""""
        """
```

This is a string with `"""` as the delimiter. Several community members have stated this is confusing and we should require in a case like this that the delimiter always have more characters. That would then be:

```
var v = """""""
        contents""""""
        """""""
```

Design decision: Yes, the delimiter must be longer than any sequence of quotes in the string itself.

Allow new-lines in all interpolations

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

- [x] Proposed
- [x] Implementation: https://github.com/dotnet/roslyn/pull/56853
- [x] Specification: this file.

Summary

The language today non-verbatim and verbatim interpolated strings (`$""` and `$@""` respectively). The primary *sensible* difference for these is that a non-verbatim interpolated string works like a normal string and cannot contain newlines in its text segments, and must instead use escapes (like `\r\n`). Conversely, a verbatim interpolated string can contain newlines in its text segments (like a verbatim string), and doesn't escape newlines or other character (except for `""` to escape a quote itself).

This is all reasonable and will not change with this proposal.

What is unreasonable today is that we extend the restriction on 'no newlines' in a non-verbatim interpolated string *beyond* its text segments into the *interpolations* themselves. This means, for example, that you cannot write the following:

```c#
var v = $"Count is\t: { this.Is.A.Really(long(expr))
                            .That.I.Should(
                                be + able)[
                                    to.Wrap()] }.";
```

Ultimately, the 'interpolation must be on a single line itself' rule is just a restriction of the current implementation. That restriction really isn't necessary, and can be annoying, and would be fairly trivial to remove (see work https://github.com/dotnet/roslyn/pull/54875 to show how). In the end, all it does is force the dev to place things on a single line, or force them into a verbatim interpolated string (both of which may be unpalatable).

The interpolation expressions themselves are not text, and shouldn't be beholden to any escaping/newline rules therin.

Specification change

diff

```
single_regular_balanced_text_character
-    : '<Any character except / (U+002F), @ (U+0040), \" (U+0022), $
(U+0024), ( (U+0028), ) (U+0029), [ (U+005B), ] (U+005D), { (U+007B), }
(U+007D) and new_line_character>'
-    | '</ (U+002F), if not directly followed by / (U+002F) or * (U+002A)>'
+    : <Any character except @ (U+0040), \" (U+0022), $ (U+0024), (
(U+0028), ) (U+0029), [ (U+005B), ] (U+005D), { (U+007B), } (U+007D)>
+    | comment
    ;
```

LDM Discussions

https://github.com/dotnet/csharplang/blob/main/meetings/2021/LDM-2021-09-20.md

Utf8 Strings Literals

Article • 09/30/2022

> **Note**
>
> This article is a feature specification. The specification serves as the design document for the feature. It includes proposed specification changes, along with information needed during the design and development of the feature. These articles are published until the proposed spec changes are finalized and incorporated in the current ECMA specification.
>
> There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent language design meeting (LDM) notes.
>
> You can learn more about the process for adopting feature speclets into the C# language standard in the article on the specifications.

Summary

This proposal adds the ability to write UTF8 string literals in C# and have them automatically encoded into their UTF-8 `byte` representation.

Motivation

UTF8 is the language of the web and its use is necessary in significant portions of the .NET stack. While much of data comes in the form of `byte[]` off the network stack there is still significant uses of constants in the code. For example networking stack has to commonly write constants like `"HTTP/1.0\r\n"`, `" AUTH"` or . `"Content-Length: "`.

Today there is no efficient syntax for doing this as C# represents all strings using UTF16 encoding. That means developers have to choose between the convenience of encoding at runtime which incurs overhead, including the time spent at startup actually performing the encoding operation (and allocations if targeting a type that doesn't actually require them), or manually translating the bytes and storing in a `byte[]`.

```
c#

// Efficient but verbose and error prone
static ReadOnlySpan<byte> AuthWithTrailingSpace => new byte[] { 0x41, 0x55,
```

```
0x54, 0x48, 0x20 };
WriteBytes(AuthWithTrailingSpace);

// Incurs allocation and startup costs performing an encoding that could
have been done at compile-time
static readonly byte[] s_authWithTrailingSpace =
Encoding.UTF8.GetBytes("AUTH ");
WriteBytes(s_authWithTrailingSpace);

// Simplest / most convenient but terribly inefficient
WriteBytes(Encoding.UTF8.GetBytes("AUTH "));
```

This trade off is a pain point that comes up frequently for our partners in the runtime, ASP.NET and Azure. Often times it causes them to leave performance on the table because they don't want to go through the hassle of writing out the `byte[]` encoding by hand.

To fix this we will allow for UTF8 literals in the language and encode them into the UTF8 `byte[]` at compile time.

Detailed design

`u8` suffix on string literals

The language will provide the `u8` suffix on string literals to force the type to be UTF8. The suffix is case-insensitive, `U8` suffix will be supported and will have the same meaning as `u8` suffix.

When the `u8` suffix is used, the value of the literal is a `ReadOnlySpan<byte>` containing a UTF-8 byte representation of the string. A null terminator is placed beyond the last byte in memory (and outside the length of the `ReadOnlySpan<byte>`) in order to handle some interop scenarios where the call expects null terminated strings.

c#

```
string s1 = "hello"u8;             // Error
var s2 = "hello"u8;                // Okay and type is ReadOnlySpan<byte>
ReadOnlySpan<byte> s3 = "hello"u8; // Okay.
byte[] s4 = "hello"u8;             // Error - Cannot implicitly convert type
'System.ReadOnlySpan<byte>' to 'byte[]'.
byte[] s5 = "hello"u8.ToArray();   // Okay.
Span<byte> s6 = "hello"u8;         // Error - Cannot implicitly convert type
'System.ReadOnlySpan<byte>' to 'System.Span<byte>'.
```

Since the literals would be allocated as global constants, the lifetime of the resulting `ReadOnlySpan<byte>` would not prevent it from being returned or passed around to elsewhere. However, certain contexts, most notably within async functions, do not allow locals of ref struct types, so there would be a usage penalty in those situations, with a `ToArray()` call or similar being required.

A `u8` literal doesn't have a constant value. That is because `ReadOnlySpan<byte>` cannot be type of a constant today. If the definition of `const` is expanded in the future to consider `ReadOnlySpan<byte>`, then this value should also be considered a constant. Practically though this means a `u8` literal cannot be used as the default value of an optional parameter.

```
c#

// Error: The argument is not constant
void Write(ReadOnlySpan<byte> message = "missing"u8) { ... }
```

When the input text for the literal is a malformed UTF16 string, then the language will emit an error:

```
c#

var bytes = "hello \uD801\uD802"u8; // Error: the input string is not valid
UTF16
```

Addition operator

A new bullet point will be added to §11.9.5 Addition operator as follows.

- UTF8 byte representation concatenation:

  ```
  C#

  ReadOnlySpan<byte> operator +(ReadOnlySpan<byte> x, ReadOnlySpan<byte>
  y);
  ```

  This binary `+` operator performs byte sequences concatenation and is applicable if and only if both operands are semantically UTF8 byte representations. An operand is semantically a UTF8 byte representation when it is either a value of a `u8` literal, or a value produced by the UTF8 byte representation concatenation operator.

> The result of the UTF8 byte representation concatenation is a `ReadOnlySpan<byte>` that consists of the bytes of the left operand followed by the bytes of the right operand. A null terminator is placed beyond the last byte in memory (and outside the length of the `ReadOnlySpan<byte>`) in order to handle some interop scenarios where the call expects null terminated strings.

Lowering

The language will lower the UTF8 encoded strings exactly as if the developer had typed the resulting `byte[]` literal in code. For example:

```
c#
```

```
ReadOnlySpan<byte> span = "hello"u8;

// Equivalent to

ReadOnlySpan<byte> span = new ReadOnlySpan<byte>(new byte[] { 0x68, 0x65,
0x6c, 0x6c, 0x6f, 0x00 }).
                                Slice(0,5); // The `Slice` call will be
optimized away by the compiler.
```

That means all optimizations that apply to the `new byte[] { ... }` form will apply to utf8 literals as well. This means the call site will be allocation free as C# will optimize this to be stored in the `.data` section of the PE file.

Multiple consecutive applications of UTF8 byte representation concatenation operators are collapsed into a single creation of `ReadOnlySpan<byte>` with byte array containing the final byte sequence.

```
c#
```

```
ReadOnlySpan<byte> span = "h"u8 + "el"u8 + "lo"u8;

// Equivalent to

ReadOnlySpan<byte> span = new ReadOnlySpan<byte>(new byte[] { 0x68, 0x65,
0x6c, 0x6c, 0x6f, 0x00 }).
                                Slice(0,5); // The `Slice` call will be
optimized away by the compiler.
```

Drawbacks

Relying on core APIs

The compiler implementation will use `UTF8Encoding` for both invalid string detection as well as translation to `byte[]`. The exact APIs will possibly depend on which target framework the compiler is using. But `UTF8Encoding` will be the workhorse of the implementation.

Historically the compiler has avoided using runtime APIs for literal processing. That is because it takes control of how constants are processed away from the language and into the runtime. Concretely it means items like bug fixes can change constant encoding and mean that the outcome of C# compilation depends on which runtime the compiler is executing on.

This is not a hypothetical problem. Early versions of Roslyn used `double.Parse` to handle floating point constant parsing. That caused a number of problems. First it meant that some floating point values had different representations between the native compiler and Roslyn. Second as .NET core evolved and fixed long standing bugs in the `double.Parse` code it meant that the meaning of those constants changed in the language depending on what runtime the compiler executed on. As a result the compiler ended up writing its own version of floating point parsing code and removing the dependency on `double.Parse`.

This scenario was discussed with the runtime team and we do not feel it has the same problems we've hit before. The UTF8 parsing is stable across runtimes and there are no known issues in this area that are areas for future compat concerns. If one does come up we can re-evaluate the strategy.

Alternatives

Target type only

The design could rely on target typing only and remove the `u8` suffix on `string` literals. In the majority of cases today the `string` literal is being assigned directly to a `ReadOnlySpan<byte>` hence it's unnecessary.

```
c#
```

```
ReadOnlySpan<byte> span = "Hello World;"
```

The `u8` suffix exists primarily to support two scenarios: `var` and overload resolution. For the latter consider the following use case:

```c#
void Write(ReadOnlySpan<byte> span) { ... }
void Write(string s) {
    var bytes = Encoding.Utf8.GetBytes(s);
    Write(bytes.AsSpan());
}
```

Given the implementation it is better to call `Write(ReadOnlySpan<byte>)` and the `u8` suffix makes this convenient: `Write("hello"u8)`. Lacking that developers need to resort to awkward casting `Write((ReadOnlySpan<byte>)"hello")`.

Still this is a convenience item, the feature can exist without it and it is non-breaking to add it at a later time.

Wait for Utf8String type

While the .NET ecosystem is standardizing on `ReadOnlySpan<byte>` as the defacto Utf8 string type today it's possible the runtime will introduce an actual `Utf8String` type is the future.

We should evaluate our design here in the face of this possible change and reflect on whether we'd regret the decisions we've made. This should be weighed though against the realistic probability we'll introduce `Utf8String`, a probability which seems to decrease every day we find `ReadOnlySpan<byte>` as an acceptable alternative.

It seems unlikely that we would regret the target type conversion between string literals and `ReadOnlySpan<byte>`. The use of `ReadOnlySpan<byte>` as utf8 is embedded in our APIs now and hence there is still value in the conversion even if `Utf8String` comes along and is a "better" type. The language could simply prefer conversions to `Utf8String` over `ReadOnlySpan<byte>`.

It seems more likely that we'd regret the `u8` suffix pointing to `ReadOnlySpan<byte>` instead of `Utf8String`. It would be similar to how we regret that `stackalloc int[]` has a natural type of `int*` instead of `Span<int>`. This is not a deal breaker though, just an inconvenience.

Conversions between `string` constants and `byte` sequences

The language will allow conversions between `string` constants and `byte` sequences where the text is converted into the equivalent UTF8 byte representation. Specifically the

compiler will allow *string_constant_to_UTF8_byte_representation_conversion* - implicit conversions from `string` constants to `byte[]`, `Span<byte>`, and `ReadOnlySpan<byte>`. A new bullet point will be added to the implicit conversions §10.2 section. This conversion is not a standard conversion §10.4.

```
c#
```

```
byte[] array = "hello";             // new byte[] { 0x68, 0x65, 0x6c, 0x6c,
0x6f }
Span<byte> span = "dog";            // new byte[] { 0x64, 0x6f, 0x67 }
ReadOnlySpan<byte> span = "cat";    // new byte[] { 0x63, 0x61, 0x74 }
```

When the input text for the conversion is a malformed UTF16 string then the language will emit an error:

```
c#
```

```
const string text = "hello \uD801\uD802";
byte[] bytes = text; // Error: the input string is not valid UTF16
```

The predominant usage of this feature is expected to be with literals but it will work with any `string` constant value. A conversion from a `string` constant with `null` value will be supported as well. The result of the conversion will be `default` value of the target type.

```
c#
```

```
const string data = "dog"
ReadOnlySpan<byte> span = data;     // new byte[] { 0x64, 0x6f, 0x67 }
```

In the case of any constant operation on strings, such as `+`, the encoding to UTF8 will occur on the final `string` vs. happening for the individual parts and then concatenating the results. This ordering is important to consider because it can impact whether or not the conversion succeeds.

```
c#
```

```
const string first = "\uD83D";  // high surrogate
const string second = "\uDE00"; // low surrogate
ReadOnlySpan<byte> span = first + second;
```

The two parts here are invalid on their own as they are incomplete portions of a surrogate pair. Individually there is no correct translation to UTF8 but together they form a complete surrogate pair that can be successfully translated to UTF8.

The *string_constant_to_UTF8_byte_representation_conversion* is not allowed in Linq Expression Trees.

While the inputs to these conversions are constants and the data is fully encoded at compile time, the conversion is **not** considered constant by the language. That is because arrays are not constant today. If the definition of `const` is expanded in the future to consider arrays then these conversions should also be considered. Practically though this means a result of these conversions cannot be used as the default value of an optional parameter.

```c#
// Error: The argument is not constant
void Write(ReadOnlySpan<byte> message = "missing") { ... }
```

Once implemented string literals will have the same problem that other literals have in the language: what type they represent depends on how they are used. C# provides a literal suffix to disambiguate the meaning for other literals. For example developers can write `3.14f` to force the value to be a `float` or `1l` to force the value to be a `long`.

Unresolved questions

(Resolved) Conversions between a `string` constant with `null` value and `byte` sequences

Whether this conversion is supported and, if so, how it is performed is not specified.

Proposal:

Allow implicit conversions from a `string` constant with `null` value to `byte[]`, `Span<byte>`, and `ReadOnlySpan<byte>`. The result of the conversion is `default` value of the target type.

Resolution:

The proposal is approved - https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-01-26.md#conversions-from-null-literals .

(Resolved) Where does *string_constant_to_UTF8_byte_representation_conversion*

belong?

Is *string_constant_to_UTF8_byte_representation_conversion* a bullet point in the implicit conversions §10.2 section on its own, or is it part of §10.2.11, or does it belong to some other existing implicit conversions group?

Proposal:

It is a new bullet point in implicit conversions §10.2, similar to "Implicit interpolated string conversions" or "Method group conversions". It doesn't feel like it belongs to "Implicit constant expression conversions" because, even though the source is a constant expression, the result is never a constant expression. Also, "Implicit constant expression conversions" are considered to be "Standard implicit conversions" §10.4.2, which is likely to lead to non-trivial behavior changes involving user-defined conversions.

Resolution:

We will introduce a new conversion kind for string constant to UTF-8 bytes - https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-01-26.md#conversion-kinds

(Resolved) Is *string_constant_to_UTF8_byte_representation_conversion* a standard conversion

In addition to "pure" Standard Conversions (the standard conversions are those pre-defined conversions that can occur as part of a user-defined conversion), compiler also treats some predefined conversions as "somewhat" standard. For example, an implicit interpolated string conversion can occur as part of a user-defined conversion if there is an explicit cast to the target type in code. As if it is a Standard Explicit Conversion, even though it is an implicit conversion not explicitly included into the set of standard implicit or explicit conversions. For example:

C#

```
class C
{
    static void Main()
    {
        C1 x = $"hello"; // error CS0266: Cannot implicitly convert type
'string' to 'C1'. An explicit conversion exists (are you missing a cast?)
        var y = (C1)$"dog"; // works
    }
```

```
}

class C1
{
    public static implicit operator C1(System.FormattableString x) => new
C1();
}
```

Proposal:

The new conversion is not a standard conversion. This will avoid non-trivial behavior changes involving user-defined conversions. For example, we won't need to worry about user-defined cinversions under implicit tuple literal conversions, etc.

Resolution:

Not a standard conversion, for now - https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-01-26.md#implicit-standard-conversion .

(Resolved) Linq Expression Tree conversion

Should *string_constant_to_UTF8_byte_representation_conversion* be allowed in context of a Linq Expression Tree conversion? We can disallow it for now, or we could simply include the "lowered" form into the tree. For example:

```C#
Expression<Func<byte[]>> x = () => "hello";           // () => new [] {104,
101, 108, 108, 111}
Expression<FuncSpanOfByte> y = () => "dog";           // () => new
Span`1(new [] {100, 111, 103})
Expression<FuncReadOnlySpanOfByte> z = () => "cat";   // () => new
ReadOnlySpan`1(new [] {99, 97, 116})
```

What about string literals with `u8` suffix? We could surface those as byte array creations:

```C#
Expression<Func<byte[]>> x = () => "hello"u8;           // () => new []
{104, 101, 108, 108, 111}
```

Resolution:

Disallow in Linq Expression Trees - https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-01-

26.md#expression-tree-representation.

(Resolved) The natural type of a string literal with `u8` suffix

The "Detailed design" section says: "The natural type though will be `ReadOnlySpan<byte>`." At the same time: "When the `u8` suffix is used the literal can still be converted to any of the allowed types: `byte[]`, `Span<byte>` or `ReadOnlySpan<byte>`."

There are several disadvantages with this approach:

- `ReadOnlySpan<byte>` is not available on desktop framework;
- There are no existing conversions from `ReadOnlySpan<byte>` to `byte[]` or `Span<byte>`. In order to support them we will likely need to treat the literals as target typed. Both the language rules and implementation will become more complicated.

Proposal:

The natural type will be `byte[]`. It is readily available on all frameworks. BTW, at runtime we will always be starting with creating a byte array, even with the original proposal. We also don't need any special conversion rules to support conversions to `Span<byte>` and `ReadOnlySpan<byte>`. There are already implicit user-defined conversions from `byte[]` to `Span<byte>` and `ReadOnlySpan<byte>`. There is even implicit user-defined conversion to `ReadOnlyMemory<byte>` (see the "Depth of the conversion" question below). There is a disadvantage, language doesn't allow chaining user-defined conversions. So, the following code will not compile:

C#

```
using System;
class C
{
    static void Main()
    {
        var y = (C2)"dog"u8; // error CS0030: Cannot convert type 'byte[]'
to 'C2'
        var z = (C3)"cat"u8; // error CS0030: Cannot convert type 'byte[]'
to 'C3'
    }
}

class C2
{
    public static implicit operator C2(Span<byte> x) => new C2();
}
```

```
class C3
{
    public static explicit operator C3(ReadOnlySpan<byte> x) => new C3();
}
```

However, as with any user-defined conversion, an explicit cast can be used to make one user-defined conversion a part of another user-defined conversion.

It feels like all motivating scenarios are going to be addressed with `byte[]` as the natural type, but the language rules and implementation will be significantly simpler.

Resolution:

The proposal is approved - https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-01-26.md#natural-type-of-u8-literals. We will likely want to have a deeper debate about whether `u8` string literals should have a type of a mutable array, but we don't think that debate is necessary for now.

(Resolved) Depth of the conversion

Will it also work anywhere that a byte[] could work? Consider:

```c#
static readonly ReadOnlyMemory<byte> s_data1 = "Data"u8;
static readonly ReadOnlyMemory<byte> s_data2 = "Data";
```

The first example likely should work because of the natural type that comes from `u8`.

The second example is hard to make work because it requires conversions in both directions. That is unless we add `ReadOnlyMemory<byte>` as one of the allowed conversion types.

Proposal:

Don't do anything special.

Resolution:

No new conversion targets added for now https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-01-26.md#conversion-depth.

(Resolved) Overload resolution breaks

The following API would become ambiguous:

```c#
M("");
static void M1(ReadOnlySpan<char> charArray) => ...;
static void M1(byte[] byteArray) => ...;
```

What should we do to address this?

Proposal:

Similar to https://github.com/dotnet/csharplang/blob/main/proposals/csharp-10.0/lambda-improvements.md#overload-resolution, Better function member (§11.6.4.3) is updated to prefer members where none of the conversions involved require converting `string` constants to UTF8 `byte` sequences.

> ### Better function member
>
> ... Given an argument list `A` with a set of argument expressions `{E1, E2, ..., En}` and two applicable function members `Mp` and `Mq` with parameter types `{P1, P2, ..., Pn}` and `{Q1, Q2, ..., Qn}`, `Mp` is defined to be a ***better function member*** than `Mq` if
>
> 1. **for each argument, the implicit conversion from `Ex` to `Px` is not a *string_constant_to_UTF8_byte_representation_conversion*, and for at least one argument, the implicit conversion from `Ex` to `Qx` is a *string_constant_to_UTF8_byte_representation_conversion*, or**
> 2. for each argument, the implicit conversion from `Ex` to `Px` is not a *function_type_conversion*, and
>
>   - `Mp` is a non-generic method or `Mp` is a generic method with type parameters `{X1, X2, ..., Xp}` and for each type parameter `Xi` the type argument is inferred from an expression or from a type other than a *function_type*, and
>   - for at least one argument, the implicit conversion from `Ex` to `Qx` is a *function_type_conversion*, or `Mq` is a generic method with type parameters `{Y1, Y2, ..., Yq}` and for at least one type parameter `Yi` the type argument is inferred from a *function_type*, or

> 3. for each argument, the implicit conversion from `Ex` to `Qx` is not better than the implicit conversion from `Ex` to `Px`, and for at least one argument, the conversion from `Ex` to `Px` is better than the conversion from `Ex` to `Qx`.

Note that the addition of this rule is not going to cover scenarios with instance methods becoming applicable and "shadowing" extension methods. For example:

```C#
using System;

class Program
{
    static void Main()
    {
        var p = new Program();
        Console.WriteLine(p.M(""));
    }

    public string M(byte[] b) => "byte[]";
}

static class E
{
    public static string M(this object o, string s) => "string";
}
```

Behavior of this code will silently change from printing "string" to printing "byte[]".

Are we Ok with this behavior change? Should it be documented as a breaking change?

Note that there is no proposal to make *string_constant_to_UTF8_byte_representation_conversion* unavailable when C#10 language version is targeted. In that case, the example above becomes an error rather than returns to C#10 behavior. This follows a general principle that target language version doesn't affect semantics of the language.

Are we Ok with this behavior? Should it be documented as a breaking change?

The new rule also is not going to prevent breaks involving tuple literal conversions. For example,

```C#
class C
{
    static void Main()
    {
```

```
        System.Console.Write(Test(("s", 1)));
    }

    static string Test((object, int) a) => "object";
    static string Test((byte[], int) a) => "array";
}
```

is going to silently print "array" instead of "object".

Are we Ok with this behavior? Should it be documented as a breaking change? Perhaps we could complicate the new rule to dig into the tuple literal conversions.

Resolution:

The prototype will not adjust any rules here, so we can hopefully see what breaks in practice - https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-01-26.md#breaking-changes .

(Resolved) Should `u8` suffix be case-insensitive?

Proposal:

Support `U8` suffix as well for consistency with numeric suffixes.

Resolution:

Approved - https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-01-26.md#suffix-case-sensitivity .

Examples today

Examples of where runtime has manually encoded the UTF8 bytes today

- https://github.com/dotnet/runtime/blob/e095fde94baa480a6d65dfdee43d9cc0ad0d0b38/src/libraries/Common/src/System/Net/Http/aspnetcore/Http2/Hpack/StatusCodes.cs#L13-L78
- https://github.com/dotnet/runtime/blob/e095fde94baa480a6d65dfdee43d9cc0ad0d0b38/src/libraries/System.Memory/src/System/Buffers/Text/Base64Encoder.cs#L581-L591
- https://github.com/dotnet/runtime/blob/e095fde94baa480a6d65dfdee43d9cc0ad0d0b38/src/libraries/System.Net.HttpListener/src/System/Net/Windows/HttpResponseStream.Windows.cs#L284
- https://github.com/dotnet/runtime/blob/e095fde94baa480a6d65dfdee43d9cc0ad0d0b38/src/libraries/System.Net.Http/src/System/Net/Http/SocketsHttpHandler/Ht

tp2Stream.cs#L30

- https://github.com/dotnet/runtime/blob/e095fde94baa480a6d65dfdee43d9cc0ad0d0b38/src/libraries/System.Net.Http/src/System/Net/Http/SocketsHttpHandler/Http3RequestStream.cs#L852
- https://github.com/dotnet/runtime/blob/e095fde94baa480a6d65dfdee43d9cc0ad0d0b38/src/libraries/System.Text.Json/src/System/Text/Json/JsonConstants.cs#L35-L42

Examples where we leave perf on the table

- https://github.com/dotnet/runtime/blob/e095fde94baa480a6d65dfdee43d9cc0ad0d0b38/src/libraries/System.Net.Security/src/System/Net/Security/Pal.Managed/SafeChannelBindingHandle.cs#L16-L17
- https://github.com/dotnet/runtime/blob/e095fde94baa480a6d65dfdee43d9cc0ad0d0b38/src/libraries/System.Net.Http/src/System/Net/Http/SocketsHttpHandler/HttpConnection.cs#L37-L43
- https://github.com/dotnet/runtime/blob/e095fde94baa480a6d65dfdee43d9cc0ad0d0b38/src/libraries/System.Net.Http/src/System/Net/Http/SocketsHttpHandler/Http2Connection.cs#L78
- https://github.com/dotnet/runtime/blob/e095fde94baa480a6d65dfdee43d9cc0ad0d0b38/src/libraries/System.Net.Mail/src/System/Net/Mail/SmtpCommands.cs#L669-L687

Design meetings

https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-01-26.md https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-04-18.md https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-06-06.md

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# feature specifications feedback

C# feature specifications is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Native-sized integers

Article • 07/22/2022

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Language support for a native-sized signed and unsigned integer types.

The motivation is for interop scenarios and for low-level libraries.

Design

The identifiers `nint` and `nuint` are new contextual keywords that represent native signed and unsigned integer types. The identifiers are only treated as keywords when name lookup does not find a viable result at that program location.

```
C#
```

```
nint x = 3;
string y = nameof(nuint);
_ = nint.Equals(x, 3);
```

The types `nint` and `nuint` are represented by the underlying types `System.IntPtr` and `System.UIntPtr` with compiler surfacing additional conversions and operations for those types as native ints.

Constants

Constant expressions may be of type `nint` or `nuint`. There is no direct syntax for native int literals. Implicit or explicit casts of other integral constant values can be used instead:

`const nint i = (nint)42;`.

`nint` constants are in the range [`int.MinValue`, `int.MaxValue`].

`nuint` constants are in the range [`uint.MinValue`, `uint.MaxValue`].

There are no `MinValue` or `MaxValue` fields on `nint` or `nuint` because, other than `nuint.MinValue`, those values cannot be emitted as constants.

Constant folding is supported for all unary operators { `+`, `-`, `~` } and binary operators { `+`, `-`, `*`, `/`, `%`, `==`, `!=`, `<`, `<=`, `>`, `>=`, `&`, `|`, `^`, `<<`, `>>` }. Constant folding operations are evaluated with `Int32` and `UInt32` operands rather than native ints for consistent behavior regardless of compiler platform. If the operation results in a constant value in 32-bits, constant folding is performed at compile-time. Otherwise the operation is executed at runtime and not considered a constant.

Conversions

There is an identity conversion between `nint` and `IntPtr`, and between `nuint` and `UIntPtr`. There is an identity conversion between compound types that differ by native ints and underlying types only: arrays, `Nullable<>`, constructed types, and tuples.

The tables below cover the conversions between special types. (The IL for each conversion includes the variants for `unchecked` and `checked` contexts if different.)

General notes on the table below:

- `conv.u` is a zero-extending conversion to native integer and `conv.i` is sign-extending conversion to native integer.
- `checked` contexts for both **widening** and **narrowing** are:
  - `conv.ovf.*` for `signed to *`
  - `conv.ovf.*.un` for `unsigned to *`
- `unchecked` contexts for **widening** are:
  - `conv.i*` for `signed to *` (where * is the target width)
  - `conv.u*` for `unsigned to *` (where * is the target width)
- `unchecked` contexts for **narrowing** are:
  - `conv.i*` for `any to signed *` (where * is the target width)
  - `conv.u*` for `any to unsigned *` (where * is the target width)

Taking a few examples:

- `sbyte to nint` and `sbyte to nuint` use `conv.i` while `byte to nint` and `byte to nuint` use `conv.u` because they are all **widening**.
- `nint to byte` and `nuint to byte` use `conv.u1` while `nint to sbyte` and `nuint to sbyte` use `conv.i1`. For `byte`, `sbyte`, `short`, and `ushort` the "stack type" is `int32`. So `conv.i1` is effectively "downcast to a signed byte and then sign-extend up to int32" while `conv.u1` is effectively "downcast to an unsigned byte and then zero-extend up to int32".
- `checked void* to nint` uses `conv.ovf.i.un` the same way that `checked void* to long` uses `conv.ovf.i8.un`.

Operand	Target	Conversion	IL
`object`	`nint`	Unboxing	`unbox`
`void*`	`nint`	PointerToVoid	nop / `conv.ovf.i.un`
`sbyte`	`nint`	ImplicitNumeric	`conv.i`
`byte`	`nint`	ImplicitNumeric	`conv.u`
`short`	`nint`	ImplicitNumeric	`conv.i`
`ushort`	`nint`	ImplicitNumeric	`conv.u`
`int`	`nint`	ImplicitNumeric	`conv.i`
`uint`	`nint`	ExplicitNumeric	`conv.u` / `conv.ovf.i.un`
`long`	`nint`	ExplicitNumeric	`conv.i` / `conv.ovf.i`
`ulong`	`nint`	ExplicitNumeric	`conv.i` / `conv.ovf.i.un`
`char`	`nint`	ImplicitNumeric	`conv.u`
`float`	`nint`	ExplicitNumeric	`conv.i` / `conv.ovf.i`
`double`	`nint`	ExplicitNumeric	`conv.i` / `conv.ovf.i`
`decimal`	`nint`	ExplicitNumeric	`long decimal.op_Explicit(decimal) conv.i` / `... conv.ovf.i`
`IntPtr`	`nint`	Identity	
`UIntPtr`	`nint`	None	
`object`	`nuint`	Unboxing	`unbox`
`void*`	`nuint`	PointerToVoid	nop

Operand	Target	Conversion	IL
`sbyte`	`nuint`	ExplicitNumeric	`conv.i` / `conv.ovf.u`
`byte`	`nuint`	ImplicitNumeric	`conv.u`
`short`	`nuint`	ExplicitNumeric	`conv.i` / `conv.ovf.u`
`ushort`	`nuint`	ImplicitNumeric	`conv.u`
`int`	`nuint`	ExplicitNumeric	`conv.i` / `conv.ovf.u`
`uint`	`nuint`	ImplicitNumeric	`conv.u`
`long`	`nuint`	ExplicitNumeric	`conv.u` / `conv.ovf.u`
`ulong`	`nuint`	ExplicitNumeric	`conv.u` / `conv.ovf.u.un`
`char`	`nuint`	ImplicitNumeric	`conv.u`
`float`	`nuint`	ExplicitNumeric	`conv.u` / `conv.ovf.u`
`double`	`nuint`	ExplicitNumeric	`conv.u` / `conv.ovf.u`
`decimal`	`nuint`	ExplicitNumeric	`ulong decimal.op_Explicit(decimal) conv.u` / `... conv.ovf.u.un`
`IntPtr`	`nuint`	None	
`UIntPtr`	`nuint`	Identity	
Enumeration	`nint`	ExplicitEnumeration	
Enumeration	`nuint`	ExplicitEnumeration	

Operand	Target	Conversion	IL
`nint`	`object`	Boxing	`box`
`nint`	`void*`	PointerToVoid	nop / `conv.ovf.u`
`nint`	`nuint`	ExplicitNumeric	`conv.u` (can be omitted) / `conv.ovf.u`
`nint`	`sbyte`	ExplicitNumeric	`conv.i1` / `conv.ovf.i1`
`nint`	`byte`	ExplicitNumeric	`conv.u1` / `conv.ovf.u1`
`nint`	`short`	ExplicitNumeric	`conv.i2` / `conv.ovf.i2`
`nint`	`ushort`	ExplicitNumeric	`conv.u2` / `conv.ovf.u2`
`nint`	`int`	ExplicitNumeric	`conv.i4` / `conv.ovf.i4`

Operand	Target	Conversion	IL
`nint`	`uint`	ExplicitNumeric	`conv.u4` / `conv.ovf.u4`
`nint`	`long`	ImplicitNumeric	`conv.i8`
`nint`	`ulong`	ExplicitNumeric	`conv.i8` / `conv.ovf.u8`
`nint`	`char`	ExplicitNumeric	`conv.u2` / `conv.ovf.u2`
`nint`	`float`	ImplicitNumeric	`conv.r4`
`nint`	`double`	ImplicitNumeric	`conv.r8`
`nint`	`decimal`	ImplicitNumeric	`conv.i8 decimal decimal.op_Implicit(long)`
`nint`	`IntPtr`	Identity	
`nint`	`UIntPtr`	None	
`nint`	Enumeration	ExplicitEnumeration	
`nuint`	`object`	Boxing	`box`
`nuint`	`void*`	PointerToVoid	nop
`nuint`	`nint`	ExplicitNumeric	`conv.i`(can be omitted) / `conv.ovf.i.un`
`nuint`	`sbyte`	ExplicitNumeric	`conv.i1` / `conv.ovf.i1.un`
`nuint`	`byte`	ExplicitNumeric	`conv.u1` / `conv.ovf.u1.un`
`nuint`	`short`	ExplicitNumeric	`conv.i2` / `conv.ovf.i2.un`
`nuint`	`ushort`	ExplicitNumeric	`conv.u2` / `conv.ovf.u2.un`
`nuint`	`int`	ExplicitNumeric	`conv.i4` / `conv.ovf.i4.un`
`nuint`	`uint`	ExplicitNumeric	`conv.u4` / `conv.ovf.u4.un`
`nuint`	`long`	ExplicitNumeric	`conv.u8` / `conv.ovf.i8.un`
`nuint`	`ulong`	ImplicitNumeric	`conv.u8`
`nuint`	`char`	ExplicitNumeric	`conv.u2` / `conv.ovf.u2.un`
`nuint`	`float`	ImplicitNumeric	`conv.r.un conv.r4`
`nuint`	`double`	ImplicitNumeric	`conv.r.un conv.r8`
`nuint`	`decimal`	ImplicitNumeric	`conv.u8 decimal decimal.op_Implicit(ulong)`

Operand	Target	Conversion	IL
`nuint`	`IntPtr`	None	
`nuint`	`UIntPtr`	Identity	
`nuint`	Enumeration	ExplicitEnumeration	

Conversion from `A` to `Nullable<B>` is:

- an implicit nullable conversion if there is an identity conversion or implicit conversion from `A` to `B`;
- an explicit nullable conversion if there is an explicit conversion from `A` to `B`;
- otherwise invalid.

Conversion from `Nullable<A>` to `B` is:

- an explicit nullable conversion if there is an identity conversion or implicit or explicit numeric conversion from `A` to `B`;
- otherwise invalid.

Conversion from `Nullable<A>` to `Nullable<B>` is:

- an identity conversion if there is an identity conversion from `A` to `B`;
- an explicit nullable conversion if there is an implicit or explicit numeric conversion from `A` to `B`;
- otherwise invalid.

Operators

The predefined operators are as follows. These operators are considered during overload resolution based on normal rules for implicit conversions *if at least one of the operands is of type* `nint` *or* `nuint`.

(The IL for each operator includes the variants for `unchecked` and `checked` contexts if different.)

Unary	Operator Signature	IL
`+`	`nint operator +(nint value)`	`nop`
`+`	`nuint operator +(nuint value)`	`nop`
`-`	`nint operator -(nint value)`	`neg`

Unary	Operator Signature	IL
`~`	`nint operator ~(nint value)`	`not`
`~`	`nuint operator ~(nuint value)`	`not`

Binary	Operator Signature	IL
`+`	`nint operator +(nint left, nint right)`	`add` / `add.ovf`
`+`	`nuint operator +(nuint left, nuint right)`	`add` / `add.ovf.un`
`-`	`nint operator -(nint left, nint right)`	`sub` / `sub.ovf`
`-`	`nuint operator -(nuint left, nuint right)`	`sub` / `sub.ovf.un`
`*`	`nint operator *(nint left, nint right)`	`mul` / `mul.ovf`
`*`	`nuint operator *(nuint left, nuint right)`	`mul` / `mul.ovf.un`
`/`	`nint operator /(nint left, nint right)`	`div`
`/`	`nuint operator /(nuint left, nuint right)`	`div.un`
`%`	`nint operator %(nint left, nint right)`	`rem`
`%`	`nuint operator %(nuint left, nuint right)`	`rem.un`
`==`	`bool operator ==(nint left, nint right)`	`beq` / `ceq`
`==`	`bool operator ==(nuint left, nuint right)`	`beq` / `ceq`
`!=`	`bool operator !=(nint left, nint right)`	`bne`
`!=`	`bool operator !=(nuint left, nuint right)`	`bne`
`<`	`bool operator <(nint left, nint right)`	`blt` / `clt`
`<`	`bool operator <(nuint left, nuint right)`	`blt.un` / `clt.un`
`<=`	`bool operator <=(nint left, nint right)`	`ble`
`<=`	`bool operator <=(nuint left, nuint right)`	`ble.un`
`>`	`bool operator >(nint left, nint right)`	`bgt` / `cgt`
`>`	`bool operator >(nuint left, nuint right)`	`bgt.un` / `cgt.un`
`>=`	`bool operator >=(nint left, nint right)`	`bge`
`>=`	`bool operator >=(nuint left, nuint right)`	`bge.un`

Binary	Operator Signature	IL
`&`	`nint operator &(nint left, nint right)`	`and`
`&`	`nuint operator &(nuint left, nuint right)`	`and`
`|`	`nint operator |(nint left, nint right)`	`or`
`|`	`nuint operator |(nuint left, nuint right)`	`or`
`^`	`nint operator ^(nint left, nint right)`	`xor`
`^`	`nuint operator ^(nuint left, nuint right)`	`xor`
`<<`	`nint operator <<(nint left, int right)`	`shl`
`<<`	`nuint operator <<(nuint left, int right)`	`shl`
`>>`	`nint operator >>(nint left, int right)`	`shr`
`>>`	`nuint operator >>(nuint left, int right)`	`shr.un`

For some binary operators, the IL operators support additional operand types (see ECMA-335 III.1.5 Operand type table). But the set of operand types supported by C# is limited for simplicity and for consistency with existing operators in the language.

Lifted versions of the operators, where the arguments and return types are `nint?` and `nuint?`, are supported.

Compound assignment operations `x op= y` where `x` or `y` are native ints follow the same rules as with other primitive types with pre-defined operators. Specifically the expression is bound as `x = (T)(x op y)` where `T` is the type of `x` and where `x` is only evaluated once.

The shift operators should mask the number of bits to shift - to 5 bits if `sizeof(nint)` is 4, and to 6 bits if `sizeof(nint)` is 8. (see §11.10) in C# spec).

The C#9 compiler will report errors binding to predefined native integer operators when compiling with an earlier language version, but will allow use of predefined conversions to and from native integers.

`csc -langversion:9 -t:library A.cs`

C#

```
public class A
{
```

```
    public static nint F;
}
```

`csc -langversion:8 -r:A.dll B.cs`

C#

```
class B : A
{
    static void Main()
    {
        F = F + 1; // error: nint operator+ not available with -
langversion:8
        F = (System.IntPtr)F + 1; // ok
    }
}
```

Pointer arithmetic

There are no predefined operators in C# for pointer addition or subtraction with native integer offsets. Instead, `nint` and `nuint` values are promoted to `long` and `ulong` and pointer arithmetic uses predefined operators for those types.

C#

```
static T* AddLeftS(nint x, T* y) => x + y;   // T* operator +(long left, T*
right)
static T* AddLeftU(nuint x, T* y) => x + y;  // T* operator +(ulong left, T*
right)
static T* AddRightS(T* x, nint y) => x + y;  // T* operator +(T* left, long
right)
static T* AddRightU(T* x, nuint y) => x + y; // T* operator +(T* left, ulong
right)
static T* SubRightS(T* x, nint y) => x - y;  // T* operator -(T* left, long
right)
static T* SubRightU(T* x, nuint y) => x - y; // T* operator -(T* left, ulong
right)
```

Binary numeric promotions

The *binary numeric promotions* informative text (see §11.4.7.3) in C# spec) is updated as follows:

> - ...

- Otherwise, if either operand is of type `ulong`, the other operand is converted to type `ulong`, or a binding-time error occurs if the other operand is of type `sbyte`, `short`, `int`, `nint`, or `long`.
- **Otherwise, if either operand is of type `nuint`, the other operand is converted to type `nuint`, or a binding-time error occurs if the other operand is of type `sbyte`, `short`, `int`, `nint`, or `long`.**
- Otherwise, if either operand is of type `long`, the other operand is converted to type `long`.
- Otherwise, if either operand is of type `uint` and the other operand is of type `sbyte`, `short`, `nint`, or `int`, both operands are converted to type `long`.
- Otherwise, if either operand is of type `uint`, the other operand is converted to type `uint`.
- **Otherwise, if either operand is of type `nint`, the other operand is converted to type `nint`.**
- Otherwise, both operands are converted to type `int`.

Dynamic

The conversions and operators are synthesized by the compiler and are not part of the underlying `IntPtr` and `UIntPtr` types. As a result those conversions and operators *are not available* from the runtime binder for `dynamic`.

C#

```
nint x = 2;
nint y = x + x; // ok
dynamic d = x;
nint z = d + x; // RuntimeBinderException: '+' cannot be applied
'System.IntPtr' and 'System.IntPtr'
```

Type members

The only constructor for `nint` or `nuint` is the parameter-less constructor.

The following members of `System.IntPtr` and `System.UIntPtr` *are explicitly excluded* from `nint` or `nuint`:

C#

```
// constructors
// arithmetic operators
```

```
// implicit and explicit conversions
public static readonly IntPtr Zero; // use 0 instead
public static int Size { get; }     // use sizeof() instead
public static IntPtr Add(IntPtr pointer, int offset);
public static IntPtr Subtract(IntPtr pointer, int offset);
public int ToInt32();
public long ToInt64();
public void* ToPointer();
```

The remaining members of `System.IntPtr` and `System.UIntPtr` *are implicitly included* in `nint` and `nuint`. For .NET Framework 4.7.2:

C#

```
public override bool Equals(object obj);
public override int GetHashCode();
public override string ToString();
public string ToString(string format);
```

Interfaces implemented by `System.IntPtr` and `System.UIntPtr` *are implicitly included* in `nint` and `nuint`, with occurrences of the underlying types replaced by the corresponding native integer types. For instance if `IntPtr` implements `ISerializable, IEquatable<IntPtr>, IComparable<IntPtr>`, then `nint` implements `ISerializable, IEquatable<nint>, IComparable<nint>`.

Overriding, hiding, and implementing

`nint` and `System.IntPtr`, and `nuint` and `System.UIntPtr`, are considered equivalent for overriding, hiding, and implementing.

Overloads cannot differ by `nint` and `System.IntPtr`, and `nuint` and `System.UIntPtr`, alone. Overrides and implementations may differ by `nint` and `System.IntPtr`, or `nuint` and `System.UIntPtr`, alone. Methods hide other methods that differ by `nint` and `System.IntPtr`, or `nuint` and `System.UIntPtr`, alone.

Miscellaneous

`nint` and `nuint` expressions used as array indices are emitted without conversion.

C#

```
static object GetItem(object[] array, nint index)
{
```

```
    return array[index]; // ok
}
```

`nint` and `nuint` cannot be used as an `enum` base type from C#.

C#

```
enum E : nint // error: byte, sbyte, short, ushort, int, uint, long, or
ulong expected
{
}
```

Reads and writes are atomic for `nint` and `nuint`.

Fields may be marked `volatile` for types `nint` and `nuint`. ECMA-334 15.5.4 does not include `enum` with base type `System.IntPtr` or `System.UIntPtr` however.

`default(nint)` and `new nint()` are equivalent to `(nint)0`; `default(nuint)` and `new nuint()` are equivalent to `(nuint)0`.

`typeof(nint)` is `typeof(IntPtr)`; `typeof(nuint)` is `typeof(UIntPtr)`.

`sizeof(nint)` and `sizeof(nuint)` are supported but require compiling in an unsafe context (as required for `sizeof(IntPtr)` and `sizeof(UIntPtr)`). The values are not compile-time constants. `sizeof(nint)` is implemented as `sizeof(IntPtr)` rather than `IntPtr.Size`; `sizeof(nuint)` is implemented as `sizeof(UIntPtr)` rather than `UIntPtr.Size`.

Compiler diagnostics for type references involving `nint` or `nuint` report `nint` or `nuint` rather than `IntPtr` or `UIntPtr`.

Metadata

`nint` and `nuint` are represented in metadata as `System.IntPtr` and `System.UIntPtr`.

Type references that include `nint` or `nuint` are emitted with a `System.Runtime.CompilerServices.NativeIntegerAttribute` to indicate which parts of the type reference are native ints.

C#

```
namespace System.Runtime.CompilerServices
{
    [AttributeUsage(
```

```
        AttributeTargets.Class |
        AttributeTargets.Event |
        AttributeTargets.Field |
        AttributeTargets.GenericParameter |
        AttributeTargets.Parameter |
        AttributeTargets.Property |
        AttributeTargets.ReturnValue,
        AllowMultiple = false,
        Inherited = false)]
    public sealed class NativeIntegerAttribute : Attribute
    {
        public NativeIntegerAttribute()
        {
            TransformFlags = new[] { true };
        }
        public NativeIntegerAttribute(bool[] flags)
        {
            TransformFlags = flags;
        }
        public readonly bool[] TransformFlags;
    }
}
```

The encoding of type references with `NativeIntegerAttribute` is covered in NativeIntegerAttribute.md.

Alternatives

An alternative to the "type erasure" approach above is to introduce new types: `System.NativeInt` and `System.NativeUInt`.

C#

```
public readonly struct NativeInt
{
    public IntPtr Value;
}
```

Distinct types would allow overloading distinct from `IntPtr` and would allow distinct parsing and `ToString()`. But there would be more work for the CLR to handle these types efficiently which defeats the primary purpose of the feature - efficiency. And interop with existing native int code that uses `IntPtr` would be more difficult.

Another alternative is to add more native int support for `IntPtr` in the framework but without any specific compiler support. Any new conversions and arithmetic operations

would be supported by the compiler automatically. But the language would not provide keywords, constants, or `checked` operations.

Design meetings

- https://github.com/dotnet/csharplang/blob/master/meetings/2017/LDM-2017-05-26.md
- https://github.com/dotnet/csharplang/blob/master/meetings/2017/LDM-2017-06-13.md
- https://github.com/dotnet/csharplang/blob/master/meetings/2017/LDM-2017-07-05.md#native-int-and-intptr-operators
- https://github.com/dotnet/csharplang/blob/master/meetings/2019/LDM-2019-10-23.md
- https://github.com/dotnet/csharplang/blob/master/meetings/2020/LDM-2020-03-25.md

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# feature specification feedback

The C# feature specifications are open source. Provide feedback here.

Open a documentation issue

Provide product feedback

Numeric IntPtr

Article • 09/30/2022

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

This is a revision on the initial native integers feature (spec), where the `nint`/`nuint` types were distinct from the underlying types `System.IntPtr`/`System.UIntPtr`. In short, we now treat `nint`/`nuint` as simple types aliasing `System.IntPtr`/`System.UIntPtr`, like we do for `int` in relation to `System.Int32`. The `System.Runtime.CompilerServices.RuntimeFeature.NumericIntPtr` runtime feature flag triggers this new behavior.

Design

8.3.5 Simple types

C# provides a set of predefined `struct` types called the simple types. The simple types are identified through keywords, but these keywords are simply aliases for predefined `struct` types in the `System` namespace, as described in the table below.

Keyword	Aliased type
`sbyte`	`System.SByte`
`byte`	`System.Byte`
`short`	`System.Int16`
`ushort`	`System.UInt16`

Keyword	Aliased type
`int`	`System.Int32`
`uint`	`System.UInt32`
`nint`	**`System.IntPtr`**
`nuint`	**`System.UIntPtr`**
`long`	`System.Int64`
`ulong`	`System.UInt64`
`char`	`System.Char`
`float`	`System.Single`
`double`	`System.Double`
`bool`	`System.Boolean`
`decimal`	`System.Decimal`

[...]

8.3.6 Integral types

C# supports **eleven** integral types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, and `char`. [...]

8.8 Unmanaged types

In other words, an *unmanaged_type* is one of the following:

- `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `char`, `float`, `double`, `decimal`, or `bool`.
- Any *enum_type*.
- Any user-defined *struct_type* that is not a constructed type and contains fields of *unmanaged_type*s only.
- In unsafe code, any *pointer_type*.

10.2.3 Implicit numeric conversions

The implicit numeric conversions are:

- From `sbyte` to `short`, `int`, **`nint`**, `long`, `float`, `double`, or `decimal`.
- From `byte` to `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `float`, `double`, or `decimal`.
- From `short` to `int`, **`nint`**, `long`, `float`, `double`, or `decimal`.
- From `ushort` to `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `float`, `double`, or `decimal`.
- From `int` to **`nint`**, `long`, `float`, `double`, or `decimal`.
- From `uint` to **`nuint`**, `long`, `ulong`, `float`, `double`, or `decimal`.
- **From `nint` to `long`, `float`, `double`, or `decimal`.**
- **From `nuint` to `ulong`, `float`, `double`, or `decimal`.**
- From `long` to `float`, `double`, or `decimal`.
- From `ulong` to `float`, `double`, or `decimal`.
- From `char` to `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `float`, `double`, or `decimal`.
- From `float` to `double`.

[...]

10.2.11 Implicit constant expression conversions

An implicit constant expression conversion permits the following conversions:

- A *constant_expression* of type `int` can be converted to type `sbyte`, `byte`, `short`, `ushort`, `uint`, **`nint`**, **`nuint`**, or `ulong`, provided the value of the *constant_expression* is within the range of the destination type. [...]

10.3.2 Explicit numeric conversions

The explicit numeric conversions are the conversions from a *numeric_type* to another *numeric_type* for which an implicit numeric conversion does not already exist:

- From `sbyte` to `byte`, `ushort`, `uint`, **`nuint`**, `ulong`, or `char`.
- From `byte` to `sbyte` or `char`.
- From `short` to `sbyte`, `byte`, `ushort`, `uint`, **`nuint`**, `ulong`, or `char`.
- From `ushort` to `sbyte`, `byte`, `short`, or `char`.
- From `int` to `sbyte`, `byte`, `short`, `ushort`, `uint`, **`nuint`**, `ulong`, or `char`.
- From `uint` to `sbyte`, `byte`, `short`, `ushort`, `int`, **`nint`**, or `char`.
- From `long` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `ulong`, or `char`.
- **From `nint` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `nuint`, `ulong`, or `char`.**
- **From `nuint` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `nint`, `long`, or `char`.**

- From `ulong` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, or `char`.
- From `char` to `sbyte`, `byte`, or `short`.
- From `float` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `char`, or `decimal`.
- From `double` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `char`, `float`, or `decimal`.
- From `decimal` to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `char`, `float`, or `double`.

[...]

10.3.3 Explicit enumeration conversions

The explicit enumeration conversions are:

- From `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `char`, `float`, `double`, or `decimal` to any *enum_type*.
- From any *enum_type* to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `char`, `float`, `double`, or `decimal`.
- From any *enum_type* to any other *enum_type*.

11.6.4.6 Better conversion target

Given two types T_1 and T_2, T_1 is a ***better conversion target*** than T_2 if one of the following holds:

- An implicit conversion from T_1 to T_2 exists and no implicit conversion from T_2 to T_1 exists
- T_1 is $\texttt{Task<}S_1\texttt{>}$, T_2 is $\texttt{Task<}S_2\texttt{>}$, and S_1 is a better conversion target than S_2
- T_1 is S_1 or $S_1\texttt{?}$ where S_1 is a signed integral type, and T_2 is S_2 or $S_2\texttt{?}$ where S_2 is an unsigned integral type. Specifically: [...]

11.7.10 Element access

[...] The number of expressions in the *argument_list* shall be the same as the rank of the *array_type*, and each expression shall be of type `int`, `uint`, **`nint`**, **`nuint`**, `long`, or `ulong`, or shall be implicitly convertible to one or more of these types.

11.7.10.2 Array access

[...] The number of expressions in the *argument_list* shall be the same as the rank of the *array_type*, and each expression shall be of type `int`, `uint`, **`nint`**, **`nuint`**, `long`, or `ulong`, or shall be implicitly convertible to one or more of these types.

[...] The run-time processing of an array access of the form `P[A]`, where `P` is a *primary_no_array_creation_expression* of an *array_type* and `A` is an *argument_list*, consists of the following steps: [...]

- The index expressions of the *argument_list* are evaluated in order, from left to right. Following evaluation of each index expression, an implicit conversion to one of the following types is performed: `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`. The first type in this list for which an implicit conversion exists is chosen. [...]

11.7.14 Postfix increment and decrement operators

Unary operator overload resolution is applied to select a specific operator implementation. Predefined `++` and `--` operators exist for the following types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `char`, `float`, `double`, `decimal`, and any enum type.

11.8.2 Unary plus operator

The predefined unary plus operators are:

```
C#

...
nint operator +(nint x);
nuint operator +(nuint x);
```

11.8.3 Unary minus operator

The predefined unary minus operators are:

- Integer negation:

```
C#

...
nint operator -(nint x);
```

11.7.14 Postfix increment and decrement operators

Predefined `++` and `--` operators exist for the following types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `char`, `float`, `double`, `decimal`, and any enum type.

11.7.19 Default value expressions

In addition, a *default_value_expression* is a constant expression if the type is one of the following value types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `char`, `float`, `double`, `decimal`, `bool`, or any enumeration type.

11.8.5 Bitwise complement operator

The predefined bitwise complement operators are:

```
C#

...
nint operator ~(nint x);
nuint operator ~(nuint x);
```

11.8.6 Prefix increment and decrement operators

Predefined `++` and `--` operators exist for the following types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `char`, `float`, `double`, `decimal`, and any enum type.

11.9 Arithmetic operators

11.9.2 Multiplication operator

The predefined multiplication operators are listed below. The operators all compute the product of `x` and `y`.

- Integer multiplication:

  ```
  C#

  ...
  nint operator *(nint x, nint y);
  ```

```
nuint operator *(nuint x, nuint y);
```

11.9.3 Division operator

The predefined division operators are listed below. The operators all compute the quotient of `x` and `y`.

- Integer division:

```
C#

...
nint operator /(nint x, nint y);
nuint operator /(nuint x, nuint y);
```

11.9.4 Remainder operator

The predefined remainder operators are listed below. The operators all compute the remainder of the division between `x` and `y`.

- Integer remainder:

```
C#

...
nint operator %(nint x, nint y);
nuint operator %(nuint x, nuint y);
```

11.9.5 Addition operator

- Integer addition:

```
C#

...
nint operator +(nint x, nint y);
nuint operator +(nuint x, nuint y);
```

11.9.6 Subtraction operator

- Integer subtraction:

```
C#
```

```
...
nint operator -(nint x, nint y);
nuint operator -(nuint x, nuint y);
```

11.10 Shift operators

The predefined shift operators are listed below.

- Shift left:

  ```
  C#
  ```

  ```
  ...
  nint operator <<(nint x, int count);
  nuint operator <<(nuint x, int count);
  ```

- Shift right:

  ```
  C#
  ```

  ```
  ...
  nint operator >>(nint x, int count);
  nuint operator >>(nuint x, int count);
  ```

  The `>>` operator shifts `x` right by a number of bits computed as described below.

  When `x` is of type `int`, **`nint`** or `long`, the low-order bits of `x` are discarded, the remaining bits are shifted right, and the high-order empty bit positions are set to zero if `x` is non-negative and set to one if `x` is negative.

  When `x` is of type `uint`, **`nuint`** or `ulong`, the low-order bits of `x` are discarded, the remaining bits are shifted right, and the high-order empty bit positions are set to zero.

- Unsigned shift right:

  ```
  C#
  ```

  ```
  ...
  nint operator >>>(nint x, int count);
  nuint operator >>>(nuint x, int count);
  ```

For the predefined operators, the number of bits to shift is computed as follows: [...]

- When the type of `x` is `nint` or `nuint`, the shift count is given by the low-order five bits of `count` on a 32 bit platform, or the lower-order six bits of `count` on a 64 bit platform.

11.11 Relational and type-testing operators

11.11.2 Integer comparison operators

The predefined integer comparison operators are:

C#

```
...
bool operator ==(nint x, nint y);
bool operator ==(nuint x, nuint y);

bool operator !=(nint x, nint y);
bool operator !=(nuint x, nuint y);

bool operator <(nint x, nint y);
bool operator <(nuint x, nuint y);

bool operator >(nint x, nint y);
bool operator >(nuint x, nuint y);

bool operator <=(nint x, nint y);
bool operator <=(nuint x, nuint y);

bool operator >=(nint x, nint y);
bool operator >=(nuint x, nuint y);
```

11.12 Logical operators

11.12.2 Integer logical operators

The predefined integer logical operators are:

C#

```
...
nint operator &(nint x, nint y);
nuint operator &(nuint x, nuint y);
```

```
nint operator |(nint x, nint y);
nuint operator |(nuint x, nuint y);

nint operator ^(nint x, nint y);
nuint operator ^(nuint x, nuint y);
```

11.20 Constant expressions

A constant expression may be either a value type or a reference type. If a constant expression is a value type, it must be one of the following types: `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, `char`, `float`, `double`, `decimal`, `bool`, or any enumeration type.

[...]

An implicit constant expression conversion permits a constant expression of type `int` to be converted to `sbyte`, `byte`, `short`, `ushort`, `uint`, **`nint`**, **`nuint`**, or `ulong`, provided the value of the constant expression is within the range of the destination type.

16.4 Array element access

Array elements are accessed using *element_access* expressions of the form `A[I₁, I₂, ..., Iₓ]`, where `A` is an expression of an array type and each `Iₑ` is an expression of type `int`, `uint`, **`nint`**, **`nuint`**, `long`, `ulong`, or can be implicitly converted to one or more of these types. The result of an array element access is a variable, namely the array element selected by the indices.

22.5 Pointer conversions

22.5.1 General

[...]

Additionally, in an unsafe context, the set of available explicit conversions is extended to include the following explicit pointer conversions:

- From any *pointer_type* to any other *pointer_type*.
- From `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, or `ulong` to any *pointer_type*.

- From any *pointer_type* to `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, **`nint`**, **`nuint`**, `long`, or `ulong`.

22.6.4 Pointer element access

[...] In a pointer element access of the form `P[E]`, `P` shall be an expression of a pointer type other than `void*`, and `E` shall be an expression that can be implicitly converted to `int`, `uint`, **`nint`**, **`nuint`**, `long`, or `ulong`.

22.6.7 Pointer arithmetic

In an unsafe context, the `+` operator and `-` operator can be applied to values of all pointer types except `void*`. Thus, for every pointer type `T*`, the following operators are implicitly defined:

```
C#

[...]
T* operator +(T* x, nint y);
T* operator +(T* x, nuint y);
T* operator +(nint x, T* y);
T* operator +(nuint x, T* y);
T* operator -(T* x, nint y);
T* operator -(T* x, nuint y);
```

Given an expression `P` of a pointer type `T*` and an expression `N` of type `int`, `uint`, **`nint`**, **`nuint`**, `long`, or `ulong`, the expressions `P + N` and `N + P` compute the pointer value of type `T*` that results from adding `N * sizeof(T)` to the address given by `P`. Likewise, the expression `P - N` computes the pointer value of type `T*` that results from subtracting `N * sizeof(T)` from the address given by `P`.

Various considerations

Breaking changes

One of the main impacts of this design is that `System.IntPtr` and `System.UIntPtr` gain some built-in operators (conversions, unary and binary).
Those include `checked` operators, which means that the following operators on those types will now throw when overflowing:

- `IntPtr + int`

- `IntPtr - int`
- `IntPtr -> int`
- `long -> IntPtr`
- `void* -> IntPtr`

Metadata encoding

This design means that `nint` and `nuint` can simply be emitted as `System.IntPtr` and `System.UIntPtr`, without the use of `System.Runtime.CompilerServices.NativeIntegerAttribute`.
Similarly, when loading metadata `NativeIntegerAttribute` can be ignored.

Improved Definite Assignment Analysis

Article • 05/29/2022

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent language design meeting (LDM) notes. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the specifications.

Summary

Definite assignment §9.4 as specified has a few gaps which have caused users inconvenience. In particular, scenarios involving comparison to boolean constants, conditional-access, and null coalescing.

Related discussions and issues

csharplang discussion of this proposal:
https://github.com/dotnet/csharplang/discussions/4240

Probably a dozen or so user reports can be found via this or similar queries (i.e. search for "definite assignment" instead of "CS0165", or search in csharplang).
https://github.com/dotnet/roslyn/issues?q=is%3Aclosed+is%3Aissue+label%3A%22Resolution-By+Design%22+cs0165

I have included related issues in the scenarios below to give a sense of the relative impact of each scenario.

Scenarios

As a point of reference, let's start with a well-known "happy case" that does work in definite assignment and in nullable.

```
cs
```

```cs
#nullable enable

C c = new C();
if (c != null && c.M(out object obj0))
{
    obj0.ToString(); // ok
}

public class C
{
    public bool M(out object obj)
    {
        obj = new object();
        return true;
    }
}
```

Comparison to bool constant

- https://github.com/dotnet/csharplang/discussions/801
- https://github.com/dotnet/roslyn/issues/45582
  - Links to 4 other issues where people were affected by this.

```cs
if ((c != null && c.M(out object obj1)) == true)
{
    obj1.ToString(); // undesired error
}

if ((c != null && c.M(out object obj2)) is true)
{
    obj2.ToString(); // undesired error
}
```

Comparison between a conditional access and a constant value

- https://github.com/dotnet/roslyn/issues/33559
- https://github.com/dotnet/csharplang/discussions/4214
- https://github.com/dotnet/csharplang/issues/3659
- https://github.com/dotnet/csharplang/issues/3485
- https://github.com/dotnet/csharplang/issues/3659

This scenario is probably the biggest one. We do support this in nullable but not in definite assignment.

```cs
if (c?.M(out object obj3) == true)
{
    obj3.ToString(); // undesired error
}
```

Conditional access coalesced to a bool constant

- https://github.com/dotnet/csharplang/discussions/916
- https://github.com/dotnet/csharplang/issues/3365

This scenario is very similar to the previous one. This is also supported in nullable but not in definite assignment.

```cs
if (c?.M(out object obj4) ?? false)
{
    obj4.ToString(); // undesired error
}
```

Conditional expressions where one arm is a bool constant

- https://github.com/dotnet/roslyn/issues/4272

It's worth pointing out that we already have special behavior for when the condition expression is constant (i.e. `true ? a : b`). We just unconditionally visit the arm indicated by the constant condition and ignore the other arm.

Also note that we haven't handled this scenario in nullable.

```cs
if (c != null ? c.M(out object obj4) : false)
{
    obj4.ToString(); // undesired error
}
```

Specification

?. (null-conditional operator) expressions

We introduce a new section **?. (null-conditional operator) expressions**. See the null-conditional operator specification (§11.7.7)and Precise rules for determining definite assignment §9.4.4 for context.

As in the definite assignment rules linked above, we refer to a given initially unassigned variable as *v*.

We introduce the concept of "directly contains". An expression *E* is said to "directly contain" a subexpression E_1 if it is not subject to a user-defined conversion §10.5 whose parameter is not of a non-nullable value type, and one of the following conditions holds:

- *E* is E_1. For example, `a?.b()` directly contains the expression `a?.b()`.
- If *E* is a parenthesized expression `(E2)`, and E_2 directly contains E_1.
- If *E* is a null-forgiving operator expression `E2!`, and E_2 directly contains E_1.
- If *E* is a cast expression `(T)E2`, and the cast does not subject E_2 to a non-lifted user-defined conversion whose parameter is not of a non-nullable value type, and E_2 directly contains E_1.

For an expression *E* of the form `primary_expression null_conditional_operations`, let E_0 be the expression obtained by textually removing the leading ? from each of the *null_conditional_operations* of *E* that have one, as in the linked specification above.

In subsequent sections we will refer to E_0 as the *non-conditional counterpart* to the null-conditional expression. Note that some expressions in subsequent sections are subject to additional rules that only apply when one of the operands directly contains a null-conditional expression.

- The definite assignment state of *v* at any point within *E* is the same as the definite assignment state at the corresponding point within *E0*.
- The definite assignment state of *v* after *E* is the same as the definite assignment state of *v* after *primary_expression*.

Remarks

We use the concept of "directly contains" to allow us to skip over relatively simple "wrapper" expressions when analyzing conditional accesses that are compared to other values. For example, `((a?.b(out x))!) == true` is expected to result in the same flow state as `a?.b == true` in general.

We also want to allow analysis to function in the presence of a number of possible conversions on a conditional access. Propagating out "state when not null" is not possible when the conversion is user-defined, though, since we can't count on user-defined conversions to honor the constraint that the output is non-null only if the input is non-null. The only exception to this is when the user-defined conversion's input is a non-nullable value type. For example:

```cs
public struct S1 { }
public struct S2 { public static implicit operator S2?(S1 s1) => null; }
```

This also includes lifted conversions like the following:

```cs
string x;

S1? s1 = null;
_ = s1?.M1(x = "a") ?? s1.Value.M2(x = "a");

x.ToString(); // ok

public struct S1
{
    public S1 M1(object obj) => this;
    public S2 M2(object obj) => new S2();
}
public struct S2
{
    public static implicit operator S2(S1 s1) => null;
}
```

When we consider whether a variable is assigned at a given point within a null-conditional expression, we simply assume that any preceding null-conditional operations within the same null-conditional expression succeeded.

For example, given a conditional expression `a?.b(out x)?.c(x)`, the non-conditional counterpart is `a.b(out x).c(x)`. If we want to know the definite assignment state of `x` before `?.c(x)`, for example, then we perform a "hypothetical" analysis of `a.b(out x)` and use the resulting state as an input to `?.c(x)`.

Boolean constant expressions

We introduce a new section "Boolean constant expressions":

For an expression *expr* where *expr* is a constant expression with a bool value:

- The definite assignment state of *v* after *expr* is determined by:
  - If *expr* is a constant expression with value *true*, and the state of *v* before *expr* is "not definitely assigned", then the state of *v* after *expr* is "definitely assigned when false".
  - If *expr* is a constant expression with value *false*, and the state of *v* before *expr* is "not definitely assigned", then the state of *v* after *expr* is "definitely assigned when true".

Remarks

We assume that if an expression has a constant value bool `false`, for example, it's impossible to reach any branch that requires the expression to return `true`. Therefore variables are assumed to be definitely assigned in such branches. This ends up combining nicely with the spec changes for expressions like `??` and `?:` and enabling a lot of useful scenarios.

It's also worth noting that we never expect to be in a conditional state *before* visiting a constant expression. That's why we do not account for scenarios such as "*expr* is a constant expression with value *true*, and the state of *v* before *expr* is "definitely assigned when true".

?? (null-coalescing expressions) augment

We augment section §9.4.4.29 as follows:

For an expression *expr* of the form `expr_first ?? expr_second`:

- ...
- The definite assignment state of *v* after *expr* is determined by:
  - ...
  - If *expr_first* directly contains a null-conditional expression *E*, and *v* is definitely assigned after the non-conditional counterpart E_0, then the definite assignment state of *v* after *expr* is the same as the definite assignment state of *v* after *expr_second*.

Remarks

The above rule formalizes that for an expression like `a?.M(out x) ?? (x = false)`, either the `a?.M(out x)` was fully evaluated and produced a non-null value, in which case `x` was

assigned, or the `x = false` was evaluated, in which case `x` was also assigned. Therefore `x` is always assigned after this expression.

This also handles the `dict?.TryGetValue(key, out var value) ?? false` scenario, by observing that *v* is definitely assigned after `dict.TryGetValue(key, out var value)`, and *v* is "definitely assigned when true" after `false`, and concluding that *v* must be "definitely assigned when true".

The more general formulation also allows us to handle some more unusual scenarios, such as:

- `if (x?.M(out y) ?? (b && z.M(out y))) y.ToString();`
- `if (x?.M(out y) ?? z?.M(out y) ?? false) y.ToString();`

?: (conditional) expressions

We augment section §9.4.4.30 as follows:

For an expression *expr* of the form `expr_cond ? expr_true : expr_false`:

- ...
- The definite assignment state of *v* after *expr* is determined by:
  - ...
  - If the state of *v* after *expr_true* is "definitely assigned when true", and the state of *v* after *expr_false* is "definitely assigned when true", then the state of *v* after *expr* is "definitely assigned when true".
  - If the state of *v* after *expr_true* is "definitely assigned when false", and the state of *v* after *expr_false* is "definitely assigned when false", then the state of *v* after *expr* is "definitely assigned when false".

Remarks

This makes it so when both arms of a conditional expression result in a conditional state, we join the corresponding conditional states and propagate it out instead of unsplitting the state and allowing the final state to be non-conditional. This enables scenarios like the following:

```cs
bool b = true;
object x = null;
int y;
if (b ? x != null && Set(out y) : x != null && Set(out y))
```

```
{
  y.ToString();
}

bool Set(out int x) { x = 0; return true; }
```

This is an admittedly niche scenario, that compiles without error in the native compiler, but was broken in Roslyn in order to match the specification at the time.

==/!= (relational equality operator) expressions

We introduce a new section **==/!= (relational equality operator) expressions**.

The general rules for expressions with embedded expressions §9.4.4.23 apply, except for the scenarios described below.

For an expression *expr* of the form `expr_first == expr_second`, where `==` is a[predefined comparison operator (§11.11) or a lifted operator (§11.4.8), the definite assignment state of *v* after *expr* is determined by:

- If *expr_first* directly contains a null-conditional expression *E* and *expr_second* is a constant expression with value *null*, and the state of *v* after the non-conditional counterpart E_0 is "definitely assigned", then the state of *v* after *expr* is "definitely assigned when false".
- If *expr_first* directly contains a null-conditional expression *E* and *expr_second* is an expression of a non-nullable value type, or a constant expression with a non-null value, and the state of *v* after the non-conditional counterpart E_0 is "definitely assigned", then the state of *v* after *expr* is "definitely assigned when true".
- If *expr_first* is of type *boolean*, and *expr_second* is a constant expression with value *true*, then the definite assignment state after *expr* is the same as the definite assignment state after *expr_first*.
- If *expr_first* is of type *boolean*, and *expr_second* is a constant expression with value *false*, then the definite assignment state after *expr* is the same as the definite assignment state of *v* after the logical negation expression `!expr_first`.

For an expression *expr* of the form `expr_first != expr_second`, where `!=` is a predefined comparison operator (§11.11) or a lifted operator ((§11.4.8)), the definite assignment state of *v* after *expr* is determined by:

- If *expr_first* directly contains a null-conditional expression *E* and *expr_second* is a constant expression with value *null*, and the state of *v* after the non-conditional counterpart E_0 is "definitely assigned", then the state of *v* after *expr* is "definitely assigned when true".

- If *expr_first* directly contains a null-conditional expression *E* and *expr_second* is an expression of a non-nullable value type, or a constant expression with a non-null value, and the state of *v* after the non-conditional counterpart E_0 is "definitely assigned", then the state of *v* after *expr* is "definitely assigned when false".
- If *expr_first* is of type *boolean*, and *expr_second* is a constant expression with value *true*, then the definite assignment state after *expr* is the same as the definite assignment state of *v* after the logical negation expression `!expr_first`.
- If *expr_first* is of type *boolean*, and *expr_second* is a constant expression with value *false*, then the definite assignment state after *expr* is the same as the definite assignment state after *expr_first*.

All of the above rules in this section are commutative, meaning that if a rule applies when evaluated in the form `expr_second op expr_first`, it also applies in the form `expr_first op expr_second`.

Remarks

The general idea expressed by these rules is:

- if a conditional access is compared to `null`, then we know the operations definitely occurred if the result of the comparison is `false`
- if a conditional access is compared to a non-nullable value type or a non-null constant, then we know the operations definitely occurred if the result of the comparison is `true`.
- since we can't trust user-defined operators to provide reliable answers where initialization safety is concerned, the new rules only apply when a predefined `==`/`!=` operator is in use.

We may eventually want to refine these rules to thread through conditional state which is present at the end of a member access or call. Such scenarios don't really happen in definite assignment, but they do happen in nullable in the presence of `[NotNullWhen(true)]` and similar attributes. This would require special handling for `bool` constants in addition to just handling for `null`/non-null constants.

Some consequences of these rules:

- `if (a?.b(out var x) == true)) x() else x();` will error in the 'else' branch
- `if (a?.b(out var x) == 42)) x() else x();` will error in the 'else' branch
- `if (a?.b(out var x) == false)) x() else x();` will error in the 'else' branch
- `if (a?.b(out var x) == null)) x() else x();` will error in the 'then' branch
- `if (a?.b(out var x) != true)) x() else x();` will error in the 'then' branch

- `if (a?.b(out var x) != 42)) x() else x();` will error in the 'then' branch
- `if (a?.b(out var x) != false)) x() else x();` will error in the 'then' branch
- `if (a?.b(out var x) != null)) x() else x();` will error in the 'else' branch

`is` operator and `is` pattern expressions

We introduce a new section **`is` operator and `is` pattern expressions**.

For an expression *expr* of the form `E is T`, where *T* is any type or pattern

- The definite assignment state of *v* before *E* is the same as the definite assignment state of *v* before *expr*.
- The definite assignment state of *v* after *expr* is determined by:
  - If *E* directly contains a null-conditional expression, and the state of *v* after the non-conditional counterpart E_0 is "definitely assigned", and `T` is any type or a pattern that does not match a `null` input, then the state of *v* after *expr* is "definitely assigned when true".
  - If *E* directly contains a null-conditional expression, and the state of *v* after the non-conditional counterpart E_0 is "definitely assigned", and `T` is a pattern that matches a `null` input, then the state of *v* after *expr* is "definitely assigned when false".
  - If *E* is of type boolean and `T` is a pattern which only matches a `true` input, then the definite assignment state of *v* after *expr* is the same as the definite assignment state of *v* after E.
  - If *E* is of type boolean and `T` is a pattern which only matches a `false` input, then the definite assignment state of *v* after *expr* is the same as the definite assignment state of *v* after the logical negation expression `!expr`.
  - Otherwise, if the definite assignment state of *v* after E is "definitely assigned", then the definite assignment state of *v* after *expr* is "definitely assigned".

Remarks

This section is meant to address similar scenarios as in the `==`/`!=` section above. This specification does not address recursive patterns, e.g. `(a?.b(out x), c?.d(out y)) is (object, object)`. Such support may come later if time permits.

Additional scenarios

This specification doesn't currently address scenarios involving pattern switch expressions and switch statements. For example:

```cs
_ = c?.M(out object obj4) switch
{
    not null => obj4.ToString() // undesired error
};
```

It seems like support for this could come later if time permits.

There have been several categories of bugs filed for nullable which require we essentially increase the sophistication of pattern analysis. It is likely that any ruling we make which improves definite assignment would also be carried over to nullable.

https://github.com/dotnet/roslyn/issues/49353
https://github.com/dotnet/roslyn/issues/46819
https://github.com/dotnet/roslyn/issues/44127

Drawbacks

It feels odd to have the analysis "reach down" and have special recognition of conditional accesses, when typically flow analysis state is supposed to propagate upward. We are concerned about how a solution like this could intersect painfully with possible future language features that do null checks.

Alternatives

Two alternatives to this proposal:

1. Introduce "state when null" and "state when not null" to the language and compiler. This has been judged to be too much effort for the scenarios we are trying to solve, but that we could potentially implement the above proposal and then move to a "state when null/not null" model later on without breaking people.
2. Do nothing.

Unresolved questions

There are impacts on switch expressions that should be specified:
https://github.com/dotnet/csharplang/discussions/4240#discussioncomment-343395

Design meetings

https://github.com/dotnet/csharplang/discussions/4243

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

C# feature specification feedback

C# feature specification is an open source project. Select a link to provide feedback:

- Open a documentation issue
- Provide product feedback

`ref readonly` parameters

Article • 08/16/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow parameter declaration-site modifier `ref readonly` and change callsite rules as follows:

Expand table

Callsite annotation	`ref` **parameter**	`ref readonly` **parameter**	`in` **parameter**	`out` **parameter**
`ref`	Allowed	**Allowed**	**Warning**	Error
`in`	Error	**Allowed**	Allowed	Error
`out`	Error	**Error**	Error	Allowed
No annotation	Error	**Warning**	Allowed	Error

(Note that there is one change to the existing rules: `in` parameter with `ref` callsite annotation produces a warning instead of an error.)

Change argument value rules as follows:

Expand table

Value kind	`ref` **parameter**	`ref readonly` **parameter**	`in` **parameter**	`out` **parameter**
rvalue	Error	**Warning**	Allowed	Error

Value kind	`ref` parameter	`ref readonly` parameter	`in` parameter	`out` parameter
lvalue	Allowed	**Allowed**	Allowed	Allowed

Where lvalue means a variable (i.e., a value with a location; does not have to be writable/assignable) and rvalue means any kind of value.

Motivation

C# 7.2 introduced in parameters as a way to pass readonly references. `in` parameters allow both lvalues and rvalues and can be used without any annotation at the callsite. However, APIs which capture or return references from their parameters would like to disallow rvalues and also enforce some indication at the callsite that a reference is being captured. `ref readonly` parameters are ideal in such cases as they warn if used with rvalues or without any annotation at the callsite.

Furthermore, there are APIs that need only read-only references but use

- `ref` parameters since they were introduced before `in` became available and changing to `in` would be a source and binary breaking change, e.g., `QueryInterface`, or
- `in` parameters to accept readonly references even though passing rvalues to them doesn't really make sense, e.g., `ReadOnlySpan<T>..ctor(in T value)`, or
- `ref` parameters to disallow rvalues even though they don't mutate the passed reference, e.g., `Unsafe.IsNullRef`.

These APIs could migrate to `ref readonly` parameters without breaking users. For details on binary compatibility, see the proposed metadata encoding. Specifically, changing

- `ref` → `ref readonly` would only be a binary breaking change for virtual methods,
- `ref` → `in` would also be a binary breaking change for virtual methods, but not a source breaking change (because the rules change to only warn for `ref` arguments passed to `in` parameters),
- `in` → `ref readonly` would not be a breaking change (but no callsite annotation or rvalue would result in a warning),
  - note that this would be a source breaking change for users using older compiler versions (as they interpret `ref readonly` parameters as `ref` parameters, disallowing `in` or no annotation at the callsite) and new compiler versions with `LangVersion <= 11` (for consistency with older compiler versions, an error will be

emitted that `ref readonly` parameters are not supported unless the corresponding arguments are passed with the `ref` modifier).

In the opposite direction, changing

- `ref readonly` → `ref` would be potentially a source breaking change (unless only `ref` callsite annotation was used and only readonly references used as arguments), and a binary breaking change for virtual methods,
- `ref readonly` → `in` would not be a breaking change (but `ref` callsite annotation would result in a warning).

Note that the rules outlined above apply to method signatures, but not to delegate signatures. For example, changing `ref` to `in` in a delegate signature can be a source breaking change (if a user is assigning a method with `ref` parameter to that delegate type, it would become an error after the API change).

Detailed design

In general, rules for `ref readonly` parameters are the same as specified for `in` parameters in their proposal, except where explicitly changed in this proposal.

Parameter declarations

No changes in grammar are necessary. The modifier `ref readonly` will be allowed for parameters. Apart from normal methods, `ref readonly` will be allowed for indexer parameters (like `in` but unlike `ref`), but disallowed for operator parameters (like `ref` but unlike `in`).

Default parameter values will be allowed for `ref readonly` parameters with a warning since they are equivalent to passing rvalues. This allows API authors to change `in` parameters with default values to `ref readonly` parameters without introducing a source breaking change.

Value kind checks

Note that even though `ref` argument modifier is allowed for `ref readonly` parameters, nothing changes w.r.t. value kind checks, i.e.,

- `ref` can only be used with assignable values;
- to pass readonly references, one has to use the `in` argument modifier instead;

- to pass rvalues, one has to use no modifier (which results in a warning for `ref readonly` parameters as described in the summary of this proposal).

Overload resolution

Overload resolution will allow mixing `ref`/`ref readonly`/`in`/no callsite annotations and parameter modifiers as denoted by the table in the summary of this proposal, i.e., all *allowed* and *warning* cases will be considered as possible candidates during overload resolution. Specifically, there's a change in existing behavior where methods with `in` parameter will match calls with the corresponding argument marked as `ref`—this change will be gated on LangVersion.

However, the warning for passing an argument with no callsite modifier to a `ref readonly` parameter will be suppressed if the parameter is

- the receiver in an extension method invocation,
- used implicitly as part of custom collection initializer or interpolated string handler.

By-value overloads will be preferred over `ref readonly` overloads in case there is no argument modifier (`in` parameters have the same behavior).

Method conversions

Similarly, for the purpose of anonymous function [§10.7] and method group [§10.8] conversions, these modifiers are considered compatible (but any allowed conversion between different modifiers results in a warning):

- `ref readonly` parameter of the target method is allowed to match `in` or `ref` parameter of the delegate,
- `in` parameter of the target method is allowed to match `ref readonly` or, gated on LangVersion, `ref` parameter of the delegate.
- Note: `ref` parameter of the target method is *not* allowed to match `in` nor `ref readonly` parameter of the delegate.

For example:

```cs
DIn dIn = (ref int p) => { }; // error: cannot match `ref` to `in`
DRef dRef = (in int p) => { }; // warning: mismatch between `in` and `ref`
DRR dRR = (ref int p) => { }; // error: cannot match `ref` to `ref readonly`
dRR = (in int p) => { }; // warning: mismatch between `in` and `ref
readonly`
```

```
dIn = (ref readonly int p) => { }; // warning: mismatch between `ref
readonly` and `in`
dRef = (ref readonly int p) => { }; // warning: mismatch between `ref
readonly` and `ref`
delegate void DIn(in int p);
delegate void DRef(ref int p);
delegate void DRR(ref readonly int p);
```

Note that there is no change in behavior of function pointer conversions. As a reminder, implicit function pointer conversions are disallowed if there is a mismatch between reference kind modifiers, and explicit casts are always allowed without any warnings.

Signature matching

Members declared in a single type cannot differ in signature solely by `ref`/`out`/`in`/`ref readonly`. For other purposes of signature matching (e.g., hiding or overriding), `ref readonly` can be interchanged with `in` modifier, but that results in a warning at the declaration site [§7.6]. This doesn't apply when matching `partial` declaration with its implementation and when matching interceptor signature with intercepted signature. Note that there is no change in overriding for `ref`/`in` and `ref readonly`/`ref` modifier pairs, they cannot be interchanged, because the signatures aren't binary compatible. For consistency, the same is true for other signature matching purposes (e.g., hiding).

Metadata encoding

As a reminder,

- `ref` parameters are emitted as plain byref types (`T&` in IL),
- `in` parameters are like `ref` plus they are annotated with `System.Runtime.CompilerServices.IsReadOnlyAttribute`. In C# 7.3 and later, they are also emitted with `[in]` and if virtual, `modreq(System.Runtime.InteropServices.InAttribute)`.

`ref readonly` parameters will be emitted as `[in] T&`, plus annotated with the following attribute:

cs

```
namespace System.Runtime.CompilerServices
{
    [AttributeUsage(AttributeTargets.Parameter, AllowMultiple = false,
Inherited = false)]
```

```
    public sealed class RequiresLocationAttribute : Attribute
    {
    }
}
```

Furthermore, if virtual, they will be emitted with `modreq(System.Runtime.InteropServices.InAttribute)` to ensure binary compatibility with `in` parameters. Note that unlike `in` parameters, no `[IsReadOnly]` will be emitted for `ref readonly` parameters to avoid increasing metadata size and also to make older compiler versions interpret `ref readonly` parameters as `ref` parameters (and hence `ref` → `ref readonly` won't be a source breaking change even between different compiler versions).

The `RequiresLocationAttribute` will be matched by namespace-qualified name and synthesized by the compiler if not already included in the compilation.

Specifying the attribute in source will be an error if it's applied to a parameter, similarly to `ParamArrayAttribute`.

Function pointers

In function pointers, `in` parameters are emitted with `modreq(System.Runtime.InteropServices.InAttribute)` (see function pointers proposal). `ref readonly` parameters will be emitted without that `modreq`, but instead with `modopt(System.Runtime.CompilerServices.RequiresLocationAttribute)`. Older compiler versions will ignore the `modopt` and hence interpret `ref readonly` parameters as `ref` parameters (consistent with older compiler behavior for normal methods with `ref readonly` parameters as described above) and new compiler versions aware of the `modopt` will use it to recognize `ref readonly` parameters to emit warnings during conversions and invocations. For consistency with older compiler versions, new compiler versions with `LangVersion <= 11` will report errors that `ref readonly` parameters are not supported unless the corresponding arguments are passed with the `ref` modifier.

Note that it is a binary break to change modifiers in function pointer signatures if they are part of public APIs, hence it will be a binary break when changing `ref` or `in` to `ref readonly`. However, a source break will only occur for callers with `LangVersion <= 11` when changing `in` → `ref readonly` (if invoking the pointer with `in` callsite modifier), consistent with normal methods.

Breaking changes

The `ref`/`in` mismatch relaxation in overload resolution introduces a behavior breaking change demonstrated in the following example:

```cs
class C
{
    string M(in int i) => "C";
    static void Main()
    {
        int i = 5;
        System.Console.Write(new C().M(ref i));
    }
}
static class E
{
    public static string M(this C c, ref int i) => "E";
}
```

In C# 11, the call binds to `E.M`, hence `"E"` is printed. In C# 12, `C.M` is allowed to bind (with a warning) and no extension scopes are searched since we have an applicable candidate, hence `"C"` is printed.

There is also a source breaking change due to the same reason. The example below prints `"1"` in C# 11, but fails to compile with an ambiguity error in C# 12:

```cs
var i = 5;
System.Console.Write(C.M(null, ref i));

interface I1 { }
interface I2 { }
static class C
{
    public static string M(I1 o, ref int x) => "1";
    public static string M(I2 o, in int x) => "2";
}
```

The examples above demonstrate the breaks for method invocations, but since they are caused by overload resolution changes, they can be similarly triggered for method conversions.

Alternatives

Parameter declarations

API authors could annotate `in` parameters designed to accept only lvalues with a custom attribute and provide an analyzer to flag incorrect usages. This would not allow API authors to change signatures of existing APIs that opted to use `ref` parameters to disallow rvalues. Callers of such APIs would still need to perform extra work to get a `ref` if they have access only to a `ref readonly` variable. Changing these APIs from `ref` to `[RequiresLocation] in` would be a source breaking change (and in case of virtual methods, also a binary breaking change).

Instead of allowing the modifier `ref readonly`, the compiler could recognize when a special attribute (like `[RequiresLocation]`) is applied to a parameter. This was discussed in LDM 2022-04-25, deciding this is a language feature, not an analyzer, so it should look like one.

Value kind checks

Passing lvalues without any modifiers to `ref readonly` parameters could be permitted without any warnings, similarly to C++'s implicit byref parameters. This was discussed in LDM 2022-05-11, noting that the primary motivation for `ref readonly` parameters are APIs which capture or return references from these parameters, so marker of some kind is a good thing.

Passing rvalue to a `ref readonly` could be an error, not a warning. That was initially accepted in LDM 2022-04-25, but later e-mail discussions relaxed this because we would lose the ability to change existing APIs without breaking users.

`in` could be the "natural" callsite modifier for `ref readonly` parameters and using `ref` could result in warnings. This would ensure a consistent code style and make it obvious at the callsite that the reference is readonly (unlike `ref`). It was initially accepted in LDM 2022-04-25. However, warnings could be a friction point for API authors to move from `ref` to `ref readonly`. Also, `in` has been redefined as `ref readonly` + convenience features, hence this was rejected in LDM 2022-05-11.

Pending LDM review

Parameter declarations

Inverse ordering of modifiers (`readonly ref` instead of `ref readonly`) could be allowed. This would be inconsistent with how `readonly ref` returns and fields behave (inverse ordering is disallowed or means something different, respectively) and could clash with readonly parameters if implemented in the future.

Default parameter values could be an error for `ref readonly` parameters.

Value kind checks

Errors could be emitted instead of warnings when passing rvalues to `ref readonly` parameters or mismatching callsite annotations and parameter modifiers. Similarly, special `modreq` could be used instead of an attribute to ensure `ref readonly` parameters are distinct from `in` parameters on the binary level. This would provide stronger guarantees, so it would be good for new APIs, but prevent adoption in existing runtime APIs which cannot introduce breaking changes.

Value kind checks could be relaxed to allow passing readonly references via `ref` into `in`/`ref readonly` parameters. That would be similar to how ref assignments and ref returns work today—they also allow passing references as readonly via the `ref` modifier on the source expression. However, the `ref` there is usually close to the place where the target is declared as `ref readonly`, so it is clear we are passing a reference as readonly, unlike invocations whose argument and parameter modifiers are usually far apart. Furthermore, they allow *only* the `ref` modifier unlike arguments which allow also `in`, hence `in` and `ref` would become interchangeable for arguments, or `in` would become practically obsolete if users wanted to make their code consistent (they would probably use `ref` everywhere since it's the only modifier allowed for ref assignments and ref returns).

Overload resolution

Overload resolution, overriding, and conversion could disallow interchangeability of `ref readonly` and `in` modifiers.

The overload resolution change for existing `in` parameters could be taken unconditionally (not considering LangVersion), but that would be a breaking change.

Invoking an extension method with `ref readonly` receiver could result in warning "Argument 1 should be passed with `ref` or `in` keyword" as would happen for non-extension invocations with no callsite modifiers (user could fix such warning by turning the extension method invocation into static method invocation). The same warning could be reported when using custom collection initializer or interpolated string handler with `ref readonly` parameter, although user could not work around it.

`ref readonly` overloads could be preferred over by-value overloads when there is no callsite modifier or there could be an ambiguity error.

Method conversions

We could allow `ref` parameter of the target method to match `in` and `ref readonly` parameter of the delegate. This would enable API authors to change for example `ref` to `in` in delegate signatures without breaking their users (consistently with what is allowed for normal method signatures). However, it would also result in the following violation of `readonly` guarantees with just a warning:

cs

```
class Program
{
    static readonly int f = 123;
    static void Main()
    {
        var d = (in int x) => { };
        d = (ref int x) => { x = 42; }; // warning: mismatch between `ref`
and `in`
        d(f); // changes value of `f` even though it is `readonly`!
        System.Console.WriteLine(f); // prints 42
    }
}
```

Function pointer conversions could warn on `ref readonly`/`ref`/`in` mismatch, but if we wanted to gate that on LangVersion, a significant implementation investment would be required as today type conversions do not need access to compilation. Furthermore, even though mismatch is currently an error, it is easy for users to add a cast to allow the mismatch if they want.

Metadata encoding

Specifying the `RequiresLocationAttribute` in source could be allowed, similarly to `In` and `Out` attributes. Alternatively, it could be an error when applied in other contexts than just parameters, similarly to `IsReadOnly` attribute; to preserve further design space.

Function pointer `ref readonly` parameters could be emitted with different `modopt`/`modreq` combinations (note that "source break" in this table means for callers with `LangVersion <= 11`):

Expand table

Modifiers	Can be recognized across compilations	Old compilers see them as	`ref` → `ref readonly`	`in` → `ref readonly`
`modreq(In)` `modopt(RequiresLocation)`	yes	`in`	binary, source break	binary break
`modreq(In)`	no	`in`	binary, source break	ok
`modreq(RequiresLocation)`	yes	unsupported	binary, source break	binary, source break
`modopt(RequiresLocation)`	yes	`ref`	binary break	binary, source break

We could emit both `[RequiresLocation]` and `[IsReadOnly]` attributes for `ref readonly` parameters. Then `in` → `ref readonly` would not be a breaking change even for older compiler versions, but `ref` → `ref readonly` would become a source breaking change for older compiler versions (as they would interpret `ref readonly` as `in`, disallowing `ref` modifiers) and new compiler versions with `LangVersion <= 11` (for consistency).

We could make the behavior for `LangVersion <= 11` different from the behavior for older compiler versions. For example, it could be an error whenever a `ref readonly` parameter is called (even when using the `ref` modifier at the callsite), or it could be always allowed without any errors.

Breaking changes

This proposal suggests accepting a behavior breaking change because it should be rare to hit, is gated by LangVersion, and users can work around it by calling the extension method explicitly. Instead, we could mitigate it by

- disallowing the `ref`/`in` mismatch (that would only prevent migration to `in` for old APIs that used `ref` because `in` wasn't available yet),
- modifying the overload resolution rules to continue looking for a better match (determined by betterness rules specified below) when there's a ref kind mismatch introduced in this proposal,
  - or alternatively continue only for `ref` vs. `in` mismatch, not the others (`ref readonly` vs. `ref`/`in`/by-value).

Betterness rules

The following example currently results in three ambiguity errors for the three invocations of `M`. We could add new betterness rules to resolve the ambiguities. This would also resolve the source breaking change described earlier. One way would be to make the example print `221` (where `ref readonly` parameter is matched with `in` argument since it would be a warning to call it with no modifier whereas for `in` parameter that's allowed).

cs

```
interface I1 { }
interface I2 { }
class C
{
    static string M(I1 o, in int i) => "1";
    static string M(I2 o, ref readonly int i) => "2";
    static void Main()
    {
        int i = 5;
        System.Console.Write(M(null, ref i));
        System.Console.Write(M(null, in i));
        System.Console.Write(M(null, i));
    }
}
```

New betterness rules could mark as worse the parameter whose argument could have been passed with a different argument modifier to make it better. In other words, user should be always able to turn a worse parameter into a better parameter by changing its corresponding argument modifier. For example, when an argument is passed by `in`, a `ref readonly` parameter is preferred over an `in` parameter because user could pass the argument by-value to choose the `in` parameter. This rule is just an extension of by-value/`in` preference rule that is in effect today (it's the last overload resolution rule and the whole overload is better if any of its parameter is better and none is worse than the corresponding parameter of another overload).

Expand table

argument	better parameter	worse parameter
`ref`/`in`	`ref readonly`	`in`
`ref`	`ref`	`ref readonly`/`in`
by-value	by-value/`in`	`ref readonly`

argument	better parameter	worse parameter
`in`	`in`	`ref`

We should handle method conversions similarly. The following example currently results in two ambiguity errors for the two delegate assignments. New betterness rules could prefer a method parameter whose refness modifier matches the corresponding target delegate parameter refness modifier over one that has a mismatch. Hence, the following example would print `12`.

cs

```cs
class C
{
    void M(I1 o, ref readonly int x) => System.Console.Write("1");
    void M(I2 o, ref int x) => System.Console.Write("2");
    void Run()
    {
        D1 m1 = this.M;
        D2 m2 = this.M;

        var i = 5;
        m1(null, in i);
        m2(null, ref i);
    }
    static void Main() => new C().Run();
}
interface I1 { }
interface I2 { }
class X : I1, I2 { }
delegate void D1(X s, ref readonly int x);
delegate void D2(X s, ref int x);
```

Design meetings

- LDM 2022-04-25: feature accepted
- LDM 2022-05-09: discussion split into three parts
- LDM 2022-05-11: allowed `ref` and no callsite annotation for `ref readonly` parameters

Collaborate with us on GitHub

.NET C# feature specification feedback

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

C# feature specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Recursive Pattern Matching

Article • 03/24/2022

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Pattern matching extensions for C# enable many of the benefits of algebraic data types and pattern matching from functional languages, but in a way that smoothly integrates with the feel of the underlying language. Elements of this approach are inspired by related features in the programming languages F# and Scala.

Detailed design

Is Expression

The `is` operator is extended to test an expression against a *pattern*.

antlr

```
relational_expression
    : is_pattern_expression
    ;
is_pattern_expression
    : relational_expression 'is' pattern
    ;
```

This form of *relational_expression* is in addition to the existing forms in the C# specification. It is a compile-time error if the *relational_expression* to the left of the `is` token does not designate a value or does not have a type.

Every *identifier* of the pattern introduces a new local variable that is *definitely assigned* after the `is` operator is `true` (i.e. *definitely assigned when true*).

> Note: There is technically an ambiguity between *type* in an `is-expression` and *constant_pattern*, either of which might be a valid parse of a qualified identifier. We try to bind it as a type for compatibility with previous versions of the language; only if that fails do we resolve it as we do an expression in other contexts, to the first thing found (which must be either a constant or a type). This ambiguity is only present on the right-hand-side of an `is` expression.

Patterns

Patterns are used in the *is_pattern* operator, in a *switch_statement*, and in a *switch_expression* to express the shape of data against which incoming data (which we call the input value) is to be compared. Patterns may be recursive so that parts of the data may be matched against sub-patterns.

antlr

```
pattern
    : declaration_pattern
    | constant_pattern
    | var_pattern
    | positional_pattern
    | property_pattern
    | discard_pattern
    ;
declaration_pattern
    : type simple_designation
    ;
constant_pattern
    : constant_expression
    ;
var_pattern
    : 'var' designation
    ;
positional_pattern
    : type? '(' subpatterns? ')' property_subpattern? simple_designation?
    ;
subpatterns
    : subpattern
    | subpattern ',' subpatterns
    ;
subpattern
    : pattern
    | identifier ':' pattern
    ;
property_subpattern
    : '{' '}'
```

```
    | '{' subpatterns ','? '}'
    ;
property_pattern
    : type? property_subpattern simple_designation?
    ;
simple_designation
    : single_variable_designation
    | discard_designation
    ;
discard_pattern
    : '_'
    ;
```

Declaration Pattern

antlr

```
declaration_pattern
    : type simple_designation
    ;
```

The *declaration_pattern* both tests that an expression is of a given type and casts it to that type if the test succeeds. This may introduce a local variable of the given type named by the given identifier, if the designation is a *single_variable_designation*. That local variable is *definitely assigned* when the result of the pattern-matching operation is `true`.

The runtime semantic of this expression is that it tests the runtime type of the left-hand *relational_expression* operand against the *type* in the pattern. If it is of that runtime type (or some subtype) and not `null`, the result of the `is operator` is `true`.

Certain combinations of static type of the left-hand-side and the given type are considered incompatible and result in compile-time error. A value of static type `E` is said to be *pattern-compatible* with a type `T` if there exists an identity conversion, an implicit reference conversion, a boxing conversion, an explicit reference conversion, or an unboxing conversion from `E` to `T`, or if one of those types is an open type. It is a compile-time error if an input of type `E` is not *pattern-compatible* with the *type* in a type pattern that it is matched with.

The type pattern is useful for performing run-time type tests of reference types, and replaces the idiom

C#

```
var v = expr as Type;
if (v != null) { // code using v
```

With the slightly more concise

C#

```
if (expr is Type v) { // code using v
```

It is an error if *type* is a nullable value type.

The type pattern can be used to test values of nullable types: a value of type `Nullable<T>` (or a boxed `T`) matches a type pattern `T2 id` if the value is non-null and the type of `T2` is `T`, or some base type or interface of `T`. For example, in the code fragment

C#

```
int? x = 3;
if (x is int v) { // code using v
```

The condition of the `if` statement is `true` at runtime and the variable `v` holds the value `3` of type `int` inside the block. After the block the variable `v` is in scope but not definitely assigned.

Constant Pattern

antlr

```
constant_pattern
    : constant_expression
    ;
```

A constant pattern tests the value of an expression against a constant value. The constant may be any constant expression, such as a literal, the name of a declared `const` variable, or an enumeration constant. When the input value is not an open type, the constant expression is implicitly converted to the type of the matched expression; if the type of the input value is not *pattern-compatible* with the type of the constant expression, the pattern-matching operation is an error.

The pattern *c* is considered matching the converted input value *e* if `object.Equals(c, e)` would return `true`.

We expect to see `e is null` as the most common way to test for `null` in newly written code, as it cannot invoke a user-defined `operator==`.

Var Pattern

antlr

```
var_pattern
    : 'var' designation
    ;
designation
    : simple_designation
    | tuple_designation
    ;
simple_designation
    : single_variable_designation
    | discard_designation
    ;
single_variable_designation
    : identifier
    ;
discard_designation
    : _
    ;
tuple_designation
    : '(' designations? ')'
    ;
designations
    : designation
    | designations ',' designation
    ;
```

If the *designation* is a *simple_designation*, an expression *e* matches the pattern. In other words, a match to a *var pattern* always succeeds with a *simple_designation*. If the *simple_designation* is a *single_variable_designation*, the value of *e* is bound to a newly introduced local variable. The type of the local variable is the static type of *e*.

If the *designation* is a *tuple_designation*, then the pattern is equivalent to a *positional_pattern* of the form `(var` *designation*, ... `)` where the *designation*s are those found within the *tuple_designation*. For example, the pattern `var (x, (y, z))` is equivalent to `(var x, (var y, var z))`.

It is an error if the name `var` binds to a type.

Discard Pattern

antlr

```
discard_pattern
    : '_'
    ;
```

An expression *e* matches the pattern `_` always. In other words, every expression matches the discard pattern.

A discard pattern may not be used as the pattern of an *is_pattern_expression*.

Positional Pattern

A positional pattern checks that the input value is not `null`, invokes an appropriate `Deconstruct` method, and performs further pattern matching on the resulting values. It also supports a tuple-like pattern syntax (without the type being provided) when the type of the input value is the same as the type containing `Deconstruct`, or if the type of the input value is a tuple type, or if the type of the input value is `object` or `ITuple` and the runtime type of the expression implements `ITuple`.

antlr

```
positional_pattern
    : type? '(' subpatterns? ')' property_subpattern? simple_designation?
    ;
subpatterns
    : subpattern
    | subpattern ',' subpatterns
    ;
subpattern
    : pattern
    | identifier ':' pattern
    ;
```

If the *type* is omitted, we take it to be the static type of the input value.

Given a match of an input value to the pattern *type* `(` *subpattern_list* `)`, a method is selected by searching in *type* for accessible declarations of `Deconstruct` and selecting one among them using the same rules as for the deconstruction declaration.

It is an error if a *positional_pattern* omits the type, has a single *subpattern* without an *identifier*, has no *property_subpattern* and has no *simple_designation*. This disambiguates between a *constant_pattern* that is parenthesized and a *positional_pattern*.

In order to extract the values to match against the patterns in the list,

- If *type* was omitted and the input value's type is a tuple type, then the number of subpatterns is required to be the same as the cardinality of the tuple. Each tuple element is matched against the corresponding *subpattern*, and the match succeeds if all of these succeed. If any *subpattern* has an *identifier*, then that must name a tuple element at the corresponding position in the tuple type.
- Otherwise, if a suitable `Deconstruct` exists as a member of *type*, it is a compile-time error if the type of the input value is not *pattern-compatible* with *type*. At runtime the input value is tested against *type*. If this fails then the positional pattern match fails. If it succeeds, the input value is converted to this type and `Deconstruct` is invoked with fresh compiler-generated variables to receive the `out` parameters. Each value that was received is matched against the corresponding *subpattern*, and the match succeeds if all of these succeed. If any *subpattern* has an *identifier*, then that must name a parameter at the corresponding position of `Deconstruct`.
- Otherwise if *type* was omitted, and the input value is of type `object` or `ITuple` or some type that can be converted to `ITuple` by an implicit reference conversion, and no *identifier* appears among the subpatterns, then we match using `ITuple`.
- Otherwise the pattern is a compile-time error.

The order in which subpatterns are matched at runtime is unspecified, and a failed match may not attempt to match all subpatterns.

Example

This example uses many of the features described in this specification

c#

```
    var newState = (GetState(), action, hasKey) switch {
        (DoorState.Closed, Action.Open, _) => DoorState.Opened,
        (DoorState.Opened, Action.Close, _) => DoorState.Closed,
        (DoorState.Closed, Action.Lock, true) => DoorState.Locked,
        (DoorState.Locked, Action.Unlock, true) => DoorState.Closed,
        (var state, _, _) => state };
```

Property Pattern

A property pattern checks that the input value is not `null` and recursively matches values extracted by the use of accessible properties or fields.

antlr

```
property_pattern
    : type? property_subpattern simple_designation?
    ;
property_subpattern
    : '{' '}'
    | '{' subpatterns ','? '}'
    ;
```

It is an error if any *subpattern* of a *property_pattern* does not contain an *identifier* (it must be of the second form, which has an *identifier*). A trailing comma after the last subpattern is optional.

Note that a null-checking pattern falls out of a trivial property pattern. To check if the string `s` is non-null, you can write any of the following forms

C#

```
if (s is object o) ... // o is of type object
if (s is string x) ... // x is of type string
if (s is {} x) ... // x is of type string
if (s is {}) ...
```

Given a match of an expression *e* to the pattern *type* `{` *property_pattern_list* `}`, it is a compile-time error if the expression *e* is not *pattern-compatible* with the type *T* designated by *type*. If the type is absent, we take it to be the static type of *e*. If the *identifier* is present, it declares a pattern variable of type *type*. Each of the identifiers appearing on the left-hand-side of its *property_pattern_list* must designate an accessible readable property or field of *T*. If the *simple_designation* of the *property_pattern* is present, it defines a pattern variable of type *T*.

At runtime, the expression is tested against *T*. If this fails then the property pattern match fails and the result is `false`. If it succeeds, then each *property_subpattern* field or property is read and its value matched against its corresponding pattern. The result of the whole match is `false` only if the result of any of these is `false`. The order in which subpatterns are matched is not specified, and a failed match may not match all subpatterns at runtime. If the match succeeds and the *simple_designation* of the *property_pattern* is a *single_variable_designation*, it defines a variable of type *T* that is assigned the matched value.

> Note: The property pattern can be used to pattern-match with anonymous types.

Example

C#

```
if (o is string { Length: 5 } s)
```

Switch Expression

A *switch_expression* is added to support `switch`-like semantics for an expression context.

The C# language syntax is augmented with the following syntactic productions:

antlr

```
multiplicative_expression
    : switch_expression
    | multiplicative_expression '*' switch_expression
    | multiplicative_expression '/' switch_expression
    | multiplicative_expression '%' switch_expression
    ;
switch_expression
    : range_expression 'switch' '{' '}'
    | range_expression 'switch' '{' switch_expression_arms ','? '}'
    ;
switch_expression_arms
    : switch_expression_arm
    | switch_expression_arms ',' switch_expression_arm
    ;
switch_expression_arm
    : pattern case_guard? '=>' expression
    ;
case_guard
    : 'when' null_coalescing_expression
    ;
```

The *switch_expression* is not permitted as an *expression_statement*.

> We are looking at relaxing this in a future revision.

The type of the *switch_expression* is the *best common type* (§11.6.3.15) of the expressions appearing to the right of the `=>` tokens of the *switch_expression_arm*s if such a type exists and the expression in every arm of the switch expression can be implicitly converted to that type. In addition, we add a new *switch expression conversion*, which is a predefined implicit conversion from a switch expression to every type `T` for which there exists an implicit conversion from each arm's expression to `T`.

It is an error if some *switch_expression_arm*'s pattern cannot affect the result because some previous pattern and guard will always match.

A switch expression is said to be *exhaustive* if some arm of the switch expression handles every value of its input. The compiler shall produce a warning if a switch expression is not *exhaustive*.

At runtime, the result of the *switch_expression* is the value of the *expression* of the first *switch_expression_arm* for which the expression on the left-hand-side of the *switch_expression* matches the *switch_expression_arm*'s pattern, and for which the *case_guard* of the *switch_expression_arm*, if present, evaluates to `true`. If there is no such *switch_expression_arm*, the *switch_expression* throws an instance of the exception `System.Runtime.CompilerServices.SwitchExpressionException`.

Optional parens when switching on a tuple literal

In order to switch on a tuple literal using the *switch_statement*, you have to write what appear to be redundant parens

```
C#

switch ((a, b))
{
```

To permit

```
C#

switch (a, b)
{
```

the parentheses of the switch statement are optional when the expression being switched on is a tuple literal.

Order of evaluation in pattern-matching

Giving the compiler flexibility in reordering the operations executed during pattern-matching can permit flexibility that can be used to improve the efficiency of pattern-matching. The (unenforced) requirement would be that properties accessed in a pattern, and the Deconstruct methods, are required to be "pure" (side-effect free, idempotent, etc). That doesn't mean that we would add purity as a language concept, only that we would allow the compiler flexibility in reordering operations.

Resolution 2018-04-04 LDM: confirmed: the compiler is permitted to reorder calls to `Deconstruct`, property accesses, and invocations of methods in `ITuple`, and may assume

that returned values are the same from multiple calls. The compiler should not invoke functions that cannot affect the result, and we will be very careful before making any changes to the compiler-generated order of evaluation in the future.

Some Possible Optimizations

The compilation of pattern matching can take advantage of common parts of patterns. For example, if the top-level type test of two successive patterns in a *switch_statement* is the same type, the generated code can skip the type test for the second pattern.

When some of the patterns are integers or strings, the compiler can generate the same kind of code it generates for a switch-statement in earlier versions of the language.

For more on these kinds of optimizations, see [Scott and Ramsey (2000)].

Pattern-matching changes for C# 9.0

Article • 09/21/2021

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

We are considering a small handful of enhancements to pattern-matching for C# 9.0 that have natural synergy and work well to address a number of common programming problems:

- https://github.com/dotnet/csharplang/issues/2925 Type patterns
- https://github.com/dotnet/csharplang/issues/1350 Parenthesized patterns to enforce or emphasize precedence of the new combinators
- https://github.com/dotnet/csharplang/issues/1350 Conjunctive `and` patterns that require both of two different patterns to match;
- https://github.com/dotnet/csharplang/issues/1350 Disjunctive `or` patterns that require either of two different patterns to match;
- https://github.com/dotnet/csharplang/issues/1350 Negated `not` patterns that require a given pattern *not* to match; and
- https://github.com/dotnet/csharplang/issues/812 Relational patterns that require the input value to be less than, less than or equal to, etc a given constant.

Parenthesized Patterns

Parenthesized patterns permit the programmer to put parentheses around any pattern. This is not so useful with the existing patterns in C# 8.0, however the new pattern combinators introduce a precedence that the programmer may want to override.

antlr

```
primary_pattern
    : parenthesized_pattern
    | // all of the existing forms
```

```
    ;
parenthesized_pattern
    : '(' pattern ')'
    ;
```

Type Patterns

We permit a type as a pattern:

antlr

```
primary_pattern
    : type-pattern
    | // all of the existing forms
    ;
type_pattern
    : type
    ;
```

This retcons the existing *is-type-expression* to be an *is-pattern-expression* in which the pattern is a *type-pattern*, though we would not change the syntax tree produced by the compiler.

One subtle implementation issue is that this grammar is ambiguous. A string such as `a.b` can be parsed either as a qualified name (in a type context) or a dotted expression (in an expression context). The compiler is already capable of treating a qualified name the same as a dotted expression in order to handle something like `e is Color.Red`. The compiler's semantic analysis would be further extended to be capable of binding a (syntactic) constant pattern (e.g. a dotted expression) as a type in order to treat it as a bound type pattern in order to support this construct.

After this change, you would be able to write

C#

```
void M(object o1, object o2)
{
    var t = (o1, o2);
    if (t is (int, string)) {} // test if o1 is an int and o2 is a string
    switch (o1) {
        case int: break; // test if o1 is an int
        case System.String: break; // test if o1 is a string
    }
}
```

Relational Patterns

Relational patterns permit the programmer to express that an input value must satisfy a relational constraint when compared to a constant value:

C#

```
    public static LifeStage LifeStageAtAge(int age) => age switch
    {
        < 0 =>  LifeStage.Prenatal,
        < 2 =>  LifeStage.Infant,
        < 4 =>  LifeStage.Toddler,
        < 6 =>  LifeStage.EarlyChild,
        < 12 => LifeStage.MiddleChild,
        < 20 => LifeStage.Adolescent,
        < 40 => LifeStage.EarlyAdult,
        < 65 => LifeStage.MiddleAdult,
        _ =>    LifeStage.LateAdult,
    };
```

Relational patterns support the relational operators `<`, `<=`, `>`, and `>=` on all of the built-in types that support such binary relational operators with two operands of the same type in an expression. Specifically, we support all of these relational patterns for `sbyte`, `byte`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `char`, `float`, `double`, `decimal`, `nint`, and `nuint`.

antlr

```
primary_pattern
    : relational_pattern
    ;
relational_pattern
    : '<' relational_expression
    | '<=' relational_expression
    | '>' relational_expression
    | '>=' relational_expression
    ;
```

The expression is required to evaluate to a constant value. It is an error if that constant value is `double.NaN` or `float.NaN`. It is an error if the expression is a null constant.

When the input is a type for which a suitable built-in binary relational operator is defined that is applicable with the input as its left operand and the given constant as its right operand, the evaluation of that operator is taken as the meaning of the relational pattern. Otherwise we convert the input to the type of the expression using an explicit nullable or unboxing conversion. It is a compile-time error if no such conversion exists.

The pattern is considered not to match if the conversion fails. If the conversion succeeds then the result of the pattern-matching operation is the result of evaluating the expression `e OP v` where `e` is the converted input, `OP` is the relational operator, and `v` is the constant expression.

Pattern Combinators

Pattern *combinators* permit matching both of two different patterns using `and` (this can be extended to any number of patterns by the repeated use of `and`), either of two different patterns using `or` (ditto), or the *negation* of a pattern using `not`.

A common use of a combinator will be the idiom

```c#
if (e is not null) ...
```

More readable than the current idiom `e is object`, this pattern clearly expresses that one is checking for a non-null value.

The `and` and `or` combinators will be useful for testing ranges of values

```c#
bool IsLetter(char c) => c is >= 'a' and <= 'z' or >= 'A' and <= 'Z';
```

This example illustrates that `and` will have a higher parsing priority (i.e. will bind more closely) than `or`. The programmer can use the *parenthesized pattern* to make the precedence explicit:

```c#
bool IsLetter(char c) => c is (>= 'a' and <= 'z') or (>= 'A' and <= 'Z');
```

Like all patterns, these combinators can be used in any context in which a pattern is expected, including nested patterns, the *is-pattern-expression*, the *switch-expression*, and the pattern of a switch statement's case label.

```antlr
pattern
    : disjunctive_pattern
    ;
```

```
disjunctive_pattern
    : disjunctive_pattern 'or' conjunctive_pattern
    | conjunctive_pattern
    ;
conjunctive_pattern
    : conjunctive_pattern 'and' negated_pattern
    | negated_pattern
    ;
negated_pattern
    : 'not' negated_pattern
    | primary_pattern
    ;
primary_pattern
    : // all of the patterns forms previously defined
    ;
```

Change to 7.5.4.2 Grammar Ambiguities

Due to the introduction of the *type pattern*, it is possible for a generic type to appear before the token `=>`. We therefore add `=>` to the set of tokens listed in *7.5.4.2 Grammar Ambiguities* to permit disambiguation of the `<` that begins the type argument list. See also https://github.com/dotnet/roslyn/issues/47614.

Open Issues with Proposed Changes

Syntax for relational operators

Are `and`, `or`, and `not` some kind of contextual keyword? If so, is there a breaking change (e.g. compared to their use as a designator in a *declaration-pattern*).

Semantics (e.g. type) for relational operators

We expect to support all of the primitive types that can be compared in an expression using a relational operator. The meaning in simple cases is clear

c#

```
bool IsValidPercentage(int x) => x is >= 0 and <= 100;
```

But when the input is not such a primitive type, what type do we attempt to convert it to?

c#

```
bool IsValidPercentage(object x) => x is >= 0 and <= 100;
```

We have proposed that when the input type is already a comparable primitive, that is the type of the comparison. However, when the input is not a comparable primitive, we treat the relational as including an implicit type test to the type of the constant on the right-hand-side of the relational. If the programmer intends to support more than one input type, that must be done explicitly:

c#

```
bool IsValidPercentage(object x) => x is
    >= 0 and <= 100 or    // integer tests
    >= 0F and <= 100F or  // float tests
    >= 0D and <= 100D;    // double tests
```

Flowing type information from the left to the right of `and`

It has been suggested that when you write an `and` combinator, type information learned on the left about the top-level type could flow to the right. For example

C#

```
bool isSmallByte(object o) => o is byte and < 100;
```

Here, the *input type* to the second pattern is narrowed by the *type narrowing* requirements of left of the `and`. We would define type narrowing semantics for all patterns as follows. The *narrowed type* of a pattern `P` is defined as follows:

1. If `P` is a type pattern, the *narrowed type* is the type of the type pattern's type.
2. If `P` is a declaration pattern, the *narrowed type* is the type of the declaration pattern's type.
3. If `P` is a recursive pattern that gives an explicit type, the *narrowed type* is that type.
4. If `P` is matched via the rules for `ITuple`, the *narrowed type* is the type `System.Runtime.CompilerServices.ITuple`.
5. If `P` is a constant pattern where the constant is not the null constant and where the expression has no *constant expression conversion* to the *input type*, the *narrowed type* is the type of the constant.
6. If `P` is a relational pattern where the constant expression has no *constant expression conversion* to the *input type*, the *narrowed type* is the type of the constant.

7. If `P` is an `or` pattern, the *narrowed type* is the common type of the *narrowed type* of the subpatterns if such a common type exists. For this purpose, the common type algorithm considers only identity, boxing, and implicit reference conversions, and it considers all subpatterns of a sequence of `or` patterns (ignoring parenthesized patterns).
8. If `P` is an `and` pattern, the *narrowed type* is the *narrowed type* of the right pattern. Moreover, the *narrowed type* of the left pattern is the *input type* of the right pattern.
9. Otherwise the *narrowed type* of `P` is `P`'s input type.

Variable definitions and definite assignment

The addition of `or` and `not` patterns creates some interesting new problems around pattern variables and definite assignment. Since variables can normally be declared at most once, it would seem any pattern variable declared on one side of an `or` pattern would not be definitely assigned when the pattern matches. Similarly, a variable declared inside a `not` pattern would not be expected to be definitely assigned when the pattern matches. The simplest way to address this is to forbid declaring pattern variables in these contexts. However, this may be too restrictive. There are other approaches to consider.

One scenario that is worth considering is this

```
C#

if (e is not int i) return;
M(i); // is i definitely assigned here?
```

This does not work today because, for an *is-pattern-expression*, the pattern variables are considered *definitely assigned* only where the *is-pattern-expression* is true ("definitely assigned when true").

Supporting this would be simpler (from the programmer's perspective) than also adding support for a negated-condition `if` statement. Even if we add such support, programmers would wonder why the above snippet does not work. On the other hand, the same scenario in a `switch` makes less sense, as there is no corresponding point in the program where *definitely assigned when false* would be meaningful. Would we permit this in an *is-pattern-expression* but not in other contexts where patterns are permitted? That seems irregular.

Related to this is the problem of definite assignment in a *disjunctive-pattern*.

```C#
if (e is 0 or int i)
{
    M(i); // is i definitely assigned here?
}
```

We would only expect `i` to be definitely assigned when the input is not zero. But since we don't know whether the input is zero or not inside the block, `i` is not definitely assigned. However, what if we permit `i` to be declared in different mutually exclusive patterns?

```C#
if ((e1, e2) is (0, int i) or (int i, 0))
{
    M(i);
}
```

Here, the variable `i` is definitely assigned inside the block, and takes it value from the other element of the tuple when a zero element is found.

It has also been suggested to permit variables to be (multiply) defined in every case of a case block:

```C#
    case (0, int x):
    case (int x, 0):
        Console.WriteLine(x);
```

To make any of this work, we would have to carefully define where such multiple definitions are permitted and under what conditions such a variable is considered definitely assigned.

Should we elect to defer such work until later (which I advise), we could say in C# 9

- beneath a `not` or `or`, pattern variables may not be declared.

Then, we would have time to develop some experience that would provide insight into the possible value of relaxing that later.

Diagnostics, subsumption, and exhaustiveness

These new pattern forms introduce many new opportunities for diagnosable programmer error. We will need to decide what kinds of errors we will diagnose, and how to do so. Here are some examples:

```
C#
```

```
case >= 0 and <= 100D:
```

This case can never match (because the input cannot be both an `int` and a `double`). We already have an error when we detect a case that can never match, but its wording ("The switch case has already been handled by a previous case" and "The pattern has already been handled by a previous arm of the switch expression") may be misleading in new scenarios. We may have to modify the wording to just say that the pattern will never match the input.

```
C#
```

```
case 1 and 2:
```

Similarly, this would be an error because a value cannot be both `1` and `2`.

```
C#
```

```
case 1 or 2 or 3 or 1:
```

This case is possible to match, but the `or 1` at the end adds no meaning to the pattern. I suggest we should aim to produce an error whenever some conjunct or disjunct of a compound pattern does not either define a pattern variable or affect the set of matched values.

```
C#
```

```
case < 2: break;
case 0 or 1 or 2 or 3 or 4 or 5: break;
```

Here, `0 or 1 or` adds nothing to the second case, as those values would have been handled by the first case. This too deserves an error.

```
C#
```

```
byte b = ...;
int x = b switch { <100 => 0, 100 => 1, 101 => 2, >101 => 3 };
```

A switch expression such as this should be considered *exhaustive* (it handles all possible input values).

In C# 8.0, a switch expression with an input of type `byte` is only considered exhaustive if it contains a final arm whose pattern matches everything (a *discard-pattern* or *var-pattern*). Even a switch expression that has an arm for every distinct `byte` value is not considered exhaustive in C# 8. In order to properly handle exhaustiveness of relational patterns, we will have to handle this case too. This will technically be a breaking change, but no user is likely to notice.

Extended property patterns

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow property subpatterns to reference nested members, for instance:

```C#
if (e is MethodCallExpression { Method.Name: "MethodName" })
```

Instead of:

```C#
if (e is MethodCallExpression { Method: { Name: "MethodName" } })
```

Motivation

When you want to match a child property, nesting another recursive pattern adds too much noise which will hurt readability with no real advantage.

Detailed design

The *property_pattern* syntax is modified as follow:

```diff
```

```
property_pattern
  : type? property_pattern_clause simple_designation?
  ;

property_pattern_clause
  : '{' (subpattern (',' subpattern)* ','?)? '}'
  ;

subpattern
- : identifier ':' pattern
+ : subpattern_name ':' pattern
  ;

+subpattern_name
+ : identifier
+ | subpattern_name '.' identifier
+ ;
```

The receiver for each name lookup is the type of the previous member *T0*, starting from the *input type* of the *property_pattern*. if *T* is a nullable type, *T0* is its underlying type, otherwise *T0* is equal to *T*.

For example, a pattern of the form `{ Prop1.Prop2: pattern }` is exactly equivalent to `{ Prop1: { Prop2: pattern } }`.

Note that this will include the null check when *T* is a nullable value type or a reference type. This null check means that the nested properties available will be the properties of *T0*, not of *T*.

Repeated member paths are allowed. The compilation of pattern matching can take advantage of common parts of patterns.

Pattern match `Span<char>` on a constant string

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Permit pattern matching a `Span<char>` and a `ReadOnlySpan<char>` on a constant string.

Motivation

For perfomance, usage of `Span<char>` and `ReadOnlySpan<char>` is preferred over string in many scenarios. The framework has added many new APIs to allow you to use `ReadOnlySpan<char>` in place of a `string`.

A common operation on strings is to use a switch to test if it is a particular value, and the compiler optimizes such a switch. However there is currently no way to do the same on a `ReadOnlySpan<char>` efficiently, other than implementing the switch and the optimization manually.

In order to encourage adoption of `ReadOnlySpan<char>` we allow pattern matching a `ReadOnlySpan<char>`, on a constant `string`, thus also allowing it to be used in a switch.

C#

```
static bool Is123(ReadOnlySpan<char> s)
{
    return s is "123";
}
```

```
static bool IsABC(Span<char> s)
{
    return s switch { "ABC" => true, _ => false };
}
```

Detailed design

We alter the spec for constant patterns as follows (the proposed addition is shown in bold):

> A constant pattern tests the value of an expression against a constant value. The constant may be any constant expression, such as a literal, the name of a declared `const` variable, or an enumeration constant, or a `typeof` expression etc.
>
> If both *e* and *c* are of integral types, the pattern is considered matched if the result of the expression `e == c` is `true`.
>
> **If *e* is of type `System.Span<char>` or `System.ReadOnlySpan<char>`, and *c* is a constant string, and *c* does not have a constant value of `null`, then the pattern is considered matching if `System.MemoryExtensions.SequenceEqual<char>(e, System.MemoryExtensions.AsSpan(c))` returns `true`.**
>
> Otherwise the pattern is considered matching if `object.Equals(e, c)` returns `true`. In this case it is a compile-time error if the static type of *e* is not *pattern compatible* with the type of the constant.

Well-known members

`System.Span<T>` and `System.ReadOnlySpan<T>` are matched by name, must be `ref struct`s, and can be defined outside corlib.

`System.MemoryExtensions` is matched by name and can be defined outside corlib.

The signature of `System.MemoryExtensions.SequenceEqual` overloads must match:

- `public static bool SequenceEqual<T>(System.Span<T>, System.ReadOnlySpan<T>)`
- `public static bool SequenceEqual<T>(System.ReadOnlySpan<T>, System.ReadOnlySpan<T>)`

The signature of `System.MemoryExtensions.AsSpan` must match:

- `public static System.ReadOnlySpan<char> AsSpan(string)`

Methods with optional parameters are excluded from consideration.

Drawbacks

None

Alternatives

None

Unresolved questions

1. Should matching be defined independently from `MemoryExtensions.SequenceEqual()` etc.?

   > ... the pattern is considered matching if `e.Length == c.Length` and `e[i] == c[i]` for all characters in `e`.

   Recommendation: Define in terms of `MemoryExtensions.SequenceEqual()` for performance. If `MemoryExtensions` is missing, report compile error.

2. Should matching against `(string)null` be allowed?

   If so, should `(string)null` subsume `""` since `MemoryExtensions.AsSpan(null) == MemoryExtensions.AsSpan("")`?

   C#

   ```
   static bool IsEmpty(ReadOnlySpan<char> span)
   {
       return span switch
       {
           (string)null => true, // ok?
           "" => true,           // error: unreachable?
           _ => false,
       };
   }
   ```

   Recommendation: Constant pattern `(string)null` should be reported as an error.

3. Should the constant pattern match include a runtime type test of the expression value for `Span<char>` or `ReadOnlySpan<char>`?

```C#
static bool Is123<T>(Span<T> s)
{
    return s is "123"; // test for Span<char>?
}

static bool IsABC<T>(Span<T> s)
{
    return s is Span<char> and "ABC"; // ok?
}

static bool IsEmptyString<T>(T t) where T : ref struct
{
    return t is ""; // test for ReadOnlySpan<char>, Span<char>, string?
}
```

Recommendation: No implicit runtime type test for constant pattern. (`IsABC<T>()` example is allowed because the type test is explicit.)

4. Should subsumption consider constant string patterns, list patterns, and `Length` property pattern?

```C#
static int ToNum(ReadOnlySpan<char> s)
{
    return s switch
    {
        { Length: 0 } => 0,
        "" => 1,        // error: unreachable?
        ['A',..] => 2,
        "ABC" => 3,     // error: unreachable?
        _ => 4,
    };
}
```

Recommendation: Same subsumption behavior as used when the expression value is `string`. (Does that mean no subsumption between constant strings, list patterns, and `Length`, other than treating `[..]` as matching any?)

Design meetings

https://github.com/dotnet/csharplang/blob/master/meetings/2020/LDM-2020-10-07.md#readonlyspanchar-patterns

List patterns

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Lets you to match an array or a list with a sequence of patterns e.g. `array is [1, 2, 3]` will match an integer array of the length three with 1, 2, 3 as its elements, respectively.

Detailed design

The pattern syntax is modified as follow:

antlr

```
list_pattern_clause
  : '[' (pattern (',' pattern)* ','?)? ']'
  ;

list_pattern
  : list_pattern_clause simple_designation?
  ;

slice_pattern
  : '..' pattern?
  ;

primary_pattern
  : list_pattern
  | slice_pattern
  | // all of the pattern forms previously defined
  ;
```

There are two new patterns:

- The *list_pattern* is used to match elements.
- A *slice_pattern* is only permitted once and only directly in a *list_pattern_clause* and discards ***zero or more*** elements.

Pattern compatibility

A *list_pattern* is compatible with any type that is *countable* as well as *indexable* — it has an accessible indexer that takes an `Index` as an argument or otherwise an accessible indexer with a single `int` parameter. If both indexers are present, the former is preferred.

A *slice_pattern* with a subpattern is compatible with any type that is *countable* as well as *sliceable* — it has an accessible indexer that takes a `Range` as an argument or otherwise an accessible `Slice` method with two `int` parameters. If both are present, the former is preferred.

A *slice_pattern* without a subpattern is compatible with any type that is compatible with a *list_pattern*.

This set of rules is derived from the ***range indexer pattern*** .

Subsumption checking

Subsumption checking works just like positional patterns with ITuple - corresponding subpatterns are matched by position plus an additional node for testing length.

For example, the following code produces an error because both patterns yield the same DAG:

C#

```
case [_, .., 1]: // expr.Length is >= 2 && expr[^1] is 1
case [.., _, 1]: // expr.Length is >= 2 && expr[^1] is 1
```

Unlike:

C#

```
case [_, 1, ..]: // expr.Length is >= 2 && expr[1] is 1
case [.., 1, _]: // expr.Length is >= 2 && expr[^2] is 1
```

The order in which subpatterns are matched at runtime is unspecified, and a failed match may not attempt to match all subpatterns.

Given a specific length, it's possible that two subpatterns refer to the same element, in which case a test for this value is inserted into the decision DAG.

- For instance, `[_, >0, ..] or [.., <=0, _]` becomes `length >= 2 && ([1] > 0 || length == 3 || [^2] <= 0)` where the length value of 3 implies the other test.
- Conversely, `[_, >0, ..] and [.., <=0, _]` becomes `length >= 2 && [1] > 0 && length != 3 && [^2] <= 0` where the length value of 3 disallows the other test.

As a result, an error is produced for something like `case [.., p]: case [p]:` because at runtime, we're matching the same element in the second case.

If a slice subpattern matches a list or a length value, subpatterns are treated as if they were a direct subpattern of the containing list. For instance, `[..[1, 2, 3]]` subsumes a pattern of the form `[1, 2, 3]`.

The following assumptions are made on the members being used:

- The property that makes the type *countable* is assumed to always return a non-negative value, if and only if the type is *indexable*. For instance, the pattern `{ Length: -1 }` can never match an array.
- The member that makes the type *sliceable* is assumed to be well-behaved, that is, the return value is never null and that it is a proper subslice of the containing list.

The behavior of a pattern-matching operation is undefined if any of the above assumptions doesn't hold.

Lowering

A pattern of the form `expr is [1, 2, 3]` is equivalent to the following code:

C#

```
expr.Length is 3
&& expr[new Index(0, fromEnd: false)] is 1
&& expr[new Index(1, fromEnd: false)] is 2
&& expr[new Index(2, fromEnd: false)] is 3
```

A *slice_pattern* acts like a proper discard i.e. no tests will be emitted for such pattern, rather it only affects other nodes, namely the length and indexer. For instance, a pattern

of the form `expr is [1, .. var s, 3]` is equivalent to the following code (if compatible via explicit `Index` and `Range` support):

```
C#
```

```
expr.Length is >= 2
&& expr[new Index(0, fromEnd: false)] is 1
&& expr[new Range(new Index(1, fromEnd: false), new Index(1, fromEnd:
true))] is var s
&& expr[new Index(1, fromEnd: true)] is 3
```

The *input type* for the *slice_pattern* is the return type of the underlying `this[Range]` or `Slice` method with two exceptions: For `string` and arrays, `string.Substring` and `RuntimeHelpers.GetSubArray` will be used, respectively.

Unresolved questions

1. Should we support multi-dimensional arrays? (answer [LDM 2021-05-26]: Not supported. If we want to make a general MD-array focused release, we would want to revisit all the areas they're currently lacking, not just list patterns.)
2. Should we accept a general *pattern* following `..` in a *slice_pattern*? (answer [LDM 2021-05-26]: Yes, any pattern is allowed after a slice.)
3. By this definition, the pattern `[..]` tests for `expr.Length >= 0`. Should we omit such test, assuming `Length` is always non-negative? (answer [LDM 2021-05-26]: `[..]` will not emit a Length check)

Ranges

Article • 11/09/2021

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

This feature is about delivering two new operators that allow constructing `System.Index` and `System.Range` objects, and using them to index/slice collections at runtime.

Overview

Well-known types and members

To use the new syntactic forms for `System.Index` and `System.Range`, new well-known types and members may be necessary, depending on which syntactic forms are used.

To use the "hat" operator (`^`), the following is required

```C#
namespace System
{
    public readonly struct Index
    {
        public Index(int value, bool fromEnd);
    }
}
```

To use the `System.Index` type as an argument in an array element access, the following member is required:

C#

```
int System.Index.GetOffset(int length);
```

The `..` syntax for `System.Range` will require the `System.Range` type, as well as one or more of the following members:

C#

```
namespace System
{
    public readonly struct Range
    {
        public Range(System.Index start, System.Index end);
        public static Range StartAt(System.Index start);
        public static Range EndAt(System.Index end);
        public static Range All { get; }
    }
}
```

The `..` syntax allows for either, both, or none of its arguments to be absent. Regardless of the number of arguments, the `Range` constructor is always sufficient for using the `Range` syntax. However, if any of the other members are present and one or more of the `..` arguments are missing, the appropriate member may be substituted.

Finally, for a value of type `System.Range` to be used in an array element access expression, the following member must be present:

C#

```
namespace System.Runtime.CompilerServices
{
    public static class RuntimeHelpers
    {
        public static T[] GetSubArray<T>(T[] array, System.Range range);
    }
}
```

System.Index

C# has no way of indexing a collection from the end, but rather most indexers use the "from start" notion, or do a "length - i" expression. We introduce a new Index expression that means "from the end". The feature will introduce a new unary prefix "hat" operator. Its single operand must be convertible to `System.Int32`. It will be lowered into the appropriate `System.Index` factory method call.

We augment the grammar for *unary_expression* with the following additional syntax form:

```antlr
unary_expression
    : '^' unary_expression
    ;
```

We call this the *index from end* operator. The predefined *index from end* operators are as follows:

```C#
System.Index operator ^(int fromEnd);
```

The behavior of this operator is only defined for input values greater than or equal to zero.

Examples:

```C#
var array = new int[] { 1, 2, 3, 4, 5 };
var thirdItem = array[2];    // array[2]
var lastItem = array[^1];    // array[new Index(1, fromEnd: true)]
```

System.Range

C# has no syntactic way to access "ranges" or "slices" of collections. Usually users are forced to implement complex structures to filter/operate on slices of memory, or resort to LINQ methods like `list.Skip(5).Take(2)`. With the addition of `System.Span<T>` and other similar types, it becomes more important to have this kind of operation supported on a deeper level in the language/runtime, and have the interface unified.

The language will introduce a new range operator `x..y`. It is a binary infix operator that accepts two expressions. Either operand can be omitted (examples below), and they have to be convertible to `System.Index`. It will be lowered to the appropriate `System.Range` factory method call.

We replace the C# grammar rules for *multiplicative_expression* with the following (in order to introduce a new precedence level):

antlr

```
range_expression
    : unary_expression
    | range_expression? '..' range_expression?
    ;

multiplicative_expression
    : range_expression
    | multiplicative_expression '*' range_expression
    | multiplicative_expression '/' range_expression
    | multiplicative_expression '%' range_expression
    ;
```

All forms of the *range operator* have the same precedence. This new precedence group is lower than the *unary operators* and higher than the *multiplicative arithmetic operators*.

We call the `..` operator the *range operator*. The built-in range operator can roughly be understood to correspond to the invocation of a built-in operator of this form:

C#

```
System.Range operator ..(Index start = 0, Index end = ^0);
```

Examples:

C#

```
var array = new int[] { 1, 2, 3, 4, 5 };
var slice1 = array[2..^3];    // array[new Range(2, new Index(3, fromEnd:
true))]
var slice2 = array[..^3];     // array[Range.EndAt(new Index(3, fromEnd:
true))]
var slice3 = array[2..];      // array[Range.StartAt(2)]
var slice4 = array[..];       // array[Range.All]
```

Moreover, `System.Index` should have an implicit conversion from `System.Int32`, in order to avoid the need to overload mixing integers and indexes over multi-dimensional signatures.

Adding Index and Range support to existing library types

Implicit Index support

The language will provide an instance indexer member with a single parameter of type `Index` for types which meet the following criteria:

- The type is Countable.
- The type has an accessible instance indexer which takes a single `int` as the argument.
- The type does not have an accessible instance indexer which takes an `Index` as the first parameter. The `Index` must be the only parameter or the remaining parameters must be optional.

A type is ***Countable*** if it has a property named `Length` or `Count` with an accessible getter and a return type of `int`. The language can make use of this property to convert an expression of type `Index` into an `int` at the point of the expression without the need to use the type `Index` at all. In case both `Length` and `Count` are present, `Length` will be preferred. For simplicity going forward, the proposal will use the name `Length` to represent `Count` or `Length`.

For such types, the language will act as if there is an indexer member of the form `T this[Index index]` where `T` is the return type of the `int` based indexer including any `ref` style annotations. The new member will have the same `get` and `set` members with matching accessibility as the `int` indexer.

The new indexer will be implemented by converting the argument of type `Index` into an `int` and emitting a call to the `int` based indexer. For discussion purposes, let's use the example of `receiver[expr]`. The conversion of `expr` to `int` will occur as follows:

- When the argument is of the form `^expr2` and the type of `expr2` is `int`, it will be translated to `receiver.Length - expr2`.
- Otherwise, it will be translated as `expr.GetOffset(receiver.Length)`.

Regardless of the specific conversion strategy, the order of evaluation should be equivalent to the following:

1. `receiver` is evaluated;
2. `expr` is evaluated;
3. `length` is evaluated, if needed;
4. the `int` based indexer is invoked.

This allows for developers to use the `Index` feature on existing types without the need for modification. For example:

```
C#
```

```
List<char> list = ...;
var value = list[^1];

// Gets translated to
var value = list[list.Count - 1];
```

The `receiver` and `Length` expressions will be spilled as appropriate to ensure any side effects are only executed once. For example:

C#

```
class Collection {
    private int[] _array = new[] { 1, 2, 3 };

    public int Length {
        get {
            Console.Write("Length ");
            return _array.Length;
        }
    }

    public int this[int index] => _array[index];
}

class SideEffect {
    Collection Get() {
        Console.Write("Get ");
        return new Collection();
    }

    void Use() {
        int i = Get()[^1];
        Console.WriteLine(i);
    }
}
```

This code will print "Get Length 3".

Implicit Range support

The language will provide an instance indexer member with a single parameter of type `Range` for types which meet the following criteria:

- The type is Countable.
- The type has an accessible member named `Slice` which has two parameters of type `int`.

- The type does not have an instance indexer which takes a single `Range` as the first parameter. The `Range` must be the only parameter or the remaining parameters must be optional.

For such types, the language will bind as if there is an indexer member of the form `T this[Range range]` where `T` is the return type of the `Slice` method including any `ref` style annotations. The new member will also have matching accessibility with `Slice`.

When the `Range` based indexer is bound on an expression named `receiver`, it will be lowered by converting the `Range` expression into two values that are then passed to the `Slice` method. For discussion purposes, let's use the example of `receiver[expr]`.

The first argument of `Slice` will be obtained by converting the range typed expression in the following way:

- When `expr` is of the form `expr1..expr2` (where `expr2` can be omitted) and `expr1` has type `int`, then it will be emitted as `expr1`.
- When `expr` is of the form `^expr1..expr2` (where `expr2` can be omitted), then it will be emitted as `receiver.Length - expr1`.
- When `expr` is of the form `..expr2` (where `expr2` can be omitted), then it will be emitted as `0`.
- Otherwise, it will be emitted as `expr.Start.GetOffset(receiver.Length)`.

This value will be re-used in the calculation of the second `Slice` argument. When doing so it will be referred to as `start`. The second argument of `Slice` will be obtained by converting the range typed expression in the following way:

- When `expr` is of the form `expr1..expr2` (where `expr1` can be omitted) and `expr2` has type `int`, then it will be emitted as `expr2 - start`.
- When `expr` is of the form `expr1..^expr2` (where `expr1` can be omitted), then it will be emitted as `(receiver.Length - expr2) - start`.
- When `expr` is of the form `expr1..` (where `expr1` can be omitted), then it will be emitted as `receiver.Length - start`.
- Otherwise, it will be emitted as `expr.End.GetOffset(receiver.Length) - start`.

Regardless of the specific conversion strategy, the order of evaluation should be equivalent to the following:

1. `receiver` is evaluated;
2. `expr` is evaluated;
3. `length` is evaluated, if needed;

4. the `Slice` method is invoked.

The `receiver`, `expr`, and `length` expressions will be spilled as appropriate to ensure any side effects are only executed once. For example:

```C#
class Collection {
    private int[] _array = new[] { 1, 2, 3 };

    public int Length {
        get {
            Console.Write("Length ");
            return _array.Length;
        }
    }

    public int[] Slice(int start, int length) {
        var slice = new int[length];
        Array.Copy(_array, start, slice, 0, length);
        return slice;
    }
}

class SideEffect {
    Collection Get() {
        Console.Write("Get ");
        return new Collection();
    }

    void Use() {
        var array = Get()[0..2];
        Console.WriteLine(array.Length);
    }
}
```

This code will print "Get Length 2".

The language will special case the following known types:

- `string`: the method `Substring` will be used instead of `Slice`.
- `array`: the method `System.Runtime.CompilerServices.RuntimeHelpers.GetSubArray` will be used instead of `Slice`.

Alternatives

The new operators (`^` and `..`) are syntactic sugar. The functionality can be implemented by explicit calls to `System.Index` and `System.Range` factory methods, but it

will result in a lot more boilerplate code, and the experience will be unintuitive.

IL Representation

These two operators will be lowered to regular indexer/method calls, with no change in subsequent compiler layers.

Runtime behavior

- Compiler can optimize indexers for built-in types like arrays and strings, and lower the indexing to the appropriate existing methods.
- `System.Index` will throw if constructed with a negative value.
- `^0` does not throw, but it translates to the length of the collection/enumerable it is supplied to.
- `Range.All` is semantically equivalent to `0..^0`, and can be deconstructed to these indices.

Considerations

Detect Indexable based on ICollection

The inspiration for this behavior was collection initializers. Using the structure of a type to convey that it had opted into a feature. In the case of collection initializers types can opt into the feature by implementing the interface `IEnumerable` (non generic).

This proposal initially required that types implement `ICollection` in order to qualify as Indexable. That required a number of special cases though:

- `ref struct`: these cannot implement interfaces yet types like `Span<T>` are ideal for index / range support.
- `string`: does not implement `ICollection` and adding that `interface` has a large cost.

This means to support key types special casing is already needed. The special casing of `string` is less interesting as the language does this in other areas (`foreach` lowering, constants, etc ...). The special casing of `ref struct` is more concerning as it's special casing an entire class of types. They get labeled as Indexable if they simply have a property named `Count` with a return type of `int`.

After consideration the design was normalized to say that any type which has a property `Count` / `Length` with a return type of `int` is Indexable. That removes all special casing, even for `string` and arrays.

Detect just Count

Detecting on the property names `Count` or `Length` does complicate the design a bit. Picking just one to standardize though is not sufficient as it ends up excluding a large number of types:

- Use `Length`: excludes pretty much every collection in System.Collections and sub-namespaces. Those tend to derive from `ICollection` and hence prefer `Count` over length.
- Use `Count`: excludes `string`, arrays, `Span<T>` and most `ref struct` based types

The extra complication on the initial detection of Indexable types is outweighed by its simplification in other aspects.

Choice of Slice as a name

The name `Slice` was chosen as it's the de-facto standard name for slice style operations in .NET. Starting with netcoreapp2.1 all span style types use the name `Slice` for slicing operations. Prior to netcoreapp2.1 there really aren't any examples of slicing to look to for an example. Types like `List<T>`, `ArraySegment<T>`, `SortedList<T>` would've been ideal for slicing but the concept didn't exist when types were added.

Thus, `Slice` being the sole example, it was chosen as the name.

Index target type conversion

Another way to view the `Index` transformation in an indexer expression is as a target type conversion. Instead of binding as if there is a member of the form `return_type this[Index]`, the language instead assigns a target typed conversion to `int`.

This concept could be generalized to all member access on Countable types. Whenever an expression with type `Index` is used as an argument to an instance member invocation and the receiver is Countable then the expression will have a target type conversion to `int`. The member invocations applicable for this conversion include methods, indexers, properties, extension methods, etc ... Only constructors are excluded as they have no receiver.

The target type conversion will be implemented as follows for any expression which has a type of `Index`. For discussion purposes lets use the example of `receiver[expr]`:

- When `expr` is of the form `^expr2` and the type of `expr2` is `int`, it will be translated to `receiver.Length - expr2`.
- Otherwise, it will be translated as `expr.GetOffset(receiver.Length)`.

The `receiver` and `Length` expressions will be spilled as appropriate to ensure any side effects are only executed once. For example:

C#

```
class Collection {
    private int[] _array = new[] { 1, 2, 3 };

    public int Length {
        get {
            Console.Write("Length ");
            return _array.Length;
        }
    }

    public int GetAt(int index) => _array[index];
}

class SideEffect {
    Collection Get() {
        Console.Write("Get ");
        return new Collection();
    }

    void Use() {
        int i = Get().GetAt(^1);
        Console.WriteLine(i);
    }
}
```

This code will print "Get Length 3".

This feature would be beneficial to any member which had a parameter that represented an index. For example `List<T>.InsertAt`. This also has the potential for confusion as the language can't give any guidance as to whether or not an expression is meant for indexing. All it can do is convert any `Index` expression to `int` when invoking a member on a Countable type.

Restrictions:

- This conversion is only applicable when the expression with type `Index` is directly an argument to the member. It would not apply to any nested expressions.

Decisions made during implementation

- All members in the pattern must be instance members
- If a Length method is found but it has the wrong return type, continue looking for Count
- The indexer used for the Index pattern must have exactly one int parameter
- The Slice method used for the Range pattern must have exactly two int parameters
- When looking for the pattern members, we look for original definitions, not constructed members

Design meetings

- Jan 10, 2018
- Jan 18, 2018
- Jan 22, 2018
- Dec 3, 2018
- Mar 25, 2019
- April 1st, 2019
- April 15, 2019

Null coalescing assignment

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Simplifies a common coding pattern where a variable is assigned a value if it is null.

As part of this proposal, we will also loosen the type requirements on `??` to allow an expression whose type is an unconstrained type parameter to be used on the left-hand side.

Motivation

It is common to see code of the form

```C#
if (variable == null)
{
    variable = expression;
}
```

This proposal adds a non-overloadable binary operator to the language that performs this function.

There have been at least eight separate community requests for this feature.

Detailed design

We add a new form of assignment operator

```antlr
assignment_operator
    : '??='
    ;
```

Which follows the existing semantic rules for compound assignment operators (§11.18.3), except that we elide the assignment if the left-hand side is non-null. The rules for this feature are as follows.

Given `a ??= b`, where `A` is the type of `a`, `B` is the type of `b`, and `A0` is the underlying type of `A` if `A` is a nullable value type:

1. If `A` does not exist or is a non-nullable value type, a compile-time error occurs.
2. If `B` is not implicitly convertible to `A` or `A0` (if `A0` exists), a compile-time error occurs.
3. If `A0` exists and `B` is implicitly convertible to `A0`, and `B` is not dynamic, then the type of `a ??= b` is `A0`. `a ??= b` is evaluated at runtime as:

   ```C#
   var tmp = a.GetValueOrDefault();
   if (!a.HasValue) { tmp = b; a = tmp; }
   tmp
   ```

   Except that `a` is only evaluated once.
4. Otherwise, the type of `a ??= b` is `A`. `a ??= b` is evaluated at runtime as `a ?? (a = b)`, except that `a` is only evaluated once.

For the relaxation of the type requirements of `??`, we update the spec where it currently states that, given `a ?? b`, where `A` is the type of `a`:

> 1. If A exists and is not a nullable type or a reference type, a compile-time error occurs.

We relax this requirement to:

1. If A exists and is a non-nullable value type, a compile-time error occurs.

This allows the null coalescing operator to work on unconstrained type parameters, as the unconstrained type parameter T exists, is not a nullable type, and is not a reference type.

Drawbacks

As with any language feature, we must question whether the additional complexity to the language is repaid in the additional clarity offered to the body of C# programs that would benefit from the feature.

Alternatives

The programmer can write `(x = x ?? y)`, `if (x == null) x = y;`, or `x ?? (x = y)` by hand.

Unresolved questions

- [ ] Requires LDM review
- [ ] Should we also support `&&=` and `||=` operators?

Design meetings

None.

Target-typed `new` expressions

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Do not require type specification for constructors when the type is known.

Motivation

Allow field initialization without duplicating the type.

```C#
Dictionary<string, List<int>> field = new() {
    { "item1", new() { 1, 2, 3 } }
};
```

Allow omitting the type when it can be inferred from usage.

```C#
XmlReader.Create(reader, new() { IgnoreWhitespace = true });
```

Instantiate an object without spelling out the type.

```C#
private readonly static object s_syncObj = new();
```

Specification

A new syntactic form, *target_typed_new* of the *object_creation_expression* is accepted in which the *type* is optional.

antlr

```
object_creation_expression
    : 'new' type '(' argument_list? ')' object_or_collection_initializer?
    | 'new' type object_or_collection_initializer
    | target_typed_new
    ;
target_typed_new
    : 'new' '(' argument_list? ')' object_or_collection_initializer?
    ;
```

A *target_typed_new* expression does not have a type. However, there is a new *object creation conversion* that is an implicit conversion from expression, that exists from a *target_typed_new* to every type.

Given a target type `T`, the type `T0` is `T`'s underlying type if `T` is an instance of `System.Nullable`. Otherwise `T0` is `T`. The meaning of a *target_typed_new* expression that is converted to the type `T` is the same as the meaning of a corresponding *object_creation_expression* that specifies `T0` as the type.

It is a compile-time error if a *target_typed_new* is used as an operand of a unary or binary operator, or if it is used where it is not subject to an *object creation conversion*.

> **Open Issue:** should we allow delegates and tuples as the target-type?

The above rules include delegates (a reference type) and tuples (a struct type). Although both types are constructible, if the type is inferable, an anonymous function or a tuple literal can already be used.

C#

```
(int a, int b) t = new(1, 2); // "new" is redundant
Action a = new(() => {}); // "new" is redundant

(int a, int b) t = new(); // OK; same as (0, 0)
Action a = new(); // no constructor found
```

Miscellaneous

The following are consequences of the specification:

- `throw new()` is allowed (the target type is `System.Exception`)
- Target-typed `new` is not allowed with binary operators.
- It is disallowed when there is no type to target: unary operators, collection of a `foreach`, in a `using`, in a deconstruction, in an `await` expression, as an anonymous type property (`new { Prop = new() }`), in a `lock` statement, in a `sizeof`, in a `fixed` statement, in a member access (`new().field`), in a dynamically dispatched operation (`someDynamic.Method(new())`), in a LINQ query, as the operand of the `is` operator, as the left operand of the `??` operator, ...
- It is also disallowed as a `ref`.
- The following kinds of types are not permitted as targets of the conversion
  - **Enum types:** `new()` will work (as `new Enum()` works to give the default value), but `new(1)` will not work as enum types do not have a constructor.
  - **Interface types:** This would work the same as the corresponding creation expression for COM types.
  - **Array types:** arrays need a special syntax to provide the length.
  - **dynamic:** we don't allow `new dynamic()`, so we don't allow `new()` with `dynamic` as a target type.
  - **tuples:** These have the same meaning as an object creation using the underlying type.
  - All the other types that are not permitted in the *object_creation_expression* are excluded as well, for instance, pointer types.

Drawbacks

There were some concerns with target-typed `new` creating new categories of breaking changes, but we already have that with `null` and `default`, and that has not been a significant problem.

Alternatives

Most of complaints about types being too long to duplicate in field initialization is about *type arguments* not the type itself, we could infer only type arguments like `new Dictionary(...)` (or similar) and infer type arguments locally from arguments or the collection initializer.

Questions

- Should we forbid usages in expression trees? (no)
- How the feature interacts with `dynamic` arguments? (no special treatment)
- How IntelliSense should work with `new()`? (only when there is a single target-type)

Design meetings

- LDM-2017-10-18
- LDM-2018-05-21
- LDM-2018-06-25
- LDM-2018-08-22
- LDM-2018-10-17
- LDM-2020-03-25

Target-Typed Conditional Expression

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Conditional Expression Conversion

For a conditional expression `c ? e1 : e2`, when

1. there is no common type for `e1` and `e2`, or
2. for which a common type exists but one of the expressions `e1` or `e2` has no implicit conversion to that type

we define a new implicit *conditional expression conversion* that permits an implicit conversion from the conditional expression to any type `T` for which there is a conversion-from-expression from `e1` to `T` and also from `e2` to `T`. It is an error if a conditional expression neither has a common type between `e1` and `e2` nor is subject to a *conditional expression conversion*.

Better Conversion from Expression

We change

> ### Better conversion from expression
>
> Given an implicit conversion `C1` that converts from an expression `E` to a type `T1`, and an implicit conversion `C2` that converts from an expression `E` to a type `T2`, `C1` is a ***better conversion*** than `C2` if `E` does not exactly match `T2` and at least one of the following holds:

> - `E` exactly matches `T1` (§11.6.4.4)
> - `T1` is a better conversion target than `T2` (§11.6.4.6)

to

> ### Better conversion from expression
>
> Given an implicit conversion `C1` that converts from an expression `E` to a type `T1`, and an implicit conversion `C2` that converts from an expression `E` to a type `T2`, `C1` is a ***better conversion*** than `C2` if `E` does not exactly match `T2` and at least one of the following holds:
>
> - `E` exactly matches `T1` (§11.6.4.4)
> - **`C1` is not a *conditional expression conversion* and `C2` is a *conditional expression conversion***.
> - `T1` is a better conversion target than `T2` (§11.6.4.6) **and either `C1` and `C2` are both *conditional expression conversions* or neither is a *conditional expression conversion***.

Cast Expression

The current C# language specification says

> A *cast_expression* of the form `(T)E`, where `T` is a *type* and `E` is a *unary_expression*, performs an explicit conversion (§10.3) of the value of `E` to type `T`.

In the presence of the *conditional expression conversion* there may be more than one possible conversion from `E` to `T`. With the addition of *conditional expression conversion*, we prefer any other conversion to a *conditional expression conversion*, and use the *conditional expression conversion* only as a last resort.

Design Notes

The reason for the change to *Better conversion from expression* is to handle a case such as this:

```
C#
```

```
M(b ? 1 : 2);
```

```
void M(short);
void M(long);
```

This approach does have two small downsides. First, it is not quite the same as the switch expression:

C#

```
M(b ? 1 : 2); // calls M(long)
M(b switch { true => 1, false => 2 }); // calls M(short)
```

This is still a breaking change, but its scope is less likely to affect real programs:

C#

```
M(b ? 1 : 2, 1); // calls M(long, long) without this feature; ambiguous with
this feature.

M(short, short);
M(long, long);
```

This becomes ambiguous because the conversion to `long` is better for the first argument (because it does not use the *conditional expression conversion*), but the conversion to `short` is better for the second argument (because `short` is a *better conversion target* than `long`). This breaking change seems less serious because it does not silently change the behavior of an existing program.

The reason for the notes on the cast expression is to handle a case such as this:

C#

```
_ = (short)(b ? 1 : 2);
```

This program currently uses the explicit conversion from `int` to `short`, and we want to preserve the current language meaning of this program. The change would be unobservable at runtime, but with the following program the change would be observable:

C#

```
_ = (A)(b ? c : d);
```

where `c` is of type `C`, `d` is of type `D`, and there is an implicit user-defined conversion from `C` to `D`, and an implicit user-defined conversion from `D` to `A`, and an implicit user-defined conversion from `C` to `A`. If this code is compiled before C# 9.0, when `b` is true we convert from `c` to `D` then to `A`. If we use the *conditional expression conversion*, then when `b` is true we convert from `c` to `A` directly, which executes a different sequence of user code. Therefore we treat the *conditional expression conversion* as a last resort in a cast, to preserve existing behavior.

Extension `GetEnumerator` support for `foreach` loops.

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow foreach loops to recognize an extension method GetEnumerator method that otherwise satisfies the foreach pattern, and loop over the expression when it would otherwise be an error.

Motivation

This will bring foreach inline with how other features in C# are implemented, including async and pattern-based deconstruction.

Detailed design

The spec change is relatively straightforward. We modify `The foreach statement` section to this text:

> The compile-time processing of a foreach statement first determines the ***collection type***, ***enumerator type*** and ***element type*** of the expression. This determination proceeds as follows:
>
> - If the type `X` of *expression* is an array type then there is an implicit reference conversion from `X` to the `IEnumerable` interface (since `System.Array` implements this interface). The ***collection type*** is the `IEnumerable` interface, the

enumerator type is the `IEnumerator` interface and the ***element type*** is the element type of the array type `X`.

- If the type `X` of *expression* is `dynamic` then there is an implicit conversion from *expression* to the `IEnumerable` interface (§10.2.10). The ***collection type*** is the `IEnumerable` interface and the ***enumerator type*** is the `IEnumerator` interface. If the `var` identifier is given as the *local_variable_type* then the ***element type*** is `dynamic`, otherwise it is `object`.

- Otherwise, determine whether the type `X` has an appropriate `GetEnumerator` method:
  - Perform member lookup on the type `X` with identifier `GetEnumerator` and no type arguments. If the member lookup does not produce a match, or it produces an ambiguity, or produces a match that is not a method group, check for an enumerable interface as described below. It is recommended that a warning be issued if member lookup produces anything except a method group or no match.
  - Perform overload resolution using the resulting method group and an empty argument list. If overload resolution results in no applicable methods, results in an ambiguity, or results in a single best method but that method is either static or not public, check for an enumerable interface as described below. It is recommended that a warning be issued if overload resolution produces anything except an unambiguous public instance method or no applicable methods.
  - If the return type `E` of the `GetEnumerator` method is not a class, struct or interface type, an error is produced and no further steps are taken.
  - Member lookup is performed on `E` with the identifier `Current` and no type arguments. If the member lookup produces no match, the result is an error, or the result is anything except a public instance property that permits reading, an error is produced and no further steps are taken.
  - Member lookup is performed on `E` with the identifier `MoveNext` and no type arguments. If the member lookup produces no match, the result is an error, or the result is anything except a method group, an error is produced and no further steps are taken.
  - Overload resolution is performed on the method group with an empty argument list. If overload resolution results in no applicable methods, results in an ambiguity, or results in a single best method but that method is either static or not public, or its return type is not `bool`, an error is produced and no further steps are taken.

    - The ***collection type*** is `X`, the ***enumerator type*** is `E`, and the ***element type*** is the type of the `Current` property.

- Otherwise, check for an enumerable interface:
    - If among all the types `Ti` for which there is an implicit conversion from `X` to `IEnumerable<Ti>`, there is a unique type `T` such that `T` is not `dynamic` and for all the other `Ti` there is an implicit conversion from `IEnumerable<T>` to `IEnumerable<Ti>`, then the ***collection type*** is the interface `IEnumerable<T>`, the ***enumerator type*** is the interface `IEnumerator<T>`, and the ***element type*** is `T`.
    - Otherwise, if there is more than one such type `T`, then an error is produced and no further steps are taken.
    - Otherwise, if there is an implicit conversion from `X` to the `System.Collections.IEnumerable` interface, then the ***collection type*** is this interface, the ***enumerator type*** is the interface `System.Collections.IEnumerator`, and the ***element type*** is `object`.

- Otherwise, determine whether the type 'X' has an appropriate `GetEnumerator` extension method:
    - Perform extension method lookup on the type `X` with identifier `GetEnumerator`. If the member lookup does not produce a match, or it produces an ambiguity, or produces a match which is not a method group, an error is produced and no further steps are taken. It is recommended that a warning be issues if member lookup produces anything except a method group or no match.
    - Perform overload resolution using the resulting method group and a single argument of type `X`. If overload resolution produces no applicable methods, results in an ambiguity, or results in a single best method but that method is not accessible, an error is produced an no further steps are taken.
        - This resolution permits the first argument to be passed by ref if `X` is a struct type, and the ref kind is `in`.
    - If the return type `E` of the `GetEnumerator` method is not a class, struct or interface type, an error is produced and no further steps are taken.
    - Member lookup is performed on `E` with the identifier `Current` and no type arguments. If the member lookup produces no match, the result is an error, or the result is anything except a public instance property that permits reading, an error is produced and no further steps are taken.
    - Member lookup is performed on `E` with the identifier `MoveNext` and no type arguments. If the member lookup produces no match, the result is an error,

or the result is anything except a method group, an error is produced and no further steps are taken.
    - Overload resolution is performed on the method group with an empty argument list. If overload resolution results in no applicable methods, results in an ambiguity, or results in a single best method but that method is either static or not public, or its return type is not `bool`, an error is produced and no further steps are taken.
    - The ***collection type*** is `X`, the ***enumerator type*** is `E`, and the ***element type*** is the type of the `Current` property.

- Otherwise, an error is produced and no further steps are taken.

For `await foreach`, the rules are similarly modified. The only change that is required to that spec is removing the `Extension methods do not contribute.` line from the description, as the rest of that spec is based on the above rules with different names substituted for the pattern methods.

Drawbacks

Every change adds additional complexity to the language, and this potentially allows things that weren't designed to be `foreach`ed to be `foreach`ed, like `Range`.

Alternatives

Doing nothing.

Unresolved questions

None at this point.

Collection expressions

Article • 03/30/2024

> **Note**
>
> This article is a feature specification. The specification serves as the design document for the feature. It includes proposed specification changes, along with information needed during the design and development of the feature. These articles are published until the proposed spec changes are finalized and incorporated in the current ECMA specification.
>
> There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent language design meeting (LDM) notes.
>
> You can learn more about the process for adopting feature speclets into the C# language standard in the article on the specifications.

Summary

Collection expressions introduce a new terse syntax, `[e1, e2, e3, etc]`, to create common collection values. Inlining other collections into these values is possible using a spread element `..e` like so: `[e1, ..c2, e2, ..c2]`.

Several collection-like types can be created without requiring external BCL support. These types are:

- Array types, such as `int[]`.
- Span<T> and ReadOnlySpan<T>.
- Types that support collection initializers, such as List<T>.

Further support is present for collection-like types not covered under the above through a new attribute and API pattern that can be adopted directly on the type itself.

Motivation

- Collection-like values are hugely present in programming, algorithms, and especially in the C#/.NET ecosystem. Nearly all programs will utilize these values to store data and send or receive data from other components. Currently, almost all C# programs must use many different and unfortunately verbose approaches to

create instances of such values. Some approaches also have performance drawbacks. Here are some common examples:
    - Arrays, which require either `new Type[]` or `new[]` before the `{ ... }` values.
    - Spans, which may use `stackalloc` and other cumbersome constructs.
    - Collection initializers, which require syntax like `new List<T>` (lacking inference of a possibly verbose `T`) prior to their values, and which can cause multiple reallocations of memory because they use N `.Add` invocations without supplying an initial capacity.
    - Immutable collections, which require syntax like `ImmutableArray.Create(...)` to initialize the values, and which can cause intermediary allocations and data copying. More efficient construction forms (like `ImmutableArray.CreateBuilder`) are unwieldy and still produce unavoidable garbage.

- Looking at the surrounding ecosystem, we also find examples everywhere of list creation being more convenient and pleasant to use. TypeScript, Dart, Swift, Elm, Python, and more opt for a succinct syntax for this purpose, with widespread usage, and to great effect. Cursory investigations have revealed no substantive problems arising in those ecosystems with having these literals built in.

- C# has also added list patterns in C# 11. This pattern allows matching and deconstruction of list-like values using a clean and intuitive syntax. However, unlike almost all other pattern constructs, this matching/deconstruction syntax lacks the corresponding construction syntax.

- Getting the best performance for constructing each collection type can be tricky. Simple solutions often waste both CPU and memory. Having a literal form allows for maximum flexibility from the compiler implementation to optimize the literal to produce at least as good a result as a user could provide, but with simple code. Very often the compiler will be able to do better, and the specification aims to allow the implementation large amounts of leeway in terms of implementation strategy to ensure this.

An inclusive solution is needed for C#. It should meet the vast majority of casse for customers in terms of the collection-like types and values they already have. It should also feel natural in the language and mirror the work done in pattern matching.

This leads to a natural conclusion that the syntax should be like `[e1, e2, e3, e-etc]` or `[e1, ..c2, e2]`, which correspond to the pattern equivalents of `[p1, p2, p3, p-etc]` and `[p1, ..p2, p3]`.

Detailed design

The following grammar productions are added:

```diff
primary_no_array_creation_expression
  ...
+ | collection_expression
  ;

+ collection_expression
  : '[' ']'
  | '[' collection_element ( ',' collection_element )* ']'
  ;

+ collection_element
  : expression_element
  | spread_element
  ;

+ expression_element
  : expression
  ;

+ spread_element
  : '..' expression
  ;
```

Collection literals are target-typed.

Spec clarifications

- For brevity, `collection_expression` will be referred to as "literal" in the following sections.
- `expression_element` instances will commonly be referred to as `e1`, `e_n`, etc.
- `spread_element` instances will commonly be referred to as `..s1`, `..s_n`, etc.
- *span type* means either `Span<T>` or `ReadOnlySpan<T>`.
- Literals will commonly be shown as `[e1, ..s1, e2, ..s2, etc]` to convey any number of elements in any order. Importantly, this form will be used to represent all cases such as:
  - Empty literals `[]`
  - Literals with no `expression_element` in them.
  - Literals with no `spread_element` in them.
  - Literals with arbitrary ordering of any element type.

- The *iteration type* of `..s_n` is the type of the *iteration variable* determined as if `s_n` were used as the expression being iterated over in a foreach_statement.
- Variables starting with `__name` are used to represent the results of the evaluation of `name`, stored in a location so that it is only evaluated once. For example `__e1` is the evaluation of `e1`.
- `List<T>`, `IEnumerable<T>`, etc. refer to the respective types in the `System.Collections.Generic` namespace.
- The specification defines a translation of the literal to existing C# constructs. Similar to the *query expression translation*, the literal is itself only legal if the translation would result in legal code. The purpose of this rule is to avoid having to repeat other rules of the language that are implied (for example, about convertibility of expressions when assigned to storage locations).
- An implementation is not required to translate literals exactly as specified below. Any translation is legal if the same result is produced and there are no observable differences in the production of the result.
  - For example, an implementation could translate literals like `[1, 2, 3]` directly to a `new int[] { 1, 2, 3 }` expression that itself bakes the raw data into the assembly, eliding the need for `__index` or a sequence of instructions to assign each value. Importantly, this does mean if any step of the translation might cause an exception at runtime that the program state is still left in the state indicated by the translation.
- References to 'stack allocation' refer to any strategy to allocate on the stack and not the heap. Importantly, it does not imply or require that that strategy be through the actual `stackalloc` mechanism. For example, the use of inline arrays is also an allowed and desirable approach to accomplish stack allocation where available.
- Collections are assumed to be well-behaved. For example:
  - It is assumed that the value of `Count` on a collection will produce that same value as the number of elements when enumerated.
  - The types used in this spec defined in the `System.Collections.Generic` namespace are presumed to be side-effect free. As such, the compiler can optimize scenarios where such types might be used as intermediary values, but otherwise not be exposed.
  - It is assumed that a call to some applicable `.AddRange(x)` member on a collection will result in the same final value as iterating over `x` and adding all of its enumerated values individually to the collection with `.Add`.

- The behavior of collection literals with collections that are not well-behaved is undefined.

Conversions

A *collection expression conversion* allows a collection expression to be converted to a type.

An implicit *collection expression conversion* exists from a collection expression to the following types:

- A single dimensional *array type* `T[]`, in which case the *element type* is `T`
- A *span type*:
  - `System.Span<T>`
  - `System.ReadOnlySpan<T>`

    in which cases the *element type* is `T`
- A *type* with an appropriate *create method*, in which case the *element type* is the *iteration type* determined from a `GetEnumerator` instance method or enumerable interface, not from an extension method
- A *struct* or *class type* that implements `System.Collections.IEnumerable` where:
  - The *type* has an *applicable* constructor that can be invoked with no arguments, and the constructor is accessible at the location of the collection expression.
  - If the collection expression has any elements, the *type* has an instance or extension method `Add` where:
    - The method can be invoked with a single value argument.
    - If the method is generic, the type arguments can be inferred from the collection and argument.
    - The method is accessible at the location of the collection expression.

    In which case the *element type* is the *iteration type* of the *type*.
- An *interface type*:
  - `System.Collections.Generic.IEnumerable<T>`
  - `System.Collections.Generic.IReadOnlyCollection<T>`
  - `System.Collections.Generic.IReadOnlyList<T>`
  - `System.Collections.Generic.ICollection<T>`
  - `System.Collections.Generic.IList<T>`

    in which cases the *element type* is `T`

The implicit conversion exists if the type has an *element type* `U` where for each *element* E_i in the collection expression:

- If E_i is an *expression element*, there is an implicit conversion from E_i to `U`.
- If E_i is a *spread element* $..S_i$, there is an implicit conversion from the *iteration type* of S_i to `U`.

There is no *collection expression conversion* from a collection expression to a multi dimensional *array type*.

Types for which there is an implicit collection expression conversion from a collection expression are the valid *target types* for that collection expression.

The following additional implicit conversions exist from a *collection expression*:

- To a *nullable value type* `T?` where there is a *collection expression conversion* from the collection expression to a value type `T`. The conversion is a *collection expression conversion* to `T` followed by an *implicit nullable conversion* from `T` to `T?`.

- To a reference type `T` where there is a *create method* associated with `T` that returns a type `U` and an *implicit reference conversion* from `U` to `T`. The conversion is a *collection expression conversion* to `U` followed by an *implicit reference conversion* from `U` to `T`.

- To an interface type `I` where there is a *create method* associated with `I` that returns a type `V` and an *implicit boxing conversion* from `V` to `I`. The conversion is a *collection expression conversion* to `V` followed by an *implicit boxing conversion* from `V` to `I`.

Create methods

A *create method* is indicated with a `[CollectionBuilder(...)]` attribute on the *collection type*. The attribute specifies the *builder type* and *method name* of a method to be invoked to construct an instance of the collection type.

```
c#

namespace System.Runtime.CompilerServices
{
    [AttributeUsage(
        AttributeTargets.Class | AttributeTargets.Struct |
AttributeTargets.Interface,
```

```
        Inherited = false,
        AllowMultiple = false)]
    public sealed class CollectionBuilderAttribute : System.Attribute
    {
        public CollectionBuilderAttribute(Type builderType, string
methodName);
        public Type BuilderType { get; }
        public string MethodName { get; }
    }
}
```

The attribute can be applied to a `class`, `struct`, `ref struct`, or `interface`. The attribute is not inherited although the attribute can be applied to a base `class` or an `abstract class`.

The *builder type* must be a non-generic `class` or `struct`.

First, the set of applicable *create methods* `CM` is determined.
It consists of methods that meet the following requirements:

- The method must have the name specified in the `[CollectionBuilder(...)]` attribute.
- The method must be defined on the *builder type* directly.
- The method must be `static`.
- The method must be accessible where the collection expression is used.
- The *arity* of the method must match the *arity* of the collection type.
- The method must have a single parameter of type `System.ReadOnlySpan<E>`, passed by value.
- There is an *identity conversion*, *implicit reference conversion*, or *boxing conversion* from the method return type to the *collection type*.

Methods declared on base types or interfaces are ignored and not part of the `CM` set.

If the `CM` set is empty, then the *collection type* doesn't have an *element type* and doesn't have a *create method*. None of the following steps apply.

If only one method among those in the `CM` set has an *identity conversion* from `E` to the *element type* of the *collection type*, that is the *create method* for the *collection type*. Otherwise, the *collection type* doesn't have a *create method*.

An error is reported if the `[CollectionBuilder]` attribute does not refer to an invokable method with the expected signature.

For a *collection expression* with a target type $C<S_0, S_1, \ldots>$ where the *type declaration* $C<T_0, T_1, \ldots>$ has an associated *builder method* $B.M<U_0, U_1, \ldots>()$, the *generic type*

arguments from the target type are applied in order — and from outermost containing type to innermost — to the *builder method*.

The span parameter for the *create method* can be explicitly marked `scoped` or `[UnscopedRef]`. If the parameter is implicitly or explicitly `scoped`, the compiler *may* allocate the storage for the span on the stack rather than the heap.

For example, a possible *create method* for `ImmutableArray<T>`:

C#

```
[CollectionBuilder(typeof(ImmutableArray), "Create")]
public struct ImmutableArray<T> { ... }

public static class ImmutableArray
{
    public static ImmutableArray<T> Create<T>(ReadOnlySpan<T> items) { ... }
}
```

With the *create method* above, `ImmutableArray<int> ia = [1, 2, 3];` could be emitted as:

C#

```
[InlineArray(3)] struct __InlineArray3<T> { private T _element0; }

Span<int> __tmp = new __InlineArray3<int>();
__tmp[0] = 1;
__tmp[1] = 2;
__tmp[2] = 3;
ImmutableArray<int> ia =
    ImmutableArray.Create((ReadOnlySpan<int>)__tmp);
```

Construction

The elements of a collection expression are *evaluated* in order, left to right. Each element is evaluated exactly once, and any further references to the elements refer to the results of this initial evaluation.

A spread element may be *iterated* before or after the subsequent elements in the collection expression are *evaluated*.

An unhandled exception thrown from any of the methods used during construction will be uncaught and will prevent further steps in the construction.

`Length`, `Count`, and `GetEnumerator` are assumed to have no side effects.

If the target type is a *struct* or *class type* that implements `System.Collections.IEnumerable`, and the target type does not have a *create method*, the construction of the collection instance is as follows:

- The elements are evaluated in order. Some or all elements may be evaluated *during* the steps below rather than before.
- The compiler *may* determine the *known length* of the collection expression by invoking *countable* properties — or equivalent properties from well-known interfaces or types — on each *spread element expression*.
- The constructor that is applicable with no arguments is invoked.
- For each element in order:
  - If the element is an *expression element*, the applicable `Add` instance or extension method is invoked with the element *expression* as the argument. (Unlike classic *collection initializer behavior*, element evaluation and `Add` calls are not necessarily interleaved.)
  - If the element is a *spread element* then one of the following is used:
    - An applicable `GetEnumerator` instance or extension method is invoked on the *spread element expression* and for each item from the enumerator the applicable `Add` instance or extension method is invoked on the *collection instance* with the item as the argument. If the enumerator implements `IDisposable`, then `Dispose` will be called after enumeration, regardless of exceptions.
    - An applicable `AddRange` instance or extension method is invoked on the *collection instance* with the spread element *expression* as the argument.
    - An applicable `CopyTo` instance or extension method is invoked on the *spread element expression* with the collection instance and `int` index as arguments.
- During the construction steps above, an applicable `EnsureCapacity` instance or extension method *may* be invoked one or more times on the *collection instance* with an `int` capacity argument.

If the target type is an *array*, a *span*, a type with a *create method*, or an *interface*, the construction of the collection instance is as follows:

- The elements are evaluated in order. Some or all elements may be evaluated *during* the steps below rather than before.

- The compiler *may* determine the *known length* of the collection expression by invoking *countable* properties — or equivalent properties from well-known interfaces or types — on each *spread element expression*.

- An *initialization instance* is created as follows:
  - If the target type is an *array* and the collection expression has a *known length*, an array is allocated with the expected length.
  - If the target type is a *span* or a type with a *create method*, and the collection has a *known length*, a span with the expected length is created referring to contiguous storage.
  - Otherwise intermediate storage is allocated.

- For each element in order:
  - If the element is an *expression element*, the initialization instance *indexer* is invoked to add the evaluated expression at the current index.
  - If the element is a *spread element* then one of the following is used:
    - A member of a well-known interface or type is invoked to copy items from the spread element expression to the initialization instance.
    - An applicable `GetEnumerator` instance or extension method is invoked on the *spread element expression* and for each item from the enumerator, the initialization instance *indexer* is invoked to add the item at the current index. If the enumerator implements `IDisposable`, then `Dispose` will be called after enumeration, regardless of exceptions.
    - An applicable `CopyTo` instance or extension method is invoked on the *spread element expression* with the initialization instance and `int` index as arguments.

- If intermediate storage was allocated for the collection, a collection instance is allocated with the actual collection length and the values from the initialization instance are copied to the collection instance, or if a span is required the compiler *may* use a span of the actual collection length from the intermediate storage. Otherwise the initialization instance is the collection instance.

- If the target type has a *create method*, the create method is invoked with the span instance.

> *Note:* The compiler may *delay* adding elements to the collection — or *delay* iterating through spread elements — until after evaluating subsequent elements. (When subsequent spread elements have *countable* properties that would allow calculating the expected length of the collection before allocating the collection.) Conversely,

the compiler may *eagerly* add elements to the collection — and *eagerly* iterate through spread elements — when there is no advantage to delaying.

Consider the following collection expression:

```
c#
int[] x = [a, ..b, ..c, d];
```

If spread elements `b` and `c` are *countable*, the compiler could delay adding items from `a` and `b` until after `c` is evaluated, to allow allocating the resulting array at the expected length. After that, the compiler could eagerly add items from `c`, before evaluating `d`.

```
c#
var __tmp1 = a;
var __tmp2 = b;
var __tmp3 = c;
var __result = new int[2 + __tmp2.Length + __tmp3.Length];
int __index = 0;
__result[__index++] = __tmp1;
foreach (var __i in __tmp2) __result[__index++] = __i;
foreach (var __i in __tmp3) __result[__index++] = __i;
__result[__index++] = d;
x = __result;
```

Empty collection literal

- The empty literal `[]` has no type. However, similar to the *null-literal*, this literal can be implicitly converted to any *constructible* collection type.

  For example, the following is not legal as there is no *target type* and there are no other conversions involved:

  ```
  c#
  var v = []; // illegal
  ```

- Spreading an empty literal is permitted to be elided. For example:

  ```
  c#
  bool b = ...
  ```

```
List<int> l = [x, y, .. b ? [1, 2, 3] : []];
```

Here, if `b` is false, it is not required that any value actually be constructed for the empty collection expression since it would immediately be spread into zero values in the final literal.

- The empty collection expression is permitted to be a singleton if used to construct a final collection value that is known to not be mutable. For example:

```
c#
```

```
// Can be a singleton, like Array.Empty<int>()
int[] x = [];

// Can be a singleton. Allowed to use Array.Empty<int>(),
Enumerable.Empty<int>(),
// or any other implementation that can not be mutated.
IEnumerable<int> y = [];

// Must not be a singleton.  Value must be allowed to mutate, and
should not mutate
// other references elsewhere.
List<int> z = [];
```

Ref safety

See *safe context constraint* for definitions of the *safe-context* values: *declaration-block*, *function-member*, and *caller-context*.

The *safe-context* of a collection expression is:

- The safe-context of an empty collection expression `[]` is the *caller-context*.
- If the target type is a *span type* `System.ReadOnlySpan<T>`, and `T` is one of the *primitive types* `bool`, `sbyte`, `byte`, `short`, `ushort`, `char`, `int`, `uint`, `long`, `ulong`, `float`, or `double`, and the collection expression contains *constant values only*, the safe-context of the collection expression is the *caller-context*.
- If the target type is a *span type* `System.Span<T>` or `System.ReadOnlySpan<T>`, the safe-context of the collection expression is the *declaration-block*.
- If the target type is a *ref struct type* with a *create method*, the safe-context of the collection expression is the *safe-context of an invocation* of the create method where the collection expression is the span argument to the method.

- Otherwise the safe-context of the collection expression is the *caller-context*.

A collection expression with a safe-context of *declaration-block* cannot escape the enclosing scope, and the compiler *may* store the collection on the stack rather than the heap.

To allow a collection expression for a ref struct type to escape the *declaration-block*, it may be necessary to cast the expression to another type.

```
C#
```

```
static ReadOnlySpan<int> AsSpanConstants()
{
    return [1, 2, 3]; // ok: span refers to assembly data section
}

static ReadOnlySpan<T> AsSpan2<T>(T x, T y)
{
    return [x, y];    // error: span may refer to stack data
}

static ReadOnlySpan<T> AsSpan3<T>(T x, T y, T z)
{
    return (T[])[x, y, z]; // ok: span refers to T[] on heap
}
```

Type inference

```
c#
```

```
var a = AsArray([1, 2, 3]);          // AsArray<int>(int[])
var b = AsListOfArray([[4, 5], []]); // AsListOfArray<int>(List<int[]>)

static T[] AsArray<T>(T[] arg) => arg;
static List<T[]> AsListOfArray<T>(List<T[]> arg) => arg;
```

The *type inference* rules are updated as follows.

The existing rules for the *first phase* are extracted to a new *input type inference* section, and a rule is added to *input type inference* and *output type inference* for collection expression expressions.

> 11.6.3.2 The first phase
>
> For each of the method arguments `Eᵢ`:

- An *input type inference* is made *from* `Eᵢ` *to* the corresponding *parameter type* `Tᵢ`.

An *input type inference* is made *from* an expression `E` *to* a type `T` in the following way:

- If `E` is a *collection expression* with elements `Eᵢ`, and `T` is a type with an *element type* `Tₑ` or `T` is a *nullable value type* `T0?` and `T0` has an *element type* `Tₑ`, then for each `Eᵢ`:
  - If `Eᵢ` is an *expression element*, then an *input type inference* is made *from* `Eᵢ` *to* `Tₑ`.
  - If `Eᵢ` is a *spread element* with an *iteration type* `Sᵢ`, then a *lower-bound inference* is made *from* `Sᵢ` *to* `Tₑ`.
- *[existing rules from first phase]* ...

11.6.3.7 Output type inferences

An *output type inference* is made *from* an expression `E` *to* a type `T` in the following way:

- If `E` is a *collection expression* with elements `Eᵢ`, and `T` is a type with an *element type* `Tₑ` or `T` is a *nullable value type* `T0?` and `T0` has an *element type* `Tₑ`, then for each `Eᵢ`:
  - If `Eᵢ` is an *expression element*, then an *output type inference* is made *from* `Eᵢ` *to* `Tₑ`.
  - If `Eᵢ` is a *spread element*, no inference is made from `Eᵢ`.
- *[existing rules from output type inferences]* ...

Extension methods

No changes to *extension method invocation* rules.

11.7.8.3 Extension method invocations

An extension method `Cᵢ.Mₑ` is *eligible* if:

- ...
- An implicit identity, reference, or boxing conversion exists from *expr* to the type of the first parameter of `Mₑ`.

A collection expression does not have a natural type so the existing conversions from *type* are not applicable. As a result, a collection expression cannot be used directly as the first parameter for an extension method invocation.

```c#
static class Extensions
{
    public static ImmutableArray<T> AsImmutableArray<T>(this
ImmutableArray<T> arg) => arg;
}

var x = [1].AsImmutableArray();           // error: collection expression
has no target type
var y = [2].AsImmutableArray<int>();      // error: ...
var z = Extensions.AsImmutableArray([3]); // ok
```

Overload resolution

Better conversion from expression is updated to prefer certain target types in collection expression conversions.

In the updated rules:

- A *span_type* is one of:
  - `System.Span<T>`
  - `System.ReadOnlySpan<T>`.
- An *array_or_array_interface* is one of:
  - an *array type*
  - one of the following *interface types* implemented by an *array type*:
    - `System.Collections.Generic.IEnumerable<T>`
    - `System.Collections.Generic.IReadOnlyCollection<T>`
    - `System.Collections.Generic.IReadOnlyList<T>`
    - `System.Collections.Generic.ICollection<T>`
    - `System.Collections.Generic.IList<T>`

> Given an implicit conversion C_1 that converts from an expression E to a type T_1, and an implicit conversion C_2 that converts from an expression E to a type T_2, C_1 is a ***better conversion*** than C_2 if one of the following holds:
>
> - **E is a *collection expression* and one of the following holds:**
>   - **T_1 is $\texttt{System.ReadOnlySpan<}E_1\texttt{>}$, and T_2 is $\texttt{System.Span<}E_2\texttt{>}$, and an implicit conversion exists from E_1 to E_2**

    - `T₁` is `System.ReadOnlySpan<E₁>` or `System.Span<E₁>`, and `T₂` is an *array_or_array_interface* with *element type* `E₂`, and an implicit conversion exists from `E₁` to `E₂`
    - `T₁` is not a *span_type*, and `T₂` is not a *span_type*, and an implicit conversion exists from `T₁` to `T₂`
- `E` is not a *collection expression* and one of the following holds:
    - `E` exactly matches `T₁` and `E` does not exactly match `T₂`
    - `E` exactly matches both or neither of `T₁` and `T₂`, and `T₁` is a *better conversion target* than `T₂`
- `E` is a method group, ...

Examples of differences with overload resolution between array initializers and collection expressions:

```c#
static void Generic<T>(Span<T> value) { }
static void Generic<T>(T[] value) { }

static void SpanDerived(Span<string> value) { }
static void SpanDerived(object[] value) { }

static void ArrayDerived(Span<object> value) { }
static void ArrayDerived(string[] value) { }

// Array initializers
Generic(new[] { "" });      // string[]
SpanDerived(new[] { "" });  // ambiguous
ArrayDerived(new[] { "" }); // string[]

// Collection expressions
Generic([""]);              // Span<string>
SpanDerived([""]);          // Span<string>
ArrayDerived([""]);         // ambiguous
```

Span types

The span types `ReadOnlySpan<T>` and `Span<T>` are both *constructible collection types*. Support for them follows the design for params Span<T>. Specifically, constructing either of those spans will result in an array T[] created on the stack if the params array is within limits (if any) set by the compiler. Otherwise the array will be allocated on the heap.

If the compiler chooses to allocate on the stack, it is not required to translate a literal directly to a `stackalloc` at that specific point. For example, given:

```
c#

foreach (var x in y)
{
    Span<int> span = [a, b, c];
    // do things with span
}
```

The compiler is allowed to translate that using `stackalloc` as long as the `Span` meaning stays the same and *span-safety* is maintained. For example, it can translate the above to:

```
c#

Span<int> __buffer = stackalloc int[3];
foreach (var x in y)
{
    __buffer[0] = a
    __buffer[1] = b
    __buffer[2] = c;
    Span<int> span = __buffer;
    // do things with span
}
```

The compiler can also inline arrays, if available, when choosing to allocate on the stack.

If the compiler decides to allocate on the heap, the translation for `Span<T>` is simply:

```
c#

T[] __array = [...]; // using existing rules
Span<T> __result = __array;
```

Collection literal translation

A collection expression has a *known length* if the compile-time type of each *spread element* in the collection expression is *countable*.

Interface translation

Given a target type `IEnumerable<T>`, `IReadOnlyCollection<T>`, `IReadOnlyList<T>`, `ICollection<T>`, or `IList<T>` a compliant implementation is only required to produce a value that implements that interface. A compliant implementation is free to:

1. Use an existing type that implements that interface.
2. Synthesize a type that implements the interface.

In either case, the type used is allowed to implement a larger set of interfaces than those strictly required.

Synthesized types are free to employ any strategy they want to implement the required interfaces properly. For example, a synthesized type might inline the elements directly within itself, avoiding the need for additional internal collection allocations. A synthesized type could also not use any storage whatsoever, opting to compute the values directly. For example, using `Enumerable.Range(1, 10)` for `[1, 2, 3, 4, 5, 6, 7, 8, 9, 10]`.

Non-mutable interface translation

Given a target type or `IEnumerable<T>`, `IReadOnlyCollection<T>`, `IReadOnlyList<T>`, the value generated is allowed to implement more interfaces than required. For example, implementing the mutable interfaces as well (specifically, implementing `ICollection<T>` or `IList<T>`). However, in that case:

1. The value must return `true` when queried for `ICollection<T>.IsReadOnly`. This ensures consumers can appropriately tell that the collection is non-mutable, despite implementing the mutable views.
2. The value must throw on any call to a mutation method (like `IList<T>.Add`). This ensures safety, preventing a non-mutable collection from being accidentally mutated.

It is recommended that any type that is synthesized implement all these interfaces. This ensures that maximal compatibility with existing libraries, including those that introspect the interfaces implemented by a value in order to light up performance optimizations.

Mutable interface translation

Given a target type or `ICollection<T>` or `IList<T>`:

1. The value must return `false` when queried for `ICollection<T>.IsReadOnly`.

The value generated is allowed to implement more interfaces than required. Specifically, implementing `IList<T>` even when only targeting `ICollection<T>`. However, in that case:

1. The value must support all mutation methods (like `IList<T>.RemoveAt`).

Known length translation

Having a *known length* allows for efficient construction of a result with the potential for no copying of data and no unnecessary slack space in a result.

Not having a *known length* does not prevent any result from being created. However, it may result in extra CPU and memory costs producing the data, then moving to the final destination.

- For a *known length* literal `[e1, ..s1, etc]`, the translation first starts with the following:

  ```
  c#
  int __len = count_of_expression_elements +
              __s1.Count;
              ...
              __s_n.Count;
  ```

- Given a target type `T` for that literal:

  - If `T` is some `T1[]`, then the literal is translated as:

    ```
    c#
    T1[] __result = new T1[__len];
    int __index = 0;

    __result[__index++] = __e1;
    foreach (T1 __t in __s1)
        __result[__index++] = __t;

    // further assignments of the remaining elements
    ```

    The implementation is allowed to utilize other means to populate the array. For example, utilizing efficient bulk-copy methods like `.CopyTo()`.

  - If `T` is some `Span<T1>`, then the literal is translated as the same as above, except that the `__result` initialization is translated as:

```
c#

Span<T1> __result = new T1[__len];

// same assignments as the array translation
```

The translation may use `stackalloc T1[]` or an *inline array* rather than `new T1[]` if *span-safety* is maintained.

- If `T` is some `ReadOnlySpan<T1>`, then the literal is translated the same as for the `Span<T1>` case except that the final result will be that `Span<T1>` implicitly converted to a `ReadOnlySpan<T1>`.

  A `ReadOnlySpan<T1>` where `T1` is some primitive type, and all collection elements are constant does not need its data to be on the heap, or on the stack. For example, an implementation could construct this span directly as a reference to portion of the data segment of the program.

  The above forms (for arrays and spans) are the base representations of the collection expression and are used for the following translation rules:

  - If `T` is some `C<S0, S1, …>` which has a corresponding create-method `B.M<U0, U1, …>()`, then the literal is translated as:

    ```
    c#

    // Collection literal is passed as is as the single B.M<...>(...)
    argument
    C<S0, S1, …> __result = B.M<S0, S1, …>([...])
    ```

    As the *create method* must have an argument type of some instantiated `ReadOnlySpan<T>`, the translation rule for spans applies when passing the collection expression to the create method.

  - If `T` supports collection initializers, then:

    - if the type `T` contains an accessible constructor with a single parameter `int capacity`, then the literal is translated as:

      ```
      c#

      T __result = new T(capacity: __len);
      __result.Add(__e1);
      foreach (var __t in __s1)
          __result.Add(__t);
      ```

```
// further additions of the remaining elements
```

Note: the name of the parameter is required to be `capacity`.

This form allows for a literal to inform the newly constructed type of the count of elements to allow for efficient allocation of internal storage. This avoids wasteful reallocations as the elements are added.

- otherwise, the literal is translated as:

```c#
T __result = new T();

__result.Add(__e1);
foreach (var __t in __s1)
    __result.Add(__t);

// further additions of the remaining elements
```

This allows creating the target type, albeit with no capacity optimization to prevent internal reallocation of storage.

Unknown length translation

- Given a target type `T` for an *unknown length* literal:
  - If `T` supports collection initializers, then the literal is translated as:

```c#
T __result = new T();

__result.Add(__e1);
foreach (var __t in __s1)
    __result.Add(__t);

// further additions of the remaining elements
```

This allows spreading of any iterable type, albeit with the least amount of optimization possible.

  - If `T` is some `T1[]`, then the literal has the same semantics as:

```c#
```

```
List<T1> __list = [...]; /* initialized using predefined rules */
T1[] __result = __list.ToArray();
```

The above is inefficient though; it creates the intermediary list, and then creates a copy of the final array from it. Implementations are free to optimize this away, for example producing code like so:

```c#
T1[] __result = <private_details>.CreateArray<T1>(
    count_of_expression_elements);
int __index = 0;

<private_details>.Add(ref __result, __index++, __e1);
foreach (var __t in __s1)
    <private_details>.Add(ref __result, __index++, __t);

// further additions of the remaining elements

<private_details>.Resize(ref __result, __index);
```

This allows for minimal waste and copying, without additional overhead that library collections might incur.

The counts passed to `CreateArray` are used to provide a starting size hint to prevent wasteful resizes.

- If `T` is some *span type*, an implementation may follow the above `T[]` strategy, or any other strategy with the same semantics, but better performance. For example, instead of allocating the array as a copy of the list elements, `CollectionsMarshal.AsSpan(__list)` could be used to obtain a span value directly.

Unsupported scenarios

While collection literals can be used for many scenarios, there are a few that they are not capable of replacing. These include:

- Multi-dimensional arrays (e.g. `new int[5, 10] { ... }`). There is no facility to include the dimensions, and all collection literals are either linear or map structures only.
- Collections which pass special values to their constructors. There is no facility to access the constructor being used.

- Nested collection initializers, e.g. `new Widget { Children = { w1, w2, w3 } }`. This form needs to stay since it has very different semantics from `Children = [w1, w2, w3]`. The former calls `.Add` repeatedly on `.Children` while the latter would assign a new collection over `.Children`. We could consider having the latter form fall back to adding to an existing collection if `.Children` can't be assigned, but that seems like it could be extremely confusing.

Syntax ambiguities

- There are two "true" syntactic ambiguities where there are multiple legal syntactic interpretations of code that uses a `collection_literal_expression`.
  - The `spread_element` is ambiguous with a range_expression. One could technically have:

    ```c#
    Range[] ranges = [range1, ..e, range2];
    ```

    To resolve this, we can either:
    - Require users to parenthesize `(..e)` or include a start index `0..e` if they want a range.
    - Choose a different syntax (like `...`) for spread. This would be unfortunate for the lack of consistency with slice patterns.
- There are two cases where there isn't a true ambiguity but where the syntax greatly increases parsing complexity. While not a problem given engineering time, this does still increase cognitive overhead for users when looking at code.
  - Ambiguity between `collection_literal_expression` and `attributes` on statements or local functions. Consider:

    ```c#
    [X(), Y, Z()]
    ```

    This could be one of:

    ```c#
    // A list literal inside some expression statement
    [X(), Y, Z()].ForEach(() => ...);
    ```

```
// The attributes for a statement or local function
[X(), Y, Z()] void LocalFunc() { }
```

Without complex lookahead, it would be impossible to tell without consuming the entirety of the literal.

Options to address this include:

- Allow this, doing the parsing work to determine which of these cases this is.
- Disallow this, and require the user wrap the literal in parentheses like `([X(), Y, Z()]).ForEach(...)`.
- Ambiguity between a `collection_literal_expression` in a `conditional_expression` and a `null_conditional_operations`. Consider:

```
c#

M(x ? [a, b, c]
```

This could be one of:

```
c#

// A ternary conditional picking between two collections
M(x ? [a, b, c] : [d, e, f]);

// A null conditional safely indexing into 'x':
M(x ? [a, b, c]);
```

Without complex lookahead, it would be impossible to tell without consuming the entirety of the literal.

Note: this is a problem even without a *natural type* because target typing applies through `conditional_expressions`.

As with the others, we could require parentheses to disambiguate. In other words, presume the `null_conditional_operation` interpretation unless written like so: `x ? ([1, 2, 3]) :`. However, that seems rather unfortunate. This sort of code does not seem unreasonable to write and will likely trip people up.

Drawbacks

- This introduces yet another form for collection expressions on top of the myriad ways we already have. This is extra complexity for the language. That said, this also

makes it possible to unify on one ~~ring~~ syntax to rule them all, which means existing codebases can be simplified and moved to a uniform look everywhere.
- Using `[...]` instead of `{...}` moves away from the syntax we've generally used for arrays and collection initializers already. Specifically that it uses `[...]` instead of `{...}`. However, this was already settled on by the language team when we did list patterns. We attempted to make `{...}` work with list patterns and ran into insurmountable issues. Because of this, we moved to `[...]` which, while new for C#, feels natural in many programming languages and allowed us to start fresh with no ambiguity. Using `[...]` as the corresponding literal form is complementary with our latest decisions, and gives us a clean place to work without problem.

This does introduce warts into the language. For example, the following are both legal and (fortunately) mean the exact same thing:

```c#
int[] x = { 1, 2, 3 };
int[] x = [ 1, 2, 3 ];
```

However, given the breadth and consistency brought by the new literal syntax, we should consider recommending that people move to the new form. IDE suggestions and fixes could help in that regard.

Alternatives

- What other designs have been considered? What is the impact of not doing this?

Resolved questions

- Should the compiler use `stackalloc` for stack allocation when *inline arrays* are not available and the *iteration type* is a primitive type?

  Resolution: No. Managing a `stackalloc` buffer requires additional effort over an *inline array* to ensure the buffer is not allocated repeatedly when the collection expression is within a loop. The additional complexity in the compiler and in the generated code outweighs the benefit of stack allocation on older platforms.

- In what order should we evaluate literal elements compared with Length/Count property evaluation? Should we evaluate all elements first, then all lengths? Or should we evaluate an element, then its length, then the next element, and so on?

Resolution: We evaluate all elements first, then everything else follows that.

- Can an *unknown length* literal create a collection type that needs a *known length*, like an array, span, or Construct(array/span) collection? This would be harder to do efficiently, but it might be possible through clever use of pooled arrays and/or builders.

  Resolution: Yes, we allow creating a fixes-length collection from an *unknown length* literal. The compiler is permitted to implement this in as efficient a manner as possible.

  The following text exists to record the original discussion of this topic.
  ▼ Details

  Users could always make an *unknown length* literal into a *known length* one with code like this:

  ```c#
  ImmutableArray<int> x = [a, ..unknownLength.ToArray(), b];
  ```

  However, this is unfortunate due to the need to force allocations of temporary storage. We could potentially be more efficient if we controlled how this was emitted.

- Can a `collection_expression` be target-typed to an `IEnumerable<T>` or other collection interfaces?

  For example:

  ```c#
  void DoWork(IEnumerable<long> values) { ... }
  // Needs to produce `longs` not `ints` for this to work.
  DoWork([1, 2, 3]);
  ```

  Resolution: Yes, a literal can be target-typed to any interface type `I<T>` that `List<T>` implements. For example, `IEnumerable<long>`. This is the same as target-typing to `List<long>` and then assigning that result to the specified interface type.
  The following text exists to record the original discussion of this topic.
  ▼ Details

  The open question here is determining what underlying type to actually create. One option is to look at the proposal for params IEnumerable<T> . There, we

would generate an array to pass the values along, similar to what happens with `params T[]`.

- Can/should the compiler emit `Array.Empty<T>()` for `[]`? Should we mandate that it does this, to avoid allocations whenever possible?

  Yes. The compiler should emit `Array.Empty<T>()` for any case where this is legal and the final result is non-mutable. For example, targeting `T[]`, `IEnumerable<T>`, `IReadOnlyCollection<T>` or `IReadOnlyList<T>`. It should not use `Array.Empty<T>` when the target is mutable (`ICollection<T>` or `IList<T>`).

- Should we expand on collection initializers to look for the very common `AddRange` method? It could be used by the underlying constructed type to perform adding of spread elements potentially more efficiently. We might also want to look for things like `.CopyTo` as well. There may be drawbacks here as those methods might end up causing excess allocations/dispatches versus directly enumerating in the translated code.

  Yes. An implementation is allowed to utilize other methods to initialize a collection value, under the presumption that these methods have well-defined semantics, and that collection types should be "well behaved". In practice though, an implementation should be cautious as benefits in one way (bulk copying) may come with negative consequences as well (for example, boxing a struct collection).

  An implementation should take advantage in the cases where there are no downsides. For example, with an `.AddRange(ReadOnlySpan<T>)` method.

Unresolved questions

- Should we allow inferring the *element type* when the *iteration type* is "ambiguous" (by some definition)? For example:

C#

```
Collection x = [1L, 2L];

// error CS1640: foreach statement cannot operate on variables of type
'Collection' because it implements multiple instantiations of
'IEnumerable<T>'; try casting to a specific interface instantiation
foreach (var x in new Collection) { }

static class Builder
{
    public Collection Create(ReadOnlySpan<long> items) => throw null;
}
```

```
[CollectionBuilder(...)]
class Collection : IEnumerable<int>, IEnumerable<string>
{
    IEnumerator<int> IEnumerable<int>.GetEnumerator() => throw null;
    IEnumerator<string> IEnumerable<string>.GetEnumerator() => throw null;
    IEnumerator IEnumerable.GetEnumerator() => throw null;
}
```

- Should it be legal to create and immediately index into a collection literal? Note: this requires an answer to the unresolved question below of whether collection literals have a *natural type*.
- Stack allocations for huge collections might blow the stack. Should the compiler have a heuristic for placing this data on the heap? Should the language be unspecified to allow for this flexibility? We should follow the spec for params Span<T>.
- Do we need to target-type `spread_element`? Consider, for example:

  ```
  c#

  Span<int> span = [a, ..b ? [c] : [d, e], f];
  ```

  Note: this may commonly come up in the following form to allow conditional inclusion of some set of elements, or nothing if the condition is false:

  ```
  c#

  Span<int> span = [a, ..b ? [c, d, e] : [], f];
  ```

  In order to evaluate this full literal, we need to evaluate the element expressions within. That means being able to evaluate `b ? [c] : [d, e]`. However, absent a target type to evaluate this expression in the context of, and absent any sort of *natural type*, this would we would be unable to determine what to do with either `[c]` or `[d, e]` here.

  To resolve this, we could say that when evaluating a literal's `spread_element` expression, there was an implicit target type equivalent to the target type of the literal itself. So, in the above, that would be rewritten as:

  ```
  c#

  int __e1 = a;
  Span<int> __s1 = b ? [c] : [d, e];
  ```

```
int __e2 = f;

Span<int> __result = stackalloc int[2 + __s1.Length];
int __index = 0;

__result[__index++] = a;
foreach (int __t in __s1)
  __result[index++] = __t;
__result[__index++] = f;

Span<int> span = __result;
```

Specification of a *constructible* collection type utilizing a *create method* is sensitive to the context at which conversion is classified

An existence of the conversion in this case depends on the notion of an *iteration type* of the *collection type*. If there is a *create method* that takes a `ReadOnlySpan<T>` where `T` is the *iteration type*, the conversion exists. Otherwise, it doesn't.

However, an *iteration type* is sensitive to the context at which `foreach` is performed. For the same *collection type* it can be different based on what extension methods are in scope, and it can also be undefined.

That feels fine for the purpose of `foreach` when the type isn't designed to be foreach-able on itself. If it is, extension methods cannot change how the type is foreach-ed over, no matter what the context is.

However, that feels somewhat strange for a conversion to be context sensitive like that. Effectively the conversion is "unstable". A *collection type* explicitly designed to be *constructible* is allowed to leave out a definition of a very important detail - its *iteration type*. Leaving the type "unconvertible" on itself.

Here is an example:

C#

```
using System;
using System.Collections.Generic;
using System.Runtime.CompilerServices;

[CollectionBuilder(typeof(MyCollectionBuilder),
nameof(MyCollectionBuilder.Create))]
class MyCollection
{
}
```

```
class MyCollectionBuilder
{
    public static MyCollection Create(ReadOnlySpan<long> items) => throw
null;
    public static MyCollection Create(ReadOnlySpan<string> items) => throw
null;
}

namespace Ns1
{
    static class Ext
    {
        public static IEnumerator<long> GetEnumerator(this MyCollection x)
=> throw null;
    }

    class Program
    {
        static void Main()
        {
            foreach (var l in new MyCollection())
            {
                long s = l;
            }

            MyCollection x1 = ["a", // error CS0029: Cannot implicitly
convert type 'string' to 'long'
                               2];
        }
    }
}

namespace Ns2
{
    static class Ext
    {
        public static IEnumerator<string> GetEnumerator(this MyCollection x)
=> throw null;
    }

    class Program
    {
        static void Main()
        {
            foreach (var l in new MyCollection())
            {
                string s = l;
            }

            MyCollection x1 = ["a",
                               2]; // error CS0029: Cannot implicitly
convert type 'int' to 'string'
        }
    }
}
```

```
namespace Ns3
{
    class Program
    {
        static void Main()
        {
            // error CS1579: foreach statement cannot operate on variables
of type 'MyCollection' because 'MyCollection' does not contain a public
instance or extension definition for 'GetEnumerator'
            foreach (var l in new MyCollection())
            {
            }

            MyCollection x1 = ["a", 2]; // error CS9188: 'MyCollection' has
a CollectionBuilderAttribute but no element type.
        }
    }
}
```

Given the current design, if the type doesn't define *iteration type* itself, compiler is unable to reliably validate an application of a `CollectionBuilder` attribute. If we don't know the *iteration type*, we don't know what the signature of the *create method* should be. If the *iteration type* comes from context, there is no guarantee that the type is always going to be used in a similar context.

Params Collections feature is also affected by this. It feels strange to be unable to reliably predict element type of a `params` parameter at the declaration point. The current proposal also requires to ensure that the *create method* is at least as accessible as the `params` *collection type*. It is impossible to perform this check in a reliable fashion, unless the *collection type* defines its *iteration type* itself.

Note, that we also have https://github.com/dotnet/roslyn/issues/69676 opened for compiler, which basically observes the same issue, but talks about it from the perspective of optimization.

Proposal

Require a type utilizing `CollectionBuilder` attribute to define its *iteration type* on itself. In other words this means, that the type should either implement `IEnumarable`/`IEnumerable<T>`, or it should have public `GetEnumerator` method with the right signature (this excludes any extension methods).

Also, right now *create method* is required to "be accessible where the collection expression is used". This is another point of context dependency based on accessibility. The purpose of this method is very similar to the purpose of a user-defined conversion

method, and that one must be public. Therefore, we should consider requiring the *create method* to be public as well.

Conclusion

Approved with modifications LDM-2024-01-08

The notion of *iteration type* is not applied consistently throughout *conversions*

> - To a *struct* or *class type* that implements `System.Collections.Generic.IEnumerable<T>` where:
>   - For each *element* `Ei` there is an *implicit conversion* to `T`.

It looks like an assumption is made that `T` is necessary the *iteration type* of the *struct* or *class type* in this case. However, that assumption is incorrect. Which can lead to a very strange behavior. For example:

C#

```
using System.Collections;
using System.Collections.Generic;

class MyCollection : IEnumerable<long>
{
    IEnumerator<long> IEnumerable<long>.GetEnumerator() => throw null;
    IEnumerator IEnumerable.GetEnumerator() => throw null;

    public void Add(string l) => throw null;

    public IEnumerator<string> GetEnumerator() => throw null;
}

class Program
{
    static void Main()
    {
        foreach (var l in new MyCollection())
        {
            string s = l; // Iteration type is string
        }

        MyCollection x1 = ["a", // error CS0029: Cannot implicitly convert
type 'string' to 'long'
                           2];
        MyCollection x2 = new MyCollection() { "b" };
```

```
    }
}
```

- To a *struct* or *class type* that implements `System.Collections.IEnumerable` and *does not implement* `System.Collections.Generic.IEnumerable<T>`.

It looks like implementation assumes that the *iteration type* is `object`, but the specification leaves this fact unspecified, and simply doesn't require each *element* to convert to anything. In general, however, the *iteration type* is not necessary the `object` type. Which can be observed in the following example:

C#

```
using System.Collections;
using System.Collections.Generic;

class MyCollection : IEnumerable
{
    public IEnumerator<string> GetEnumerator() => throw null;
    IEnumerator IEnumerable.GetEnumerator() => throw null;
}

class Program
{
    static void Main()
    {
        foreach (var l in new MyCollection())
        {
            string s = l; // Iteration type is string
        }
    }
}
```

The notion of *iteration type* is fundamental to Params Collections feature. And this issue leads to a strange discrepancy between the two features. For Example:

C#

```
using System.Collections;
using System.Collections.Generic;

class MyCollection : IEnumerable<long>
{
    IEnumerator<long> IEnumerable<long>.GetEnumerator() => throw null;
    IEnumerator IEnumerable.GetEnumerator() => throw null;

    public IEnumerator<string> GetEnumerator() => throw null;
```

```
    public void Add(long l) => throw null;
    public void Add(string l) => throw null;
}

class Program
{
    static void Main()
    {
        Test("2"); // error CS0029: Cannot implicitly convert type 'string'
to 'long'
        Test(["2"]); // error CS1503: Argument 1: cannot convert from
'collection expressions' to 'string'
        Test(3); // error CS1503: Argument 1: cannot convert from 'int' to
'string'
        Test([3]); // Ok

        MyCollection x1 = ["2"]; // error CS0029: Cannot implicitly convert
type 'string' to 'long'
        MyCollection x2 = [3];
    }

    static void Test(params MyCollection a)
    {
    }
}
```

C#

```
using System.Collections;
using System.Collections.Generic;

class MyCollection : IEnumerable
{
    IEnumerator IEnumerable.GetEnumerator() => throw null;

    public IEnumerator<string> GetEnumerator() => throw null;
    public void Add(object l) => throw null;
}

class Program
{
    static void Main()
    {
        Test("2", 3); // error CS1503: Argument 2: cannot convert from 'int'
to 'string'
        Test(["2", 3]); // Ok
    }

    static void Test(params MyCollection a)
    {
    }
}
```

It will probably be good to align one way or the other.

Proposal

Specify convertibility of *struct* or *class type* that implements `System.Collections.Generic.IEnumerable<T>` or `System.Collections.IEnumerable` in terms of *iteration type* and require an *implicit conversion* for each *element* `Ei` to the *iteration type*.

Conclusion

Approved LDM-2024-01-08

Should *collection expression conversion* require availability of a minimal set of APIs for construction?

A *constructible* collection type according to *conversions* can actually be not constructible, which is likely to lead to some unexpected overload resolution behavior. For example:

C#

```
class C1
{
    public static void M1(string x)
    {
    }
    public static void M1(char[] x)
    {
    }

    void Test()
    {
        M1(['a', 'b']); // error CS0121: The call is ambiguous between the
following methods or properties: 'C1.M1(string)' and 'C1.M1(char[])'
    }
}
```

However, the 'C1.M1(string)' is not a candidate that can be used because:

```
error CS1729: 'string' does not contain a constructor that takes 0 arguments
error CS1061: 'string' does not contain a definition for 'Add' and no
accessible extension method 'Add' accepting a first argument of type
```

```
'string' could be found (are you missing a using directive or an assembly
reference?)
```

Here is another example with a user-defined type and a stronger error that doesn't even mention a valid candidate:

C#

```
using System.Collections;
using System.Collections.Generic;

class C1 : IEnumerable<char>
{
    public static void M1(C1 x)
    {
    }
    public static void M1(char[] x)
    {
    }

    void Test()
    {
        M1(['a', 'b']); // error CS1061: 'C1' does not contain a definition
for 'Add' and no accessible extension method 'Add' accepting a first
argument of type 'C1' could be found (are you missing a using directive or
an assembly reference?)
    }

    public static implicit operator char[](C1 x) => throw null;
    IEnumerator<char> IEnumerable<char>.GetEnumerator() => throw null;
    IEnumerator IEnumerable.GetEnumerator() => throw null;
}
```

It looks like the situation is very similar to what we used to have with method group to delegate conversions. I.e. there were scenarios where the conversion existed, but was erroneous. We decided to improve that by ensuring that, if conversion is erroneous, then it doesn't exist.

Note, that with "Params Collections" feature we will be running into a similar issue. It might be good to disallow usage of `params` modifier for not constructible collections. However in the current proposal that check is based on *conversions* section. Here is an example:

C#

```
using System.Collections;
using System.Collections.Generic;

class C1 : IEnumerable<char>
```

```
{
    public static void M1(params C1 x) // It is probably better to report an
error about an invalid `params` modifier
    {
    }
    public static void M1(params ushort[] x)
    {
    }

    void Test()
    {
        M1('a', 'b'); // error CS1061: 'C1' does not contain a definition
for 'Add' and no accessible extension method 'Add' accepting a first
argument of type 'C1' could be found (are you missing a using directive or
an assembly reference?)
        M2('a', 'b'); // Ok
    }

    public static void M2(params ushort[] x)
    {
    }

    IEnumerator<char> IEnumerable<char>.GetEnumerator() => throw null;
    IEnumerator IEnumerable.GetEnumerator() => throw null;
}
```

It looks like the issue was somewhat discussed previously, see https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-10-02.md#collection-expressions. At that time an argument was made that the rules, as specified right now, are consistent with how interpolated string handlers are specified. Here is a quote:

> In particular, interpolated string handlers were originally specified this way, but we revised the specification after considering this issue.

While there is some similarity, there is also an important distinction worth considering. Here is a quote from https://github.com/dotnet/csharplang/blob/main/proposals/csharp-10.0/improved-interpolated-strings.md#interpolated-string-handler-conversion:

> Type `T` is said to be an *applicable_interpolated_string_handler_type* if it is attributed with `System.Runtime.CompilerServices.InterpolatedStringHandlerAttribute`. There exists an implicit *interpolated_string_handler_conversion* to `T` from an *interpolated_string_expression*, or an *additive_expression* composed entirely of _interpolated_string_expression_s and using only `+` operators.

The target type must have a special attribute which is a strong indicator of author's intent for the type to be an interpolated string handler. It is fair to assume that presence of the attribute is not a coincidence. In contrast, the fact that a type is "enumerable", doesn't necessary mean that there was author's intent for the type to be constructible. A presence of a *create method*, however, which is indicated with a `[CollectionBuilder(...)]` attribute on the *collection type*, feels like a strong indicator of author's intent for the type to be constructible.

Proposal

For a *struct* or *class type* that implements `System.Collections.IEnumerable` and that does not have a *create method conversions* section should require presence of at least the following APIs:

- An accessible constructor that is applicable with no arguments.
- An accessible `Add` instance or extension method that can be invoked with value of *iteration type* as the argument.

For the purpose of Params Collectons feature, such types are valid `params` types when these APIs are declared public and are instance (vs. extension) methods.

Conclusion

Approved with modifications LDM-2024-01-10

Design meetings

https://github.com/dotnet/csharplang/blob/main/meetings/2021/LDM-2021-11-01.md#collection-literals
https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-03-09.md#ambiguity-of--in-collection-expressions
https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-09-28.md#collection-literals
https://github.com/dotnet/csharplang/blob/main/meetings/2024/LDM-2024-01-08.md https://github.com/dotnet/csharplang/blob/main/meetings/2024/LDM-2024-01-10.md

Working group meetings

https://github.com/dotnet/csharplang/blob/main/meetings/working-groups/collection-literals/CL-2022-10-06.md
https://github.com/dotnet/csharplang/blob/main/meetings/working-groups/collection-literals/CL-2022-10-14.md
https://github.com/dotnet/csharplang/blob/main/meetings/working-groups/collection-literals/CL-2022-10-21.md
https://github.com/dotnet/csharplang/blob/main/meetings/working-groups/collection-literals/CL-2023-04-05.md
https://github.com/dotnet/csharplang/blob/main/meetings/working-groups/collection-literals/CL-2023-04-28.md
https://github.com/dotnet/csharplang/blob/main/meetings/working-groups/collection-literals/CL-2023-05-26.md
https://github.com/dotnet/csharplang/blob/main/meetings/working-groups/collection-literals/CL-2023-06-12.md
https://github.com/dotnet/csharplang/blob/main/meetings/working-groups/collection-literals/CL-2023-06-26.md
https://github.com/dotnet/csharplang/blob/main/meetings/working-groups/collection-literals/CL-2023-08-03.md
https://github.com/dotnet/csharplang/blob/main/meetings/working-groups/collection-literals/CL-2023-08-10.md

Upcoming agenda items

- Stack allocations for huge collections might blow the stack. Should the compiler have a heuristic for placing this data on the heap? Should the language be unspecified to allow for this flexibility? We should follow what the spec/impl does for params Span<T>. Options are:
  - Always stackalloc. Teach people to be careful with Span. This allows things like `Span<T> span = [1, 2, ..s]` to work, and be fine as long as `s` is small. If this could blow the stack, users could always create an array instead, and then get a span around this. This seems like the most in line with what people might want, but with extreme danger.
  - Only stackalloc when the literal has a *fixed* number of elements (i.e. no spread elements). This then likely makes things always safe, with fixed stack usage, and the compiler (hopefully) able to reuse that fixed buffer. However, it means things like `[1, 2, ..s]` would never be possible, even if the user knows it is completely safe at runtime.

- How does overload resolution work? If an API has:

```
C#
```

```
public void M(T[] values);
public void M(List<T> values);
```

What happens with `M([1, 2, 3])`? We likely need to define 'betterness' for these conversions.

- Should we expand on collection initializers to look for the very common `AddRange` method? It could be used by the underlying constructed type to perform adding of spread elements potentially more efficiently. We might also want to look for things like `.CopyTo` as well. There may be drawbacks here as those methods might end up causing excess allocations/dispatches versus directly enumerating in the translated code.

- Generic type inference should be updated to flow type information to/from collection literals. For example:

```
C#

void M<T>(T[] values);
M([1, 2, 3]);
```

It seems natural that this should be something the inference algorithm can be made aware of. Once this is supported for the 'base' constructible collection type cases (`T[]`, `I<T>`, `Span<T>` `new T()`), then it should also fall out of the `Collect(constructible_type)` case. For example:

```
C#

void M<T>(ImmutableArray<T> values);
M([1, 2, 3]);
```

Here, `Immutable<T>` is constructible through an `init void Construct(T[] values)` method. So the `T[] values` type would be used with inference against `[1, 2, 3]` leading to an inference of `int` for `T`.

- Cast/Index ambiguity.

  Today the following is an expression that is indexed into

```
c#

var v = (Expr)[1, 2, 3];
```

But it would be nice to be able to do things like:

```c#
var v = (ImmutableArray<int>)[1, 2, 3];
```

Can/should we take a break here?

- Syntactic ambiguities with `?[`.

  It might be worthwhile to change the rules for `nullable index access` to state that no space can occur between `?` and `[`. That would be a breaking change (but likely minor as VS already forces those together if you type them with a space). If we do this, then we can have `x?[y]` be parsed differently than `x ? [y]`.

  A similar thing occurs if we want to go with https://github.com/dotnet/csharplang/issues/2926. In that world `x?.y` is ambiguous with `x ? .y`. If we require the `?.` to abut, we can syntactically distinguish the two cases trivially.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# feature specifications feedback

C# feature specifications is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Lambda discard parameters

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow discards (`_`) to be used as parameters of lambdas and anonymous methods. For example:

- lambdas: `(_, _) => 0`, `(int _, int _) => 0`
- anonymous methods: `delegate(int _, int _) { return 0; }`

Motivation

Unused parameters do not need to be named. The intent of discards is clear, i.e. they are unused/discarded.

Detailed design

Method parameters - §14.6.2 In the parameter list of a lambda or anonymous method with more than one parameter named `_`, such parameters are discard parameters. Note: if a single parameter is named `_` then it is a regular parameter for backwards compatibility reasons.

Discard parameters do not introduce any names to any scopes. Note this implies they do not cause any `_` (underscore) names to be hidden.

Simple names (§11.7.4) If `K` is zero and the *simple_name* appears within a *block* and if the *block*'s (or an enclosing *block*'s) local variable declaration space (Declarations -

§7.3) contains a local variable, parameter (with the exception of discard parameters) or constant with name `I`, then the *simple_name* refers to that local variable, parameter or constant and is classified as a variable or value.

Scopes - §7.7 With the exception of discard parameters, the scope of a parameter declared in a *lambda_expression* (§11.16) is the *anonymous_function_body* of that *lambda_expression* With the exception of discard parameters, the scope of a parameter declared in an *anonymous_method_expression* (§11.16) is the *block* of that *anonymous_method_expression*.

Related spec sections

- Corresponding parameters - §11.6.2.2

Static anonymous functions

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow a 'static' modifier on lambdas and anonymous methods, which disallows capture of locals or instance state from containing scopes.

Motivation

Avoid unintentionally capturing state from the enclosing context, which can result in unexpected retention of captured objects or unexpected additional allocations.

Detailed design

A lambda or anonymous method may have a `static` modifier. The `static` modifier indicates that the lambda or anonymous method is a *static anonymous function*.

A *static anonymous function* cannot capture state from the enclosing scope. As a result, locals, parameters, and `this` from the enclosing scope are not available within a *static anonymous function*.

A *static anonymous function* cannot reference instance members from an implicit or explicit `this` or `base` reference.

A *static anonymous function* may reference `static` members from the enclosing scope.

A *static anonymous function* may reference `constant` definitions from the enclosing scope.

`nameof()` in a *static anonymous function* may reference locals, parameters, or `this` or `base` from the enclosing scope.

Accessibility rules for `private` members in the enclosing scope are the same for `static` and non-`static` anonymous functions.

No guarantee is made as to whether a *static anonymous function* definition is emitted as a `static` method in metadata. This is left up to the compiler implementation to optimize.

A non-`static` local function or anonymous function can capture state from an enclosing *static anonymous function* but cannot capture state outside the enclosing *static anonymous function*.

Removing the `static` modifier from an anonymous function in a valid program does not change the meaning of the program.

Function Pointers

Article • 03/24/2022

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

This proposal provides language constructs that expose IL opcodes that cannot currently be accessed efficiently, or at all, in C# today: `ldftn` and `calli`. These IL opcodes can be important in high performance code and developers need an efficient way to access them.

Motivation

The motivations and background for this feature are described in the following issue (as is a potential implementation of the feature):

https://github.com/dotnet/csharplang/issues/191

This is an alternate design proposal to compiler intrinsics

Detailed Design

Function pointers

The language will allow for the declaration of function pointers using the `delegate*` syntax. The full syntax is described in detail in the next section but it is meant to resemble the syntax used by `Func` and `Action` type declarations.

```
C#
```

```
unsafe class Example {
    void Example(Action<int> a, delegate*<int, void> f) {
        a(42);
        f(42);
    }
}
```

These types are represented using the function pointer type as outlined in ECMA-335. This means invocation of a `delegate*` will use `calli` where invocation of a `delegate` will use `callvirt` on the `Invoke` method. Syntactically though invocation is identical for both constructs.

The ECMA-335 definition of method pointers includes the calling convention as part of the type signature (section 7.1). The default calling convention will be `managed`. Unmanaged calling conventions can be specified by putting an `unmanaged` keyword afer the `delegate*` syntax, which will use the runtime platform default. Specific unmanaged conventions can then be specified in brackets to the `unmanaged` keyword by specifying any type starting with `CallConv` in the `System.Runtime.CompilerServices` namespace, leaving off the `CallConv` prefix. These types must come from the program's core library, and the set of valid combinations is platform-dependent.

C#

```
//This method has a managed calling convention. This is the same as leaving
the managed keyword off.
delegate* managed<int, int>;

// This method will be invoked using whatever the default unmanaged calling
convention on the runtime
// platform is. This is platform and architecture dependent and is
determined by the CLR at runtime.
delegate* unmanaged<int, int>;

// This method will be invoked using the cdecl calling convention
// Cdecl maps to System.Runtime.CompilerServices.CallConvCdecl
delegate* unmanaged[Cdecl] <int, int>;

// This method will be invoked using the stdcall calling convention, and
suppresses GC transition
// Stdcall maps to System.Runtime.CompilerServices.CallConvStdcall
// SuppressGCTransition maps to
System.Runtime.CompilerServices.CallConvSuppressGCTransition
delegate* unmanaged[Stdcall, SuppressGCTransition] <int, int>;
```

Conversions between `delegate*` types is done based on their signature including the calling convention.

C#

```
unsafe class Example {
    void Conversions() {
        delegate*<int, int, int> p1 = ...;
        delegate* managed<int, int, int> p2 = ...;
        delegate* unmanaged<int, int, int> p3 = ...;

        p1 = p2; // okay p1 and p2 have compatible signatures
        Console.WriteLine(p2 == p1); // True
        p2 = p3; // error: calling conventions are incompatible
    }
}
```

A `delegate*` type is a pointer type which means it has all of the capabilities and restrictions of a standard pointer type:

- Only valid in an `unsafe` context.
- Methods which contain a `delegate*` parameter or return type can only be called from an `unsafe` context.
- Cannot be converted to `object`.
- Cannot be used as a generic argument.
- Can implicitly convert `delegate*` to `void*`.
- Can explicitly convert from `void*` to `delegate*`.

Restrictions:

- Custom attributes cannot be applied to a `delegate*` or any of its elements.
- A `delegate*` parameter cannot be marked as `params`
- A `delegate*` type has all of the restrictions of a normal pointer type.
- Pointer arithmetic cannot be performed directly on function pointer types.

Function pointer syntax

The full function pointer syntax is represented by the following grammar:

antlr

```
pointer_type
    : ...
    | funcptr_type
    ;

funcptr_type
    : 'delegate' '*' calling_convention_specifier? '<'
funcptr_parameter_list funcptr_return_type '>'
```

```
    ;

calling_convention_specifier
    : 'managed'
    | 'unmanaged' ('[' unmanaged_calling_convention ']')?
    ;

unmanaged_calling_convention
    : 'Cdecl'
    | 'Stdcall'
    | 'Thiscall'
    | 'Fastcall'
    | identifier (',' identifier)*
    ;

funptr_parameter_list
    : (funcptr_parameter ',')*
    ;

funcptr_parameter
    : funcptr_parameter_modifier? type
    ;

funcptr_return_type
    : funcptr_return_modifier? return_type
    ;

funcptr_parameter_modifier
    : 'ref'
    | 'out'
    | 'in'
    ;

funcptr_return_modifier
    : 'ref'
    | 'ref readonly'
    ;
```

If no `calling_convention_specifier` is provided, the default is `managed`. The precise metadata encoding of the `calling_convention_specifier` and what `identifier`s are valid in the `unmanaged_calling_convention` is covered in Metadata Representation of Calling Conventions.

C#

```
delegate int Func1(string s);
delegate Func1 Func2(Func1 f);

// Function pointer equivalent without calling convention
delegate*<string, int>;
delegate*<delegate*<string, int>, delegate*<string, int>>;
```

```
// Function pointer equivalent with calling convention
delegate* managed<string, int>;
delegate*<delegate* managed<string, int>, delegate*<string, int>>;
```

Function pointer conversions

In an unsafe context, the set of available implicit conversions (Implicit conversions) is extended to include the following implicit pointer conversions:

- *Existing conversions* - (§22.5)
- From *funcptr_type* `F0` to another *funcptr_type* `F1`, provided all of the following are true:
  - `F0` and `F1` have the same number of parameters, and each parameter `D0n` in `F0` has the same `ref`, `out`, or `in` modifiers as the corresponding parameter `D1n` in `F1`.
  - For each value parameter (a parameter with no `ref`, `out`, or `in` modifier), an identity conversion, implicit reference conversion, or implicit pointer conversion exists from the parameter type in `F0` to the corresponding parameter type in `F1`.
  - For each `ref`, `out`, or `in` parameter, the parameter type in `F0` is the same as the corresponding parameter type in `F1`.
  - If the return type is by value (no `ref` or `ref readonly`), an identity, implicit reference, or implicit pointer conversion exists from the return type of `F1` to the return type of `F0`.
  - If the return type is by reference (`ref` or `ref readonly`), the return type and `ref` modifiers of `F1` are the same as the return type and `ref` modifiers of `F0`.
  - The calling convention of `F0` is the same as the calling convention of `F1`.

Allow address-of to target methods

Method groups will now be allowed as arguments to an address-of expression. The type of such an expression will be a `delegate*` which has the equivalent signature of the target method and a managed calling convention:

C#

```
unsafe class Util {
    public static void Log() { }

    void Use() {
        delegate*<void> ptr1 = &Util.Log;
```

```
        // Error: type "delegate*<void>" not compatible with "delegate*
<int>";
        delegate*<int> ptr2 = &Util.Log;
    }
}
```

In an unsafe context, a method `M` is compatible with a function pointer type `F` if all of the following are true:

- `M` and `F` have the same number of parameters, and each parameter in `M` has the same `ref`, `out`, or `in` modifiers as the corresponding parameter in `F`.
- For each value parameter (a parameter with no `ref`, `out`, or `in` modifier), an identity conversion, implicit reference conversion, or implicit pointer conversion exists from the parameter type in `M` to the corresponding parameter type in `F`.
- For each `ref`, `out`, or `in` parameter, the parameter type in `M` is the same as the corresponding parameter type in `F`.
- If the return type is by value (no `ref` or `ref readonly`), an identity, implicit reference, or implicit pointer conversion exists from the return type of `F` to the return type of `M`.
- If the return type is by reference (`ref` or `ref readonly`), the return type and `ref` modifiers of `F` are the same as the return type and `ref` modifiers of `M`.
- The calling convention of `M` is the same as the calling convention of `F`. This includes both the calling convention bit, as well as any calling convention flags specified in the unmanaged identifier.
- `M` is a static method.

In an unsafe context, an implicit conversion exists from an address-of expression whose target is a method group `E` to a compatible function pointer type `F` if `E` contains at least one method that is applicable in its normal form to an argument list constructed by use of the parameter types and modifiers of `F`, as described in the following.

- A single method `M` is selected corresponding to a method invocation of the form `E(A)` with the following modifications:
  - The arguments list `A` is a list of expressions, each classified as a variable and with the type and modifier (`ref`, `out`, or `in`) of the corresponding *funcptr_parameter_list* of `F`.
  - The candidate methods are only those methods that are applicable in their normal form, not those applicable in their expanded form.
  - The candidate methods are only those methods that are static.
- If the algorithm of overload resolution produces an error, then a compile-time error occurs. Otherwise, the algorithm produces a single best method `M` having the

same number of parameters as `F` and the conversion is considered to exist.
- The selected method `M` must be compatible (as defined above) with the function pointer type `F`. Otherwise, a compile-time error occurs.
- The result of the conversion is a function pointer of type `F`.

This means developers can depend on overload resolution rules to work in conjunction with the address-of operator:

C#

```
unsafe class Util {
    public static void Log() { }
    public static void Log(string p1) { }
    public static void Log(int i) { };

    void Use() {
        delegate*<void> a1 = &Log; // Log()
        delegate*<int, void> a2 = &Log; // Log(int i)

        // Error: ambiguous conversion from method group Log to "void*"
        void* v = &Log;
    }
```

The address-of operator will be implemented using the `ldftn` instruction.

Restrictions of this feature:

- Only applies to methods marked as `static`.
- Non-`static` local functions cannot be used in `&`. The implementation details of these methods are deliberately not specified by the language. This includes whether they are static vs. instance or exactly what signature they are emitted with.

Operators on Function Pointer Types

The section in unsafe code on operators is modified as such:

> In an unsafe context, several constructs are available for operating on all _pointer_type_s that are not _funcptr_type_s:
>
> - The `*` operator may be used to perform pointer indirection (§22.6.2).
> - The `->` operator may be used to access a member of a struct through a pointer (§22.6.3).
> - The `[]` operator may be used to index a pointer (§22.6.4).
> - The `&` operator may be used to obtain the address of a variable (§22.6.5).

> - The `++` and `--` operators may be used to increment and decrement pointers (§22.6.6).
> - The `+` and `-` operators may be used to perform pointer arithmetic (§22.6.7).
> - The `==`, `!=`, `<`, `>`, `<=`, and `=>` operators may be used to compare pointers (§22.6.8).
> - The `stackalloc` operator may be used to allocate memory from the call stack (§22.8).
> - The `fixed` statement may be used to temporarily fix a variable so its address can be obtained (§22.7).
>
> In an unsafe context, several constructs are available for operating on all _funcptr_type_s:
>
> - The `&` operator may be used to obtain the address of static methods (Allow address-of to target methods)
> - The `==`, `!=`, `<`, `>`, `<=`, and `=>` operators may be used to compare pointers (§22.6.8).

Additionally, we modify all the sections in `Pointers in expressions` to forbid function pointer types, except `Pointer comparison` and `The sizeof operator`.

Better function member

The better function member specification will be changed to include the following line:

> A `delegate*` is more specific than `void*`

This means that it is possible to overload on `void*` and a `delegate*` and still sensibly use the address-of operator.

Type Inference

In unsafe code, the following changes are made to the type inference algorithms:

Input types

§11.6.3.4

The following is added:

> If `E` is an address-of method group and `T` is a function pointer type then all the parameter types of `T` are input types of `E` with type `T`.

Output types

§11.6.3.5

The following is added:

> If `E` is an address-of method group and `T` is a function pointer type then the return type of `T` is an output type of `E` with type `T`.

Output type inferences

§11.6.3.7

The following bullet is added between bullets 2 and 3:

> - If `E` is an address-of method group and `T` is a function pointer type with parameter types `T1...Tk` and return type `Tb`, and overload resolution of `E` with the types `T1..Tk` yields a single method with return type `U`, then a *lower-bound inference* is made from `U` to `Tb`.

Better conversion from expression

§11.6.4.4

The following sub-bullet is added as a case to bullet 2:

> - `V` is a function pointer type `delegate*<V2..Vk, V1>` and `U` is a function pointer type `delegate*<U2..Uk, U1>`, and the calling convention of `V` is identical to `U`, and the refness of `Vi` is identical to `Ui`.

Lower-bound inferences

§11.6.3.10

The following case is added to bullet 3:

> - `V` is a function pointer type `delegate*<V2..Vk, V1>` and there is a function pointer type `delegate*<U2..Uk, U1>` such that `U` is identical to `delegate*`

> `<U2..Uk, U1>`, and the calling convention of `V` is identical to `U`, and the refness of `Vi` is identical to `Ui`.

The first bullet of inference from `Ui` to `Vi` is modified to:

> - If `U` is not a function pointer type and `Ui` is not known to be a reference type, or if `U` is a function pointer type and `Ui` is not known to be a function pointer type or a reference type, then an *exact inference* is made

Then, added after the 3rd bullet of inference from `Ui` to `Vi`:

> - Otherwise, if `V` is `delegate*<V2..Vk, V1>` then inference depends on the i-th parameter of `delegate*<V2..Vk, V1>`:
>   - If V1:
>   - If the return is by value, then a *lower-bound inference* is made.
>   - If the return is by reference, then an *exact inference* is made.
>   - If V2..Vk:
>   - If the parameter is by value, then an *upper-bound inference* is made.
>   - If the parameter is by reference, then an *exact inference* is made.

Upper-bound inferences

§11.6.3.11

The following case is added to bullet 2:

> - `U` is a function pointer type `delegate*<U2..Uk, U1>` and `V` is a function pointer type which is identical to `delegate*<V2..Vk, V1>`, and the calling convention of `U` is identical to `V`, and the refness of `Ui` is identical to `Vi`.

The first bullet of inference from `Ui` to `Vi` is modified to:

> - If `U` is not a function pointer type and `Ui` is not known to be a reference type, or if `U` is a function pointer type and `Ui` is not known to be a function pointer type or a reference type, then an *exact inference* is made

Then added after the 3rd bullet of inference from `Ui` to `Vi`:

> - Otherwise, if `U` is `delegate*<U2..Uk, U1>` then inference depends on the i-th parameter of `delegate*<U2..Uk, U1>`:

- If U1:
  - If the return is by value, then an *upper-bound inference* is made.
  - If the return is by reference, then an *exact inference* is made.
- If U2..Uk:
  - If the parameter is by value, then a *lower-bound inference* is made.
  - If the parameter is by reference, then an *exact inference* is made.

Metadata representation of `in`, `out`, and `ref readonly` parameters and return types

Function pointer signatures have no parameter flags location, so we must encode whether parameters and the return type are `in`, `out`, or `ref readonly` by using modreqs.

`in`

We reuse `System.Runtime.InteropServices.InAttribute`, applied as a `modreq` to the ref specifier on a parameter or return type, to mean the following:

- If applied to a parameter ref specifier, this parameter is treated as `in`.
- If applied to the return type ref specifier, the return type is treated as `ref readonly`.

`out`

We use `System.Runtime.InteropServices.OutAttribute`, applied as a `modreq` to the ref specifier on a parameter type, to mean that the parameter is an `out` parameter.

Errors

- It is an error to apply `OutAttribute` as a modreq to a return type.
- It is an error to apply both `InAttribute` and `OutAttribute` as a modreq to a parameter type.
- If either are specified via modopt, they are ignored.

Metadata Representation of Calling Conventions

Calling conventions are encoded in a method signature in metadata by a combination of the `CallKind` flag in the signature and zero or more `modopt`s at the start of the

signature. ECMA-335 currently declares the following elements in the `CallKind` flag:

```antlr
CallKind
   : default
   | unmanaged cdecl
   | unmanaged fastcall
   | unmanaged thiscall
   | unmanaged stdcall
   | varargs
   ;
```

Of these, function pointers in C# will support all but `varargs`.

In addition, the runtime (and eventually 335) will be updated to include a new `CallKind` on new platforms. This does not have a formal name currently, but this document will use `unmanaged ext` as a placeholder to stand for the new extensible calling convention format. With no `modopt`s, `unmanaged ext` is the platform default calling convention, `unmanaged` without the square brackets.

Mapping the `calling_convention_specifier` to a `CallKind`

A `calling_convention_specifier` that is omitted, or specified as `managed`, maps to the `default` `CallKind`. This is default `CallKind` of any method not attributed with `UnmanagedCallersOnly`.

C# recognizes 4 special identifiers that map to specific existing unmanaged `CallKind`s from ECMA 335. In order for this mapping to occur, these identifiers must be specified on their own, with no other identifiers, and this requirement is encoded into the spec for `unmanaged_calling_convention`s. These identifiers are `Cdecl`, `Thiscall`, `Stdcall`, and `Fastcall`, which correspond to `unmanaged cdecl`, `unmanaged thiscall`, `unmanaged stdcall`, and `unmanaged fastcall`, respectively. If more than one `identifer` is specified, or the single `identifier` is not of the specially recognized identifiers, we perform special name lookup on the identifier with the following rules:

- We prepend the `identifier` with the string `CallConv`
- We look only at types defined in the `System.Runtime.CompilerServices` namespace.
- We look only at types defined in the core library of the application, which is the library that defines `System.Object` and has no dependencies.
- We look only at public types.

If lookup succeeds on all of the `identifier`s specified in an `unmanaged_calling_convention`, we encode the `CallKind` as `unmanaged ext`, and encode each of the resolved types in the set of `modopt`s at the beginning of the function pointer signature. As a note, these rules mean that users cannot prefix these `identifier`s with `CallConv`, as that will result in looking up `CallConvCallConvVectorCall`.

When interpreting metadata, we first look at the `CallKind`. If it is anything other than `unmanaged ext`, we ignore all `modopt`s on the return type for the purposes of determining the calling convention, and use only the `CallKind`. If the `CallKind` is `unmanaged ext`, we look at the modopts at the start of the function pointer type, taking the union of all types that meet the following requirements:

- The is defined in the core library, which is the library that references no other libraries and defines `System.Object`.
- The type is defined in the `System.Runtime.CompilerServices` namespace.
- The type starts with the prefix `CallConv`.
- The type is public.

These represent the types that must be found when performing lookup on the `identifier`s in an `unmanaged_calling_convention` when defining a function pointer type in source.

It is an error to attempt to use a function pointer with a `CallKind` of `unmanaged ext` if the target runtime does not support the feature. This will be determined by looking for the presence of the `System.Runtime.CompilerServices.RuntimeFeature.UnmanagedCallKind` constant. If this constant is present, the runtime is considered to support the feature.

`System.Runtime.InteropServices.UnmanagedCallersOnlyAttribute`

`System.Runtime.InteropServices.UnmanagedCallersOnlyAttribute` is an attribute used by the CLR to indicate that a method should be called with a specific calling convention. Because of this, we introduce the following support for working with the attribute:

- It is an error to directly call a method annotated with this attribute from C#. Users must obtain a function pointer to the method and then invoke that pointer.
- It is an error to apply the attribute to anything other than an ordinary static method or ordinary static local function. The C# compiler will mark any non-static or static non-ordinary methods imported from metadata with this attribute as unsupported by the language.

- It is an error for a method marked with the attribute to have a parameter or return type that is not an `unmanaged_type`.
- It is an error for a method marked with the attribute to have type parameters, even if those type parameters are constrained to `unmanaged`.
- It is an error for a method in a generic type to be marked with the attribute.
- It is an error to convert a method marked with the attribute to a delegate type.
- It is an error to specify any types for `UnmanagedCallersOnly.CallConvs` that do not meet the requirements for calling convention `modopt`s in metadata.

When determining the calling convention of a method marked with a valid `UnmanagedCallersOnly` attribute, the compiler performs the following checks on the types specified in the `CallConvs` property to determine the effective `CallKind` and `modopt`s that should be used to determine the calling convention:

- If no types are specified, the `CallKind` is treated as `unmanaged ext`, with no calling convention `modopt`s at the start of the function pointer type.
- If there is one type specified, and that type is named `CallConvCdecl`, `CallConvThiscall`, `CallConvStdcall`, or `CallConvFastcall`, the `CallKind` is treated as `unmanaged cdecl`, `unmanaged thiscall`, `unmanaged stdcall`, or `unmanaged fastcall`, respectively, with no calling convention `modopt`s at the start of the function pointer type.
- If multiple types are specified or the single type is not named one of the specially called out types above, the `CallKind` is treated as `unmanaged ext`, with the union of the types specified treated as `modopt`s at the start of the function pointer type.

The compiler then looks at this effective `CallKind` and `modopt` collection and uses normal metadata rules to determine the final calling convention of the function pointer type.

Open Questions

Detecting runtime support for `unmanaged ext`

https://github.com/dotnet/runtime/issues/38135 tracks adding this flag. Depending on the feedback from review, we will either use the property specified in the issue, or use the presence of `UnmanagedCallersOnlyAttribute` as the flag that determines whether the runtimes supports `unmanaged ext`.

Considerations

Allow instance methods

The proposal could be extended to support instance methods by taking advantage of the `EXPLICITTHIS` CLI calling convention (named `instance` in C# code). This form of CLI function pointers puts the `this` parameter as an explicit first parameter of the function pointer syntax.

C#

```
unsafe class Instance {
    void Use() {
        delegate* instance<Instance, string> f = &ToString;
        f(this);
    }
}
```

This is sound but adds some complication to the proposal. Particularly because function pointers which differed by the calling convention `instance` and `managed` would be incompatible even though both cases are used to invoke managed methods with the same C# signature. Also in every case considered where this would be valuable to have there was a simple work around: use a `static` local function.

C#

```
unsafe class Instance {
    void Use() {
        static string toString(Instance i) => i.ToString();
        delegate*<Instance, string> f = &toString;
        f(this);
    }
}
```

Don't require unsafe at declaration

Instead of requiring `unsafe` at every use of a `delegate*`, only require it at the point where a method group is converted to a `delegate*`. This is where the core safety issues come into play (knowing that the containing assembly cannot be unloaded while the value is alive). Requiring `unsafe` on the other locations can be seen as excessive.

This is how the design was originally intended. But the resulting language rules felt very awkward. It's impossible to hide the fact that this is a pointer value and it kept peeking through even without the `unsafe` keyword. For example the conversion to `object` can't

be allowed, it can't be a member of a `class`, etc ... The C# design is to require `unsafe` for all pointer uses and hence this design follows that.

Developers will still be capable of presenting a *safe* wrapper on top of `delegate*` values the same way that they do for normal pointer types today. Consider:

```
C#

unsafe struct Action {
    delegate*<void> _ptr;

    Action(delegate*<void> ptr) => _ptr = ptr;
    public void Invoke() => _ptr();
}
```

Using delegates

Instead of using a new syntax element, `delegate*`, simply use existing `delegate` types with a `*` following the type:

```
C#

Func<object, object, bool>* ptr = &object.ReferenceEquals;
```

Handling calling convention can be done by annotating the `delegate` types with an attribute that specifies a `CallingConvention` value. The lack of an attribute would signify the managed calling convention.

Encoding this in IL is problematic. The underlying value needs to be represented as a pointer yet it also must:

1. Have a unique type to allow for overloads with different function pointer types.
2. Be equivalent for OHI purposes across assembly boundaries.

The last point is particularly problematic. This mean that every assembly which uses `Func<int>*` must encode an equivalent type in metadata even though `Func<int>*` is defined in an assembly though don't control. Additionally any other type which is defined with the name `System.Func<T>` in an assembly that is not mscorlib must be different than the version defined in mscorlib.

One option that was explored was emitting such a pointer as `mod_req(Func<int>) void*`. This doesn't work though as a `mod_req` cannot bind to a `TypeSpec` and hence cannot target generic instantiations.

Named function pointers

The function pointer syntax can be cumbersome, particularly in complex cases like nested function pointers. Rather than have developers type out the signature every time the language could allow for named declarations of function pointers as is done with `delegate`.

C#

```
func* void Action();

unsafe class NamedExample {
    void M(Action a) {
        a();
    }
}
```

Part of the problem here is the underlying CLI primitive doesn't have names hence this would be purely a C# invention and require a bit of metadata work to enable. That is doable but is a significant about of work. It essentially requires C# to have a companion to the type def table purely for these names.

Also when the arguments for named function pointers were examined we found they could apply equally well to a number of other scenarios. For example it would be just as convenient to declare named tuples to reduce the need to type out the full signature in all cases.

C#

```
(int x, int y) Point;

class NamedTupleExample {
    void M(Point p) {
        Console.WriteLine(p.x);
    }
}
```

After discussion we decided to not allow named declaration of `delegate*` types. If we find there is significant need for this based on customer usage feedback then we will investigate a naming solution that works for function pointers, tuples, generics, etc ... This is likely to be similar in form to other suggestions like full `typedef` support in the language.

Future Considerations

static delegates

This refers to the proposal to allow for the declaration of `delegate` types which can only refer to `static` members. The advantage being that such `delegate` instances can be allocation free and better in performance sensitive scenarios.

If the function pointer feature is implemented the `static delegate` proposal will likely be closed out. The proposed advantage of that feature is the allocation free nature. However recent investigations have found that is not possible to achieve due to assembly unloading. There must be a strong handle from the `static delegate` to the method it refers to in order to keep the assembly from being unloaded out from under it.

To maintain every `static delegate` instance would be required to allocate a new handle which runs counter to the goals of the proposal. There were some designs where the allocation could be amortized to a single allocation per call-site but that was a bit complex and didn't seem worth the trade off.

That means developers essentially have to decide between the following trade offs:

1. Safety in the face of assembly unloading: this requires allocations and hence `delegate` is already a sufficient option.
2. No safety in face of assembly unloading: use a `delegate*`. This can be wrapped in a `struct` to allow usage outside an `unsafe` context in the rest of the code.

Lambda improvements

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Proposed changes:

1. Allow lambdas with attributes
2. Allow lambdas with explicit return type
3. Infer a natural delegate type for lambdas and method groups

Motivation

Support for attributes on lambdas would provide parity with methods and local functions.

Support for explicit return types would provide symmetry with lambda parameters where explicit types can be specified. Allowing explicit return types would also provide control over compiler performance in nested lambdas where overload resolution must bind the lambda body currently to determine the signature.

A natural type for lambda expressions and method groups will allow more scenarios where lambdas and method groups may be used without an explicit delegate type, including as initializers in `var` declarations.

Requiring explicit delegate types for lambdas and method groups has been a friction point for customers, and has become an impediment to progress in ASP.NET with recent work on MapAction.

ASP.NET MapAction without proposed changes (`MapAction()` takes a `System.Delegate` argument):

C#

```
[HttpGet("/")] Todo GetTodo() => new(Id: 0, Name: "Name");
app.MapAction((Func<Todo>)GetTodo);

[HttpPost("/")] Todo PostTodo([FromBody] Todo todo) => todo;
app.MapAction((Func<Todo, Todo>)PostTodo);
```

ASP.NET MapAction with natural types for method groups:

C#

```
[HttpGet("/")] Todo GetTodo() => new(Id: 0, Name: "Name");
app.MapAction(GetTodo);

[HttpPost("/")] Todo PostTodo([FromBody] Todo todo) => todo;
app.MapAction(PostTodo);
```

ASP.NET MapAction with attributes and natural types for lambda expressions:

C#

```
app.MapAction([HttpGet("/")] () => new Todo(Id: 0, Name: "Name"));
app.MapAction([HttpPost("/")] ([FromBody] Todo todo) => todo);
```

Attributes

Attributes may be added to lambda expressions and lambda parameters. To avoid ambiguity between method attributes and parameter attributes, a lambda expression with attributes must use a parenthesized parameter list. Parameter types are not required.

C#

```
f = [A] () => { };        // [A] lambda
f = [return:A] x => x;    // syntax error at '=>'
f = [return:A] (x) => x;  // [A] lambda
f = [A] static x => x;    // syntax error at '=>'

f = ([A] x) => x;         // [A] x
f = ([A] ref int x) => x; // [A] x
```

Multiple attributes may be specified, either comma-separated within the same attribute list or as separate attribute lists.

```
C#
```

```
var f = [A1, A2][A3] () => { };    // ok
var g = ([A1][A2, A3] int x) => x; // ok
```

Attributes are not supported for *anonymous methods* declared with `delegate { }` syntax.

```
C#
```

```
f = [A] delegate { return 1; };         // syntax error at 'delegate'
f = delegate ([A] int x) { return x; }; // syntax error at '['
```

The parser will look ahead to differentiate a collection initializer with an element assignment from a collection initializer with a lambda expression.

```
C#
```

```
var y = new C { [A] = x };    // ok: y[A] = x
var z = new C { [A] x => x }; // ok: z[0] = [A] x => x
```

The parser will treat `?[` as the start of a conditional element access.

```
C#
```

```
x = b ? [A];               // ok
y = b ? [A] () => { } : z; // syntax error at '('
```

Attributes on the lambda expression or lambda parameters will be emitted to metadata on the method that maps to the lambda.

In general, customers should not depend on how lambda expressions and local functions map from source to metadata. How lambdas and local functions are emitted can, and has, changed between compiler versions.

The changes proposed here are targeted at the `Delegate` driven scenario. It should be valid to inspect the `MethodInfo` associated with a `Delegate` instance to determine the signature of the lambda expression or local function including any explicit attributes and additional metadata emitted by the compiler such as default parameters. This allows teams such as ASP.NET to make available the same behaviors for lambdas and local functions as ordinary methods.

Explicit return type

An explicit return type may be specified before the parenthesized parameter list.

```
C#
```

```
f = T () => default;                    // ok
f = short x => 1;                       // syntax error at '=>'
f = ref int (ref int x) => ref x;       // ok
f = static void (_) => { };             // ok
f = async async (async async) => async; // ok?
```

The parser will look ahead to differentiate a method call `T()` from a lambda expression `T () => e`.

Explicit return types are not supported for anonymous methods declared with `delegate { }` syntax.

```
C#
```

```
f = delegate int { return 1; };         // syntax error
f = delegate int (int x) { return x; }; // syntax error
```

Method type inference should make an exact inference from an explicit lambda return type.

```
C#
```

```
static void F<T>(Func<T, T> f) { ... }
F(int (i) => i); // Func<int, int>
```

Variance conversions are not allowed from lambda return type to delegate return type (matching similar behavior for parameter types).

```
C#
```

```
Func<object> f1 = string () => null; // error
Func<object?> f2 = object () => x;   // warning
```

The parser allows lambda expressions with `ref` return types within expressions without additional parentheses.

```
C#
```

```
d = ref int () => x; // d = (ref int () => x)
F(ref int () => x);  // F((ref int () => x))
```

`var` cannot be used as an explicit return type for lambda expressions.

C#

```
class var { }

d = var (var v) => v;              // error: contextual keyword 'var' cannot
be used as explicit lambda return type
d = @var (var v) => v;             // ok
d = ref var (ref var v) => ref v;  // error: contextual keyword 'var' cannot
be used as explicit lambda return type
d = ref @var (ref var v) => ref v; // ok
```

Natural (function) type

An *anonymous function* expression (§11.16) (a *lambda expression* or an *anonymous method*) has a natural type if the parameters types are explicit and the return type is either explicit or can be inferred (see §11.6.3.13).

A *method group* has a natural type if all candidate methods in the method group have a common signature. (If the method group may include extension methods, the candidates include the containing type and all extension method scopes.)

The natural type of an anonymous function expression or method group is a *function_type*. A *function_type* represents a method signature: the parameter types and ref kinds, and return type and ref kind. Anonymous function expressions or method groups with the same signature have the same *function_type*.

Function_types are used in a few specific contexts only:

- implicit and explicit conversions
- method type inference (§11.6.3) and best common type (§11.6.3.15)
- `var` initializers

A *function_type* exists at compile time only: *function_types* do not appear in source or metadata.

Conversions

From a *function_type* `F` there are implicit *function_type* conversions:

- To a *function_type* `G` if the parameters and return types of `F` are variance-convertible to the parameters and return type of `G`
- To `System.MulticastDelegate` or base classes or interfaces of `System.MulticastDelegate`
- To `System.Linq.Expressions.Expression` or `System.Linq.Expressions.LambdaExpression`

Anonymous function expressions and method groups already have *conversions from expression* to delegate types and expression tree types (see anonymous function conversions §10.7 and method group conversions §10.8). Those conversions are sufficient for converting to strongly-typed delegate types and expression tree types. The *function_type* conversions above add *conversions from type* to the base types only: `System.MulticastDelegate`, `System.Linq.Expressions.Expression`, etc.

There are no conversions to a *function_type* from a type other than a *function_type*. There are no explicit conversions for *function_types* since *function_types* cannot be referenced in source.

A conversion to `System.MulticastDelegate` or base type or interface realizes the anonymous function or method group as an instance of an appropriate delegate type. A conversion to `System.Linq.Expressions.Expression<TDelegate>` or base type realizes the lambda expression as an expression tree with an appropriate delegate type.

C#

```
Delegate d = delegate (object obj) { }; // Action<object>
Expression e = () => "";                // Expression<Func<string>>
object o = "".Clone;                    // Func<object>
```

Function_type conversions are not implicit or explicit standard conversions §10.4 and are not considered when determining whether a user-defined conversion operator is applicable to an anonymous function or method group. From evaluation of user defined conversions §10.5.3:

> For a conversion operator to be applicable, it must be possible to perform a standard conversion (§10.4) from the source type to the operand type of the operator, and it must be possible to perform a standard conversion from the result type of the operator to the target type.

C#

```
class C
{
    public static implicit operator C(Delegate d) { ... }
}

C c;
c = () => 1;      // error: cannot convert lambda expression to type 'C'
c = (C)(() => 2); // error: cannot convert lambda expression to type 'C'
```

A warning is reported for an implicit conversion of a method group to `object`, since the conversion is valid but perhaps unintentional.

C#

```
Random r = new Random();
object obj;
obj = r.NextDouble;         // warning: Converting method group to 'object'.
Did you intend to invoke the method?
obj = (object)r.NextDouble; // ok
```

Type inference

The existing rules for type inference are mostly unchanged (see §11.6.3). There are however a **couple of changes** below to specific phases of type inference.

First phase

The first phase (§11.6.3.2) allows an anonymous function to bind to `Ti` even if `Ti` is not a delegate or expression tree type (perhaps a type parameter constrained to `System.Delegate` for instance).

> For each of the method arguments `Ei`:
>
> - If `Ei` is an anonymous function **and `Ti` is a delegate type or expression tree type**, an *explicit parameter type inference* is made from `Ei` to `Ti` **and an *explicit return type inference* is made from `Ei` to `Ti`**.
> - Otherwise, if `Ei` has a type `U` and `xi` is a value parameter then a *lower-bound inference* is made *from* `U` *to* `Ti`.
> - Otherwise, if `Ei` has a type `U` and `xi` is a `ref` or `out` parameter then an *exact inference* is made *from* `U` *to* `Ti`.
> - Otherwise, no inference is made for this argument.

Explicit return type inference

An *explicit return type inference* is made *from* an expression `E` *to* a type `T` in the following way:

- **If `E` is an anonymous function with explicit return type `Ur` and `T` is a delegate type or expression tree type with return type `Vr` then an *exact inference* (§11.6.3.9) is made *from* `Ur` *to* `Vr`.**

Fixing

Fixing (§11.6.3.12) ensures other conversions are preferred over *function_type* conversions. (Lambda expressions and method group expressions only contribute to lower bounds so handling of *function_types* is needed for lower bounds only.)

An *unfixed* type variable `Xi` with a set of bounds is *fixed* as follows:

- The set of *candidate types* `Uj` starts out as the set of all types in the set of bounds for `Xi` **where function types are ignored in lower bounds if there any types that are not function types**.
- We then examine each bound for `Xi` in turn: For each exact bound `U` of `Xi` all types `Uj` which are not identical to `U` are removed from the candidate set. For each lower bound `U` of `Xi` all types `Uj` to which there is *not* an implicit conversion from `U` are removed from the candidate set. For each upper bound `U` of `Xi` all types `Uj` from which there is *not* an implicit conversion to `U` are removed from the candidate set.
- If among the remaining candidate types `Uj` there is a unique type `V` from which there is an implicit conversion to all the other candidate types, then `Xi` is fixed to `V`.
- Otherwise, type inference fails.

Best common type

Best common type (§11.6.3.15) is defined in terms of type inference so the type inference changes above apply to best common type as well.

C#

```
var fs = new[] { (string s) => s.Length, (string s) => int.Parse(s) }; //
Func<string, int>[]
```

`var`

Anonymous functions and method groups with function types can be used as initializers in `var` declarations.

```
C#
```

```
var f1 = () => default;           // error: cannot infer type
var f2 = x => x;                  // error: cannot infer type
var f3 = () => 1;                 // System.Func<int>
var f4 = string () => null;       // System.Func<string>
var f5 = delegate (object o) { }; // System.Action<object>

static void F1() { }
static void F1<T>(this T t) { }
static void F2(this string s) { }

var f6 = F1;    // error: multiple methods
var f7 = "".F1; // System.Action
var f8 = F2;    // System.Action<string>
```

Function types are not used in assignments to discards.

```
C#
```

```
d = () => 0; // ok
_ = () => 1; // error
```

Delegate types

The delegate type for the anonymous function or method group with parameter types `P1, ..., Pn` and return type `R` is:

- if any parameter or return value is not by value, or there are more than 16 parameters, or any of the parameter types or return are not valid type arguments (say, `(int* p) => { }`), then the delegate is a synthesized `internal` anonymous delegate type with signature that matches the anonymous function or method group, and with parameter names `arg1, ..., argn` or `arg` if a single parameter;
- if `R` is `void`, then the delegate type is `System.Action<P1, ..., Pn>`;
- otherwise the delegate type is `System.Func<P1, ..., Pn, R>`.

The compiler may allow more signatures to bind to `System.Action<>` and `System.Func<>` types in the future (if `ref struct` types are allowed type arguments for instance).

`modopt()` or `modreq()` in the method group signature are ignored in the corresponding delegate type.

If two anonymous functions or method groups in the same compilation require synthesized delegate types with the same parameter types and modifiers and the same return type and modifiers, the compiler will use the same synthesized delegate type.

Overload resolution

Better function member (§11.6.4.3) is updated to prefer members where none of the conversions and none of the type arguments involved inferred types from lambda expressions or method groups.

> ### Better function member
>
> ... Given an argument list `A` with a set of argument expressions `{E1, E2, ..., En}` and two applicable function members `Mp` and `Mq` with parameter types `{P1, P2, ..., Pn}` and `{Q1, Q2, ..., Qn}`, `Mp` is defined to be a ***better function member*** than `Mq` if
>
> 1. **for each argument, the implicit conversion from `Ex` to `Px` is not a *function_type_conversion*, and**
>
>   - **`Mp` is a non-generic method or `Mp` is a generic method with type parameters `{X1, X2, ..., Xp}` and for each type parameter `Xi` the type argument is inferred from an expression or from a type other than a *function_type*, and**
>   - **for at least one argument, the implicit conversion from `Ex` to `Qx` is a *function_type_conversion*, or `Mq` is a generic method with type parameters `{Y1, Y2, ..., Yq}` and for at least one type parameter `Yi` the type argument is inferred from a *function_type*, or**
>
> 2. for each argument, the implicit conversion from `Ex` to `Qx` is not better than the implicit conversion from `Ex` to `Px`, and for at least one argument, the conversion from `Ex` to `Px` is better than the conversion from `Ex` to `Qx`.

Better conversion from expression (§11.6.4.4) is updated to prefer conversions that did not involve inferred types from lambda expressions or method groups.

> ### Better conversion from expression

> Given an implicit conversion `C1` that converts from an expression `E` to a type `T1`, and an implicit conversion `C2` that converts from an expression `E` to a type `T2`, `C1` is a ***better conversion*** than `C2` if:
>
> 1. **`C1` is not a *function_type_conversion* and `C2` is a *function_type_conversion*, or**
> 2. `E` is a non-constant *interpolated_string_expression*, `C1` is an *implicit_string_handler_conversion*, `T1` is an *applicable_interpolated_string_handler_type*, and `C2` is not an *implicit_string_handler_conversion*, or
> 3. `E` does not exactly match `T2` and at least one of the following holds:
>
>   - `E` exactly matches `T1` (§11.6.4.4)
>   - `T1` is a better conversion target than `T2` (§11.6.4.6)

Syntax

antlr

```
lambda_expression
  : modifier* identifier '=>' (block | expression)
  | attribute_list* modifier* type? lambda_parameters '=>' (block |
expression)
  ;

lambda_parameters
  : lambda_parameter
  | '(' (lambda_parameter (',' lambda_parameter)*)? ')'
  ;

lambda_parameter
  : identifier
  | attribute_list* modifier* type? identifier equals_value_clause?
  ;
```

Open issues

Should default values be supported for lambda expression parameters for completeness?

Should `System.Diagnostics.ConditionalAttribute` be disallowed on lambda expressions since there are few scenarios where a lambda expression could be used conditionally?

C#

```
([Conditional("DEBUG")] static (x, y) => Assert(x == y))(a, b); // ok?
```

Should the *function_type* be available from the compiler API, in addition to the resulting delegate type?

Currently, the inferred delegate type uses `System.Action<>` or `System.Func<>` when parameter and return types are valid type arguments *and* there are no more than 16 parameters, and if the expected `Action<>` or `Func<>` type is missing, an error is reported. Instead, should the compiler use `System.Action<>` or `System.Func<>` regardless of arity? And if the expected type is missing, synthesize a delegate type otherwise?

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# feature specification feedback

C# feature specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Method group natural type improvements

Article • 05/15/2024

> **Note**
>
> This article is a feature specification. The specification serves as the design document for the feature. It includes proposed specification changes, along with information needed during the design and development of the feature. These articles are published until the proposed spec changes are finalized and incorporated in the current ECMA specification.
>
> There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent language design meeting (LDM) notes.
>
> You can learn more about the process for adopting feature speclets into the C# language standard in the article on the specifications.

Summary

This proposal refines the determination of the natural type of a method group in a few ways:

1. Consider candidate methods scope-by-scope (instance methods first, then each scope subsequent scope of extension methods)
2. Prune candidates that have no chance of succeeding, so they don't interfere with determining a unique signature:
    - Prune generic instance methods when no type arguments are provided (`var x = M;`)
    - Prune generic extension methods based on being able to reduce the extension and on constraints

Context on method group natural type

In C# 10, method groups gained a weak natural type.
That type is a "weak type" in that it only comes into play when the method group is not

target-typed (ie. it plays no role in `System.Action a = MethodGroup;`).
That weak natural type allows scenarios like `var x = MethodGroup;`.

For reference: https://github.com/dotnet/csharplang/blob/main/proposals/csharp-10.0/lambda-improvements.md#natural-function-type

> A method group has a natural type if all candidate methods in the method group have a common signature. (If the method group may include extension methods, the candidates include the containing type and all extension method scopes.)

In practice, this means that we:

1. Construct the set of all candidate methods:

- methods on the relevant type are in the set if they are static and the receiver is a type, or if they are non-static and the receiver is a value
- extension methods (across all scopes) that can be reduced are in the set

3. If the signatures of all the candidates do not match, then the method group doesn't have a natural type
4. If the arity of the resulting signature doesn't match the number of provided type arguments, then the method group doesn't have a natural type
5. Otherwise, the resulting signature is used as the natural type

Proposal

The principle is to go scope-by-scope and prune candidates that we know cannot succeed as early as possible (same principle used in overload resolution).

1. For each scope, we construct the set of all candidate methods:

- for the initial scope, methods on the relevant type with arity matching the provided type arguments and satisfying constraints with the provided type arguments are in the set if they are static and the receiver is a type, or if they are non-static and the receiver is a value
- for subsequent scopes, extension methods in that scope that can be substituted with the provided type arguments and reduced using the value of the receiver while satisfying contstraints are in the set

1. If we have no candidates in the given scope, proceed to the next scope.
2. If the signatures of all the candidates do not match, then the method group doesn't have a natural type
3. Otherwise, resulting signature is used as the natural type

4. If the scopes are exhausted, then the method group doesn't have a natural type

---

Relates to scope-by-scope proposal: https://github.com/dotnet/csharplang/issues/7364 Relates to proposal to better handle generic extension methods: https://github.com/dotnet/roslyn/issues/69222

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# feature specifications feedback

C# feature specifications is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Optional and parameter array parameters for lambdas and method groups

Article • 08/16/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

To build on top of the lambda improvements introduced in C# 10 (see relevant background), we propose adding support for default parameter values and `params` arrays in lambdas. This would enable users to implement the following lambdas:

```C#
var addWithDefault = (int addTo = 2) => addTo + 1;
addWithDefault(); // 3
addWithDefault(5); // 6

var counter = (params int[] xs) => xs.Length;
counter(); // 0
counter(1, 2, 3); // 3
```

Similarly, we will allow the same kind of behavior for method groups:

```C#
var addWithDefault = AddWithDefaultMethod;
addWithDefault(); // 3
addWithDefault(5) // 6

var counter = CountMethod;
counter(); // 0
```

```
counter(1, 2); // 2

int AddWithDefaultMethod(int addTo = 2) {
  return addTo + 1;
}
int CountMethod(params int[] xs) {
  return xs.Length;
}
```

Relevant background

Lambda Improvements in C# 10

Method group conversion specification §10.8

Motivation

App frameworks in the .NET ecosystem leverage lambdas heavily to allow users to quickly write business logic associated with an endpoint.

C#

```
var app = WebApplication.Create(args);

app.MapPost("/todos/{id}", (TodoService todoService, int id, string task) =>
{
  var todo = todoService.Create(id, task);
  return Results.Created(todo);
});
```

Lambdas don't currently support setting default values on parameters, so if a developer wanted to build an application that was resilient to scenarios where users didn't provide data, they're left to either use local functions or set the default values within the lambda body, as opposed to the more succinct proposed syntax.

C#

```
var app = WebApplication.Create(args);

app.MapPost("/todos/{id}", (TodoService todoService, int id, string task =
"foo") => {
  var todo = todoService.Create(id, task);
  return Results.Created(todo);
});
```

The proposed syntax also has the benefit of reducing confusing differences between lambdas and local functions, making it easier to reason about constructs and "grow up" lambdas to functions without compromising features, particularly in other scenarios where lambdas are used in APIs where method groups can also be provided as references. This is also the main motivation for supporting the `params` array which is not covered by the aforementioned use-case scenario.

For example:

C#

```
var app = WebApplication.Create(args);

Result TodoHandler(TodoService todoService, int id, string task = "foo") {
  var todo = todoService.Create(id, task);
  return Results.Created(todo);
}

app.MapPost("/todos/{id}", TodoHandler);
```

Current behavior

Currently, when a user implements a lambda with an optional or `params` parameter, the compiler raises an error.

C#

```
var addWithDefault = (int addTo = 2) => addTo + 1; // error CS1065: Default
values are not valid in this context.
var counter = (params int[] xs) => xs.Length; // error CS1670: params is not
valid in this context
```

When a user attempts to use a method group where the underlying method has an optional or `params` parameter, this information isn't propagated, so the call to the method doesn't typecheck due to a mismatch in the number of expected arguments.

C#

```
void M1(int i = 1) { }
var m1 = M1; // Infers Action<int>
m1(); // error CS7036: There is no argument given that corresponds to the
required parameter 'obj' of 'Action<int>'

void M2(params int[] xs) { }
var m2 = M2; // Infers Action<int[]>
```

```
m2(); // error CS7036: There is no argument given that corresponds to the
required parameter 'obj' of 'Action<int[]>'
```

New behavior

Following this proposal, default values and `params` can be applied to lambda parameters with the following behavior:

C#

```
var addWithDefault = (int addTo = 2) => addTo + 1;
addWithDefault(); // 3
addWithDefault(5); // 6

var counter = (params int[] xs) => xs.Length;
counter(); // 0
counter(1, 2, 3); // 3
```

Default values and `params` can be applied to method group parameters by specifically defining such method group:

C#

```
int AddWithDefault(int addTo = 2) {
  return addTo + 1;
}

var add1 = AddWithDefault;
add1(); // ok, default parameter value will be used

int Counter(params int[] xs) {
  return xs.Length;
}

var counter1 = Counter;
counter1(1, 2, 3); // ok, `params` will be used
```

Breaking change

Currently, the inferred type of a method group is `Action` or `Func` so the following code compiles:

C#

```
void WriteInt(int i = 0) {
  Console.Write(i);
}

var writeInt = WriteInt; // Inferred as Action<int>
DoAction(writeInt, 3); // Ok, writeInt is an Action<int>

void DoAction(Action<int> a, int p) {
  a(p);
}

int Count(params int[] xs) {
  return xs.Length;
}
var counter = Count; // Inferred as Func<int[], int>
DoFunction(counter, 3); // Ok, counter is a Func<int[], int>

int DoFunction(Func<int[], int> f, int p) {
  return f(new[] { p });
}
```

Following this change, code of this nature would cease to compile in .NET SDK 7.0.200 or later.

C#

```
void WriteInt(int i = 0) {
  Console.Write(i);
}

var writeInt = WriteInt; // Inferred as anonymous delegate type
DoAction(writeInt, 3); // Error, cannot convert from anonymous delegate type
to Action

void DoAction(Action<int> a, int p) {
  a(p);
}

int Count(params int[] xs) {
  return xs.Length;
}
var counter = Count; // Inferred as anonymous delegate type
DoFunction(counter, 3); // Error, cannot convert from anonymous delegate
type to Func

int DoFunction(Func<int[], int> f, int p) {
  return f(new[] { p });
}
```

The impact of this breaking change needs to be considered. Fortunately, the use of `var` to infer the type of a method group has only been supported since C# 10, so only code which has been written since then which explicitly relies on this behavior would break.

Detailed design

Grammar and parser changes

This enhancement requires the following changes to the grammar for lambda expressions.

diff

```
 lambda_expression
   : modifier* identifier '=>' (block | expression)
-  | attribute_list* modifier* type? lambda_parameters '=>' (block |
expression)
+  | attribute_list* modifier* type? lambda_parameter_list '=>' (block |
expression)
   ;

+lambda_parameter_list
+  : lambda_parameters (',' parameter_array)?
+  | parameter_array
+  ;

 lambda_parameter
   : identifier
-  | attribute_list* modifier* type? identifier
+  | attribute_list* modifier* type? identifier default_argument?
   ;
```

Note that this allows default parameter values and `params` arrays only for lambdas, not for anonymous methods declared with `delegate { }` syntax.

Same rules as for method parameters (§14.6.2) apply for lambda parameters:

- A parameter with a `ref`, `out` or `this` modifier cannot have a *default_argument*.
- A *parameter_array* may occur after an optional parameter, but cannot have a default value – the omission of arguments for a *parameter_array* would instead result in the creation of an empty array.

No changes to the grammar are necessary for method groups since this proposal would only change their semantics.

The following addition (in bold) is required to anonymous function conversions (§10.7):

> Specifically, an anonymous function `F` is compatible with a delegate type `D` provided:
>
> - [...]
> - If `F` has an explicitly typed parameter list, each parameter in `D` has the same type and modifiers as the corresponding parameter in `F` **ignoring `params` modifiers and default values**.

Updates of prior proposals

The following addition (in bold) is required to the function types specification in a prior proposal:

> The natural type of an anonymous function expression or method group is a *function_type*. A *function_type* represents a method signature: the parameter types, **default values, ref kinds, `params` modifiers**, and return type and ref kind. Anonymous function expressions or method groups with the same signature have the same *function_type*.

The following addition (in bold) is required to the delegate types specification in a prior proposal:

> The delegate type for the anonymous function or method group with parameter types `P1, ..., Pn` and return type `R` is:
>
> - if any parameter or return value is not by value, **or any parameter is optional or `params`**, or there are more than 16 parameters, or any of the parameter types or return are not valid type arguments (say, `(int* p) => { }`), then the delegate is a synthesized `internal` anonymous delegate type with signature that matches the anonymous function or method group, and with parameter names `arg1, ..., argn` or `arg` if a single parameter; [...]

Binder changes

Synthesizing new delegate types

As with the behavior for delegates with `ref` or `out` parameters, delegate types are synthesized for lambdas or method groups defined with optional or `params` parameters. Note that in the below examples, the notation `a'`, `b'`, etc. is used to represent these anonymous delegate types.

C#

```
var addWithDefault = (int addTo = 2) => addTo + 1;
// internal delegate int a'(int arg = 2);
var printString = (string toPrint = "defaultString") =>
Console.WriteLine(toPrint);
// internal delegate void b'(string arg = "defaultString");
var counter = (params int[] xs) => xs.Length;
// internal delegate int c'(params int[] arg);
string PathJoin(string s1, string s2, string sep = "/") { return $"{s1}{sep}
{s2}"; }
var joinFunc = pathJoin;
// internal delegate string d'(string arg1, string arg2, string arg3 = " ");
```

Conversion and unification behavior

Anonymous delegates with optional parameters will be unified when the same parameter (based on position) has the same default value, regardless of parameter name.

C#

```
int E(int j = 13) {
  return 11;
}

int F(int k = 0) {
  return 3;
}

int G(int x = 13) {
  return 4;
}

var a = (int i = 13) => 1;
// internal delegate int b'(int arg = 13);
var b = (int i = 0) => 2;
// internal delegate int c'(int arg = 0);
var c = (int i = 13) => 3;
// internal delegate int b'(int arg = 13);
var d = (int c = 13) => 1;
// internal delegate int b'(int arg = 13);

var e = E;
```

```
// internal delegate int b'(int arg = 13);
var f = F;
// internal delegate int c'(int arg = 0);
var g = G;
// internal delegate int b'(int arg = 13);

a = b; // Not allowed
a = c; // Allowed
a = d; // Allowed
c = e; // Allowed
e = f; // Not Allowed
b = f; // Allowed
e = g; // Allowed

d = (int c = 10) => 2; // Error: default parameter value is different
between new lambda
                       // and synthesized delegate b'. We won't do implicit
conversion
```

Anonymous delegates with an array as the last parameter will be unified when the last parameter has the same `params` modifier and array type, regardless of parameter name.

C#

```
int C(int[] xs) {
  return xs.Length;
}

int D(params int[] xs) {
  return xs.Length;
}

var a = (int[] xs) => xs.Length;
// internal delegate int a'(int[] xs);
var b = (params int[] xs) => xs.Length;
// internal delegate int b'(params int[] xs);

var c = C;
// internal delegate int a'(int[] xs);
var d = D;
// internal delegate int b'(params int[] xs);

a = b; // Not allowed
a = c; // Allowed
b = c; // Not allowed
b = d; // Allowed

c = (params int[] xs) => xs.Length; // Error: different delegate types; no
implicit conversion
d = (int[] xs) => xs.Length; // Error: different delegate types; no implicit
conversion
```

Similarly, there is of course compatibility with named delegates that already support optional and `params` parameters. When default values or `params` modifiers differ in a conversion, the source one will be unused if it's in a lambda expression, since the lambda cannot be called in any other way. That might seem counter-intuitive to users, hence a warning will be emitted when the source default value or `params` modifier is present and different from the target one. If the source is a method group, it can be called on its own, hence no warning will be emitted.

C#

```
delegate int DelegateNoDefault(int x);
delegate int DelegateWithDefault(int x = 1);
int MethodNoDefault(int x) => x;
int MethodWithDefault(int x = 2) => x;
DelegateNoDefault d1 = MethodWithDefault; // no warning: source is a method
group
DelegateWithDefault d2 = MethodWithDefault; // no warning: source is a
method group
DelegateWithDefault d3 = MethodNoDefault; // no warning: source is a method
group
DelegateNoDefault d4 = (int x = 1) => x; // warning: source present, target
missing
DelegateWithDefault d5 = (int x = 2) => x; // warning: source present,
target different
DelegateWithDefault d6 = (int x) => x; // no warning: source missing, target
present

delegate int DelegateNoParams(int[] xs);
delegate int DelegateWithParams(params int[] xs);
int MethodNoParams(int[] xs) => xs.Length;
int MethodWithParams(params int[] xs) => xs.Length;
DelegateNoParams d7 = MethodWithParams; // no warning: source is a method
group
DelegateWithParams d8 = MethodNoParams; // no warning: source is a method
group
DelegateNoParams d9 = (params int[] xs) => xs.Length; // warning: source
present, target missing
DelegateWithParams d10 = (int[] xs) => xs.Length; // no warning: source
missing, target present
```

IL/runtime behavior

The default parameter values will be emitted to metadata. The IL for this feature will be very similar in nature to the IL emitted for lambdas with `ref` and `out` parameters. A class which inherits from `System.Delegate` or similar will be generated, and the `Invoke` method will include `.param` directives to set default parameter values or

`System.ParamArrayAttribute` – just as would be the case for a standard named delegate with optional or `params` parameters.

These delegate types can be inspected at runtime, as normal. In code, users can introspect the `DefaultValue` in the `ParameterInfo` associated with the lambda or method group by using the associated `MethodInfo`.

C#

```
var addWithDefault = (int addTo = 2) => addTo + 1;
void AddWithDefaultMethod(int addTo = 2) {
  return addTo + 1;
}

addWithDefault.Method.GetParameters()[0].DefaultValue; // 2

var add1 = AddWithDefaultMethod;
add1.Method.GetParameters()[0].DefaultValue; // 2
```

Open questions

Open question: how does this interact with the existing `DefaultParameterValue` attribute?

Proposed answer: For parity, permit the `DefaultParameterValue` attribute on lambdas and ensure that the delegate generation behavior matches for default parameter values supported via the syntax.

C#

```
var a = (int i = 13) => 1;
// same as
var b = ([DefaultParameterValue(13)] int i) => 1;
b = a; // Allowed
```

Open question: First, note that this is outside the scope of the current proposal but it might be worth discussing in the future. Do we want to support defaults with implicitly typed lambda parameters? I.e.,

C#

```
delegate void M1(int i = 3);
M1 m = (x = 3) => x + x; // Ok

delegate void M2(long i = 2);
```

```
M2 m = (x = 3.0) => ...; //Error: cannot convert implicitly from long to
double
```

This inference leads to some tricky conversion issues which would require more discussion.

There are also parsing performance considerations here. For instance, today the term `(x =` could never be the start of a lambda expression. If this syntax was allowed for lambda defaults, then the parser would need a larger lookahead (scanning all the way until a `=>` token) in order to determine whether a term is a lambda or not.

Design meetings

- LDM 2022-10-10: decision to add support for `params` in the same way as default parameter values.

Checked user-defined operators

Article • 09/30/2022

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

C# should support defining `checked` variants of the following user-defined operators so that users can opt into or out of overflow behavior as appropriate:

- The `++` and `--` unary operators §11.7.14 and §11.8.6.
- The `-` unary operator §11.8.3.
- The `+`, `-`, `*`, and `/` binary operators §11.9.
- Explicit conversion operators.

Motivation

There is no way for a user to declare a type and support both checked and unchecked versions of an operator. This will make it hard to port various algorithms to use the proposed `generic math` interfaces exposed by the libraries team. Likewise, this makes it impossible to expose a type such as `Int128` or `UInt128` without the language simultaneously shipping its own support to avoid breaking changes.

Detailed design

Syntax

Grammar at operators (§14.10) will be adjusted to allow `checked` keyword after the `operator` keyword right before the operator token:

antlr

```
overloadable_unary_operator
    : '+' | 'checked'? '-' | '!' | '~' | 'checked'? '++' | 'checked'? '--' |
'true' | 'false'
    ;

overloadable_binary_operator
    : 'checked'? '+'   | 'checked'? '-'   | 'checked'? '*'   | 'checked'?
'/'   | '%'   | '&'   | '|'   | '^'   | '<<'
    | right_shift | '=='  | '!='  | '>'   | '<'   | '>='  | '<='
    ;

conversion_operator_declarator
    : 'implicit' 'operator' type '(' type identifier ')'
    | 'explicit' 'operator' 'checked'? type '(' type identifier ')'
    ;
```

For example:

C#

```
public static T operator checked ++(T x) {...}
public static T operator checked --(T x) {...}
public static T operator checked -(T x) {...}
public static T operator checked +(T lhs, T rhs) {...}
public static T operator checked -(T lhs, T rhs) {...}
public static T operator checked *(T lhs, T rhs) {...}
public static T operator checked /(T lhs, T rhs) {...}
public static explicit operator checked U(T x) {...}
```

C#

```
public static T I1.operator checked ++(T x) {...}
public static T I1.operator checked --(T x) {...}
public static T I1.operator checked -(T x) {...}
public static T I1.operator checked +(T lhs, T rhs) {...}
public static T I1.operator checked -(T lhs, T rhs) {...}
public static T I1.operator checked *(T lhs, T rhs) {...}
public static T I1.operator checked /(T lhs, T rhs) {...}
public static explicit I1.operator checked U(T x) {...}
```

For brevity below, an operator with the `checked` keyword is referred to as a `checked operator` and an operator without it is referred to as a `regular operator`. These terms are not applicable to operators that don't have a `checked` form.

Semantics

A user-defined `checked operator` is expected to throw an exception when the result of an operation is too large to represent in the destination type. What does it mean to be too large actually depends on the nature of the destination type and is not prescribed by the language. Typically the exception thrown is a `System.OverflowException`, but the language doesn't have any specific requirements regarding this.

A user-defined `regular operator` is expected to not throw an exception when the result of an operation is too large to represent in the destination type. Instead, it is expected to return an instance representing a truncated result. What does it mean to be too large and to be truncated actually depends on the nature of the destination type and is not prescribed by the language.

All existing user-defined operators out there that will have `checked` form supported fall into the category of `regular operators`. It is understood that many of them are likely to not follow the semantics specified above, but for the purpose of semantic analysis, compiler will assume that they are.

Checked vs. unchecked context within a `checked operator`

Checked/unchecked context within the body of a `checked operator` is not affected by the presence of the `checked` keyword. In other words, the context is the same as immediately at the beginning of the operator declaration. The developer would need to explicitly switch the context if part of their algorithm cannot rely on default context.

Names in metadata

Section "I.10.3.1 Unary operators" of ECMA-335 will be adjusted to include *op_CheckedIncrement*, *op_CheckedDecrement*, *op_CheckedUnaryNegation* as the names for methods implementing checked `++`, `--` and `-` unary operators.

Section "I.10.3.2 Binary operators" of ECMA-335 will be adjusted to include *op_CheckedAddition*, *op_CheckedSubtraction*, *op_CheckedMultiply*, *op_CheckedDivision* as the names for methods implementing checked `+`, `-`, `*`, and `/` binary operators.

Section "I.10.3.3 Conversion operators" of ECMA-335 will be adjusted to include *op_CheckedExplicit* as the name for a method implementing checked explicit conversion operator.

Unary operators

Unary `checked operators` follow the rules from §14.10.2.

Also, a `checked operator` declaration requires a pair-wise declaration of a `regular operator` (the return type should match as well). A compile-time error occurs otherwise.

```
C#
```

```
public struct Int128
{
    // This is fine, both a checked and regular operator are defined
    public static Int128 operator checked -(Int128 lhs);
    public static Int128 operator -(Int128 lhs);

    // This is fine, only a regular operator is defined
    public static Int128 operator --(Int128 lhs);

    // This should error, a regular operator must also be defined
    public static Int128 operator checked ++(Int128 lhs);
}
```

Binary operators

Binary `checked operators` follow the rules from §14.10.3 .

Also, a `checked operator` declaration requires a pair-wise declaration of a `regular operator` (the return type should match as well). A compile-time error occurs otherwise.

```
C#
```

```
public struct Int128
{
    // This is fine, both a checked and regular operator are defined
    public static Int128 operator checked +(Int128 lhs, Int128 rhs);
    public static Int128 operator +(Int128 lhs, Int128 rhs);

    // This is fine, only a regular operator is defined
    public static Int128 operator -(Int128 lhs, Int128 rhs);

    // This should error, a regular operator must also be defined
    public static Int128 operator checked *(Int128 lhs, Int128 rhs);
}
```

Candidate user-defined operators

The https://github.com/dotnet/csharplang/blob/main/spec/expressions.md#candidate-user-defined-operators section will be adjusted as follows (additions/changes are in bold).

Given a type `T` and an operation `operator op(A)`, where `op` is an overloadable operator and `A` is an argument list, the set of candidate user-defined operators provided by `T` for `operator op(A)` is determined as follows:

- Determine the type `T0`. If `T` is a nullable type, `T0` is its underlying type, otherwise `T0` is equal to `T`.
- **Find the set of user-defined operators, `U`. This set consists of:**
  - **In `unchecked` evaluation context, all regular `operator op` declarations in `T0`.**
  - **In `checked` evaluation context, all checked and regular `operator op` declarations in `T0` except regular declarations that have pair-wise matching `checked operator` declaration.**
- For all `operator op` declarations in `U` and all lifted forms of such operators, if at least one operator is applicable (Applicable function member) with respect to the argument list `A`, then the set of candidate operators consists of all such applicable operators in `T0`.
- Otherwise, if `T0` is `object`, the set of candidate operators is empty.
- Otherwise, the set of candidate operators provided by `T0` is the set of candidate operators provided by the direct base class of `T0`, or the effective base class of `T0` if `T0` is a type parameter.

Similar rules will be applied while determining the set of candidate operators in interfaces https://github.com/dotnet/csharplang/blob/main/meetings/2017/LDM-2017-06-27.md#shadowing-within-interfaces.

The section §11.7.18 will be adjusted to reflect the effect that the checked/unchecked context has on unary and binary operator overload resolution.

Example #1:

C#

```
public class MyClass
{
    public static void Add(Int128 lhs, Int128 rhs)
    {
        // Resolves to `op_CheckedAddition`
        Int128 r1 = checked(lhs + rhs);

        // Resolves to `op_Addition`
        Int128 r2 = unchecked(lhs + rhs);

        // Resolve to `op_Subtraction`
        Int128 r3 = checked(lhs - rhs);
```

```
        // Resolve to `op_Subtraction`
        Int128 r4 = unchecked(lhs - rhs);

        // Resolves to `op_CheckedMultiply`
        Int128 r5 = checked(lhs * rhs);

        // Error: Operator '*' cannot be applied to operands of type
'Int128' and 'Int128'
        Int128 r6 = unchecked(lhs * rhs);
    }

    public static void Divide(Int128 lhs, byte rhs)
    {
        // Resolves to `op_Division` - it is a better match than
`op_CheckedDivision`
        Int128 r4 = checked(lhs / rhs);
    }
}

public struct Int128
{
    public static Int128 operator checked +(Int128 lhs, Int128 rhs);
    public static Int128 operator +(Int128 lhs, Int128 rhs);

    public static Int128 operator -(Int128 lhs, Int128 rhs);

    // Cannot be declared in C#, but could be declared by some other
language
    public static Int128 operator checked *(Int128 lhs, Int128 rhs);

    // Cannot be declared in C#, but could be declared by some other
language
    public static Int128 operator checked /(Int128 lhs, int rhs);

    public static Int128 operator /(Int128 lhs, byte rhs);
}
```

Example #2:

C#

```
class C
{
    static void Add(C2 x, C3 y)
    {
        object o;

        // error CS0034: Operator '+' is ambiguous on operands of type 'C2'
and 'C3'
        o = checked(x + y);

        // C2.op_Addition
```

```
        o = unchecked(x + y);
    }
}

class C1
{
    // Cannot be declared in C#, but could be declared by some other
language
    public static C1 operator checked + (C1 x, C3 y) => new C3();
}

class C2 : C1
{
    public static C2 operator + (C2 x, C1 y) => new C2();
}

class C3 : C1
{
}
```

Example #3:

C#

```
class C
{
    static void Add(C2 x, C3 y)
    {
        object o;

        // error CS0034: Operator '+' is ambiguous on operands of type 'C2'
and 'C3'
        o = checked(x + y);

        // C1.op_Addition
        o = unchecked(x + y);
    }
}

class C1
{
    public static C1 operator + (C1 x, C3 y) => new C3();
}

class C2 : C1
{
    // Cannot be declared in C#, but could be declared by some other
language
    public static C2 operator checked + (C2 x, C1 y) => new C2();
}

class C3 : C1
```

```
{
}
```

Conversion operators

Conversion `checked operators` follow the rules from §14.10.4.

However, a `checked operator` declaration requires a pair-wise declaration of a `regular operator`. A compile-time error occurs otherwise.

The following paragraph

> The signature of a conversion operator consists of the source type and the target type. (This is the only form of member for which the return type participates in the signature.) The implicit or explicit classification of a conversion operator is not part of the operator's signature. Thus, a class or struct cannot declare both an implicit and an explicit conversion operator with the same source and target types.

will be adjusted to allow a type to declare checked and regular forms of explicit conversions with the same source and target types. A type will not be allowed to declare both an implicit and a checked explicit conversion operator with the same source and target types.

Processing of user-defined explicit conversions

The third bullet in §10.5.5:

> - Find the set of applicable user-defined and lifted conversion operators, `U`. This set consists of the user-defined and lifted implicit or explicit conversion operators declared by the classes or structs in `D` that convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`. If `U` is empty, the conversion is undefined and a compile-time error occurs.

will be replaced with the following bullet points:

- **Find the set of conversion operators, `U0`. This set consists of:**
  - **In `unchecked` evaluation context, the user-defined implicit or regular explicit conversion operators declared by the classes or structs in `D`.**

    - **In `checked` evaluation context, the user-defined implicit or regular/checked explicit conversion operators declared by the classes or structs in `D` except regular explicit conversion operators that have pair-wise matching `checked operator` declaration within the same declaring type.**
- Find the set of applicable user-defined and lifted conversion operators, `U`. This set consists of the user-defined and lifted implicit or explicit conversion operators **in `U0`** that convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`. If `U` is empty, the conversion is undefined and a compile-time error occurs.

The Checked and unchecked operators §11.7.18 section will be adjusted to reflect the effect that the checked/unchecked context has on processing of user-defined explicit conversions.

Implementing operators

A `checked operator` does not implement a `regular operator` and vice versa.

Linq Expression Trees

`Checked operators` will be supported in Linq Expression Trees. A UnaryExpression/BinaryExpression node will be created with corresponding `MethodInfo`. The following factory methods will be used:

```
C#
```

```
public static UnaryExpression NegateChecked (Expression expression,
MethodInfo? method);

public static BinaryExpression AddChecked (Expression left, Expression
right, MethodInfo? method);
public static BinaryExpression SubtractChecked (Expression left, Expression
right, MethodInfo? method);
public static BinaryExpression MultiplyChecked (Expression left, Expression
right, MethodInfo? method);

public static UnaryExpression ConvertChecked (Expression expression, Type
type, MethodInfo? method);
```

Note, that C# doesn't support assignments in expression trees, therefore checked increment/decrement will not be supported as well.

There is no factory method for checked divide. There is an open question regarding this - Checked division in Linq Expression Trees.

Dynamic

We will investigate the cost of adding support for checked operators in dynamic invocation in CoreCLR and pursue an implementation if the cost is not too high. This is a quote from https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-09.md .

Drawbacks

This adds additional complexity to the language and allows users to introduce more kinds of breaking changes to their types.

Alternatives

The generic math interfaces that the libraries plans to expose could expose named methods (such as `AddChecked`). The primary drawback is that this is less readable/maintainable and doesn't get the benefit of the language precedence rules around operators.

Placement of the `checked` keyword

Alternatively the `checked` keyword could be moved to the place right before the `operator` keyword:

C#

```
public static T checked operator ++(T x) {...}
public static T checked operator --(T x) {...}
public static T checked operator -(T x) {...}
public static T checked operator +(T lhs, T rhs) {...}
public static T checked operator -(T lhs, T rhs) {...}
public static T checked operator *(T lhs, T rhs) {...}
public static T checked operator /(T lhs, T rhs) {...}
public static explicit checked operator U(T x) {...}
```

C#

```
public static T checked I1.operator ++(T x) {...}
public static T checked I1.operator --(T x) {...}
public static T checked I1.operator -(T x) {...}
public static T checked I1.operator +(T lhs, T rhs) {...}
public static T checked I1.operator -(T lhs, T rhs) {...}
public static T checked I1.operator *(T lhs, T rhs) {...}
```

```
public static T checked I1.operator /(T lhs, T rhs) {...}
public static explicit checked I1.operator U(T x) {...}
```

Or it could be moved into the set of operator modifiers:

antlr

```
operator_modifier
    : 'public'
    | 'static'
    | 'extern'
    | 'checked'
    | operator_modifier_unsafe
    ;
```

C#

```
public static checked T operator ++(T x) {...}
public static checked T operator --(T x) {...}
public static checked T operator -(T x) {...}
public static checked T operator +(T lhs, T rhs) {...}
public static checked T operator -(T lhs, T rhs) {...}
public static checked T operator *(T lhs, T rhs) {...}
public static checked T operator /(T lhs, T rhs) {...}
public static checked explicit operator U(T x) {...}
```

C#

```
public static checked T I1.operator ++(T x) {...}
public static checked T I1.operator --(T x) {...}
public static checked T I1.operator -(T x) {...}
public static checked T I1.operator +(T lhs, T rhs) {...}
public static checked T I1.operator -(T lhs, T rhs) {...}
public static checked T I1.operator *(T lhs, T rhs) {...}
public static checked T I1.operator /(T lhs, T rhs) {...}
public static checked explicit I1.operator U(T x) {...}
```

`unchecked` keyword

There were suggestions to support `unchecked` keyword at the same position as the `checked` keyword with the following possible meanings:

- Simply to explicitly reflect the regular nature of the operator, or
- Perhaps to designate a distinct flavor of an operator that is supposed to be used in an `unchecked` context. The language could support `op_Addition`, `op_CheckedAddition`, and `op_UncheckedAddition` to help limit the number of

breaking changes. This adds another layer of complexity that is likely not necessary in most code.

Operator names in ECMA-335

Alternatively the operator names could be *op_UnaryNegationChecked*, *op_AdditionChecked*, *op_SubtractionChecked*, *op_MultiplyChecked*, *op_DivisionChecked*, with *Checked* at the end. However, it looks like there is already a pattern established to end the names with the operator word. For example, there is a *op_UnsignedRightShift* operator rather than *op_RightShiftUnsigned* operator.

`Checked operators` are inapplicable in an `unchecked` context

The compiler, when performing member lookup to find candidate user-defined operators within an `unchecked` context, could ignore `checked operators`. If metadata is encountered that only defines a `checked operator`, then a compilation error will occur.

C#

```
public class MyClass
{
    public static void Add(Int128 lhs, Int128 rhs)
    {
        // Resolves to `op_CheckedMultiply`
        Int128 r5 = checked(lhs * rhs);

        // Error: Operator '*' cannot be applied to operands of type
'Int128' and 'Int128'
        Int128 r5 = unchecked(lhs * rhs);
    }
}

public struct Int128
{
    public static Int128 operator checked *(Int128 lhs, Int128 rhs);
}
```

More complicated operator lookup and overload resolution rules in a `checked` context

The compiler, when performing member lookup to find candidate user-defined operators within a `checked` context will also consider applicable operators ending with `Checked`. That is, if the compiler was attempting to find applicable function members for

the binary addition operator, it would look for both `op_Addition` and `op_AdditionChecked`. If the only applicable function member is a `checked operator`, it will be used. If both a `regular operator` and `checked operator` exist and are equally applicable the `checked operator` will be preferred. If both a `regular operator` and a `checked operator` exist but the `regular operator` is an exact match while the `checked operator` is not, the compiler will prefer the `regular operator`.

C#

```
public class MyClass
{
    public static void Add(Int128 lhs, Int128 rhs)
    {
        // Resolves to `op_CheckedAddition`
        Int128 r1 = checked(lhs + rhs);

        // Resolves to `op_Addition`
        Int128 r2 = unchecked(lhs + rhs);

        // Resolve to `op_Subtraction`
        Int128 r3 = checked(lhs - rhs);

        // Resolve to `op_Subtraction`
        Int128 r4 = unchecked(lhs - rhs);
    }

    public static void Multiply(Int128 lhs, byte rhs)
    {
        // Resolves to `op_Multiply` even though `op_CheckedMultiply` is
also applicable
        Int128 r4 = checked(lhs * rhs);
    }
}

public struct Int128
{
    public static Int128 operator checked +(Int128 lhs, Int128 rhs);
    public static Int128 operator +(Int128 lhs, Int128 rhs);

    public static Int128 operator -(Int128 lhs, Int128 rhs);

    public static Int128 operator checked *(Int128 lhs, int rhs);
    public static Int128 operator *(Int128 lhs, byte rhs);
}
```

Yet another way to build the set of candidate user-defined operators

Unary operator overload resolution

Assuming that `regular operator` matches `unchecked` evaluation context, `checked operator` matches `checked` evaluation context and an operator that doesn't have `checked` form (for example, `+`) matches either context, the first bullet in https://github.com/dotnet/csharplang/blob/main/spec/expressions.md#unary-operator-overload-resolution:

> - The set of candidate user-defined operators provided by `X` for the operation `operator op(x)` is determined using the rules of Candidate user-defined operators.

will be replaced with the following two bullet points:

- The set of candidate user-defined operators provided by `X` for the operation `operator op(x)` **matching the current checked/unchecked context** is determined using the rules of Candidate user-defined operators.
- If the set of candidate user-defined operators is not empty, then this becomes the set of candidate operators for the operation. Otherwise, the set of candidate user-defined operators provided by `X` for the operation `operator op(x)` **matching the opposite checked/unchecked context** is determined using the rules of Candidate user-defined operators.

Binary operator overload resolution

Assuming that `regular operator` matches `unchecked` evaluation context, `checked operator` matches `checked` evaluation context and an operator that doesn't have a `checked` form (for example, `%`) matches either context, the first bullet in https://github.com/dotnet/csharplang/blob/main/spec/expressions.md#binary-operator-overload-resolution:

> - The set of candidate user-defined operators provided by `X` and `Y` for the operation `operator op(x,y)` is determined. The set consists of the union of the candidate operators provided by `X` and the candidate operators provided by `Y`, each determined using the rules of Candidate user-defined operators. If `X` and `Y` are the same type, or if `X` and `Y` are derived from a common base type, then shared candidate operators only occur in the combined set once.

will be replaced with the following two bullet points:

- The set of candidate user-defined operators provided by `X` and `Y` for the operation `operator op(x,y)` **matching the current checked/unchecked context** is determined. The set consists of the union of the candidate operators provided by `X` and the candidate operators provided by `Y`, each determined using the rules of Candidate user-defined operators. If `X` and `Y` are the same type, or if `X` and `Y` are derived from a common base type, then shared candidate operators only occur in the combined set once.
- If the set of candidate user-defined operators is not empty, then this becomes the set of candidate operators for the operation. Otherwise, the set of candidate user-defined operators provided by `X` and `Y` for the operation `operator op(x,y)` **matching the opposite checked/unchecked context** is determined. The set consists of the union of the candidate operators provided by `X` and the candidate operators provided by `Y`, each determined using the rules of Candidate user-defined operators. If `X` and `Y` are the same type, or if `X` and `Y` are derived from a common base type, then shared candidate operators only occur in the combined set once.

Example #1:

C#

```
public class MyClass
{
    public static void Add(Int128 lhs, Int128 rhs)
    {
        // Resolves to `op_CheckedAddition`
        Int128 r1 = checked(lhs + rhs);

        // Resolves to `op_Addition`
        Int128 r2 = unchecked(lhs + rhs);

        // Resolve to `op_Subtraction`
        Int128 r3 = checked(lhs - rhs);

        // Resolve to `op_Subtraction`
        Int128 r4 = unchecked(lhs - rhs);

        // Resolves to `op_CheckedMultiply`
        Int128 r5 = checked(lhs * rhs);

        // Resolves to `op_CheckedMultiply`
        Int128 r5 = unchecked(lhs * rhs);
    }

    public static void Divide(Int128 lhs, byte rhs)
    {
        // Resolves to `op_CheckedDivision`
```

```
        Int128 r4 = checked(lhs / rhs);
    }
}

public struct Int128
{
    public static Int128 operator checked +(Int128 lhs, Int128 rhs);
    public static Int128 operator +(Int128 lhs, Int128 rhs);

    public static Int128 operator -(Int128 lhs, Int128 rhs);

    public static Int128 operator checked *(Int128 lhs, Int128 rhs);

    public static Int128 operator checked /(Int128 lhs, int rhs);
    public static Int128 operator /(Int128 lhs, byte rhs);
}
```

Example #2:

C#

```
class C
{
    static void Add(C2 x, C3 y)
    {
        object o;

        // C1.op_CheckedAddition
        o = checked(x + y);

        // C2.op_Addition
        o = unchecked(x + y);
    }
}

class C1
{
    public static C1 operator checked + (C1 x, C3 y) => new C3();
}

class C2 : C1
{
    public static C2 operator + (C2 x, C1 y) => new C2();
}

class C3 : C1
{
}
```

Example #3:

C#

```
class C
{
    static void Add(C2 x, C3 y)
    {
        object o;

        // C2.op_CheckedAddition
        o = checked(x + y);

        // C1.op_Addition
        o = unchecked(x + y);
    }
}

class C1
{
    public static C1 operator + (C1 x, C3 y) => new C3();
}

class C2 : C1
{
    public static C2 operator checked + (C2 x, C1 y) => new C2();
}

class C3 : C1
{
}
```

Example #4:

C#

```
class C
{
    static void Add(C2 x, byte y)
    {
        object o;

        // C1.op_CheckedAddition
        o = checked(x + y);

        // C2.op_Addition
        o = unchecked(x + y);
    }

    static void Add2(C2 x, int y)
    {
        object o;

        // C2.op_Addition
```

```
        o = checked(x + y);

        // C2.op_Addition
        o = unchecked(x + y);
    }
}

class C1
{
    public static C1 operator checked + (C1 x, byte y) => new C1();
}

class C2 : C1
{
    public static C2 operator + (C2 x, int y) => new C2();
}
```

Example #5:

C#

```
class C
{
    static void Add(C2 x, byte y)
    {
        object o;

        // C2.op_CheckedAddition
        o = checked(x + y);

        // C1.op_Addition
        o = unchecked(x + y);
    }

    static void Add2(C2 x, int y)
    {
        object o;

        // C1.op_Addition
        o = checked(x + y);

        // C1.op_Addition
        o = unchecked(x + y);
    }
}

class C1
{
    public static C1 operator + (C1 x, int y) => new C1();
}

class C2 : C1
```

```
{
    public static C2 operator checked + (C2 x, byte y) => new C2();
}
```

Processing of user-defined explicit conversions

Assuming that `regular operator` matches `unchecked` evaluation context and `checked operator` matches `checked` evaluation context, the third bullet in https://github.com/dotnet/csharplang/blob/main/spec/conversions.md#processing-of-user-defined-explicit-conversions :

> - Find the set of applicable user-defined and lifted conversion operators, `U`. This set consists of the user-defined and lifted implicit or explicit conversion operators declared by the classes or structs in `D` that convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`. If `U` is empty, the conversion is undefined and a compile-time error occurs.

will be replaced with the following bullet points:

- Find the set of applicable user-defined and lifted explicit conversion operators **matching the current checked/unchecked context**, `U0`. This set consists of the user-defined and lifted explicit conversion operators declared by the classes or structs in `D` that **match the current checked/unchecked context** and convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`.
- Find the set of applicable user-defined and lifted explicit conversion operators **matching the opposite checked/unchecked context**, `U1`. If `U0` is not empty, `U1` is empty. Otherwise, this set consists of the user-defined and lifted explicit conversion operators declared by the classes or structs in `D` that **match the opposite checked/unchecked context** and convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`.
- Find the set of applicable user-defined and lifted conversion operators, `U`. This set consists of operators from `U0`, `U1`, and the user-defined and lifted implicit conversion operators declared by the classes or structs in `D` that convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`. If `U` is empty, the conversion is undefined and a compile-time error occurs.

Yet another another way to build the set of candidate user-defined operators

Unary operator overload resolution

The first bullet in section §11.4.4 will be adjusted as follows (additions are in bold).

- The set of candidate user-defined operators provided by `X` for the operation `operator op(x)` is determined using the rules of "Candidate user-defined operators" section below. **If the set contains at least one operator in checked form, all operators in regular form are removed from the set.**

The section §11.7.18 will be adjusted to reflect the effect that the checked/unchecked context has on unary operator overload resolution.

Binary operator overload resolution

The first bullet in section §11.4.5 will be adjusted as follows (additions are in bold).

- The set of candidate user-defined operators provided by `X` and `Y` for the operation `operator op(x,y)` is determined. The set consists of the union of the candidate operators provided by `X` and the candidate operators provided by `Y`, each determined using the rules of "Candidate user-defined operators" section below. If `X` and `Y` are the same type, or if `X` and `Y` are derived from a common base type, then shared candidate operators only occur in the combined set once. **If the set contains at least one operator in checked form, all operators in regular form are removed from the set.**

The Checked and unchecked operators §11.7.18 section will be adjusted to reflect the effect that the checked/unchecked context has on binary operator overload resolution.

Candidate user-defined operators

The https://github.com/dotnet/csharplang/blob/main/spec/expressions.md#candidate-user-defined-operators section will be adjusted as follows (additions are in bold).

Given a type `T` and an operation `operator op(A)`, where `op` is an overloadable operator and `A` is an argument list, the set of candidate user-defined operators provided by `T` for `operator op(A)` is determined as follows:

- Determine the type `T0`. If `T` is a nullable type, `T0` is its underlying type, otherwise `T0` is equal to `T`.
- For all `operator op` declarations **in their checked and regular forms in `checked` evaluation context and only in their regular form in `unchecked` evaluation context** in `T0` and all lifted forms of such operators, if at least one operator is applicable (Applicable function member) with respect to the argument list `A`, then the set of candidate operators consists of all such applicable operators in `T0`.
- Otherwise, if `T0` is `object`, the set of candidate operators is empty.
- Otherwise, the set of candidate operators provided by `T0` is the set of candidate operators provided by the direct base class of `T0`, or the effective base class of `T0` if `T0` is a type parameter.

Similar filtering will be applied while determining the set of candidate operators in interfaces https://github.com/dotnet/csharplang/blob/main/meetings/2017/LDM-2017-06-27.md#shadowing-within-interfaces.

The https://github.com/dotnet/csharplang/blob/main/spec/expressions.md#the-checked-and-unchecked-operators section will be adjusted to reflect the effect that the checked/unchecked context has on unary and binary operator overload resolution.

Example #1:

C#

```
public class MyClass
{
    public static void Add(Int128 lhs, Int128 rhs)
    {
        // Resolves to `op_CheckedAddition`
        Int128 r1 = checked(lhs + rhs);

        // Resolves to `op_Addition`
        Int128 r2 = unchecked(lhs + rhs);

        // Resolve to `op_Subtraction`
        Int128 r3 = checked(lhs - rhs);

        // Resolve to `op_Subtraction`
        Int128 r4 = unchecked(lhs - rhs);

        // Resolves to `op_CheckedMultiply`
        Int128 r5 = checked(lhs * rhs);

        // Error: Operator '*' cannot be applied to operands of type
'Int128' and 'Int128'
        Int128 r5 = unchecked(lhs * rhs);
```

```
    }

    public static void Divide(Int128 lhs, byte rhs)
    {
        // Resolves to `op_CheckedDivision`
        Int128 r4 = checked(lhs / rhs);
    }
}

public struct Int128
{
    public static Int128 operator checked +(Int128 lhs, Int128 rhs);
    public static Int128 operator +(Int128 lhs, Int128 rhs);

    public static Int128 operator -(Int128 lhs, Int128 rhs);

    public static Int128 operator checked *(Int128 lhs, Int128 rhs);

    public static Int128 operator checked /(Int128 lhs, int rhs);
    public static Int128 operator /(Int128 lhs, byte rhs);
}
```

Example #2:

C#

```
class C
{
    static void Add(C2 x, C3 y)
    {
        object o;

        // C1.op_CheckedAddition
        o = checked(x + y);

        // C2.op_Addition
        o = unchecked(x + y);
    }
}

class C1
{
    public static C1 operator checked + (C1 x, C3 y) => new C3();
}

class C2 : C1
{
    public static C2 operator + (C2 x, C1 y) => new C2();
}

class C3 : C1
```

```
{
}
```

Example #3:

C#

```
class C
{
    static void Add(C2 x, C3 y)
    {
        object o;

        // C2.op_CheckedAddition
        o = checked(x + y);

        // C1.op_Addition
        o = unchecked(x + y);
    }
}

class C1
{
    public static C1 operator + (C1 x, C3 y) => new C3();
}

class C2 : C1
{
    public static C2 operator checked + (C2 x, C1 y) => new C2();
}

class C3 : C1
{
}
```

Example #4:

C#

```
class C
{
    static void Add(C2 x, byte y)
    {
        object o;

        // C2.op_Addition
        o = checked(x + y);

        // C2.op_Addition
        o = unchecked(x + y);
```

```
    }

    static void Add2(C2 x, int y)
    {
        object o;

        // C2.op_Addition
        o = checked(x + y);

        // C2.op_Addition
        o = unchecked(x + y);
    }
}

class C1
{
    public static C1 operator checked + (C1 x, byte y) => new C1();
}

class C2 : C1
{
    public static C2 operator + (C2 x, int y) => new C2();
}
```

Example #5:

C#

```
class C
{
    static void Add(C2 x, byte y)
    {
        object o;

        // C2.op_CheckedAddition
        o = checked(x + y);

        // C1.op_Addition
        o = unchecked(x + y);
    }

    static void Add2(C2 x, int y)
    {
        object o;

        // C1.op_Addition
        o = checked(x + y);

        // C1.op_Addition
        o = unchecked(x + y);
    }
}
```

```
class C1
{
    public static C1 operator + (C1 x, int y) => new C1();
}

class C2 : C1
{
    public static C2 operator checked + (C2 x, byte y) => new C2();
}
```

Processing of user-defined explicit conversions

The third bullet in §10.5.5:

> - Find the set of applicable user-defined and lifted conversion operators, `U`. This set consists of the user-defined and lifted implicit or explicit conversion operators declared by the classes or structs in `D` that convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`. If `U` is empty, the conversion is undefined and a compile-time error occurs.

will be replaced with the following bullet points:

- Find the set of applicable user-defined and lifted explicit conversion operators, `U0`. This set consists of the user-defined and lifted explicit conversion operators declared by the classes or structs in `D` **in their checked and regular forms in `checked` evaluation context and only in their regular form in `unchecked` evaluation context** and convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`.
- If `U0` contains at least one operator in checked form, all operators in regular form are removed from the set.
- Find the set of applicable user-defined and lifted conversion operators, `U`. This set consists of operators from `U0`, and the user-defined and lifted implicit conversion operators declared by the classes or structs in `D` that convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`. If `U` is empty, the conversion is undefined and a compile-time error occurs.

The Checked and unchecked operators §11.7.18 section will be adjusted to reflect the effect that the checked/unchecked context has on processing of user-defined explicit conversions.

Checked vs. unchecked context within a `checked operator`

The compiler could treat the default context of a `checked operator` as checked. The developer would need to explicitly use `unchecked` if part of their algorithm should not participate in the `checked context`. However, this might not work well in the future if we start allowing `checked`/`unchecked` tokens as modifiers on operators to set the context within the body. The modifier and the keyword could contradict each other. Also, we wouldn't be able to do the same (treat default context as unchecked) for a `regular operator` because that would be a breaking change.

Unresolved questions

Should the language allow `checked` and `unchecked` modifiers on methods (e.g. `static checked void M()`)? This would allow removing nesting levels for methods that require it.

Checked division in Linq Expression Trees

There is no factory method to create a checked division node and there is no `ExpressionType.DivideChecked` member. We could still use the following factory method to create regular divide node with `MethodInfo` pointing to the `op_CheckedDivision` method. Consumers will have to check the name to infer the context.

```
C#
```

```
public static BinaryExpression Divide (Expression left, Expression right,
MethodInfo? method);
```

Note, even though §11.7.18 section lists `/` operator as one of the operators affected by checked/unchecked evaluation context, IL doesn't have a special op code to perform checked division. Compiler always uses the factory method reardless of the context today.

Proposal: Checked user-defined devision will not be supported in Linq Expression Trees.

(Resolved) Should we support implicit checked conversion operators?

In general, implicit conversion operators are not supposed to throw.

Proposal: No.

Resolution: Approved - https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-07.md#checked-implicit-conversions

Design meetings

https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-07.md https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-09.md https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-14.md https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-23.md

Unsigned right shift operator

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

An unsigned right shift operator will be supported by C# as a built-in operator (for primitive integral types) and as a user-defined operator.

Motivation

When working with signed integral value, it is not uncommon that you need to shift bits right without replicating the high order bit on each shift. While this can be achieved for primitive integral types with a regular shift operator, a cast to an unsigned type before the shift operation and a cast back after it is required. Within the context of the generic math interfaces the libraries are planning to expose, this is potentially more problematic as the type might not necessary have an unsigned counterpart defined or known upfront by the generic math code, yet an algorithm might rely on ability to perform an unsigned right shift operation.

Detailed design

Operators and punctuators

Section §6.4.6 will be adjusted to include `>>>` operator - the unsigned right shift operator:

```
antlr
```

```
unsigned_right_shift
    : '>>>'
    ;

unsigned_right_shift_assignment
    : '>>>='
    ;
```

No characters of any kind (not even whitespace) are allowed between the tokens in *unsigned_right_shift* and *unsigned_right_shift_assignment* productions. These productions are treated specially in order to enable the correct handling of *type_parameter_lists*.

Shift operators

Section §11.10 will be adjusted to include `>>>` operator - the unsigned right shift operator:

The `<<`, `>>` and `>>>` operators are used to perform bit shifting operations.

antlr

```
shift_expression
    : additive_expression
    | shift_expression '<<' additive_expression
    | shift_expression right_shift additive_expression
    | shift_expression unsigned_right_shift additive_expression
    ;
```

For an operation of the form `x << count` or `x >> count` or `x >>> count`, binary operator overload resolution (§11.4.5) is applied to select a specific operator implementation. The operands are converted to the parameter types of the selected operator, and the type of the result is the return type of the operator.

The predefined unsigned shift operators will support the same set of signatures that predefined signed shift operators support today in the current implementation.

- Shift right:

  C#

  ```
  int operator >>>(int x, int count);
  uint operator >>>(uint x, int count);
  long operator >>>(long x, int count);
  ulong operator >>>(ulong x, int count);
  nint operator >>>(nint x, int count);
  nuint operator >>>(nuint x, int count);
  ```

The `>>>` operator shifts `x` right by a number of bits computed as described below.

The low-order bits of `x` are discarded, the remaining bits are shifted right, and the high-order empty bit positions are set to zero.

For the predefined operators, the number of bits to shift is computed as follows:

- When the type of `x` is `int` or `uint`, the shift count is given by the low-order five bits of `count`. In other words, the shift count is computed from `count & 0x1F`.
- When the type of `x` is `long` or `ulong`, the shift count is given by the low-order six bits of `count`. In other words, the shift count is computed from `count & 0x3F`.

If the resulting shift count is zero, the shift operators simply return the value of `x`.

Shift operations never cause overflows and produce the same results in `checked` and `unchecked` contexts.

Assignment operators

Section §11.18 will be adjusted to include *unsigned_right_shift_assignment* as follows:

antlr

```
assignment_operator
    : '='
    | '+='
    | '-='
    | '*='
    | '/='
    | '%='
    | '&='
    | '|='
    | '^='
    | '<<='
    | right_shift_assignment
    | unsigned_right_shift_assignment
    ;
```

Integral types

The Integral types §8.3.6 section will be adjusted to include information about `>>>` operator. The relevant bullet point is the following:

- For the binary `<<`, `>>` and `>>>` operators, the left operand is converted to type `T`, where `T` is the first of `int`, `uint`, `long`, and `ulong` that can fully represent all

possible values of the operand. The operation is then performed using the precision of type `T`, and the type of the result is `T`.

Constant expressions

Operator `>>>` will be added to the set of constructs permitted in constant expressions at §11.20.

Operator overloading

Operator `>>>` will be added to the set of overloadable binary operators at §11.4.3.

Lifted operators

Operator `>>>` will be added to the set of binary operators permitting a lifted form at §11.4.8.

Operator precedence and associativity

Section §11.4.2 will be adjusted to add `>>>` operator to the "Shift" category and `>>>=` operator to the "Assignment and lambda expression" category.

Grammar ambiguities

The `>>>` operator is subject to the same grammar ambiguities described at §6.2.5 as a regular `>>` operator.

Operators

The §14.10 section will be adjusted to include `>>>` operator.

antlr

```
overloadable_binary_operator
    : '+'   | '-'   | '*'   | '/'   | '%'   | '&'   | '|'   | '^'   | '<<'
    | right_shift | unsigned_right_shift | '=='  | '!='  | '>'   | '<'   |
'>='  | '<='
    ;
```

Binary operators

The signature of a `>>>` operator is subject to the same rules as those at §14.10.3 for the signature of a `>>` operator.

Metadata name

Section "I.10.3.2 Binary operators" of ECMA-335 already reserved the name for an unsigned right shift operator - op_UnsignedRightShift.

Linq Expression Trees

The `>>>` operator will not be supported in Linq Expression Trees because semantics of predefined `>>>` operators on signed types cannot be accurately represented without adding conversions to an unsigned type and back. See https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-09.md#unsigned-right-shift-operator for more information.

Dynamic Binding

It looks like dynamic binding uses values of System.Linq.Expressions.ExpressionType enum to communicate binary operator kind to the runtime binder. Since we don't have a member specifically representing an unsigned right shift operator, dynamic binding for `>>>` operator will not be supported and the static and dynamic binding (§11.3) section will be adjusted to reflect that.

Drawbacks

Alternatives

Linq Expression Trees

The `>>>` operator will be supported in Linq Expressioin Trees.

- For a user-defined operator, a BinaryExpression node pointing to the operator method will be created.
- For predefined operators
  - when the first operand is an ansigned type, a BinaryExpression node will be created.
  - when the first operand is a signed type, a conversion for the first operand to an unsigned type will be added, a BinaryExpression node will be created and

conversion for the result back to the signed type will be added.

For example:

```C#
Expression<System.Func<int, int, int>> z = (x, y) => x >>> y; // (x, y) =>
Convert((Convert(x, UInt32) >> y), Int32)
```

`Resolution:`

Rejected, see https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-09.md#unsigned-right-shift-operator for more information.

Unresolved questions

Design meetings

https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-09.md

Relaxing shift operator requirements

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

The shift operator requirements will be relaxed so that the right-hand side operand is no longer restricted to only be `int`.

Motivation

When working with types other than `int`, it is not uncommon that you shift using the result of another computation, such as shifting based on the `leading zero count`. The natural type of something like a `leading zero count` is the same as the input type (`TSelf`) and so in many cases, this requires you to convert that result to `int` before shifting, even if that result is already within range.

Within the context of the generic math interfaces the libraries are planning to expose, this is potentially problematic as the type is not well known and so the conversion to `int` may not be possible or even well-defined.

Detailed design

Shift operators

§11.10 should be reworded as follows:

```
diff
```

```
- When declaring an overloaded shift operator, the type of the first operand
must always be the class or struct containing the operator declaration,
and the type of the second operand must always be int.
+ When declaring an overloaded shift operator, the type of the first operand
must always be the class or struct containing the operator declaration.
```

That is, the restriction that the first operand be the class or struct containing the operator declaration remains. While the restriction that the second operand must be `int` is removed.

Binary operators

§14.10.3 should be reworded as follows:

diff

```
-*  A binary `<<` or `>>` operator must take two parameters, the first of
which must have type `T` or `T?` and the second of which must have type
`int` or `int?`, and can return any type.
+*  A binary `<<` or `>>` operator must take two parameters, the first of
which must have type `T` or `T?`, and can return any type.
```

That is, the restriction that the first parameter be `T` or `T?` remains. While the restriction that the second operand must be `int` or `int?` is removed.

Binary operator overload resolution

The first bullet point at §11.4.5 should be reworded as follows:

- The set of candidate user-defined operators provided by `X` and `Y` for the operation `operator op(x,y)` is determined. The set consists of the union of the candidate operators provided by `X` and **, unless the operator is a shift operator,** the candidate operators provided by `Y`, each determined using the rules of Candidate user-defined operators §11.4.6. If `X` and `Y` are the same type, or if `X` and `Y` are derived from a common base type, then shared candidate operators only occur in the combined set once.

That is, for shift operators, candidate operators are only those provided by type `X`.

Drawbacks

Users will be able to define operators that do not follow the recommended guidelines, such as implementing `cout << "string"` in C#.

Alternatives

The generic math interfaces being exposed by the libraries could expose explicitly named methods instead. This may make code more difficult to read/maintain.

The generic math interfaces could require the shift take `int` and that a conversion be performed. This conversion may be expensive or may be not possible depending on the type in question.

Unresolved questions

Is there concern around preserving the "intent" around why the second operand was restricted to `int`?

Design meetings

https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-09.md

CallerArgumentExpression

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow developers to capture the expressions passed to a method, to enable better error messages in diagnostic/testing APIs and reduce keystrokes.

Motivation

When an assertion or argument validation fails, the developer wants to know as much as possible about where and why it failed. However, today's diagnostic APIs do not fully facilitate this. Consider the following method:

C#

```
T Single<T>(this T[] array)
{
    Debug.Assert(array != null);
    Debug.Assert(array.Length == 1);

    return array[0];
}
```

When one of the asserts fail, only the filename, line number, and method name will be provided in the stack trace. The developer will not be able to tell which assert failed from this information-- they will have to open the file and navigate to the provided line number to see what went wrong.

This is also the reason testing frameworks have to provide a variety of assert methods. With xUnit, `Assert.True` and `Assert.False` are not frequently used because they do not provide enough context about what failed.

While the situation is a bit better for argument validation because the names of invalid arguments are shown to the developer, the developer must pass these names to exceptions manually. If the above example were rewritten to use traditional argument validation instead of `Debug.Assert`, it would look like

C#

```
T Single<T>(this T[] array)
{
    if (array == null)
    {
        throw new ArgumentNullException(nameof(array));
    }

    if (array.Length != 1)
    {
        throw new ArgumentException("Array must contain a single element.",
nameof(array));
    }

    return array[0];
}
```

Notice that `nameof(array)` must be passed to each exception, although it's already clear from context which argument is invalid.

Detailed design

In the above examples, including the string `"array != null"` or `"array.Length == 1"` in the assert message would help the developer determine what failed. Enter `CallerArgumentExpression`: it's an attribute the framework can use to obtain the string associated with a particular method argument. We would add it to `Debug.Assert` like so

C#

```
public static class Debug
{
    public static void Assert(bool condition,
[CallerArgumentExpression("condition")] string message = null);
}
```

The source code in the above example would stay the same. However, the code the compiler actually emits would correspond to

C#

```
T Single<T>(this T[] array)
{
    Debug.Assert(array != null, "array != null");
    Debug.Assert(array.Length == 1, "array.Length == 1");

    return array[0];
}
```

The compiler specially recognizes the attribute on `Debug.Assert`. It passes the string associated with the argument referred to in the attribute's constructor (in this case, `condition`) at the call site. When either assert fails, the developer will be shown the condition that was false and will know which one failed.

For argument validation, the attribute cannot be used directly, but can be made use of through a helper class:

C#

```
public static class Verify
{
    public static void Argument(bool condition, string message,
[CallerArgumentExpression("condition")] string conditionExpression = null)
    {
        if (!condition) throw new ArgumentException(message: message,
paramName: conditionExpression);
    }

    public static void InRange(int argument, int low, int high,
        [CallerArgumentExpression("argument")] string argumentExpression =
null,
        [CallerArgumentExpression("low")] string lowExpression = null,
        [CallerArgumentExpression("high")] string highExpression = null)
    {
        if (argument < low)
        {
            throw new ArgumentOutOfRangeException(paramName:
argumentExpression,
                message: $"{argumentExpression} ({argument}) cannot be less
than {lowExpression} ({low}).");
        }

        if (argument > high)
        {
            throw new ArgumentOutOfRangeException(paramName:
argumentExpression,
```

```
                message: $"{argumentExpression} ({argument}) cannot be
greater than {highExpression} ({high}).");
        }
    }

    public static void NotNull<T>(T argument,
[CallerArgumentExpression("argument")] string argumentExpression = null)
        where T : class
    {
        if (argument == null) throw new ArgumentNullException(paramName:
argumentExpression);
    }
}

T Single<T>(this T[] array)
{
    Verify.NotNull(array); // paramName: "array"
    Verify.Argument(array.Length == 1, "Array must contain a single
element."); // paramName: "array.Length == 1"

    return array[0];
}

T ElementAt(this T[] array, int index)
{
    Verify.NotNull(array); // paramName: "array"
    // paramName: "index"
    // message: "index (-1) cannot be less than 0 (0).", or
    //          "index (6) cannot be greater than array.Length - 1 (5)."
    Verify.InRange(index, 0, array.Length - 1);

    return array[index];
}
```

A proposal to add such a helper class to the framework is underway at https://github.com/dotnet/corefx/issues/17068. If this language feature was implemented, the proposal could be updated to take advantage of this feature.

Extension methods

The `this` parameter in an extension method may be referenced by `CallerArgumentExpression`. For example:

C#

```
public static void ShouldBe<T>(this T @this, T expected,
[CallerArgumentExpression("this")] string thisExpression = null) {}

contestant.Points.ShouldBe(1337); // thisExpression: "contestant.Points"
```

`thisExpression` will receive the expression corresponding to the object before the dot. If it's called with static method syntax, e.g. `Ext.ShouldBe(contestant.Points, 1337)`, it will behave as if first parameter wasn't marked `this`.

There should always be an expression corresponding to the `this` parameter. Even if an instance of a class calls an extension method on itself, e.g. `this.Single()` from inside a collection type, the `this` is mandated by the compiler so `"this"` will get passed. If this rule is changed in the future, we can consider passing `null` or the empty string.

Extra details

- Like the other `Caller*` attributes, such as `CallerMemberName`, this attribute may only be used on parameters with default values.
- Multiple parameters marked with `CallerArgumentExpression` are permitted, as shown above.
- The attribute's namespace will be `System.Runtime.CompilerServices`.
- If `null` or a string that is not a parameter name (e.g. `"notAParameterName"`) is provided, the compiler will pass in an empty string.
- The type of the parameter `CallerArgumentExpressionAttribute` is applied to must have a standard conversion from `string`. This means no user-defined conversions from `string` are allowed, and in practice means the type of such a parameter must be `string`, `object`, or an interface implemented by `string`.

Drawbacks

- People who know how to use decompilers will be able to see some of the source code at call sites for methods marked with this attribute. This may be undesirable/unexpected for closed-source software.

- Although this is not a flaw in the feature itself, a source of concern may be that there exists a `Debug.Assert` API today that only takes a `bool`. Even if the overload taking a message had its second parameter marked with this attribute and made optional, the compiler would still pick the no-message one in overload resolution. Therefore, the no-message overload would have to be removed to take advantage of this feature, which would be a binary (although not source) breaking change.

Alternatives

- If being able to see source code at call sites for methods that use this attribute proves to be a problem, we can make the attribute's effects opt-in. Developers will enable it through an assembly-wide `[assembly: EnableCallerArgumentExpression]` attribute they put in `AssemblyInfo.cs`.
  - In the case the attribute's effects are not enabled, calling methods marked with the attribute would not be an error, to allow existing methods to use the attribute and maintain source compatibility. However, the attribute would be ignored and the method would be called with whatever default value was provided.

C#

```
// Assembly1

void Foo(string bar); // V1
void Foo(string bar, string barExpression = "not provided"); // V2
void Foo(string bar, [CallerArgumentExpression("bar")] string barExpression
= "not provided"); // V3

// Assembly2

Foo(a); // V1: Compiles to Foo(a), V2, V3: Compiles to Foo(a, "not
provided")
Foo(a, "provided"); // V2, V3: Compiles to Foo(a, "provided")

// Assembly3

[assembly: EnableCallerArgumentExpression]

Foo(a); // V1: Compiles to Foo(a), V2: Compiles to Foo(a, "not provided"),
V3: Compiles to Foo(a, "a")
Foo(a, "provided"); // V2, V3: Compiles to Foo(a, "provided")
```

- To prevent the binary compatibility problem from occurring every time we want to add new caller info to `Debug.Assert`, an alternative solution would be to add a `CallerInfo` struct to the framework that contains all the necessary information about the caller.

C#

```
struct CallerInfo
{
    public string MemberName { get; set; }
    public string TypeName { get; set; }
    public string Namespace { get; set; }
    public string FullTypeName { get; set; }
    public string FilePath { get; set; }
    public int LineNumber { get; set; }
```

```
    public int ColumnNumber { get; set; }
    public Type Type { get; set; }
    public MethodBase Method { get; set; }
    public string[] ArgumentExpressions { get; set; }
}

[Flags]
enum CallerInfoOptions
{
    MemberName = 1, TypeName = 2, ...
}

public static class Debug
{
    public static void Assert(bool condition,
        // If a flag is not set here, the corresponding CallerInfo member is
not populated by the caller, so it's
        // pay-for-play friendly.
        [CallerInfo(CallerInfoOptions.FilePath | CallerInfoOptions.Method |
CallerInfoOptions.ArgumentExpressions)] CallerInfo callerInfo =
default(CallerInfo))
    {
        string filePath = callerInfo.FilePath;
        MethodBase method = callerInfo.Method;
        string conditionExpression = callerInfo.ArgumentExpressions[0];
        //...
    }
}

class Bar
{
    void Foo()
    {
        Debug.Assert(false);

        // Translates to:

        var callerInfo = new CallerInfo();
        callerInfo.FilePath = @"C:\Bar.cs";
        callerInfo.Method = MethodBase.GetCurrentMethod();
        callerInfo.ArgumentExpressions = new string[] { "false" };
        Debug.Assert(false, callerInfo);
    }
}
```

This was originally proposed at https://github.com/dotnet/csharplang/issues/87 .

There are a few disadvantages of this approach:

- Despite being pay-for-play friendly by allowing you to specify which properties you need, it could still hurt perf significantly by allocating an array for the expressions/calling `MethodBase.GetCurrentMethod` even when the assert passes.

- Additionally, while passing a new flag to the `CallerInfo` attribute won't be a breaking change, `Debug.Assert` won't be guaranteed to actually receive that new parameter from call sites that compiled against an old version of the method.

Unresolved questions

TBD

Design meetings

N/A

Extended `nameof` scope

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow `nameof(parameter)` inside an attribute on a method or parameter. For example:

- `[MyAttribute(nameof(parameter))] void M(int parameter) { }`
- `[MyAttribute(nameof(TParameter))] void M<TParameter>() { }`
- `void M(int parameter, [MyAttribute(nameof(parameter))] int other) { }`

Motivation

Attributes like `NotNullWhen` or `CallerExpression` need to refer to parameters, but those parameters are currently not in scope.

Detailed design

Methods

The method's *type_parameters* are in scope throughout the *method_declaration*, and can be used to form types throughout that scope in *return_type*, *method_body*, and *type_parameter_constraints_clauses* but not in *attributes*, **except within a `nameof` expression in *attributes***.

Method parameters

A method declaration creates a separate declaration space for parameters, type parameters and local variables. Names are introduced into this declaration space by the

type parameter list and the formal parameter list of the method and by local variable declarations in the block of the method. **Names are introduced into this declaration space by the type parameter list and the formal parameter list of the method in `nameof` expressions in attributes placed on the method or its parameters.**

[...]
Within the block of a method, formal parameters can be referenced by their identifiers in simple_name expressions (Simple names). **Within a `nameof` expression in attributes placed on the method or its parameters, formal parameters can be referenced by their identifiers in *simple_name* expressions.**

Anonymous function signatures

The scope of the parameters of the anonymous function is the anonymous_function_body (§7.7) **and `nameof` expressions in attributes placed on the anonymous function or its parameters**.

Delegate declarations

The scope of the parameters of the delegate is `nameof` expressions in attributes placed on the declaration, its type parameters or its parameters.

Simple names

A *simple_name* is either of the form `I` or of the form `I<A1,...,Ak>`, where `I` is a single identifier and `<A1,...,Ak>` is an optional *type_argument_list*. When no *type_argument_list* is specified, consider `K` to be zero. The *simple_name* is evaluated and classified as follows:

- If `K` is zero and the *simple_name* appears within a block and if the block's (or an enclosing block's) local variable declaration space (Declarations) contains a local variable, parameter or constant with name `I`, then the *simple_name* refers to that local variable, parameter or constant and is classified as a variable or value.
- If `K` is zero and the *simple_name* appears within the body of a generic method declaration and if that declaration includes a type parameter with name `I`, then the *simple_name* refers to that type parameter.
- **If `K` is zero and the *simple_name* appears within a `nameof` expression in an attribute on the method declaration or its parameters and if that declaration includes a parameter or type parameter with name `I`, then the *simple_name* refers to that parameter or type parameter.**
- Otherwise, for each instance type `T` (The instance type), starting with the instance type of the immediately enclosing type declaration and continuing with the

instance type of each enclosing class or struct declaration (if any):
[...]

- Otherwise, for each namespace `N`, starting with the namespace in which the *simple_name* occurs, continuing with each enclosing namespace (if any), and ending with the global namespace, the following steps are evaluated until an entity is located:
  [...]
- Otherwise, the simple_name is undefined and a compile-time error occurs.

Scopes

- The scope of a type parameter declared by a type_parameter_list on a method_declaration is [...] **and `nameof` expressions in an attribute on the method declaration or its parameters.**
- The scope of a parameter declared in a method_declaration (Methods) is the *method_body* of that method_declaration **and `nameof` expressions in an attribute on the method declaration or its parameters.**

Related spec sections

- Declarations

Global Using Directive

Article • 03/24/2022

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Syntax for a using directive is extended with an optional `global` keyword that can precede the `using` keyword:

antlr

```
compilation_unit
    : extern_alias_directive* global_using_directive* using_directive*
global_attributes? namespace_member_declaration*
    ;

global_using_directive
    : global_using_alias_directive
    | global_using_namespace_directive
    | global_using_static_directive
    ;

global_using_alias_directive
    : 'global' 'using' identifier '=' namespace_or_type_name ';'
    ;

global_using_namespace_directive
    : 'global' 'using' namespace_name ';'
    ;

global_using_static_directive
    : 'global' 'using' 'static' type_name ';'
    ;
```

- The *global_using_directive*s are allowed only on the Compilation Unit level (cannot be used inside a *namespace_declaration*).
- The *global_using_directive*s, if any, must precede any *using_directive*s.

- The scope of a *global_using_directive*s extends over the *namespace_member_declaration*s of all compilation units within the program. The scope of a *global_using_directive* specifically does not include other *global_using_directive*s. Thus, peer *global_using_directive*s or those from a different compilation unit do not affect each other, and the order in which they are written is insignificant. The scope of a *global_using_directive* specifically does not include *using_directive*s immediately contained in any compilation unit of the program.

The effect of adding a *global_using_directive* to a program can be thought of as the effect of adding a similar *using_directive* that resolves to the same target namespace or type to every compilation unit of the program. However, the target of a *global_using_directive* is resolved in context of the compilation unit that contains it.

§7.7 Scopes

These are the relevant bullet points with proposed additions (which are **in bold**):

- The scope of name defined by an *extern_alias_directive* extends over the ***global_using_directives*, ***using_directive*s, *global_attributes* and *namespace_member_declaration*s of its immediately containing compilation unit or namespace body. An *extern_alias_directive* does not contribute any new members to the underlying declaration space. In other words, an *extern_alias_directive* is not transitive, but, rather, affects only the compilation unit or namespace body in which it occurs.
- **The scope of a name defined or imported by a *global_using_directive* extends over the *global_attributes* and *namespace_member_declaration*s of all the *compilation_unit*s in the program.**

§7.8 Namespace and type names

Changes are made to the algorithm determining the meaning of a *namespace_or_type_name* as follows.

This is the relevant bullet point with proposed additions (which are **in bold**):

- If the *namespace_or_type_name* is of the form `I` or of the form `I<A1, ..., Ak>`:
  - If `K` is zero and the *namespace_or_type_name* appears within a generic method declaration (§14.6) and if that declaration includes a type parameter (§14.2.3) with name `I`, then the *namespace_or_type_name* refers to that type parameter.

- Otherwise, if the *namespace_or_type_name* appears within a type declaration, then for each instance type `T` (§14.3.2), starting with the instance type of that type declaration and continuing with the instance type of each enclosing class or struct declaration (if any):
  - If `K` is zero and the declaration of `T` includes a type parameter with name `I`, then the *namespace_or_type_name* refers to that type parameter.
  - Otherwise, if the *namespace_or_type_name* appears within the body of the type declaration, and `T` or any of its base types contain a nested accessible type having name `I` and `K` type parameters, then the *namespace_or_type_name* refers to that type constructed with the given type arguments. If there is more than one such type, the type declared within the more derived type is selected. Note that non-type members (constants, fields, methods, properties, indexers, operators, instance constructors, destructors, and static constructors) and type members with a different number of type parameters are ignored when determining the meaning of the *namespace_or_type_name*.
- If the previous steps were unsuccessful then, for each namespace `N`, starting with the namespace in which the *namespace_or_type_name* occurs, continuing with each enclosing namespace (if any), and ending with the global namespace, the following steps are evaluated until an entity is located:
  - If `K` is zero and `I` is the name of a namespace in `N`, then:
    - If the location where the *namespace_or_type_name* occurs is enclosed by a namespace declaration for `N` and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I` with a namespace or type, **or any namespace declaration for `N` in the program contains a *global_using_alias_directive* that associates the name `I` with a namespace or type,** then the *namespace_or_type_name* is ambiguous and a compile-time error occurs.
    - Otherwise, the *namespace_or_type_name* refers to the namespace named `I` in `N`.
  - Otherwise, if `N` contains an accessible type having name `I` and `K` type parameters, then:
    - If `K` is zero and the location where the *namespace_or_type_name* occurs is enclosed by a namespace declaration for `N` and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I` with a namespace or type, **or any namespace declaration for `N` in the program contains a *global_using_alias_directive* that associates the name `I` with a namespace or type,** then the *namespace_or_type_name* is ambiguous and a compile-time error occurs.

        - Otherwise, the *namespace_or_type_name* refers to the type constructed with the given type arguments.
    - Otherwise, if the location where the *namespace_or_type_name* occurs is enclosed by a namespace declaration for `N`:
        - If `K` is zero and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I` with an imported namespace or type, **or any namespace declaration for `N` in the program contains a *global_using_alias_directive* that associates the name `I` with an imported namespace or type,** then the *namespace_or_type_name* refers to that namespace or type.
        - Otherwise, if the namespaces and type declarations imported by the *using_namespace_directive*s and *using_alias_directive*s of the namespace declaration **and the namespaces and type declarations imported by the *global_using_namespace_directive*s and *global_using_static_directive*s of any namespace declaration for `N` in the program** contain exactly one accessible type having name `I` and `K` type parameters, then the *namespace_or_type_name* refers to that type constructed with the given type arguments.
        - Otherwise, if the namespaces and type declarations imported by the *using_namespace_directive*s and *using_alias_directive*s of the namespace declaration **and the namespaces and type declarations imported by the *global_using_namespace_directive*s and *global_using_static_directive*s of any namespace declaration for `N` in the program** contain more than one accessible type having name `I` and `K` type parameters, then the *namespace_or_type_name* is ambiguous and an error occurs.
    - Otherwise, the *namespace_or_type_name* is undefined and a compile-time error occurs.

Simple names §11.7.4

Changes are made to the *simple_name* evaluation rules as follows.

This is the relevant bullet point with proposed additions (which are **in bold**):

- Otherwise, for each namespace `N`, starting with the namespace in which the *simple_name* occurs, continuing with each enclosing namespace (if any), and ending with the global namespace, the following steps are evaluated until an entity is located:
    - If `K` is zero and `I` is the name of a namespace in `N`, then:

    - If the location where the *simple_name* occurs is enclosed by a namespace declaration for `N` and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I` with a namespace or type, **or any namespace declaration for `N` in the program contains a *global_using_alias_directive* that associates the name `I` with a namespace or type,** then the *simple_name* is ambiguous and a compile-time error occurs.
    - Otherwise, the *simple_name* refers to the namespace named `I` in `N`.
- Otherwise, if `N` contains an accessible type having name `I` and `K` type parameters, then:
    - If `K` is zero and the location where the *simple_name* occurs is enclosed by a namespace declaration for `N` and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I` with a namespace or type, **or any namespace declaration for `N` in the program contains a *global_using_alias_directive* that associates the name `I` with a namespace or type,** then the *simple_name* is ambiguous and a compile-time error occurs.
    - Otherwise, the *namespace_or_type_name* refers to the type constructed with the given type arguments.
- Otherwise, if the location where the *simple_name* occurs is enclosed by a namespace declaration for `N`:
    - If `K` is zero and the namespace declaration contains an *extern_alias_directive* or *using_alias_directive* that associates the name `I` with an imported namespace or type, **or any namespace declaration for `N` in the program contains a *global_using_alias_directive* that associates the name `I` with an imported namespace or type,** then the *simple_name* refers to that namespace or type.
    - Otherwise, if the namespaces and type declarations imported by the *using_namespace_directive*s and *using_static_directive*s of the namespace declaration **and the namespaces and type declarations imported by the *global_using_namespace_directive*s and *global_using_static_directive*s of any namespace declaration for `N` in the program** contain exactly one accessible type or non-extension static member having name `I` and `K` type parameters, then the *simple_name* refers to that type or member constructed with the given type arguments.
    - Otherwise, if the namespaces and types imported by the *using_namespace_directive*s of the namespace declaration **and the namespaces and type declarations imported by the *global_using_namespace_directive*s and *global_using_static_directive*s of**

any namespace declaration for `N` in the program contain more than one accessible type or non-extension-method static member having name `I` and `K` type parameters, then the *simple_name* is ambiguous and an error occurs.

Extension method invocations §11.7.8.3

Changes are made to the algorithm to find the best *type_name* `C` as follows. This is the relevant bullet point with proposed additions (which are **in bold**):

- Starting with the closest enclosing namespace declaration, continuing with each enclosing namespace declaration, and ending with the containing compilation unit, successive attempts are made to find a candidate set of extension methods:
  - If the given namespace or compilation unit directly contains non-generic type declarations `Ci` with eligible extension methods `Mj`, then the set of those extension methods is the candidate set.
  - If types `Ci` imported by *using_static_declarations* and directly declared in namespaces imported by *using_namespace_directive*s in the given namespace or compilation unit **and, if containing compilation unit is reached, imported by *global_using_static_declarations* and directly declared in namespaces imported by *global_using_namespace_directive*s in the program** directly contain eligible extension methods `Mj`, then the set of those extension methods is the candidate set.

Compilation units §13.2

A *compilation_unit* defines the overall structure of a source file. A compilation unit consists of **zero or more *global_using_directive*s followed by** zero or more *using_directive*s followed by zero or more *global_attributes* followed by zero or more *namespace_member_declaration*s.

antlr

```
compilation_unit
    : extern_alias_directive* global_using_directive* using_directive*
global_attributes? namespace_member_declaration*
    ;
```

A C# program consists of one or more compilation units, each contained in a separate source file. When a C# program is compiled, all of the compilation units are processed together. Thus, compilation units can depend on each other, possibly in a circular fashion.

The *global_using_directives* of a compilation unit affect the *global_attributes* and *namespace_member_declarations* of all compilation units in the program.

Extern aliases §13.4

The scope of an *extern_alias_directive* extends over the ***global_using_directives*, *using_directives*, *global_attributes*** and *namespace_member_declarations* of its immediately containing compilation unit or namespace body.

Using alias directives §13.5.2

The order in which *using_alias_directives* are written has no significance, and resolution of the *namespace_or_type_name* referenced by a *using_alias_directive* is not affected by the *using_alias_directive* itself or by other *using_directives* in the immediately containing compilation unit or namespace body, **and, if the *using_alias_directive* is immediately contained in a compilation unit, is not affected by the *global_using_directives* in the program**. In other words, the *namespace_or_type_name* of a *using_alias_directive* is resolved as if the immediately containing compilation unit or namespace body had no *using_directives* **and, if the *using_alias_directive* is immediately contained in a compilation unit, the program had no *global_using_directives***. A *using_alias_directive* may however be affected by *extern_alias_directives* in the immediately containing compilation unit or namespace body.

Global Using alias directives

A *global_using_alias_directive* introduces an identifier that serves as an alias for a namespace or type within the program.

antlr

```
global_using_alias_directive
    : 'global' 'using' identifier '=' namespace_or_type_name ';'
    ;
```

Within member declarations in any compilation unit of a program that contains a *global_using_alias_directive*, the identifier introduced by the *global_using_alias_directive* can be used to reference the given namespace or type.

The *identifier* of a *global_using_alias_directive* must be unique within the declaration space of any compilation unit of a program that contains the *global_using_alias_directive*.

Just like regular members, names introduced by *global_using_alias_directive*s are hidden by similarly named members in nested scopes.

The order in which *global_using_alias_directive*s are written has no significance, and resolution of the *namespace_or_type_name* referenced by a *global_using_alias_directive* is not affected by the *global_using_alias_directive* itself or by other *global_using_directive*s or *using_directive*s in the program. In other words, the *namespace_or_type_name* of a *global_using_alias_directive* is resolved as if the immediately containing compilation unit had no *using_directive*s and the entire containing program had no *global_using_directive*s. A *global_using_alias_directive* may however be affected by *extern_alias_directive*s in the immediately containing compilation unit.

A *global_using_alias_directive* can create an alias for any namespace or type.

Accessing a namespace or type through an alias yields exactly the same result as accessing that namespace or type through its declared name.

Using aliases can name a closed constructed type, but cannot name an unbound generic type declaration without supplying type arguments.

Global Using namespace directives

A *global_using_namespace_directive* imports the types contained in a namespace into the program, enabling the identifier of each type to be used without qualification.

antlr

```
global_using_namespace_directive
    : 'global' 'using' namespace_name ';'
    ;
```

Within member declarations in a program that contains a *global_using_namespace_directive*, the types contained in the given namespace can be referenced directly.

A *global_using_namespace_directive* imports the types contained in the given namespace, but specifically does not import nested namespaces.

Unlike a *global_using_alias_directive*, a *global_using_namespace_directive* may import types whose identifiers are already defined within a compilation unit of the program. In effect, in a given compilation unit, names imported by any

global_using_namespace_directive in the program are hidden by similarly named members in the compilation unit.

When more than one namespace or type imported by *global_using_namespace_directive*s or *global_using_static_directive*s in the same program contain types by the same name, references to that name as a *type_name* are considered ambiguous.

Furthermore, when more than one namespace or type imported by *global_using_namespace_directive*s or *global_using_static_directive*s in the same program contain types or members by the same name, references to that name as a *simple_name* are considered ambiguous.

The *namespace_name* referenced by a *global_using_namespace_directive* is resolved in the same way as the *namespace_or_type_name* referenced by a *global_using_alias_directive*. Thus, *global_using_namespace_directive*s in the same program do not affect each other and can be written in any order.

Global Using static directives

A *global_using_static_directive* imports the nested types and static members contained directly in a type declaration into the containing program, enabling the identifier of each member and type to be used without qualification.

antlr

```
global_using_static_directive
    : 'global' 'using' 'static' type_name ';'
    ;
```

Within member declarations in a program that contains a *global_using_static_directive*, the accessible nested types and static members (except extension methods) contained directly in the declaration of the given type can be referenced directly.

A *global_using_static_directive* specifically does not import extension methods directly as static methods, but makes them available for extension method invocation.

A *global_using_static_directive* only imports members and types declared directly in the given type, not members and types declared in base classes.

Ambiguities between multiple *global_using_namespace_directive*s and *global_using_static_directives* are discussed in the section for *global_using_namespace_directive*s (above).

Qualified alias member §13.8

Changes are made to the algorithm determining the meaning of a *qualified_alias_member* as follows.

This is the relevant bullet point with proposed additions (which are **in bold**):

- Otherwise, starting with the namespace declaration (§13.3) immediately containing the *qualified_alias_member* (if any), continuing with each enclosing namespace declaration (if any), and ending with the compilation unit containing the *qualified_alias_member*, the following steps are evaluated until an entity is located:
  - If the namespace declaration or compilation unit contains a *using_alias_directive* that associates `N` with a type, **or, when a compilation unit is reached, the program contains a *global_using_alias_directive* that associates `N` with a type,** then the *qualified_alias_member* is undefined and a compile-time error occurs.
  - Otherwise, if the namespace declaration or compilation unit contains an *extern_alias_directive* or *using_alias_directive* that associates `N` with a namespace, ***or, when a compilation unit is reached, the program contains a *global_using_alias_directive* that associates `N` with a namespace,** then:
    - If the namespace associated with `N` contains a namespace named `I` and `K` is zero, then the *qualified_alias_member* refers to that namespace.
    - Otherwise, if the namespace associated with `N` contains a non-generic type named `I` and `K` is zero, then the *qualified_alias_member* refers to that type.
    - Otherwise, if the namespace associated with `N` contains a type named `I` that has `K` type parameters, then the *qualified_alias_member* refers to that type constructed with the given type arguments.
    - Otherwise, the *qualified_alias_member* is undefined and a compile-time error occurs.

Allow using alias directive to reference any kind of Type

Article • 08/16/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Relax the using_alias_directive (§13.5.2) to allow it to point at any sort of type, not just named types. This would support types not allowed today, like: tuple types, pointer types, array types, etc. For example, this would now be allowed:

```c#
using Point = (int x, int y);
```

Motivation

For ages, C# has had the ability to introduce aliases for namespaces and named types (classes, delegated, interfaces, records and structs). This worked acceptably well as it provided a means to introduce non-conflicting names in cases where a normal named pulled in from `using_directive`s might be ambiguous, and it allowed a way to provide a simpler name when dealing with complex generic types. However, the rise of additional complex type symbols in the language has caused more use to arise where aliases would be valuable but are currently not allowed. For example, both tuples and function-pointers often can have large and complex regular textual forms that can be painful to continually write out, and a burden to try to read. Aliases would help in these cases by giving a short, developer-provided, name that can then be used in place of those full structural forms.

Detailed design

We will change the grammar of `using_alias_directive` thusly:

```
using_alias_directive
-    : 'using' identifier '=' namespace_or_type_name ';'
+    : 'using' identifier '=' (namespace_name | type) ';'
    ;
```

Top-level reference type nullability annotations are disallowed.

Interestingly, most of the spec language in §13.5.2 does not need to change. Most language in it already refers to 'namespace or type', for example:

> A using_alias_directive introduces an identifier that serves as an alias for a namespace or type within the immediately enclosing compilation unit or namespace body.

This remains true, just that the grammar now allows the 'type' to be any arbitrary type, not the limited set allowed for by `namespace_or_type_name` previously.

The sections that do need updating are:

diff

```
- The order in which using_alias_directives are written has no significance, and resolution of the namespace_or_type_name referenced by a using_alias_directive is not affected by the using_alias_directive itself or by other using_directives in the immediately containing compilation unit or namespace body. In other words, the namespace_or_type_name of a using_alias_directive is resolved as if the immediately containing compilation unit or namespace body had no using_directives. A using_alias_directive may however be affected by extern_alias_directives in the immediately containing compilation unit or namespace body. In the example
+ The order in which using_alias_directives are written has no significance, and resolution of the `(namespace_name | type)` referenced by a using_alias_directive is not affected by the using_alias_directive itself or by other using_directives in the immediately containing compilation unit or namespace body. In other words, the `(namespace_name | type)` of a using_alias_directive is resolved as if the immediately containing compilation unit or namespace body had no using_directives. A using_alias_directive may however be affected by extern_alias_directives in the immediately containing compilation unit or namespace body. In the example
```

diff

```
- The namespace_name referenced by a using_namespace_directive is resolved
in the same way as the namespace_or_type_name referenced by a
using_alias_directive. Thus, using_namespace_directives in the same
compilation unit or namespace body do not affect each other and can be
written in any order.
+ The namespace_name referenced by a using_namespace_directive is resolved
in the same way as the namespace_or_type_name referenced by a
using_alias_directive. Thus, using_namespace_directives in the same
compilation unit or namespace body do not affect each other and can be
written in any order.
```

diff

```
+ It is illegal for a using alias type to be a nullable reference type.

    1. `using X = string?;` is not legal.
    2. `using X = List<string?>;` is legal.  The alias is to `List<...>`
which is itself not a nullable reference type itself, even though it
contains one as a type argument.
    3. `using X = int?;` is legal.  This is a nullable *value* type, not a
nullable *reference* type.
```

Supporting aliases to types containing pointers.

A new unsafe context is added through an optional 'unsafe' keyword in the using_alias_directive production:

diff

```
using_alias_directive
+    : 'using' 'unsafe'? identifier '=' (namespace_name | type) ';'
    ;

using_static_directive
+    : 'using' 'static' 'unsafe'? type_name ';'
    ;

+ 'unsafe' can only be used with an using_alias_directive or
using_static_directive, not a using_directive.
+ The 'unsafe' keyword present in a 'using_alias_directive' causes the
entire textual extent of the 'type' portion (not the 'namespace_name'
portion) to become an unsafe context.
+ The 'unsafe' keyword present in a 'using_static_directive' causes the
```

```
entire textual extent of the 'type_name' portion to become an unsafe
context.
```

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

C# feature specification feedback

The C# feature specifications are open source. Provide feedback here.

- Open a documentation issue
- Provide product feedback

Attributes on local functions

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Attributes

Local function declarations are now permitted to have attributes (§21). Parameters and type parameters on local functions are also allowed to have attributes.

Attributes with a specified meaning when applied to a method, its parameters, or its type parameters will have the same meaning when applied to a local function, its parameters, or its type parameters, respectively.

A local function can be made conditional in the same sense as a conditional method (§21.5.3) by decorating it with a `[ConditionalAttribute]`. A conditional local function must also be `static`. All restrictions on conditional methods also apply to conditional local functions, including that the return type must be `void`.

Extern

The `extern` modifier is now permitted on local functions. This makes the local function external in the same sense as an external method (§14.6.8).

Similarly to an external method, the *local-function-body* of an external local function must be a semicolon. A semicolon *local-function-body* is only permitted on an external local function.

An external local function must also be `static`.

Syntax

The local functions grammar is modified as follows:

```
local-function-header
    : attributes? local-function-modifiers? return-type identifier type-
parameter-list?
        ( formal-parameter-list? ) type-parameter-constraints-clauses
    ;

local-function-modifiers
    : (async | unsafe | static | extern)*
    ;

local-function-body
    : block
    | arrow-expression-body
    | ';'
    ;
```

"pattern-based using" and "using declarations"

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

The language will add two new capabilities around the `using` statement in order to make resource management simpler: `using` should recognize a disposable pattern in addition to `IDisposable` and add a `using` declaration to the language.

Motivation

The `using` statement is an effective tool for resource management today but it requires quite a bit of ceremony. Methods that have a number of resources to manage can get syntactically bogged down with a series of `using` statements. This syntax burden is enough that most coding style guidelines explicitly have an exception around braces for this scenario.

The `using` declaration removes much of the ceremony here and gets C# on par with other languages that include resource management blocks. Additionally the pattern-based `using` lets developers expand the set of types that can participate here. In many cases removing the need to create wrapper types that only exist to allow for a values use in a `using` statement.

Together these features allow developers to simplify and expand the scenarios where `using` can be applied.

Detailed Design

using declaration

The language will allow for `using` to be added to a local variable declaration. Such a declaration will have the same effect as declaring the variable in a `using` statement at the same location.

C#

```
if (...)
{
   using FileStream f = new FileStream(@"C:\users\jaredpar\using.md");
   // statements
}

// Equivalent to
if (...)
{
   using (FileStream f = new FileStream(@"C:\users\jaredpar\using.md"))
   {
    // statements
   }
}
```

The lifetime of a `using` local will extend to the end of the scope in which it is declared. The `using` locals will then be disposed in the reverse order in which they are declared.

C#

```
{
    using var f1 = new FileStream("...");
    using var f2 = new FileStream("..."), f3 = new FileStream("...");
    ...
    // Dispose f3
    // Dispose f2
    // Dispose f1
}
```

There are no restrictions around `goto`, or any other control flow construct in the face of a `using` declaration. Instead the code acts just as it would for the equivalent `using` statement:

C#

```
{
    using var f1 = new FileStream("...");
  target:
    using var f2 = new FileStream("...");
    if (someCondition)
    {
        // Causes f2 to be disposed but has no effect on f1
        goto target;
    }
}
```

A local declared in a `using` local declaration will be implicitly read-only. This matches the behavior of locals declared in a `using` statement.

The language grammar for `using` declarations will be the following:

antlr

```
local-using-declaration:
  'using' type using-declarators

using-declarators:
  using-declarator
  using-declarators , using-declarator

using-declarator:
  identifier = expression
```

Restrictions around `using` declaration:

- May not appear directly inside a `case` label but instead must be within a block inside the `case` label.
- May not appear as part of an `out` variable declaration.
- Must have an initializer for each declarator.
- The local type must be implicitly convertible to `IDisposable` or fulfill the `using` pattern.

pattern-based using

The language will add the notion of a disposable pattern for `ref struct` types: that is a `ref struct` which has an accessible `Dispose` instance method. Types which fit the disposable pattern can participate in a `using` statement or declaration without being required to implement `IDisposable`.

```

```

C#

```
ref struct Resource
{
    public void Dispose() { ... }
}

using (var r = new Resource())
{
    // statements
}
```

This will allow developers to leverage `using` for `ref struct` types. These types can't implement interfaces today and hence can't participate in `using` statements.

The same restrictions from a traditional `using` statement apply here as well: local variables declared in the `using` are read-only, a `null` value will not cause an exception to be thrown, etc ... The code generation will be different only in that there will not be a cast to `IDisposable` before calling Dispose:

C#

```
{
      Resource r = new Resource();
      try {
            // statements
      }
      finally {
            if (r != null) r.Dispose();
      }
}
```

In order to fit the disposable pattern the `Dispose` method must be an accessible instance member, parameterless and have a `void` return type. It cannot be an extension method.

Considerations

case labels without blocks

A `using declaration` is illegal directly inside a `case` label due to complications around its actual lifetime. One potential solution is to simply give it the same lifetime as an `out var` in the same location. It was deemed the extra complexity to the feature

implementation and the ease of the work around (just add a block to the `case` label) didn't justify taking this route.

Future Expansions

fixed locals

A `fixed` statement has all of the properties of `using` statements that motivated the ability to have `using` locals. Consideration should be given to extending this feature to `fixed` locals as well. The lifetime and ordering rules should apply equally well for `using` and `fixed` here.

Async Streams

Article • 10/11/2022

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

C# has support for iterator methods and async methods, but no support for a method that is both an iterator and an async method. We should rectify this by allowing for `await` to be used in a new form of `async` iterator, one that returns an `IAsyncEnumerable<T>` or `IAsyncEnumerator<T>` rather than an `IEnumerable<T>` or `IEnumerator<T>`, with `IAsyncEnumerable<T>` consumable in a new `await foreach`. An `IAsyncDisposable` interface is also used to enable asynchronous cleanup.

Related discussion

- https://github.com/dotnet/roslyn/issues/261
- https://github.com/dotnet/roslyn/issues/114

Detailed design

Interfaces

IAsyncDisposable

There has been much discussion of `IAsyncDisposable` (e.g. https://github.com/dotnet/roslyn/issues/114) and whether it's a good idea. However, it's a required concept to add in support of async iterators. Since `finally` blocks may

contain `await`s, and since `finally` blocks need to be run as part of disposing of iterators, we need async disposal. It's also just generally useful any time cleaning up of resources might take any period of time, e.g. closing files (requiring flushes), deregistering callbacks and providing a way to know when deregistration has completed, etc.

The following interface is added to the core .NET libraries (e.g. System.Private.CoreLib / System.Runtime):

```
C#
```

```
namespace System
{
    public interface IAsyncDisposable
    {
        ValueTask DisposeAsync();
    }
}
```

As with `Dispose`, invoking `DisposeAsync` multiple times is acceptable, and subsequent invocations after the first should be treated as nops, returning a synchronously completed successful task (`DisposeAsync` need not be thread-safe, though, and need not support concurrent invocation). Further, types may implement both `IDisposable` and `IAsyncDisposable`, and if they do, it's similarly acceptable to invoke `Dispose` and then `DisposeAsync` or vice versa, but only the first should be meaningful and subsequent invocations of either should be a nop. As such, if a type does implement both, consumers are encouraged to call once and only once the more relevant method based on the context, `Dispose` in synchronous contexts and `DisposeAsync` in asynchronous ones.

(I'm leaving discussion of how `IAsyncDisposable` interacts with `using` to a separate discussion. And coverage of how it interacts with `foreach` is handled later in this proposal.)

Alternatives considered:

- *`DisposeAsync` accepting a `CancellationToken`*: while in theory it makes sense that anything async can be canceled, disposal is about cleanup, closing things out, free'ing resources, etc., which is generally not something that should be canceled; cleanup is still important for work that's canceled. The same `CancellationToken` that caused the actual work to be canceled would typically be the same token passed to `DisposeAsync`, making `DisposeAsync` worthless because cancellation of the work would cause `DisposeAsync` to be a nop. If someone wants to avoid being

blocked waiting for disposal, they can avoid waiting on the resulting `ValueTask`, or wait on it only for some period of time.

- *`DisposeAsync` returning a `Task`*: Now that a non-generic `ValueTask` exists and can be constructed from an `IValueTaskSource`, returning `ValueTask` from `DisposeAsync` allows an existing object to be reused as the promise representing the eventual async completion of `DisposeAsync`, saving a `Task` allocation in the case where `DisposeAsync` completes asynchronously.
- *Configuring `DisposeAsync` with a `bool continueOnCapturedContext` (`ConfigureAwait`)*: While there may be issues related to how such a concept is exposed to `using`, `foreach`, and other language constructs that consume this, from an interface perspective it's not actually doing any `await`'ing and there's nothing to configure... consumers of the `ValueTask` can consume it however they wish.
- *`IAsyncDisposable` inheriting `IDisposable`*: Since only one or the other should be used, it doesn't make sense to force types to implement both.
- *`IDisposableAsync` instead of `IAsyncDisposable`*: We've been following the naming that things/types are an "async something" whereas operations are "done async", so types have "Async" as a prefix and methods have "Async" as a suffix.

IAsyncEnumerable / IAsyncEnumerator

Two interfaces are added to the core .NET libraries:

C#

```
namespace System.Collections.Generic
{
    public interface IAsyncEnumerable<out T>
    {
        IAsyncEnumerator<T> GetAsyncEnumerator(CancellationToken
cancellationToken = default);
    }

    public interface IAsyncEnumerator<out T> : IAsyncDisposable
    {
        ValueTask<bool> MoveNextAsync();
        T Current { get; }
    }
}
```

Typical consumption (without additional language features) would look like:

C#

```
IAsyncEnumerator<T> enumerator = enumerable.GetAsyncEnumerator();
try
{
    while (await enumerator.MoveNextAsync())
    {
        Use(enumerator.Current);
    }
}
finally { await enumerator.DisposeAsync(); }
```

Discarded options considered:

- *`Task<bool> MoveNextAsync(); T current { get; }`*: Using `Task<bool>` would support using a cached task object to represent synchronous, successful `MoveNextAsync` calls, but an allocation would still be required for asynchronous completion. By returning `ValueTask<bool>`, we enable the enumerator object to itself implement `IValueTaskSource<bool>` and be used as the backing for the `ValueTask<bool>` returned from `MoveNextAsync`, which in turn allows for significantly reduced overheads.
- *`ValueTask<(bool, T)> MoveNextAsync();`*: It's not only harder to consume, but it means that `T` can no longer be covariant.
- *`ValueTask<T?> TryMoveNextAsync();`*: Not covariant.
- *`Task<T?> TryMoveNextAsync();`*: Not covariant, allocations on every call, etc.
- *`ITask<T?> TryMoveNextAsync();`*: Not covariant, allocations on every call, etc.
- *`ITask<(bool,T)> TryMoveNextAsync();`*: Not covariant, allocations on every call, etc.
- *`Task<bool> TryMoveNextAsync(out T result);`*: The `out` result would need to be set when the operation returns synchronously, not when it asynchronously completes the task potentially sometime long in the future, at which point there'd be no way to communicate the result.
- *`IAsyncEnumerator<T>` not implementing `IAsyncDisposable`*: We could choose to separate these. However, doing so complicates certain other areas of the proposal, as code must then be able to deal with the possibility that an enumerator doesn't provide disposal, which makes it difficult to write pattern-based helpers. Further, it will be common for enumerators to have a need for disposal (e.g. any C# async iterator that has a finally block, most things enumerating data from a network connection, etc.), and if one doesn't, it is simple to implement the method purely as `public ValueTask DisposeAsync() => default(ValueTask);` with minimal additional overhead.
- _ `IAsyncEnumerator<T> GetAsyncEnumerator()`: No cancellation token parameter.

Viable alternative:

C#

```
namespace System.Collections.Generic
{
    public interface IAsyncEnumerable<out T>
    {
        IAsyncEnumerator<T> GetAsyncEnumerator();
    }

    public interface IAsyncEnumerator<out T> : IAsyncDisposable
    {
        ValueTask<bool> WaitForNextAsync();
        T TryGetNext(out bool success);
    }
}
```

`TryGetNext` is used in an inner loop to consume items with a single interface call as long as they're available synchronously. When the next item can't be retrieved synchronously, it returns false, and any time it returns false, a caller must subsequently invoke `WaitForNextAsync` to either wait for the next item to be available or to determine that there will never be another item. Typical consumption (without additional language features) would look like:

C#

```
IAsyncEnumerator<T> enumerator = enumerable.GetAsyncEnumerator();
try
{
    while (await enumerator.WaitForNextAsync())
    {
        while (true)
        {
            int item = enumerator.TryGetNext(out bool success);
            if (!success) break;
            Use(item);
        }
    }
}
finally { await enumerator.DisposeAsync(); }
```

The advantage of this is two-fold, one minor and one major:

- *Minor: Allows for an enumerator to support multiple consumers*. There may be scenarios where it's valuable for an enumerator to support multiple concurrent consumers. That can't be achieved when `MoveNextAsync` and `Current` are separate such that an implementation can't make their usage atomic. In contrast, this approach provides a single method `TryGetNext` that supports pushing the

enumerator forward and getting the next item, so the enumerator can enable atomicity if desired. However, it's likely that such scenarios could also be enabled by giving each consumer its own enumerator from a shared enumerable. Further, we don't want to enforce that every enumerator support concurrent usage, as that would add non-trivial overheads to the majority case that doesn't require it, which means a consumer of the interface generally couldn't rely on this any way.
- *Major: Performance*. The `MoveNextAsync`/`Current` approach requires two interface calls per operation, whereas the best case for `WaitForNextAsync`/`TryGetNext` is that most iterations complete synchronously, enabling a tight inner loop with `TryGetNext`, such that we only have one interface call per operation. This can have a measurable impact in situations where the interface calls dominate the computation.

However, there are non-trivial downsides, including significantly increased complexity when consuming these manually, and an increased chance of introducing bugs when using them. And while the performance benefits show up in microbenchmarks, we don't believe they'll be impactful in the vast majority of real usage. If it turns out they are, we can introduce a second set of interfaces in a light-up fashion.

Discarded options considered:

- `ValueTask<bool> WaitForNextAsync(); bool TryGetNext(out T result);`: `out` parameters can't be covariant. There's also a small impact here (an issue with the try pattern in general) that this likely incurs a runtime write barrier for reference type results.

Cancellation

There are several possible approaches to supporting cancellation:

1. `IAsyncEnumerable<T>`/`IAsyncEnumerator<T>` are cancellation-agnostic: `CancellationToken` doesn't appear anywhere. Cancellation is achieved by logically baking the `CancellationToken` into the enumerable and/or enumerator in whatever manner is appropriate, e.g. when calling an iterator, passing the `CancellationToken` as an argument to the iterator method and using it in the body of the iterator, as is done with any other parameter.
2. `IAsyncEnumerator<T>.GetAsyncEnumerator(CancellationToken)`: You pass a `CancellationToken` to `GetAsyncEnumerator`, and subsequent `MoveNextAsync` operations respect it however it can.
3. `IAsyncEnumerator<T>.MoveNextAsync(CancellationToken)`: You pass a `CancellationToken` to each individual `MoveNextAsync` call.

4. 1 && 2: You both embed `CancellationToken`s into your enumerable/enumerator and pass `CancellationToken`s into `GetAsyncEnumerator`.
5. 1 && 3: You both embed `CancellationToken`s into your enumerable/enumerator and pass `CancellationToken`s into `MoveNextAsync`.

From a purely theoretical perspective, (5) is the most robust, in that (a) `MoveNextAsync` accepting a `CancellationToken` enables the most fine-grained control over what's canceled, and (b) `CancellationToken` is just any other type that can passed as an argument into iterators, embedded in arbitrary types, etc.

However, there are multiple problems with that approach:

- How does a `CancellationToken` passed to `GetAsyncEnumerator` make it into the body of the iterator? We could expose a new `iterator` keyword that you could dot off of to get access to the `CancellationToken` passed to `GetEnumerator`, but a) that's a lot of additional machinery, b) we're making it a very first-class citizen, and c) the 99% case would seem to be the same code both calling an iterator and calling `GetAsyncEnumerator` on it, in which case it can just pass the `CancellationToken` as an argument into the method.
- How does a `CancellationToken` passed to `MoveNextAsync` get into the body of the method? This is even worse, as if it's exposed off of an `iterator` local object, its value could change across awaits, which means any code that registered with the token would need to unregister from it prior to awaits and then re-register after; it's also potentially quite expensive to need to do such registering and unregistering in every `MoveNextAsync` call, regardless of whether implemented by the compiler in an iterator or by a developer manually.
- How does a developer cancel a `foreach` loop? If it's done by giving a `CancellationToken` to an enumerable/enumerator, then either a) we need to support `foreach`'ing over enumerators, which raises them to being first-class citizens, and now you need to start thinking about an ecosystem built up around enumerators (e.g. LINQ methods) or b) we need to embed the `CancellationToken` in the enumerable anyway by having some `WithCancellation` extension method off of `IAsyncEnumerable<T>` that would store the provided token and then pass it into the wrapped enumerable's `GetAsyncEnumerator` when the `GetAsyncEnumerator` on the returned struct is invoked (ignoring that token). Or, you can just use the `CancellationToken` you have in the body of the foreach.
- If/when query comprehensions are supported, how would the `CancellationToken` supplied to `GetEnumerator` or `MoveNextAsync` be passed into each clause? The easiest way would simply be for the clause to capture it, at which point whatever token is passed to `GetAsyncEnumerator`/`MoveNextAsync` is ignored.

An earlier version of this document recommended (1), but we since switched to (4).

The two main problems with (1):

- producers of cancellable enumerables have to implement some boilerplate, and can only leverage the compiler's support for async-iterators to implement a `IAsyncEnumerator<T> GetAsyncEnumerator(CancellationToken)` method.
- it is likely that many producers would be tempted to just add a `CancellationToken` parameter to their async-enumerable signature instead, which will prevent consumers from passing the cancellation token they want when they are given an `IAsyncEnumerable` type.

There are two main consumption scenarios:

1. `await foreach (var i in GetData(token)) ...` where the consumer calls the async-iterator method,
2. `await foreach (var i in givenIAsyncEnumerable.WithCancellation(token)) ...` where the consumer deals with a given `IAsyncEnumerable` instance.

We find that a reasonable compromise to support both scenarios in a way that is convenient for both producers and consumers of async-streams is to use a specially annotated parameter in the async-iterator method. The `[EnumeratorCancellation]` attribute is used for this purpose. Placing this attribute on a parameter tells the compiler that if a token is passed to the `GetAsyncEnumerator` method, that token should be used instead of the value originally passed for the parameter.

Consider `IAsyncEnumerable<int> GetData([EnumeratorCancellation] CancellationToken token = default)`. The implementer of this method can simply use the parameter in the method body. The consumer can use either consumption patterns above:

1. if you use `GetData(token)`, then the token is saved into the async-enumerable and will be used in iteration,
2. if you use `givenIAsyncEnumerable.WithCancellation(token)`, then the token passed to `GetAsyncEnumerator` will supersede any token saved in the async-enumerable.

foreach

`foreach` will be augmented to support `IAsyncEnumerable<T>` in addition to its existing support for `IEnumerable<T>`. And it will support the equivalent of `IAsyncEnumerable<T>` as a pattern if the relevant members are exposed publicly, falling back to using the interface directly if not, in order to enable struct-based extensions that avoid allocating

as well as using alternative awaitables as the return type of `MoveNextAsync` and `DisposeAsync`.

Syntax

Using the syntax:

```
C#
```

```
foreach (var i in enumerable)
```

C# will continue to treat `enumerable` as a synchronous enumerable, such that even if it exposes the relevant APIs for async enumerables (exposing the pattern or implementing the interface), it will only consider the synchronous APIs.

To force `foreach` to instead only consider the asynchronous APIs, `await` is inserted as follows:

```
C#
```

```
await foreach (var i in enumerable)
```

No syntax would be provided that would support using either the async or the sync APIs; the developer must choose based on the syntax used.

Semantics

The compile-time processing of an `await foreach` statement first determines the ***collection type***, ***enumerator type*** and ***iteration type*** of the expression (very similar to https://github.com/dotnet/csharpstandard/blob/draft-v6/standard/statements.md#1295-the-foreach-statement). This determination proceeds as follows:

- If the type `X` of *expression* is `dynamic` or an array type, then an error is produced and no further steps are taken.
- Otherwise, determine whether the type `X` has an appropriate `GetAsyncEnumerator` method:
  - Perform member lookup on the type `X` with identifier `GetAsyncEnumerator` and no type arguments. If the member lookup does not produce a match, or it produces an ambiguity, or produces a match that is not a method group, check for an enumerable interface as described below.

    - Perform overload resolution using the resulting method group and an empty argument list. If overload resolution results in no applicable methods, results in an ambiguity, or results in a single best method but that method is either static or not public, check for an enumerable interface as described below.
    - If the return type `E` of the `GetAsyncEnumerator` method is not a class, struct or interface type, an error is produced and no further steps are taken.
    - Member lookup is performed on `E` with the identifier `Current` and no type arguments. If the member lookup produces no match, the result is an error, or the result is anything except a public instance property that permits reading, an error is produced and no further steps are taken.
    - Member lookup is performed on `E` with the identifier `MoveNextAsync` and no type arguments. If the member lookup produces no match, the result is an error, or the result is anything except a method group, an error is produced and no further steps are taken.
    - Overload resolution is performed on the method group with an empty argument list. If overload resolution results in no applicable methods, results in an ambiguity, or results in a single best method but that method is either static or not public, or its return type is not awaitable into `bool`, an error is produced and no further steps are taken.
    - The collection type is `X`, the enumerator type is `E`, and the iteration type is the type of the `Current` property.
- Otherwise, check for an enumerable interface:
    - If among all the types T_i for which there is an implicit conversion from `X` to $\texttt{IAsyncEnumerable<}_i\texttt{>}$, there is a unique type `T` such that `T` is not dynamic and for all the other T_i there is an implicit conversion from `IAsyncEnumerable<T>` to $\texttt{IAsyncEnumerable<}T_i\texttt{>}$, then the collection type is the interface `IAsyncEnumerable<T>`, the enumerator type is the interface `IAsyncEnumerator<T>`, and the iteration type is `T`.
    - Otherwise, if there is more than one such type `T`, then an error is produced and no further steps are taken.
- Otherwise, an error is produced and no further steps are taken.

The above steps, if successful, unambiguously produce a collection type `C`, enumerator type `E` and iteration type `T`.

```
C#
```

```
await foreach (V v in x) «embedded_statement»
```

is then expanded to:

C#

```
{
    E e = ((C)(x)).GetAsyncEnumerator();
    try {
        while (await e.MoveNextAsync()) {
            V v = (V)(T)e.Current;
            «embedded_statement»
        }
    }
    finally {
        ... // Dispose e
    }
}
```

The body of the `finally` block is constructed according to the following steps:

- If the type `E` has an appropriate `DisposeAsync` method:
  - Perform member lookup on the type `E` with identifier `DisposeAsync` and no type arguments. If the member lookup does not produce a match, or it produces an ambiguity, or produces a match that is not a method group, check for the disposal interface as described below.
  - Perform overload resolution using the resulting method group and an empty argument list. If overload resolution results in no applicable methods, results in an ambiguity, or results in a single best method but that method is either static or not public, check for the disposal interface as described below.
  - If the return type of the `DisposeAsync` method is not awaitable, an error is produced and no further steps are taken.
  - The `finally` clause is expanded to the semantic equivalent of:

    C#

    ```
    finally {
        await e.DisposeAsync();
    }
    ```

- Otherwise, if there is an implicit conversion from `E` to the `System.IAsyncDisposable` interface, then
  - If `E` is a non-nullable value type then the `finally` clause is expanded to the semantic equivalent of:

    C#

    ```
    finally {
        await ((System.IAsyncDisposable)e).DisposeAsync();
    }
    ```

- Otherwise the `finally` clause is expanded to the semantic equivalent of:

  ```C#
  finally {
      System.IAsyncDisposable d = e as System.IAsyncDisposable;
      if (d != null) await d.DisposeAsync();
  }
  ```

  except that if `E` is a value type, or a type parameter instantiated to a value type, then the conversion of `e` to `System.IAsyncDisposable` shall not cause boxing to occur.

- Otherwise, the `finally` clause is expanded to an empty block:

  ```C#
  finally {
  }
  ```

ConfigureAwait

This pattern-based compilation will allow `ConfigureAwait` to be used on all of the awaits, via a `ConfigureAwait` extension method:

```C#
await foreach (T item in enumerable.ConfigureAwait(false))
{
   ...
}
```

This will be based on types we'll add to .NET as well, likely to System.Threading.Tasks.Extensions.dll:

```C#
// Approximate implementation, omitting arg validation and the like
namespace System.Threading.Tasks
{
    public static class AsyncEnumerableExtensions
    {
        public static ConfiguredAsyncEnumerable<T> ConfigureAwait<T>(this
IAsyncEnumerable<T> enumerable, bool continueOnCapturedContext) =>
            new ConfiguredAsyncEnumerable<T>(enumerable,
continueOnCapturedContext);
```

```
        public struct ConfiguredAsyncEnumerable<T>
        {
            private readonly IAsyncEnumerable<T> _enumerable;
            private readonly bool _continueOnCapturedContext;

            internal ConfiguredAsyncEnumerable(IAsyncEnumerable<T>
enumerable, bool continueOnCapturedContext)
            {
                _enumerable = enumerable;
                _continueOnCapturedContext = continueOnCapturedContext;
            }

            public ConfiguredAsyncEnumerator<T> GetAsyncEnumerator() =>
                new ConfiguredAsyncEnumerator<T>
(_enumerable.GetAsyncEnumerator(), _continueOnCapturedContext);

            public struct ConfiguredAsyncEnumerator<T>
            {
                private readonly IAsyncEnumerator<T> _enumerator;
                private readonly bool _continueOnCapturedContext;

                internal Enumerator(IAsyncEnumerator<T> enumerator, bool
continueOnCapturedContext)
                {
                    _enumerator = enumerator;
                    _continueOnCapturedContext = continueOnCapturedContext;
                }

                public ConfiguredValueTaskAwaitable<bool> MoveNextAsync() =>

_enumerator.MoveNextAsync().ConfigureAwait(_continueOnCapturedContext);

                public T Current => _enumerator.Current;

                public ConfiguredValueTaskAwaitable DisposeAsync() =>

_enumerator.DisposeAsync().ConfigureAwait(_continueOnCapturedContext);
            }
        }
    }
}
```

Note that this approach will not enable `ConfigureAwait` to be used with pattern-based enumerables, but then again it's already the case that the `ConfigureAwait` is only exposed as an extension on `Task`/`Task<T>`/`ValueTask`/`ValueTask<T>` and can't be applied to arbitrary awaitable things, as it only makes sense when applied to Tasks (it controls a behavior implemented in Task's continuation support), and thus doesn't make sense when using a pattern where the awaitable things may not be tasks. Anyone returning awaitable things can provide their own custom behavior in such advanced scenarios.

(If we can come up with some way to support a scope- or assembly-level `ConfigureAwait` solution, then this won't be necessary.)

Async Iterators

The language / compiler will support producing `IAsyncEnumerable<T>`s and `IAsyncEnumerator<T>`s in addition to consuming them. Today the language supports writing an iterator like:

C#

```
static IEnumerable<int> MyIterator()
{
    try
    {
        for (int i = 0; i < 100; i++)
        {
            Thread.Sleep(1000);
            yield return i;
        }
    }
    finally
    {
        Thread.Sleep(200);
        Console.WriteLine("finally");
    }
}
```

but `await` can't be used in the body of these iterators. We will add that support.

Syntax

The existing language support for iterators infers the iterator nature of the method based on whether it contains any `yield`s. The same will be true for async iterators. Such async iterators will be demarcated and differentiated from synchronous iterators via adding `async` to the signature, and must then also have either `IAsyncEnumerable<T>` or `IAsyncEnumerator<T>` as its return type. For example, the above example could be written as an async iterator as follows:

C#

```
static async IAsyncEnumerable<int> MyIterator()
{
    try
    {
        for (int i = 0; i < 100; i++)
```

```
        {
            await Task.Delay(1000);
            yield return i;
        }
    }
    finally
    {
        await Task.Delay(200);
        Console.WriteLine("finally");
    }
}
```

Alternatives considered:

- *Not using `async` in the signature*: Using `async` is likely technically required by the compiler, as it uses it to determine whether `await` is valid in that context. But even if it's not required, we've established that `await` may only be used in methods marked as `async`, and it seems important to keep the consistency.
- *Enabling custom builders for `IAsyncEnumerable<T>`*: That's something we could look at for the future, but the machinery is complicated and we don't support that for the synchronous counterparts.
- *Having an `iterator` keyword in the signature*: Async iterators would use `async iterator` in the signature, and `yield` could only be used in `async` methods that included `iterator`; `iterator` would then be made optional on synchronous iterators. Depending on your perspective, this has the benefit of making it very clear by the signature of the method whether `yield` is allowed and whether the method is actually meant to return instances of type `IAsyncEnumerable<T>` rather than the compiler manufacturing one based on whether the code uses `yield` or not. But it is different from synchronous iterators, which don't and can't be made to require one. Plus some developers don't like the extra syntax. If we were designing it from scratch, we'd probably make this required, but at this point there's much more value in keeping async iterators close to sync iterators.

LINQ

There are over ~200 overloads of methods on the `System.Linq.Enumerable` class, all of which work in terms of `IEnumerable<T>`; some of these accept `IEnumerable<T>`, some of them produce `IEnumerable<T>`, and many do both. Adding LINQ support for `IAsyncEnumerable<T>` would likely entail duplicating all of these overloads for it, for another ~200. And since `IAsyncEnumerator<T>` is likely to be more common as a standalone entity in the asynchronous world than `IEnumerator<T>` is in the synchronous world, we could potentially need another ~200 overloads that work with

`IAsyncEnumerator<T>`. Plus, a large number of the overloads deal with predicates (e.g. `Where` that takes a `Func<T, bool>`), and it may be desirable to have `IAsyncEnumerable<T>`-based overloads that deal with both synchronous and asynchronous predicates (e.g. `Func<T, ValueTask<bool>>` in addition to `Func<T, bool>`). While this isn't applicable to all of the now ~400 new overloads, a rough calculation is that it'd be applicable to half, which means another ~200 overloads, for a total of ~600 new methods.

That is a staggering number of APIs, with the potential for even more when extension libraries like Interactive Extensions (Ix) are considered. But Ix already has an implementation of many of these, and there doesn't seem to be a great reason to duplicate that work; we should instead help the community improve Ix and recommend it for when developers want to use LINQ with `IAsyncEnumerable<T>`.

There is also the issue of query comprehension syntax. The pattern-based nature of query comprehensions would allow them to "just work" with some operators, e.g. if Ix provides the following methods:

C#

```
public static IAsyncEnumerable<TResult> Select<TSource, TResult>(this
IAsyncEnumerable<TSource> source, Func<TSource, TResult> func);
public static IAsyncEnumerable<T> Where(this IAsyncEnumerable<T> source,
Func<T, bool> func);
```

then this C# code will "just work":

C#

```
IAsyncEnumerable<int> enumerable = ...;
IAsyncEnumerable<int> result = from item in enumerable
                               where item % 2 == 0
                               select item * 2;
```

However, there is no query comprehension syntax that supports using `await` in the clauses, so if Ix added, for example:

C#

```
public static IAsyncEnumerable<TResult> Select<TSource, TResult>(this
IAsyncEnumerable<TSource> source, Func<TSource, ValueTask<TResult>> func);
```

then this would "just work":

```C#
IAsyncEnumerable<string> result = from url in urls
                                  where item % 2 == 0
                                  select SomeAsyncMethod(item);

async ValueTask<int> SomeAsyncMethod(int item)
{
    await Task.Yield();
    return item * 2;
}
```

but there'd be no way to write it with the `await` inline in the `select` clause. As a separate effort, we could look into adding `async { ... }` expressions to the language, at which point we could allow them to be used in query comprehensions and the above could instead be written as:

```C#
IAsyncEnumerable<int> result = from item in enumerable
                               where item % 2 == 0
                               select async
                               {
                                   await Task.Yield();
                                   return item * 2;
                               };
```

or to enabling `await` to be used directly in expressions, such as by supporting `async from`. However, it's unlikely a design here would impact the rest of the feature set one way or the other, and this isn't a particularly high-value thing to invest in right now, so the proposal is to do nothing additional here right now.

Integration with other asynchronous frameworks

Integration with `IObservable<T>` and other asynchronous frameworks (e.g. reactive streams) would be done at the library level rather than at the language level. For example, all of the data from an `IAsyncEnumerator<T>` can be published to an `IObserver<T>` simply by `await foreach`'ing over the enumerator and `OnNext`'ing the data to the observer, so an `AsObservable<T>` extension method is possible. Consuming an `IObservable<T>` in a `await foreach` requires buffering the data (in case another item is pushed while the previous item is still being processing), but such a push-pull adapter can easily be implemented to enable an `IObservable<T>` to be pulled from with an

`IAsyncEnumerator<T>`. Etc. Rx/Ix already provide prototypes of such implementations, and libraries like https://github.com/dotnet/corefx/tree/master/src/System.Threading.Channels provide various kinds of buffering data structures. The language need not be involved at this stage.

Permit `stackalloc` in nested contexts

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Stack allocation

We modify the section *Stack allocation* (§22.9) of the C# language specification to relax the places when a `stackalloc` expression may appear. We delete

```antlr
local_variable_initializer_unsafe
    : stackalloc_initializer
    ;

stackalloc_initializer
    : 'stackalloc' unmanaged_type '[' expression ']'
    ;
```

and replace them with

```antlr
primary_no_array_creation_expression
    : stackalloc_initializer
    ;

stackalloc_initializer
    : 'stackalloc' unmanaged_type '[' expression? ']' array_initializer?
    | 'stackalloc' '[' expression? ']' array_initializer
    ;
```

Note that the addition of an *array_initializer* to *stackalloc_initializer* (and making the index expression optional) was an extension in C# 7.3 and is not described here.

The *element type* of the `stackalloc` expression is the *unmanaged_type* named in the stackalloc expression, if any, or the common type among the elements of the *array_initializer* otherwise.

The type of the *stackalloc_initializer* with *element type* `K` depends on its syntactic context:

- If the *stackalloc_initializer* appears directly as the *local_variable_initializer* of a *local_variable_declaration* statement or a *for_initializer*, then its type is `K*`.
- Otherwise its type is `System.Span<K>`.

Stackalloc Conversion

The *stackalloc conversion* is a new built-in implicit conversion from expression. When the type of a *stackalloc_initializer* is `K*`, there is an implicit *stackalloc conversion* from the *stackalloc_initializer* to the type `System.Span<K>`.

`Lock` object

Article • 05/16/2024

> **Note**
>
> This article is a feature specification. The specification serves as the design document for the feature. It includes proposed specification changes, along with information needed during the design and development of the feature. These articles are published until the proposed spec changes are finalized and incorporated in the current ECMA specification.
>
> There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent language design meeting (LDM) notes.
>
> You can learn more about the process for adopting feature speclets into the C# language standard in the article on the specifications.

Summary

Special-case how `System.Threading.Lock` interacts with the `lock` keyword (calling its `EnterScope` method under the hood). Add static analysis warnings to prevent accidental misuse of the type where possible.

Motivation

.NET 9 is introducing a new System.Threading.Lock type as a better alternative to existing monitor-based locking. The presence of the `lock` keyword in C# might lead developers to think they can use it with this new type. Doing so wouldn't lock according to the semantics of this type but would instead treat it as any other object and would use monitor-based locking.

cs

```cs
namespace System.Threading
{
    public sealed class Lock
    {
        public void Enter();
        public void Exit();
        public Scope EnterScope();
```

```
        public ref struct Scope
        {
            public void Dispose();
        }
    }
}
```

Detailed design

Semantics of the lock statement (§13.13) are changed to special-case the `System.Threading.Lock` type:

> A `lock` statement of the form `lock (x) { ... }`
>
> 1. **where `x` is an expression of type `System.Threading.Lock`, is precisely equivalent to:**
>
>    ```cs
>    using (x.EnterScope())
>    {
>        ...
>    }
>    ```
>
>   **and `System.Threading.Lock` must have the following shape:**
>
>    ```cs
>    namespace System.Threading
>    {
>        public sealed class Lock
>        {
>            public Scope EnterScope();
>
>            public ref struct Scope
>            {
>                public void Dispose();
>            }
>        }
>    }
>    ```
>
> 2. where `x` is an expression of a *reference_type*, is precisely equivalent to: [...]

Note that the shape might not be fully checked (e.g., there will be no errors nor warnings if the `Lock` type is not `sealed`), but the feature might not work as expected

(e.g., there will be no warnings when converting `Lock` to a derived type, since the feature assumes there are no derived types).

Additionally, new warnings are added to implicit reference conversions (§10.2.8) when upcasting the `System.Threading.Lock` type:

> The implicit reference conversions are:
>
> - From any *reference_type* to `object` and `dynamic`.
>   - **A warning is reported when the *reference_type* is known to be `System.Threading.Lock`.**
> - From any *class_type* `S` to any *class_type* `T`, provided `S` is derived from `T`.
>   - **A warning is reported when `S` is known to be `System.Threading.Lock`.**
> - From any *class_type* `S` to any *interface_type* `T`, provided `S` implements `T`.
>   - **A warning is reported when `S` is known to be `System.Threading.Lock`.**
> - [...]

cs

```cs
object l = new System.Threading.Lock(); // warning
lock (l) { } // monitor-based locking is used here
```

Note that this warning occurs even for equivalent explicit conversions.

The compiler avoids reporting the warning in some cases when the instance cannot be locked after converting to `object`:

- when the conversion is implicit and part of an object equality operator invocation.

cs

```cs
var l = new System.Threading.Lock();
if (l != null) // no warning even though `l` is implicitly converted to
`object` for `operator!=(object, object)`
    // ...
```

To escape out of the warning and force use of monitor-based locking, one can use

- the usual warning suppression means (`#pragma warning disable`),
- `Monitor` APIs directly,
- indirect casting like `object AsObject<T>(T l) => (object)l;`.

Alternatives

- Support a general pattern that other types can also use to interact with the `lock` keyword. This is a future work that might be implemented when `ref struct`s can participate in generics. Discussed in LDM 2023-12-04.
- To avoid ambiguity between the existing monitor-based locking and the new `Lock` (or pattern in the future), we could:
  - Introduce a new syntax instead of reusing the existing `lock` statement.
  - Require the new lock types to be `struct`s (since the existing `lock` disallows value types). There could be problems with default constructors and copying if the structs have lazy initialization.
- The codegen could be hardened against thread aborts (which are themselves obsoleted).
- We could warn also when `Lock` is passed as a type parameter, because locking on a type parameter always uses monitor-based locking:

  ```cs
  M(new Lock()); // could warn here

  void M<T>(T x) // (specifying `where T : Lock` makes no difference)
  {
      lock (x) { } // because this uses Monitor
  }
  ```

  However, that would cause warnings when storing `Lock`s in a list which is undesirable:

  ```cs
  List<Lock> list = new();
  list.Add(new Lock()); // would warn here
  ```

- We could include static analysis to prevent usage of `System.Threading.Lock` in `using`s with `await`s. I.e., we could emit either an error or a warning for code like `using (lockVar.EnterScope()) { await ... }`. Currently, this is not needed since `Lock.Scope` is a `ref struct`, so that code is illegal anyway. However, if we ever allowed `ref struct`s in `async` methods or changed `Lock.Scope` to not be a `ref struct`, this analysis would become beneficial. (We would also likely need to consider for this all lock types matching the general pattern if implemented in the

future. Although there might need to be an opt-out mechanism as some lock types might be allowed to be used with `await`.) Alternatively, this could be implemented as an analyzer shipped as part of the runtime.

- We could relax the restriction that value types cannot be `lock`ed
  - for the new `Lock` type (only needed if the API proposal changed it from `class` to `struct`),
  - for the general pattern where any type can participate when implemented in the future.

- We could allow the new `lock` in `async` methods where `await` is not used inside the `lock`.
  - Currently, since `lock` is lowered to `using` with a `ref struct` as the resource, this results in a compile-time error. The workaround is to extract the `lock` into a separate non-`async` method.
  - Instead of using the `ref struct Scope`, we could emit `Lock.Enter` and `Lock.Exit` methods in `try`/`finally`. However, the `Exit` method must throw when it's called from a different thread than `Enter`, hence it contains a thread lookup which is avoided when using the `Scope`.
  - Best would be to allow compiling `using` on a `ref struct` in `async` methods if there is no `await` inside the `using` body.

Design meetings

- LDM 2023-05-01: initial decision to support a `lock` pattern
- LDM 2023-10-16: triaged into the working set for .NET 9
- LDM 2023-12-04: rejected the general pattern, accepted only special-casing the `Lock` type + adding static analysis warnings

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# feature specifications feedback

C# feature specifications is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

File Scoped Namespaces

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

File scoped namespaces use a less verbose format for the typical case of files containing only one namespace. The file scoped namespace format is `namespace X.Y.Z;` (note the semicolon and lack of braces). This allows for files like the following:

```
c#

namespace X.Y.Z;

using System;

class X
{
}
```

The semantics are that using the `namespace X.Y.Z;` form is equivalent to writing `namespace X.Y.Z { ... }` where the remainder of the file following the file-scoped namespace is in the `...` section of a standard namespace declaration.

Motivation

Analysis of the C# ecosystem shows that approximately 99.7% files are all of either one of these forms:

```
c#
```

```
namespace X.Y.Z
{
    // usings

    // types
}
```

or

```
c#
```

```
// usings

namespace X.Y.Z
{
    // types
}
```

However, both these forms force the user to indent the majority of their code and add a fair amount of ceremony for what is effectively a very basic concept. This affects clarity, uses horizontal and vertical space, and is often unsatisfying for users both used to C# and coming from other languages (which commonly have less ceremony here).

The primary goal of the feature therefore is to meet the needs of the majority of the ecosystem with less unnecessary boilerplate.

Detailed design

This proposal takes the form of a diff to the existing compilation units (§13.2) section of the specification.

Diff

~~A *compilation_unit* defines the overall structure of a source file. A compilation unit consists of zero or more *using_directives* followed by zero or more *global_attributes* followed by zero or more *namespace_member_declarations*.~~

A *compilation_unit* defines the overall structure of a source file. A compilation unit consists of zero or more *using_directive*s followed by zero or more *global_attributes* followed by a *compilation_unit_body*. A *compilation_unit_body* can either be a *file_scoped_namespace_declaration* or zero or more *statement*s and *namespace_member_declaration*s.

```
```

antlr

```
compilation_unit
~~    : extern_alias_directive* using_directive* global_attributes?
namespace_member_declaration*~~
    : extern_alias_directive* using_directive* global_attributes?
compilation_unit_body
    ;

compilation_unit_body
    : statement* namespace_member_declaration*
    | file_scoped_namespace_declaration
    ;
```

... unchanged ...

A *file_scoped_namespace_declaration* will contribute members corresponding to the *namespace_declaration* it is semantically equivalent to. See (Namespace Declarations) for more details.

Namespace declarations

A *namespace_declaration* consists of the keyword `namespace`, followed by a namespace name and body, optionally followed by a semicolon. A *file_scoped_namespace_declaration* consists of the keyword `namespace`, followed by a namespace name, a semicolon and an optional list of *extern_alias_directive*s, *using_directive*s and *type_declaration*s.

antlr

```
namespace_declaration
    : 'namespace' qualified_identifier namespace_body ';'?
    ;

file_scoped_namespace_declaration
    : 'namespace' qualified_identifier ';' extern_alias_directive*
using_directive* type_declaration*
    ;

... unchanged ...
```

... unchanged ...

the two namespace declarations above contribute to the same declaration space, in this case declaring two classes with the fully qualified names `N1.N2.A` and `N1.N2.B`. Because

the two declarations contribute to the same declaration space, it would have been an error if each contained a declaration of a member with the same name.

A *file_scoped_namespace_declaration* permits a namespace declaration to be written without the `{ ... }` block. For example:

```
C#
```

```
extern alias A;
namespace Name;
using B;
class C
{
}
```

is semantically equivalent to

```
C#
```

```
extern alias A;
namespace Name
{
    using B;
    class C
    {
    }
}
```

Specifically, a *file_scoped_namespace_declaration* is treated the same as a *namespace_declaration* at the same location in the *compilation_unit* with the same *qualified_identifier*. The *extern_alias_directive*s, *using_directive*s and *type_declaration*s of that *file_scoped_namespace_declaration* act as if they were declared in the same order inside the *namespace_body* of that *namespace_declaration*.

A source file cannot contain both a *file_scoped_namespace_declaration* and a *namespace_declaration*. A source file cannot contain multiple *file_scoped_namespace_declaration*s. A *compilation_unit* cannot contain both a *file_scoped_namespace_declaration* and any top level *statement*s. *type_declaration*s cannot precede a *file_scoped_namespace_declaration*.

Extern aliases

... unchanged ...

Extending Partial Methods

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

This proposal aims to remove all restrictions around the signatures of `partial` methods in C#. The goal being to expand the set of scenarios in which these methods can work with source generators as well as being a more general declaration form for C# methods.

See also the original partial methods specification (§14.6.9).

Motivation

C# has limited support for developers splitting methods into declarations and definitions / implementations.

C#

```
partial class C
{
    // The declaration of C.M
    partial void M(string message);
}

partial class C
{
    // The definition of C.M
    partial void M(string message) => Console.WriteLine(message);
}
```

One behavior of `partial` methods is that when the definition is absent then the language will simply erase any calls to the `partial` method. Essentially it behaves like a call to a `[Conditional]` method where the condition was evaluated to false.

C#

```
partial class D
{
    partial void M(string message);

    void Example()
    {
        M(GetIt()); // Call to M and GetIt erased at compile time
    }

    string GetIt() => "Hello World";
}
```

The original motivation for this feature was source generation in the form of designer generated code. Users were constantly editing the generated code because they wanted to hook some aspect of the generated code. Most notably parts of the Windows Forms startup process, after components were initialized.

Editing the generated code was error prone because any action which caused the designer to regenerate the code would cause the user edit to be erased. The `partial` method feature eased this tension because it allowed designers to emit hooks in the form of `partial` methods.

Designers could emit hooks like `partial void OnComponentInit()` and developers could define declarations for them or not define them. In either case though the generated code would compile and developers who were interested in the process could hook in as needed.

This does mean that partial methods have several restrictions:

1. Must have a `void` return type.
2. Cannot have `out` parameters.
3. Cannot have any accessibility (implicitly `private`).

These restrictions exist because the language must be able to emit code when the call site is erased. Given they can be erased `private` is the only possible accessibility because the member can't be exposed in assembly metadata. These restrictions also serve to limit the set of scenarios in which `partial` methods can be applied.

The proposal here is to remove all of the existing restrictions around `partial` methods. Essentially let them have `out`, non-void return types or any type of accessibility. Such `partial` declarations would then have the added requirement that a definition must exist. That means the language does not have to consider the impact of erasing the call sites.

This would expand the set of generator scenarios that `partial` methods could participate in and hence link in nicely with our source generators feature. For example a regex could be defined using the following pattern:

C#

```
[RegexGenerated("(dog|cat|fish)")]
partial bool IsPetMatch(string input);
```

This gives both the developer a simple declarative way of opting into generators as well as giving generators a very easy set of declarations to look through in the source code to drive their generated output.

Compare that with the difficulty that a generator would have hooking up the following snippet of code.

C#

```
var regex = new RegularExpression("(dog|cat|fish)");
if (regex.IsMatch(someInput))
{

}
```

Given that the compiler doesn't allow generators to modify code hooking up this pattern would be pretty much impossible for generators. They would need to resort to reflection in the `IsMatch` implementation, or asking users to change their call sites to a new method + refactor the regex to pass the string literal as an argument. It's pretty messy.

Detailed Design

The language will change to allow `partial` methods to be annotated with an explicit accessibility modifier. This means they can be labeled as `private`, `public`, etc ...

When a `partial` method has an explicit accessibility modifier though the language will require that the declaration has a matching definition even when the accessibility is

`private`:

```C#
partial class C
{
    // Okay because no definition is required here
    partial void M1();

    // Okay because M2 has a definition
    private partial void M2();

    // Error: partial method M3 must have a definition
    private partial void M3();
}

partial class C
{
    private partial void M2() { }
}
```

Further the language will remove all restrictions on what can appear on a `partial` method which has an explicit accessibility. Such declarations can contain non-void return types, `out` parameters, `extern` modifier, etc ... These signatures will have the full expressivity of the C# language.

```C#
partial class D
{
    // Okay
    internal partial bool TryParse(string s, out int i);
}

partial class D
{
    internal partial bool TryParse(string s, out int i) { }
}
```

This explicitly allows for `partial` methods to participate in `overrides` and `interface` implementations:

```C#
interface IStudent
{
    string GetName();
}
```

```
partial class C : IStudent
{
    public virtual partial string GetName();
}

partial class C
{
    public virtual partial string GetName() => "Jarde";
}
```

The compiler will change the error it emits when a `partial` method contains an illegal element to essentially say:

> Cannot use `ref` on a `partial` method that lacks explicit accessibility

This will help point developers in the right direction when using this feature.

Restrictions:

- `partial` declarations with explicit accessibility must have a definition
- `partial` declarations and definition signatures must match on all method and parameter modifiers. The only aspects which can differ are parameter names and attribute lists (this is not new but rather an existing requirement of `partial` methods).

Questions

partial on all members

Given that we're expanding `partial` to be more friendly to source generators should we also expand it to work on all class members? For example should we be able to declare `partial` constructors, operators, etc ...

Resolution The idea is sound but at this point in the C# 9 schedule we're trying to avoid unnecessary feature creep. Want to solve the immediate problem of expanding the feature to work with modern source generators.

Extending `partial` to support other members will be considered for the C# 10 release. Seems likely that we will consider this extension.

Use abstract instead of partial

The crux of this proposal is essentially ensuring that a declaration has a corresponding definition / implementation. Given that should we use `abstract` since it's already a language keyword that forces the developer to think about having an implementation?

Resolution There was a healthy discussion about this but eventually it was decided against. Yes the requirements are familiar but the concepts are significantly different. Could easily lead the developer to believe they were creating virtual slots when they were not doing so.

Covariant returns

Article • 02/23/2022

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Support *covariant return types*. Specifically, permit the override of a method to declare a more derived return type than the method it overrides, and similarly to permit the override of a read-only property to declare a more derived type. Override declarations appearing in more derived types would be required to provide a return type at least as specific as that appearing in overrides in its base types. Callers of the method or property would statically receive the more refined return type from an invocation.

Motivation

It is a common pattern in code that different method names have to be invented to work around the language constraint that overrides must return the same type as the overridden method.

This would be useful in the factory pattern. For example, in the Roslyn code base we would have

C#

```
class Compilation ...
{
    public virtual Compilation WithOptions(Options options)...
}
```

C#

```
class CSharpCompilation : Compilation
{
    public override CSharpCompilation WithOptions(Options options)...
}
```

Detailed design

This is a specification for covariant return types in C#. Our intent is to permit the override of a method to return a more derived return type than the method it overrides, and similarly to permit the override of a read-only property to return a more derived return type. Callers of the method or property would statically receive the more refined return type from an invocation, and overrides appearing in more derived types would be required to provide a return type at least as specific as that appearing in overrides in its base types.

Class Method Override

The existing constraint on class override (§14.6.5) methods

> - The override method and the overridden base method have the same return type.

is modified to

> - The override method must have a return type that is convertible by an identity conversion or (if the method has a value return - not a ref return) implicit reference conversion to the return type of the overridden base method.

And the following additional requirements are appended to that list:

> - The override method must have a return type that is convertible by an identity conversion or (if the method has a value return - not a ref return) implicit reference conversion to the return type of every override of the overridden base method that is declared in a (direct or indirect) base type of the override method.
> - The override method's return type must be at least as accessible as the override method (Accessibility domains - §7.5.3).

This constraint permits an override method in a `private` class to have a `private` return type. However it requires a `public` override method in a `public` type to have a `public` return type.

Class Property and Indexer Override

The existing constraint on class override (§14.7.6) properties

> An overriding property declaration shall specify the exact same accessibility modifiers and name as the inherited property, and there shall be an identity conversion ~~between the type of the overriding and the inherited property~~. If the inherited property has only a single accessor (i.e., if the inherited property is read-only or write-only), the overriding property shall include only that accessor. If the inherited property includes both accessors (i.e., if the inherited property is read-write), the overriding property can include either a single accessor or both accessors.

is modified to

> An overriding property declaration shall specify the exact same accessibility modifiers and name as the inherited property, and there shall be an identity conversion **or (if the inherited property is read-only and has a value return - not a ref return) implicit reference conversion from the type of the overriding property to the type of the inherited property**. If the inherited property has only a single accessor (i.e., if the inherited property is read-only or write-only), the overriding property shall include only that accessor. If the inherited property includes both accessors (i.e., if the inherited property is read-write), the overriding property can include either a single accessor or both accessors. **The overriding property's type must be at least as accessible as the overriding property (Accessibility domains - §7.5.3).**

The remainder of the draft specification below proposes a further extension to covariant returns of interface methods to be considered later.

Interface Method, Property, and Indexer Override

Adding to the kinds of members that are permitted in an interface with the addition of the DIM feature in C# 8.0, we further add support for `override` members along with covariant returns. These follow the rules of `override` members as specified for classes, with the following differences:

The following text in classes:

> The method overridden by an override declaration is known as the ***overridden base method***. For an override method `M` declared in a class `C`, the overridden base method is determined by examining each base class of `C`, starting with the direct base class of `C` and continuing with each successive direct base class, until in a given base class type at least one accessible method is located which has the same signature as `M` after substitution of type arguments.

is given the corresponding specification for interfaces:

> The method overridden by an override declaration is known as the ***overridden base method***. For an override method `M` declared in an interface `I`, the overridden base method is determined by examining each direct or indirect base interface of `I`, collecting the set of interfaces declaring an accessible method which has the same signature as `M` after substitution of type arguments. If this set of interfaces has a *most derived type*, to which there is an identity or implicit reference conversion from every type in this set, and that type contains a unique such method declaration, then that is the *overridden base method*.

We similarly permit `override` properties and indexers in interfaces as specified for classes in *15.7.6 Virtual, sealed, override, and abstract accessors*.

Name Lookup

Name lookup in the presence of class `override` declarations currently modify the result of name lookup by imposing on the found member details from the most derived `override` declaration in the class hierarchy starting from the type of the identifier's qualifier (or `this` when there is no qualifier). For example, in *12.6.2.2 Corresponding parameters* we have

> For virtual methods and indexers defined in classes, the parameter list is picked from the first declaration or override of the function member found when starting with the static type of the receiver, and searching through its base classes.

to this we add

> For virtual methods and indexers defined in interfaces, the parameter list is picked from the declaration or override of the function member found in the most derived

> type among those types containing the declaration of override of the function member. It is a compile-time error if no unique such type exists.

For the result type of a property or indexer access, the existing text

> - If I identifies an instance property, then the result is a property access with an associated instance expression of E and an associated type that is the type of the property. If T is a class type, the associated type is picked from the first declaration or override of the property found when starting with T, and searching through its base classes.

is augmented with

> If T is an interface type, the associated type is picked from the declaration or override of the property found in the most derived of T or its direct or indirect base interfaces. It is a compile-time error if no unique such type exists.

A similar change should be made in *12.7.7.3 Indexer access*

In *12.7.6 Invocation expressions* we augment the existing text

> - Otherwise, the result is a value, with an associated type of the return type of the method or delegate. If the invocation is of an instance method, and the receiver is of a class type T, the associated type is picked from the first declaration or override of the method found when starting with T and searching through its base classes.

with

> If the invocation is of an instance method, and the receiver is of an interface type T, the associated type is picked from the declaration or override of the method found in the most derived interface from among T and its direct and indirect base interfaces. It is a compile-time error if no unique such type exists.

Implicit Interface Implementations

This section of the specification

> For purposes of interface mapping, a class member `A` matches an interface member `B` when:

> - `A` and `B` are methods, and the name, type, and formal parameter lists of `A` and `B` are identical.
> - `A` and `B` are properties, the name and type of `A` and `B` are identical, and `A` has the same accessors as `B` (`A` is permitted to have additional accessors if it is not an explicit interface member implementation).
> - `A` and `B` are events, and the name and type of `A` and `B` are identical.
> - `A` and `B` are indexers, the type and formal parameter lists of `A` and `B` are identical, and `A` has the same accessors as `B` (`A` is permitted to have additional accessors if it is not an explicit interface member implementation).

is modified as follows:

> For purposes of interface mapping, a class member `A` matches an interface member `B` when:
>
> - `A` and `B` are methods, and the name and formal parameter lists of `A` and `B` are identical, and the return type of `A` is convertible to the return type of `B` via an identity of implicit reference convertion to the return type of `B`.
> - `A` and `B` are properties, the name of `A` and `B` are identical, `A` has the same accessors as `B` (`A` is permitted to have additional accessors if it is not an explicit interface member implementation), and the type of `A` is convertible to the return type of `B` via an identity conversion or, if `A` is a readonly property, an implicit reference conversion.
> - `A` and `B` are events, and the name and type of `A` and `B` are identical.
> - `A` and `B` are indexers, the formal parameter lists of `A` and `B` are identical, `A` has the same accessors as `B` (`A` is permitted to have additional accessors if it is not an explicit interface member implementation), and the type of `A` is convertible to the return type of `B` via an identity conversion or, if `A` is a readonly indexer, an implicit reference conversion.

This is technically a breaking change, as the program below prints "C1.M" today, but would print "C2.M" under the proposed revision.

```c#
using System;

interface I1 { object M(); }
class C1 : I1 { public object M() { return "C1.M"; } }
class C2 : C1, I1 { public new string M() { return "C2.M"; } }
class Program
```

```
{
    static void Main()
    {
        I1 i = new C2();
        Console.WriteLine(i.M());
    }
}
```

Due to this breaking change, we might consider not supporting covariant return types on implicit implementations.

Constraints on Interface Implementation

We will need a rule that an explicit interface implementation must declare a return type no less derived than the return type declared in any override in its base interfaces.

API Compatibility Implications

TBD

Open Issues

The specification does not say how the caller gets the more refined return type. Presumably that would be done in a way similar to the way that callers get the most derived override's parameter specifications.

If we have the following interfaces:

C#

```
interface I1 { I1 M(); }
interface I2 { I2 M(); }
interface I3: I1, I2 { override I3 M(); }
```

Note that in `I3`, the methods `I1.M()` and `I2.M()` have been “merged”. When implementing `I3`, it is necessary to implement them both together.

Generally, we require an explicit implementation to refer to the original method. The question is, in a class

C#

```
class C : I1, I2, I3
{
    C IN.M();
}
```

What does that mean here? What should *N* be?

I suggest that we permit implementing either `I1.M` or `I2.M` (but not both), and treat that as an implementation of both.

Drawbacks

- [ ] Every language change must pay for itself.
- [ ] We should ensure that the performance is reasonable, even in the case of deep inheritance hierarchies
- [ ] We should ensure that artifacts of the translation strategy do not affect language semantics, even when consuming new IL from old compilers.

Alternatives

We could relax the language rules slightly to allow, in source,

C#

```
abstract class Cloneable
{
    public abstract Cloneable Clone();
}

class Digit : Cloneable
{
    public override Cloneable Clone()
    {
        return this.Clone();
    }

    public new Digit Clone() // Error: 'Digit' already defines a member
called 'Clone' with the same parameter types
    {
        return this;
    }
}
```

Unresolved questions

- [ ] How will APIs that have been compiled to use this feature work in older versions of the language?

Design meetings

- some discussion at https://github.com/dotnet/roslyn/issues/357 .
- https://github.com/dotnet/csharplang/blob/master/meetings/2020/LDM-2020-01-08.md
- Offline discussion toward a decision to support overriding of class methods only in C# 9.0.

Primary constructors

Article • 08/16/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Classes and structs can have a parameter list, and their base class specification can have an argument list. Primary constructor parameters are in scope throughout the class or struct declaration, and if they are captured by a function member or anonymous function, they are appropriately stored (e.g. as unspeakable private fields of the declared class or struct).

The proposal "retcons" the primary constructors already available on records in terms of this more general feature with some additional members synthesized.

Motivation

The ability of a class or struct in C# to have more than one constructor provides for generality, but at the expense of some tedium in the declaration syntax, because the constructor input and the class state need to be cleanly separated.

Primary constructors put the parameters of one constructor in scope for the whole class or struct to be used for initialization or directly as object state. The trade-off is that any other constructors must call through the primary constructor.

c#

```
public class C(bool b, int i, string s) : B(b) // b passed to base
constructor
{
    public int I { get; set; } = i; // i used for initialization
```

```
    public string S // s used directly in function members
    {
        get => s;
        set => s = value ?? throw new ArgumentNullException(nameof(S));
    }
    public C(string s) : this(true, 0, s) { } // must call this(...)
}
```

Detailed design

This describes the generalized design across records and non-records, and then details how the existing primary constructors for records are specified by adding a set of synthesized members in the presence of a primary constructor.

Syntax

Class and struct declarations are augmented to allow a parameter list on the type name, an argument list on the base class, and a body consisting of just a `;`:

antlr

```
class_declaration
  : attributes? class_modifier* 'partial'? class_designator identifier
type_parameter_list?
  parameter_list? class_base? type_parameter_constraints_clause* class_body
  ;

class_designator
  : 'record' 'class'?
  | 'class'

class_base
  : ':' class_type argument_list?
  | ':' interface_type_list
  | ':' class_type  argument_list? ',' interface_type_list
  ;

class_body
  : '{' class_member_declaration* '}' ';'?
  | ';'
  ;

struct_declaration
  : attributes? struct_modifier* 'partial'? 'record'? 'struct' identifier
type_parameter_list?
    parameter_list? struct_interfaces? type_parameter_constraints_clause*
struct_body
  ;
```

```
struct_body
  : '{' struct_member_declaration* '}' ';'?
  | ';'
  ;

interface_declaration
  : attributes? interface_modifier* 'partial'? 'interface'
    identifier variant_type_parameter_list? interface_base?
    type_parameter_constraints_clause* interface_body
  ;

interface_body
  : '{' interface_member_declaration* '}' ';'?
  | ';'
  ;

enum_declaration
  : attributes? enum_modifier* 'enum' identifier enum_base? enum_body
  ;

enum_body
  : '{' enum_member_declarations? '}' ';'?
  | '{' enum_member_declarations ',' '}' ';'?
  | ';'
  ;
```

Note: These productions replace `record_declaration` in Records and `record_struct_declaration` in Record structs, which both become obsolete.

It is an error for a `class_base` to have an `argument_list` if the enclosing `class_declaration` does not contain a `parameter_list`. At most one partial type declaration of a partial class or struct may provide a `parameter_list`. The parameters in the `parameter_list` of a `record` declaration must all be value parameters.

Note, according to this proposal `class_body`, `struct_body`, `interface_body` and `enum_body` are allowed to consist of just a `;`.

A class or struct with a `parameter_list` has an implicit public constructor whose signature corresponds to the value parameters of the type declaration. This is called the ***primary constructor*** for the type, and causes the implicitly declared parameterless constructor, if present, to be suppressed. It is an error to have a primary constructor and a constructor with the same signature already present in the type declaration.

Lookup

The lookup of simple names is augmented to handle primary constructor parameters. The changes are highlighted in **bold** in the following excerpt:

- Otherwise, for each instance type `T` (§14.3.2), starting with the instance type of the immediately enclosing type declaration and continuing with the instance type of each enclosing class or struct declaration (if any):
  - **If the declaration of `T` includes a primary constructor parameter `I` and the reference occurs within the `argument_list` of `T`'s `class_base` or within an initializer of a field, property or event of `T`, the result is the primary constructor parameter `I`**
  - **Otherwise,** if `e` is zero and the declaration of `T` includes a type parameter with name `I`, then the *simple_name* refers to that type parameter.
  - Otherwise, if a member lookup (§11.5) of `I` in `T` with `e` type arguments produces a match:
    - If `T` is the instance type of the immediately enclosing class or struct type and the lookup identifies one or more methods, the result is a method group with an associated instance expression of `this`. If a type argument list was specified, it is used in calling a generic method (§11.7.8.2).
    - Otherwise, if `T` is the instance type of the immediately enclosing class or struct type, if the lookup identifies an instance member, and if the reference occurs within the *block* of an instance constructor, an instance method, or an instance accessor (§11.2.1), the result is the same as a member access (§11.7.6) of the form `this.I`. This can only happen when `e` is zero.
    - Otherwise, the result is the same as a member access (§11.7.6) of the form `T.I` or `T.I<A₁, ..., Aₑ>`.
  - **Otherwise, if the declaration of `T` includes a primary constructor parameter `I`, the result is the primary constructor parameter `I`.**

The first addition corresponds to the change incurred by primary constructors on records, and ensures that primary constructor parameters are found before any corresponding fields within initializers and base class arguments. It extends this rule to static initializers as well. However, since records always have an instance member with the same name as the parameter, the extension can only lead to a change in an error message. Illegal access to a parameter vs. illegal access to an instance member.

The second addition allows primary constructor parameters to be found elsewhere within the type body, but only if not shadowed by members.

It is an error to reference a primary constructor parameter if the reference does not occur within one of the following:

- a `nameof` argument

- an initializer of an instance field, property or event of the declaring type (type declaring primary constructor with the parameter).
- the `argument_list` of `class_base` of the declaring type.
- the body of an instance method (note that instance constructors are excluded) of the declaring type.
- the body of an instance accessor of the declaring type.

In other words, primary constructor parameters are in scope throughout the declaring type body. They shadow members of the declaring type within an initializer of a field, property or event of the declaring type, or within the `argument_list` of `class_base` of the declaring type. They are shadowed by members of the declaring type everywhere else.

Thus, in the following declaration:

```
c#

class C(int i)
{
    protected int i = i; // references parameter
    public int I => i; // references field
}
```

The initializer for the field `i` references the parameter `i`, whereas the body of the property `I` references the field `i`.

Warn on shadowing by a member from base

Compiler will produce a warning on usage of an identifier when a base member shadows a primary constructor parameter if that primary constructor parameter was not passed to the base type via its constructor.

A primary constructor parameter is considered to be passed to the base type via its constructor when all the following conditions are true for an argument in *class_base*:

- The argument represents an implicit or explicit identity conversion of a primary constructor parameter;
- The argument is not part of an expanded `params` argument;

Semantics

A primary constructor leads to the generation of an instance constructor on the enclosing type with the given parameters. If the `class_base` has an argument list, the

generated instance constructor will have a `base` initializer with the same argument list.

Primary constructor parameters in class/struct declarations can be declared `ref`, `in` or `out`. Declaring `ref` or `out` parameters remains illegal in primary constructors of record declaration.

All instance member initializers in the class body will become assignments in the generated constructor.

If a primary constructor parameter is referenced from within an instance member, and the reference is not inside of a `nameof` argument, it is captured into the state of the enclosing type, so that it remains accessible after the termination of the constructor. A likely implementation strategy is via a private field using a mangled name. In a readonly struct the capture fields will be readonly. Therefore, access to captured parameters of a readonly struct will have similar restrictions as access to readonly fields. Access to captured parameters within a readonly member will have similar restrictions as access to instance fields in the same context.

Capturing is not allowed for parameters that have ref-like type, and capturing is not allowed for `ref`, `in` or `out` parameters. This is similar to a limitation for capturing in lambdas.

If a primary constructor parameter is only referenced from within instance member initializers, those can directly reference the parameter of the generated constructor, as they are executed as part of it.

Primary Constructor will do the following sequence of operations:

1. Parameter values are stored in capture fields, if any.
2. Instance initializers are executed
3. Base constructor initializer is called

Parameter references in any user code are replaced with corresponding capture field references.

For instance this declaration:

```c#
public class C(bool b, int i, string s) : B(b) // b passed to base
constructor
{
    public int I { get; set; } = i; // i used for initialization
    public string S // s used directly in function members
    {
        get => s;
```

```
        set => s = value ?? throw new NullArgumentException(nameof(X));
    }
    public C(string s) : this(true, 0, s) { } // must call this(...)
}
```

Generates code similar to the following:

```c#
public class C : B
{
    public int I { get; set; }
    public string S
    {
        get => __s;
        set => __s = value ?? throw new NullArgumentException(nameof(X));
    }
    public C(string s) : this(0, s) { ... } // must call this(...)

    // generated members
    private string __s; // for capture of s
    public C(bool b, int i, string s)
    {
        __s = s; // capture s
        I = i; // run I's initializer
        B(b) // run B's constructor
    }
}
```

It is an error for a non-primary constructor declaration to have the same parameter list as the primary constructor. All non-primary constructor declarations must use a `this` initializer, so that the primary constructor is ultimately called.

Records produce a warning if a primary constructor parameter isn't read within the (possibly generated) instance initializers or base initializer. Similar warnings will be reported for primary constructor parameters in classes and structures:

- for a by-value parameter, if the parameter is not captured and is not read within any instance initializers or base initializer.
- for an `in` parameter, if the parameter is not read within any instance initializers or base initializer.
- for a `ref` parameter, if the parameter is not read or written to within any instance initializers or base initializer.

Identical simple names and type names

There is a special language rule for scenarios often referred to as "Color Color" scenarios - Identical simple names and type names.

> In a member access of the form `E.I`, if `E` is a single identifier, and if the meaning of `E` as a *simple_name* (§11.7.4) is a constant, field, property, local variable, or parameter with the same type as the meaning of `E` as a *type_name* (§7.8.1), then both possible meanings of `E` are permitted. The member lookup of `E.I` is never ambiguous, since `I` shall necessarily be a member of the type `E` in both cases. In other words, the rule simply permits access to the static members and nested types of `E` where a compile-time error would otherwise have occurred.

With respect to primary constructors, the rule affects whether an identifier within an instance member should be treated as a type reference, or as a primary constructor parameter reference, which, in turn, captures the parameter into the the state of the enclosing type. Even though "the member lookup of `E.I` is never ambiguous", when lookup yields a member group, in some cases it is impossible to determine whether a member access refers to a static member or an instance member without fully resolving (binding) the member access. At the same time, capturing a primary constructor parameter changes properties of enclosing type in a way that affects semantic analysis. For example, the type might become unmanaged and fail certain constraints because of that. There are even scenarios for which binding can succeed either way, depending on whether the parameter is considered captured or not. For example:

C#

```
struct S1(Color Color)
{
    public void Test()
    {
        Color.M1(this);
    }
}

class Color
{
    public void M1<T>(T x, int y = 0)
    {
        System.Console.WriteLine("instance");
    }

    public static void M1<T>(T x) where T : unmanaged
    {
        System.Console.WriteLine("static");
    }
}
```

If we treat receiver `Color` as a value, we capture the parameter and 'S1' becomes managed. Then the static method becomes inapplicable due to the constraint and we would call instance method. However, if we treat the receiver as a type, we don't capture the parameter and 'S1' remains unmanaged, then both methods are applicable, but the static method is "better" because it doesn't have an optional parameter. Neither choice leads to an error, but each would result in distinct behavior.

Given this, compiler will produce an ambiguity error for a member access `E.I` when all the following conditions are met:

- Member lookup of `E.I` yields a member group containing instance and static members at the same time. Extension methods applicable to the receiver type are treated as instance methods for the purpose of this check.
- If `E` is treated as a simple name, rather than a type name, it would refer to a primary constructor parameter and would capture the parameter into the state of the enclosing type.

Double storage warnings

If a primary constructor parameter is passed to the base and *also* captured, there's a high risk that it is inadvertently stored twice in the object.

Compiler will produce a warning for `in` or by value argument in a `class_base argument_list` when all the following conditions are true:

- The argument represents an implicit or explicit identity conversion of a primary constructor parameter;
- The argument is not part of an expanded `params` argument;
- The primary constructor parameter is captured into the state of the enclosing type.

Compiler will produce a warning for a `variable_initializer` when all the following conditions are true:

- The variable initializer represents an implicit or explicit identity conversion of a primary constructor parameter;
- The primary constructor parameter is captured into the state of the enclosing type.

For example:

```c#
public class Person(string name)
{
    public string Name { get; set; } = name;   // initialization
```

```
    public override string ToString() => name; // capture
}
```

Attributes targeting primary constructors

At https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-03-13.md we decided to embrace the https://github.com/dotnet/csharplang/issues/7047 proposal.

The "method" attribute target is allowed on a *class_declaration/struct_declaration* with *parameter_list* and results in the corresponding primary constructor having that attribute. Attributes with the `method` target on a *class_declaration/struct_declaration* without *parameter_list* are are ignored with a warning.

C#

```
[method: FooAttr] // Good
public partial record Rec(
    [property: Foo] int X,
    [field: NonSerialized] int Y
);
[method: BarAttr] // warning CS0657: 'method' is not a valid attribute
location for this declaration. Valid attribute locations for this
declaration are 'type'. All attributes in this block will be ignored.
public partial record Rec
{
    public void Frobnicate()
    {
        ...
    }
}
[method: Attr] // Good
public record MyUnit1();
[method: Attr] // warning CS0657: 'method' is not a valid attribute location
for this declaration. Valid attribute locations for this declaration are
'type'. All attributes in this block will be ignored.
public record MyUnit2;
```

Primary constructors on records

With this proposal, records no longer need to separately specify a primary constructor mechanism. Instead, record (class and struct) declarations that have primary constructors would follow the general rules, with these simple additions:

- For each primary constructor parameter, if a member with the same name already exists, it must be an instance property or field. If not, a public init-only auto-property of the same name is synthesized with a property initializer assigning from the parameter.
- A deconstructor is synthesized with out parameters to match the primary constructor parameters.
- If an explicit constructor declaration is a "copy constructor" - a constructor that takes a single parameter of the enclosing type - it is not required to call a `this` initializer, and will not execute the member initializers present in the record declaration.

Drawbacks

- The allocation size of constructed objects is less obvious, as the compiler determines whether to allocate a field for a primary constructor parameter based on the full text of the class. This risk is similar to the implicit capture of variables by lambda expressions.
- A common temptation (or accidental pattern) might be to capture the "same" parameter at multiple levels of inheritance as it is passed up the constructor chain instead of explicitly allotting it a protected field at the base class, leading to duplicated allocations for the same data in objects. This is very similar to today's risk of overriding auto-properties with auto-properties.
- As proposed here, there is no place for additional logic that might usually be expressed in constructor bodies. The "primary constructor bodies" extension below addresses that.
- As proposed, execution order semantics are subtly different from within ordinary constructors, delaying member initializers to after the base calls. This could probably be remedied, but at the cost of some of the extension proposals (notably "primary constructor bodies").
- The proposal only works for scenarios where a single constructor can be designated primary.
- There is no way to express separate accessibility of the class and the primary constructor. An example is when public constructors all delegate to one private "build-it-all" constructor. If necessary, syntax could be proposed for that later.

Alternatives

No capture

A much simpler version of the feature would prohibit primary constructor parameters from occurring in member bodies. Referencing them would be an error. Fields would have to be explicitly declared if storage is desired beyond the initialization code.

```c#
public class C(string s)
{
    public S1 => s; // Nope!
    public S2 { get; } = s; // Still allowed
}
```

This could still be evolved to the full proposal at a later time, and would avoid a number of decisions and complexities, at the cost of removing less boilerplate initially, and probably also seeming unintuitive.

Explicit generated fields

An alternative approach is for primary constructor parameters to always and visibly generate a field of the same name. Instead of closing over the parameters in the same manner as local and anonymous functions, there would explicitly be a generated member declaration, similar to the public properties generated for primary construcor parameters in records. Just like for records, if a suitable member already exists, one would not be generated.

If the generated field is private it could still be elided when it is not used as a field in member bodies. In classes, however, a private field would often not be the right choice, because of the state duplication it could cause in derived classes. An option here would be to instead generating a protected field in classes, encouraging reuse of storage across inheritance layers. However, then we would not be able to elide the declaration, and would incur allocation cost for every primary constructor parameter.

This would align non-record primary constructors more closely with record ones, in that members are always (at least conceptually) generated, albeit different kinds of members with different accessibilities. But it would also lead to surprising differences from how parameters and locals are captured elsewhere in C#. If we were ever to allow local classes, for example, they would capture enclosing parameters and locals implicitly. Visibly generating shadowing fields for them would not seem to be a reasonable behavior.

Another problem often raised with this approach is that many developers have different naming conventions for parameters and fields. Which should be used for the primary

constructor parameter? Either choice would lead to inconsistency with the rest of the code.

Finally, visibly generating member declarations is really the name of the game for records, but much more surprising and "out of character" for non-record classes and structs. All in all, those are the reasons why the main proposal opts for implicit capture, with sensible behavior (consistent with records) for explicit member declarations when they are desired.

Remove instance members from initializer scope

The lookup rules above are intended to allow for the current behavior of primary constructor parameters in records when a corresponding member is manually declared, and to explain the behavior of the generated member when it is not. This requires lookup to differ between "initialization scope" (this/base initializers, member initializers) and "body scope" (member bodies), which the above proposal achieves by changing *when* primary constructor parameters are looked for, depending on where the reference occurs.

An observation is that referencing an instance member with a simple name in initializer scope always leads to an error. Instead of merely shadowing instance members in those places, could we simply take them out of scope? That way, there wouldn't be this weird conditional ordering of scopes.

This alternative is probably possible, but it would have some consequences that are somewhat far-reaching and potentially undesirable. First of all, if we remove instance members from initializer scope then a simple name that *does* correspond to an instance member and *not* to a primary constructor parameter could accidentally bind to something outside of the type declaration! This seems like it would rarely be intentional, and an error would be better.

Furthermore, *static* members are fine to reference in initialization scope. So we would have to distinguish between static and instance members in lookup, something we don't do today. (We do distinguish in overload resolution but that is not in play here). So that would have to also be changed, leading to yet more situations where e.g. in static contexts something would bind "further out" rather than error because it found an instance member.

All in all this "simplification" would lead to quite a downstream complication that no-one asked for.

Possible extensions

These are variations or additions to the core proposal that may be considered in conjunction with it, or at a later stage if deemed useful.

Primary constructor parameter access within constructors

The rules above make it an error to reference a primary constructor parameter within another constructor. This could be allowed within the *body* of other constructors, though, since the primary constructor runs first. However it would need to remain disallowed within the argument list of the `this` initializer.

```c#
public class C(bool b, int i, string s) : B(b)
{
    public C(string s) : this(b, s) // b still disallowed
    {
        i++; // could be allowed
    }
}
```

Such access would still incur capture, as that would be the only way the constructor body could get at the variable after the primary constructor has already run.

The prohibition on primary constructor parameters in the this-initializer's arguments could be weakened to allow them, but make them not definitely assigned, but that does not seem useful.

Allow constructors without a `this` initializer

Constructors without a `this` initializer (i.e. with an implicit or explicit `base` initializer) could be allowed. Such a constructor would *not* run instance field, property and event initializers, as those would be considered to be part of the primary constructor only.

In the presence of such base-calling constructors, there are a couple of options for how primary constructor parameter capture is handled. The simplest is to completely disallow capture in this situation. Primary constructor parameters would be for initialization only when such constructors exist.

Alternatively, if combined with the previously described option to allow access to primary constructor parameters within constructors, the parameters could enter the constructor body as not definitely assigned, and ones that are captured would need to be definitely assigned by the end of the constructor body. They would essentially be implicit out parameters. That way, captured primary constructor parameters would

always have a sensible (i.e. explicitly assigned) value by the time they are consumed by other function members.

An attraction of this extension (in either form) is that it fully generalizes the current exemption for "copy constructors" in records, without leading to situations where uninitialized primary constructor parameters are observed. Essentially, constructors that initialize the object in alternative ways are fine. The capture-related restrictions would not be a breaking change for existing manually defined copy constructors in records, because records never capture their primary constructor parameters (they generate fields instead).

```
c#
```

```
public class C(bool b, int i, string s) : B(b)
{
    public int I { get; set; } = i; // i used for initialization
    public string S // s used directly in function members
    {
        get => s;
        set => s = value ?? throw new NullArgumentException(nameof(X));
    }
    public C(string s2) : base(true) // cannot use `string s` because it
would shadow
    {
        s = s2; // must initialize s because it is captured by S
    }
    protected C(C original) : base(original) // copy constructor
    {
        this.s = original.s; // assignment to b and i not required because
not captured
    }
}
```

Primary constructor bodies

Constructors themselves often contain parameter validation logic or other nontrivial initialization code that cannot be expressed as initializers.

Primary constructors could be extended to allow statement blocks to appear directly in the class body. Those statements would be inserted in the generated constructor at the point where they appear between initializing assignments, and would thus be executed interspersed with initializers. For instance:

```
c#
```

```
public class C(int i, string s) : B(s)
{
```

```
    {
        if (i < 0) throw new ArgumentOutOfRangeException(nameof(i));
    }
    int[] a = new int[i];
    public int S => s;
}
```

A lot of this scenario might be adequately be covered if we were to introduce "final initializers" which run after the constructors *and* any object/collection initializers have completed. However, argument validation is one thing that would ideally happen as early as possible.

Primary constructor bodies could also provide a place for allowing an access modifier for the primary constructor, allowing it to deviate from the accessibility of the enclosing type.

Combined parameter and member declarations

A possible and often mentioned addition could be to allow primary constructor parameters to be annotated so that they would *also* declare a member on the type. Most commonly it is proposed to allow an access specifier on the parameters to trigger the member generation:

```
c#

public class C(bool b, protected int i, string s) : B(b) // i is a field as
well as a parameter
{
    void M()
    {
        ... i ... // refers to the field i
        ... s ... // closes over the parameter s
    }
}
```

There are some problems:

- What if a property is desired, not a field? Having `{ get; set; }` syntax inline in a parameter list does not seem appetizing.
- What if different naming conventions are used for parameters and fields? Then this feature would be useless.

This is a potential future addition that can be adopted or not. The current proposal leaves the possibility open.

Open questions

Lookup order for type parameters

The https://github.com/dotnet/csharplang/blob/main/proposals/csharp-12.0/primary-constructors.md#lookup section specifies that type parameters of declaring type should come before type's primary constructor parameters in every context where those parameters are in scope. However, we already have existing behavior with records - primary constructor parameters come before type parameters in base initializer and field initializers.

What should we do about this discrepancy?

- Adjust the rules to match the behavior.
- Adjust the behavior (a possible breaking change).
- Disallow a primiry constructor parameter to use type parameter's name (a possible breaking change).
- Do nothing, accept the inconsistency between the spec and implementation.

Conclusion:

Adjust the rules to match the behavior (https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-09-25.md#primary-constructors).

Field targeting attributes for captured primary constructor parameters

Should we allow field targeting attributes for captured primary constructor parameters?

C#

```
class C1([field: Test] int x) // Parameter is captured, the attribute goes
to the capture field
{
    public int X => x;
}
class C2([field: Test] int x) // Parameter is not captured, the attribute is
ignored with a warning CS0657: 'field' is not a valid attribute location for
this declaration. Valid attribute locations for this declaration are
'param'. All attributes in this block will be ignored.
{
```

```
    public int X = x;
}
```

Right now the attributes are ignored with the warning regardless of whether the parameter is captured.

Note that for records, field targeted attributes are allowed when a property is synthesized for it. The attributes go on the backing field then.

C#

```
record R1([field: Test]int X); // Ok, the attribute goes on the backing
field
record R2([field: Test]int X) // warning CS0657: 'field' is not a valid
attribute location for this declaration. Valid attribute locations for this
declaration are 'param'. All attributes in this block will be ignored.
{
    public int X = X;
}
```

Conclusion:

Not allowed (https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-05-03.md#attributes-on-captured-parameters).

Warn on shadowing by a member from base

Should we report a warning when a member from base is shadowing a primary constructor parameter inside a member (see https://github.com/dotnet/csharplang/discussions/7109#discussioncomment-5666621)?

Conclusion:

An alternative design is approved - https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-05-08.md#primary-constructors

Capturing instance of the enclosing type in a closure

When a parameter captured into the state of the enclosing type is also referenced in a lambda inside an instance initializer or a base initializer, the lambda and the state of the enclosing type should refer to the same location for the parameter. For example:

C#

```
partial class C1
{
    public System.Func<int> F1 = Execute1(() => p1++);
}

partial class C1 (int p1)
{
    public int M1() { return p1++; }
    static System.Func<int> Execute1(System.Func<int> f)
    {
        _ = f();
        return f;
    }
}
```

Since naive implementation of capturing a parameter into the state of the type simply captures the parameter in a private instance field, the lambda needs to refer to the same field. As a result, it needs to be able to access the instance of the type. This requires capturing `this` into a closure before the base constructor is invoked. That, in turn, results in safe, but an unverifiable IL. Is this acceptable?

Alternatively we could:

- Disallow lambdas like that;
- Or, instead, capture parameters like that in an instance of a separate class (yet another closure), and share that instance between the closure and the instance of the enclosing type. Thus eliminating the need to capture `this` in a closure.

Conclusion:

We are comfortable with capturing `this` into a closure before the base constructor is invoked (https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-02-15.md). The runtime team didn't find the IL pattern problematic as well.

Assigning to `this` within a struct

C# allows to assign to `this` within a struct. If the struct captures a primary constructor parameter, the assignment is going to overwrite its value, which might be not obvious to the user. Do we want to report a warning for assignments like this?

C#

```
struct S(int x)
{
    int X => x;

    void M(S s)
    {
        this = s; // 'x' is overwritten
    }
}
```

Conclusion:

Allowed, no warning (https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-02-15.md).

Double storage warning for initialization plus capture

We have a warning if a primary constructor parameter is passed to the base and *also* captured, because there's a high risk that it is inadvertently stored twice in the object.

It seems that there's a similar risk if a parameter is used to initialize a member, and is also captured. Here's a small example:

c#

```
public class Person(string name)
{
    public string Name { get; set; } = name;   // initialization
    public override string ToString() => name; // capture
}
```

For a given instance of `Person`, changes to `Name` would not be reflected in the output of `ToString`, which is probably unintended on the developer's part.

Should we introduce a double storage warning for this situation?

This is how it would work:

The compiler will produce a warning for a `variable_initializer` when all the following conditions are true:

- The variable initializer represents an implicit or explicit identity conversion of a primary constructor parameter;

- The primary constructor parameter is captured into the state of the enclosing type.

Conclusion:

Approved, see https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-05-15.md#primary-constructors

LDM meetings

- https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-10-17.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-01-18.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-02-15.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-02-22.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-03-13.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-05-03.md#primary-constructors
- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-05-08.md#primary-constructors
- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-05-15.md#primary-constructors
- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-09-25.md#primary-constructors

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# feature specification feedback

The C# feature specifications are open source. Provide feedback here.

Open a documentation issue

Provide product feedback

Init Only Setters

Article • 02/01/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

This proposal adds the concept of init only properties and indexers to C#. These properties and indexers can be set at the point of object creation but become effectively `get` only once object creation has completed. This allows for a much more flexible immutable model in C#.

Motivation

The underlying mechanisms for building immutable data in C# haven't changed since 1.0. They remain:

1. Declaring fields as `readonly`.
2. Declaring properties that contain only a `get` accessor.

These mechanisms are effective at allowing the construction of immutable data but they do so by adding cost to the boilerplate code of types and opting such types out of features like object and collection initializers. This means developers must choose between ease of use and immutability.

A simple immutable object like `Point` requires twice as much boiler plate code to support construction as it does to declare the type. The bigger the type the bigger the cost of this boiler plate:

```
C#
```

```
struct Point
{
    public int X { get; }
    public int Y { get; }

    public Point(int x, int y)
    {
        this.X = x;
        this.Y = y;
    }
}
```

The `init` accessor makes immutable objects more flexible by allowing the caller to mutate the members during the act of construction. That means the object's immutable properties can participate in object initializers and thus removes the need for all constructor boilerplate in the type. The `Point` type is now simply:

C#

```
struct Point
{
    public int X { get; init; }
    public int Y { get; init; }
}
```

The consumer can then use object initializers to create the object

C#

```
var p = new Point() { X = 42, Y = 13 };
```

Detailed Design

init accessors

An init only property (or indexer) is declared by using the `init` accessor in place of the `set` accessor:

C#

```
class Student
{
    public string FirstName { get; init; }
```

```
    public string LastName { get; init; }
}
```

An instance property containing an `init` accessor is considered settable in the following circumstances, except when in a local function or lambda:

- During an object initializer
- During a `with` expression initializer
- Inside an instance constructor of the containing or derived type, on `this` or `base`
- Inside the `init` accessor of any property, on `this` or `base`
- Inside attribute usages with named parameters

The times above in which the `init` accessors are settable are collectively referred to in this document as the construction phase of the object.

This means the `Student` class can be used in the following ways:

C#

```
var s = new Student()
{
    FirstName = "Jared",
    LastName = "Parosns",
};
s.LastName = "Parsons"; // Error: LastName is not settable
```

The rules around when `init` accessors are settable extend across type hierarchies. If the member is accessible and the object is known to be in the construction phase then the member is settable. That specifically allows for the following:

C#

```
class Base
{
    public bool Value { get; init; }
}

class Derived : Base
{
    public Derived()
    {
        // Not allowed with get only properties but allowed with init
        Value = true;
    }
}

class Consumption
{
```

```
    void Example()
    {
        var d = new Derived() { Value = true };
    }
}
```

At the point an `init` accessor is invoked, the instance is known to be in the open construction phase. Hence an `init` accessor is allowed to take the following actions in addition to what a normal `set` accessor can do:

1. Call other `init` accessors available through `this` or `base`
2. Assign `readonly` fields declared on the same type through `this`

C#

```
class Complex
{
    readonly int Field1;
    int Field2;
    int Prop1 { get; init ; }
    int Prop2
    {
        get => 42;
        init
        {
            Field1 = 13; // okay
            Field2 = 13; // okay
            Prop1 = 13; // okay
        }
    }
}
```

The ability to assign `readonly` fields from an `init` accessor is limited to those fields declared on the same type as the accessor. It cannot be used to assign `readonly` fields in a base type. This rule ensures that type authors remain in control over the mutability behavior of their type. Developers who do not wish to utilize `init` cannot be impacted from other types choosing to do so:

C#

```
class Base
{
    internal readonly int Field;
    internal int Property
    {
        get => Field;
        init => Field = value; // Okay
```

```
    }

    internal int OtherProperty { get; init; }
}

class Derived : Base
{
    internal readonly int DerivedField;
    internal int DerivedProperty
    {
        get => DerivedField;
        init
        {
            DerivedField = 42;  // Okay
            Property = 0;       // Okay
            Field = 13;         // Error Field is readonly
        }
    }

    public Derived()
    {
        Property = 42;  // Okay
        Field = 13;     // Error Field is readonly
    }
}
```

When `init` is used in a virtual property then all the overrides must also be marked as `init`. Likewise it is not possible to override a simple `set` with `init`.

C#

```
class Base
{
    public virtual int Property { get; init; }
}

class C1 : Base
{
    public override int Property { get; init; }
}

class C2 : Base
{
    // Error: Property must have init to override Base.Property
    public override int Property { get; set; }
}
```

An `interface` declaration can also participate in `init` style initialization via the following pattern:

C#

```
interface IPerson
{
    string Name { get; init; }
}

class Init
{
    void M<T>() where T : IPerson, new()
    {
        var local = new T()
        {
            Name = "Jared"
        };
        local.Name = "Jraed"; // Error
    }
}
```

Restrictions of this feature:

- The `init` accessor can only be used on instance properties
- A property cannot contain both an `init` and `set` accessor
- All overrides of a property must have `init` if the base had `init`. This rule also applies to interface implementation.

Readonly structs

`init` accessors (both auto-implemented accessors and manually-implemented accessors) are permitted on properties of `readonly struct`s, as well as `readonly` properties. `init` accessors are not permitted to be marked `readonly` themselves, in both `readonly` and non-`readonly` `struct` types.

C#

```
readonly struct ReadonlyStruct1
{
    public int Prop1 { get; init; } // Allowed
}

struct ReadonlyStruct2
{
    public readonly int Prop2 { get; init; } // Allowed

    public int Prop3 { get; readonly init; } // Error
}
```

Metadata encoding

Property `init` accessors will be emitted as a standard `set` accessor with the return type marked with a modreq of `IsExternalInit`. This is a new type which will have the following definition:

```C#
namespace System.Runtime.CompilerServices
{
    public sealed class IsExternalInit
    {
    }
}
```

The compiler will match the type by full name. There is no requirement that it appear in the core library. If there are multiple types by this name then the compiler will tie break in the following order:

1. The one defined in the project being compiled
2. The one defined in corelib

If neither of these exist then a type ambiguity error will be issued.

The design for `IsExternalInit` is futher covered in this issue

Questions

Breaking changes

One of the main pivot points in how this feature is encoded will come down to the following question:

> Is it a binary breaking change to replace `init` with `set`?

Replacing `init` with `set` and thus making a property fully writable is never a source breaking change on a non-virtual property. It simply expands the set of scenarios where the property can be written. The only behavior in question is whether or not this remains a binary breaking change.

If we want to make the change of `init` to `set` a source and binary compatible change then it will force our hand on the modreq vs. attributes decision below because it will rule out modreqs as a solution. If on the other hand this is seen as a non-interesting then this will make the modreq vs. attribute decision less impactful.

Resolution This scenario is not seen as compelling by LDM.

Modreqs vs. attributes

The emit strategy for `init` property accessors must choose between using attributes or modreqs when emitting during metadata. These have different trade offs that need to be considered.

Annotating a property set accessor with a modreq declaration means CLI compliant compilers will ignore the accessor unless it understands the modreq. That means only compilers aware of `init` will read the member. Compilers unaware of `init` will ignore the `set` accessor and hence will not accidentally treat the property as read / write.

The downside of modreq is `init` becomes a part of the binary signature of the `set` accessor. Adding or removing `init` will break binary compatbility of the application.

Using attributes to annotate the `set` accessor means that only compilers which understand the attribute will know to limit access to it. A compiler unaware of `init` will see it as a simple read / write property and allow access.

This would seemingly mean this decision is a choice between extra safety at the expense of binary compatibility. Digging in a bit the extra safety is not exactly what it seems. It will not protect against the following circumstances:

1. Reflection over `public` members
2. The use of `dynamic`
3. Compilers that don't recognize modreqs

It should also be considered that, when we complete the IL verification rules for .NET 5, `init` will be one of those rules. That means extra enforcement will be gained from simply verifying compilers emitting verifiable IL.

The primary languages for .NET (C#, F# and VB) will all be updated to recognize these `init` accessors. Hence the only realistic scenario here is when a C# 9 compiler emits `init` properties and they are seen by an older toolset such as C# 8, VB 15, etc ... C# 8. That is the trade off to consider and weigh against binary compatibility.

Note This discussion primarily applies to members only, not to fields. While `init` fields were rejected by LDM they are still interesting to consider for the modreq vs. attribute discussion. The `init` feature for fields is a relaxation of the existing restriction of `readonly`. That means if we emit the fields as `readonly` + an attribute there is no risk of

older compilers mis-using the field because they would already recognize `readonly`. Hence using a modreq here doesn't add any extra protection.

Resolution The feature will use a modreq to encode the property `init` setter. The compelling factors were (in no particular order):

- Desire to discourage older compilers from violating `init` semantics
- Desire to make adding or removing `init` in a `virtual` declaration or `interface` both a source and binary breaking change.

Given there was also no significant support for removing `init` to be a binary compatible change it made the choice of using modreq straight forward.

init vs. initonly

There were three syntax forms which got significant consideration during our LDM meeting:

```
C#

// 1. Use init
int Option1 { get; init; }
// 2. Use init set
int Option2 { get; init set; }
// 3. Use initonly
int Option3 { get; initonly; }
```

Resolution There was no syntax which was overwhelmingly favored in LDM.

One point which got significant attention was how the choice of syntax would impact our ability to do `init` members as a general feature in the future. Choosing option 1 would mean that it would be difficult to define a property which had an `init` style `get` method in the future. Eventually it was decided that if we decided to go forward with general `init` members in future, we could allow `init` to be a modifier in the property accessor list as well as a short hand for `init set`. Essentially the following two declarations would be identical.

```
C#

int Property1 { get; init; }
int Property1 { get; init set; }
```

The decision was made to move forward with `init` as a standalone accessor in the property accessor list.

Warn on failed init

Consider the following scenario. A type declares an `init` only member which is not set in the constructor. Should the code which constructs the object get a warning if they failed to initialize the value?

At that point it is clear the field will never be set and hence has a lot of similarities with the warning around failing to initialize `private` data. Hence a warning would seemingly have some value here?

There are significant downsides to this warning though:

1. It complicates the compatibility story of changing `readonly` to `init`.
2. It requires carrying additional metadata around to denote the members which are required to be initialized by the caller.

Further if we believe there is value here in the overall scenario of forcing object creators to be warned / error'd about specific fields then this likely makes sense as a general feature. There is no reason it should be limited to just `init` members.

Resolution There will be no warning on consumption of `init` fields and properties.

LDM wants to have a broader discussion on the idea of required fields and properties. That may cause us to come back and reconsider our position on `init` members and validation.

Allow init as a field modifier

In the same way `init` can serve as a property accessor it could also serve as a designation on fields to give them similar behaviors as `init` properties. That would allow for the field to be assigned before construction was complete by the type, derived types, or object initializers.

C#

```
class Student
{
    public init string FirstName;
    public init string LastName;
}
```

```
var s = new Student()
{
    FirstName = "Jarde",
    LastName = "Parsons",
}

s.FirstName = "Jared"; // Error FirstName is readonly
```

In metadata these fields would be marked in the same way as `readonly` fields but with an additional attribute or modreq to indicate they are `init` style fields.

Resolution LDM agrees this proposal is sound but overall the scenario felt disjoint from properties. The decision was to proceed only with `init` properties for now. This has a suitable level of flexibility as an `init` property can mutate a `readonly` field on the declaring type of the property. This will be reconsidered if there is significant customer feedback that justifies the scenario.

Allow init as a type modifier

In the same way the `readonly` modifier can be applied to a `struct` to automatically declare all fields as `readonly`, the `init` only modifier can be declared on a `struct` or `class` to automatically mark all fields as `init`. This means the following two type declarations are equivalent:

C#

```
struct Point
{
    public init int X;
    public init int Y;
}

// vs.

init struct Point
{
    public int X;
    public int Y;
}
```

Resolution This feature is too *cute* here and conflicts with the `readonly struct` feature on which it is based. The `readonly struct` feature is simple in that it applies `readonly` to all members: fields, methods, etc ... The `init struct` feature would only apply to properties. This actually ends up making it confusing for users.

Given that `init` is only valid on certain aspects of a type, we rejected the idea of having it as a type modifier.

Considerations

Compatibility

The `init` feature is designed to be compatible with existing `get` only properties. Specifically it is meant to be a completely additive change for a property which is `get` only today but desires more flexbile object creation semantics.

For example consider the following type:

```C#
class Name
{
    public string First { get; }
    public string Last { get; }

    public Name(string first, string last)
    {
        First = first;
        Last = last;
    }
}
```

Adding `init` to these properties is a non-breaking change:

```C#
class Name
{
    public string First { get; init; }
    public string Last { get; init; }

    public Name(string first, string last)
    {
        First = first;
        Last = last;
    }
}
```

IL verification

When .NET Core decides to re-implement IL verification, the rules will need to be adjusted to account for `init` members. This will need to be included in the rule changes for non-mutating acess to `readonly` data.

The IL verification rules will need to be broken into two parts:

1. Allowing `init` members to set a `readonly` field.
2. Determining when an `init` member can be legally called.

The first is a simple adjustment to the existing rules. The IL verifier can be taught to recognize `init` members and from there it just needs to consider a `readonly` field to be settable on `this` in such a member.

The second rule is more complicated. In the simple case of object initializers the rule is straight forward. It should be legal to call `init` members when the result of a `new` expression is still on the stack. That is until the value has been stored in a local, array element or field or passed as an argument to another method it will still be legal to call `init` members. This ensures that once the result of the `new` expression is published to a named identifier (other than `this`) then it will no longer be legal to call `init` members.

The more complicated case though is when we mix `init` members, object initializers and `await`. That can cause the newly created object to be temporarily hoisted into a state machine and hence put into a field.

```
C#

var student = new Student()
{
    Name = await SomeMethod()
};
```

Here the result of `new Student()` will be hoised into a state machine as a field before the set of `Name` occurs. The compiler will need to mark such hoisted fields in a way that the IL verifier understands they're not user accessible and hence doesn't violate the intended semantics of `init`.

init members

The `init` modifier could be extended to apply to all instance members. This would generalize the concept of `init` during object construction and allow types to declare helper methods that could partipate in the construction process to initialize `init` fields and properties.

Such members would have all the restricions that an `init` accessor does in this design. The need is questionable though and this can be safely added in a future version of the language in a compatible manner.

Generate three accessors

One potential implementation of `init` properties is to make `init` completely separate from `set`. That means that a property can potentially have three different accessors: `get`, `set`, and `init`.

This has the potential advantage of allowing the use of modreq to enforce correctness while maintaining binary compatibility. The implementation would roughly be the following:

1. An `init` accessor is always emitted if there is a `set`. When not defined by the developer it is simply a reference to `set`.
2. The set of a property in an object initializer will always use `init` if present but fall back to `set` if it's missing.

This means that a developer can always safely delete `init` from a property.

The downside of this design is that is only useful if `init` is **always** emitted when there is a `set`. The language can't know if `init` was deleted in the past, it has to assume it was and hence the `init` must always be emitted. That would cause a significant metadata expansion and is simply not worth the cost of the compatibility here.

Required Members

Article • 03/17/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

This proposal adds a way of specifying that a property or field is required to be set during object initialization, forcing the instance creator to provide an initial value for the member in an object initializer at the creation site.

Motivation

Object hierarchies today require a lot of boilerplate to carry data across all levels of the hierarchy. Let's look at a simple hierarchy involving a `Person` as might be defined in C# 8:

C#

```
class Person
{
    public string FirstName { get; }
    public string MiddleName { get; }
    public string LastName { get; }

    public Person(string firstName, string lastName, string? middleName =
null)
    {
        FirstName = firstName;
        LastName = lastName;
        MiddleName = middleName ?? string.Empty;
    }
}
```

```
class Student : Person
{
    public int ID { get; }
    public Student(int id, string firstName, string lastName, string?
middleName = null)
        : base(firstName, lastName, middleName)
    {
        ID = id;
    }
}
```

There's lots of repetition going on here:

1. At the root of the hierarchy, the type of each property had to be repeated twice, and the name had to be repeated four times.
2. At the derived level, the type of each inherited property had to be repeated once, and the name had to be repeated twice.

This is a simple hierarchy with 3 properties and 1 level of inheritance, but many real-world examples of these types of hierarchies go many levels deeper, accumulating larger and larger numbers of properties to pass along as they do so. Roslyn is one such codebase, for example, in the various tree types that make our CSTs and ASTs. This nesting is tedious enough that we have code generators to generate the constructors and definitions of these types, and many customers take similar approaches to the problem. C# 9 introduces records, which for some scenarios can make this better:

```
C#
```

```
record Person(string FirstName, string LastName, string MiddleName = "");
record Student(int ID, string FirstName, string LastName, string MiddleName
= "") : Person(FirstName, LastName, MiddleName);
```

`record`s eliminate the first source of duplication, but the second source of duplication remains unchanged: unfortunately, this is the source of duplication that grows as the hierarchy grows, and is the most painful part of the duplication to fix up after making a change in the hierarchy as it required chasing the hierarchy through all of its locations, possibly even across projects and potentially breaking consumers.

As a workaround to avoid this duplication, we have long seen consumers embracing object initializers as a way of avoiding writing constructors. Prior to C# 9, however, this had 2 major downsides:

1. The object hierarchy has to be fully mutable, with `set` accessors on every property.
2. There is no way to ensure that every instantation of an object from the graph sets every member.

C# 9 again addressed the first issue here, by introducing the `init` accessor: with it, these properties can be set on object creation/initialization, but not subsequently. However, we again still have the second issue: properties in C# have been optional since C# 1.0. Nullable reference types, introduced in C# 8.0, addressed part of this issue: if a constructor does not initialize a non-nullable reference-type property, then the user is warned about it. However, this doesn't solve the problem: the user here wants to not repeat large parts of their type in the constructor, they want to pass the *requirement* to set properties on to their consumers. It also doesn't provide any warnings about `ID` from `Student`, as that is a value type. These scenarios are extremely common in database model ORMs, such as EF Core, which need to have a public parameterless constructor but then drive nullability of the rows based on the nullability of the properties.

This proposal seeks to address these concerns by introducing a new feature to C#: required members. Required members will be required to be initialized by consumers, rather than by the type author, with various customizations to allow flexibility for multiple constructors and other scenarios.

Detailed Design

`class`, `struct`, and `record` types gain the ability to declare a *required_member_list*. This list is the list of all the properties and fields of a type that are considered *required*, and must be initialized during the construction and initialization of an instance of the type. Types inherit these lists from their base types automatically, providing a seamless experience that removes boilerplate and repetitive code.

`required` modifier

We add `'required'` to the list of modifiers in *field_modifier* and *property_modifier*. The *required_member_list* of a type is composed of all the members that have had `required` applied to them. Thus, the `Person` type from earlier now looks like this:

C#

```
public class Person
{
    // The default constructor requires that FirstName and LastName be set
at construction time
    public required string FirstName { get; init; }
    public string MiddleName { get; init; } = "";
    public required string LastName { get; init; }
}
```

All constructors on a type that has a *required_member_list* automatically advertise a *contract* that consumers of the type must initialize all of the properties in the list. It is an error for a constructor to advertise a contract that requires a member that is not at least as accessible as the constructor itself. For example:

```
C#
```

```
public class C
{
    public required int Prop { get; protected init; }

    // Advertises that Prop is required. This is fine, because the
constructor is just as accessible as the property initer.
    protected C() {}

    // Error: ctor C(object) is more accessible than required property
Prop.init.
    public C(object otherArg) {}
}
```

`required` is only valid in `class`, `struct`, and `record` types. It is not valid in `interface` types. `required` cannot be combined with the following modifiers:

- `fixed`
- `ref readonly`
- `ref`
- `const`
- `static`

`required` is not allowed to be applied to indexers.

The compiler will issue a warning when `Obsolete` is applied to a required member of a type and:

1. The type is not marked `Obsolete`, or
2. Any constructor not attributed with `SetsRequiredMembersAttribute` is not marked `Obsolete`.

`SetsRequiredMembersAttribute`

All constructors in a type with required members, or whose base type specifies required members, must have those members set by a consumer when that constructor is called. In order to exempt constructors from this requirement, a constructor can be attributed

with `SetsRequiredMembersAttribute`, which removes these requirements. The constructor body is not validated to ensure that it definitely sets the required members of the type.

`SetsRequiredMembersAttribute` removes *all* requirements from a constructor, and those requirements are not checked for validity in any way. NB: this is the escape hatch if inheriting from a type with an invalid required members list is necessary: mark the constructor of that type with `SetsRequiredMembersAttribute`, and no errors will be reported.

If a constructor `C` chains to a `base` or `this` constructor that is attributed with `SetsRequiredMembersAttribute`, `C` must also be attributed with `SetsRequiredMembersAttribute`.

For record types, we will emit `SetsRequiredMembersAttribute` on the synthesized copy constructor of a record if the record type or any of its base types have required members.

NB: An earlier version of this proposal had a larger metalanguage around initialization, allowing adding and removing individual required members from a constructor, as well as validation that the constructor was setting all required members. This was deemed too complex for the initial release, and removed. We can look at adding more complex contracts and modifications as a later feature.

Enforcement

For every constructor `Ci` in type `T` with required members `R`, consumers calling `Ci` must do one of:

- Set all members of `R` in an *object_initializer* on the *object_creation_expression*,
- Or set all members of `R` via the *named_argument_list* section of an *attribute_target*.

unless `Ci` is attributed with `SetsRequiredMembers`.

If the current context does not permit an *object_initializer* or is not an *attribute_target*, and `Ci` is not attributed with `SetsRequiredMembers`, then it is an error to call `Ci`.

`new()` constraint

A type with a parameterless constructor that advertises a *contract* is not allowed to be substituted for a type parameter constrained to `new()`, as there is no way for the generic instantiation to ensure that the requirements are satisfied.

`struct` `default`s

Required members are not enforced on instances of `struct` types created with `default` or `default(StructType)`. They are enforced for `struct` instances created with `new StructType()`, even when `StructType` has no parameterless constructor and the default struct constructor is used.

Accessibility

It is an error to mark a member required if the member cannot be set in any context where the containing type is visible.

- If the member is a field, it cannot be `readonly`.
- If the member is a property, it must have a setter or initer at least as accessible as the member's containing type.

This means the following cases are not allowed:

C#

```
interface I
{
    int Prop1 { get; }
}
public class Base
{
    public virtual int Prop2 { get; set; }

    protected required int _field; // Error: _field is not at least as
visible as Base. Open question below about the protected constructor
scenario

    public required readonly int _field2; // Error: required fields cannot
be readonly
    protected Base() { }

    protected class Inner
    {
        protected required int PropInner { get; set; } // Error: PropInner
cannot be set inside Base or Derived
    }
}
public class Derived : Base, I
{
    required int I.Prop1 { get; } // Error: explicit interface implementions
cannot be required as they cannot be set in an object initializer

    public required override int Prop2 { get; set; } // Error: this property
is hidden by Derived.Prop2 and cannot be set in an object initializer
```

```
    public new int Prop2 { get; }

    public required int Prop3 { get; } // Error: Required member must have a
setter or initer

    public required int Prop4 { get; internal set; } // Error: Required
member setter must be at least as visible as the constructor of Derived
}
```

It is an error to hide a `required` member, as that member can no longer be set by a consumer.

When overriding a `required` member, the `required` keyword must be included on the method signature. This is done so that if we ever want to allow unrequiring a property with an override in the future, we have design space to do so.

Overrides are allowed to mark a member `required` where it was not `required` in the base type. A member so-marked is added to the required members list of the derived type.

Types are allowed to override required virtual properties. This means that if the base virtual property has storage, and the derived type tries to access the base implementation of that property, they could observe uninitialized storage. NB: This is a general C# anti-pattern, and we don't think that this proposal should attempt to address it.

Effect on nullable analysis

Members that are marked `required` are not required to be initialized to a valid nullable state at the end of a constructor. All `required` members from this type and any base types are considered by nullable analysis to be default at the beginning of any constructor in that type, unless chaining to a `this` or `base` constructor that is attributed with `SetsRequiredMembersAttribute`.

Nullable analysis will warn about all `required` members from the current and base types that do not have a valid nullable state at the end of a constructor attributed with `SetsRequiredMembersAttribute`.

C#

```
#nullable enable
public class Base
{
    public required string Prop1 { get; set; }
```

```
    public Base() {}

    [SetsRequiredMembers]
    public Base(int unused) { Prop1 = ""; }
}
public class Derived : Base
{
    public required string Prop2 { get; set; }

    [SetsRequiredMembers]
    public Derived() : base()
    {
    } // Warning: Prop1 and Prop2 are possibly null.

    [SetsRequiredMembers]
    public Derived(int unused) : base()
    {
        Prop1.ToString(); // Warning: possibly null dereference
        Prop2.ToString(); // Warning: possibly null dereference
    }

    [SetsRequiredMembers]
    public Derived(int unused, int unused2) : this()
    {
        Prop1.ToString(); // Ok
        Prop2.ToString(); // Ok
    }

    [SetsRequiredMembers]
    public Derived(int unused1, int unused2, int unused3) : base(unused1)
    {
        Prop1.ToString(); // Ok
        Prop2.ToString(); // Warning: possibly null dereference
    }
}
```

Metadata Representation

The following 2 attributes are known to the C# compiler and required for this feature to function:

C#

```
namespace System.Runtime.CompilerServices
{
    [AttributeUsage(AttributeTargets.Class | AttributeTargets.Struct |
AttributeTargets.Field | AttributeTargets.Property, AllowMultiple = false,
Inherited = false)]
    public sealed class RequiredMemberAttribute : Attribute
    {
        public RequiredMemberAttribute() {}
    }
```

```
}

namespace System.Diagnostics.CodeAnalysis
{
    [AttributeUsage(AttributeTargets.Constructor, AllowMultiple = false,
Inherited = false)]
    public sealed class SetsRequiredMembersAttribute : Attribute
    {
        public SetsRequiredMembersAttribute() {}
    }
}
```

It is an error to manually apply `RequiredMemberAttribute` to a type.

Any member that is marked `required` has a `RequiredMemberAttribute` applied to it. In addition, any type that defines such members is marked with `RequiredMemberAttribute`, as a marker to indicate that there are required members in this type. Note that if type `B` derives from `A`, and `A` defines `required` members but `B` does not add any new or override any existing `required` members, `B` will not be marked with a `RequiredMemberAttribute`. To fully determine whether there are any required members in `B`, checking the full inheritance hierarchy is necessary.

Any constructor in a type with `required` members that does not have `SetsRequiredMembersAttribute` applied to it is marked with two attributes:

1. `System.Runtime.CompilerServices.CompilerFeatureRequiredAttribute` with the feature name `"RequiredMembers"`.
2. `System.ObsoleteAttribute` with the string `"Types with required members are not supported in this version of your compiler"`, and the attribute is marked as an error, to prevent any older compilers from using these constructors.

We don't use a `modreq` here because it is a goal to maintain binary compat: if the last `required` property was removed from a type, the compiler would no longer synthesize this `modreq`, which is a binary-breaking change and all consumers would need to be recompiled. A compiler that understands `required` members will ignore this obsolete attribute. Note that members can come from base types as well: even if there are no new `required` members in the current type, if any base type has `required` members, this `Obsolete` attribute will be generated. If the constructor already has an `Obsolete` attribute, no additional `Obsolete` attribute will be generated.

We use both `ObsoleteAttribute` and `CompilerFeatureRequiredAttribute` because the latter is new this release, and older compilers don't understand it. In the future, we may

be able to drop the `ObsoleteAttribute` and/or not use it to protect new features, but for now we need both for full protection.

To build the full list of `required` members `R` for a given type `T`, including all base types, the following algorithm is run:

1. For every `Tb`, starting with `T` and working through the base type chain until `object` is reached.
2. If `Tb` is marked with `RequiredMemberAttribute`, then all members of `Tb` marked with `RequiredMemberAttribute` are gathered into `Rb`
   a. For every `Ri` in `Rb`, if `Ri` is overridden by any member of `R`, it is skipped.
   b. Otherwise, if any `Ri` is hidden by a member of `R`, then the lookup of required members fails and no further steps are taken. Calling any constructor of `T` not attributed with `SetsRequiredMembers` issues an error.
   c. Otherwise, `Ri` is added to `R`.

Open Questions

Nested member initializers

What will the enforcement mechanisms for nested member initializers be? Will they be disallowed entirely?

C#

```
class Range
{
    public required Location Start { get; init; }
    public required Location End { get; init; }
}

class Location
{
    public required int Column { get; init; }
    public required int Line { get; init; }
}

_ = new Range { Start = { Column = 0, Line = 0 }, End = { Column = 1, Line =
0 } } // Would this be allowed if Location is a struct type?
_ = new Range { Start = new Location { Column = 0, Line = 0 }, End = new
Location { Column = 1, Line = 0 } } // Or would this form be necessary
instead?
```

Discussed Questions

Level of enforcement for `init` clauses

Do we strictly enforce that members specified in a `init` clause without an initializer must initialize all members? It seems likely that we do, otherwise we create an easy pit-of-failure. However, we also run the risk of reintroducing the same problems we solved with `MemberNotNull` in C# 9. If we want to strictly enforce this, we will likely need a way for a helper method to indicate that it sets a member. Some possible syntaxes we've discussed for this:

- Allow `init` methods. These methods are only allowed to be called from a constructor or from another `init` method, and can access `this` as if it's in the constructor (ie, set `readonly` and `init` fields/properties). This can be combined with `init` clauses on such methods. A `init` clause would be considered satisfied if the member in the clause is definitely assigned in the body of the method/constructor. Calling a method with a `init` clause that includes a member counts as assigning to that member. If we do decided that this is a route we want to pursue, now or in the future, it seems likely that we should not use `init` as the keyword for the init clause on a constructor, as that would be confusing.
- Allow the `!` operator to suppress the warning/error explicitly. If initializing a member in a complicated way (such as in a shared method), the user can add a `!` to the init clause to indicate the compiler should not check for initialization.

Conclusion: After discussion we like the idea of the `!` operator. It allows the user to be intentional about more complicated scenarios while also not creating a large design hole around init methods and annotating every method as setting members X or Y. `!` was chosen because we already use it for suppressing nullable warnings, and using it to tell the compiler "I'm smarter than you" in another place is a natural extension of the syntax form.

Required interface members

This proposal does not allow interfaces to mark members as required. This protects us from having to figure out complex scenarios around `new()` and interface constraints in generics right now, and is directly related to both factories and generic construction. In order to ensure that we have design space in this area, we forbid `required` in interfaces, and forbid types with *required_member_lists* from being substituted for type parameters

constrained to `new()`. When we want to take a broader look at generic construction scenarios with factories, we can revisit this issue.

Syntax questions

- Is `init` the right word? `init` as a postfix modifier on the constructor might interfere if we ever want to reuse it for factories and also enable `init` methods with a prefix modifier. Other possibilities:
  - `set`
- Is `required` the right modifier for specifying that all members are initialized? Others suggested:
  - `default`
  - `all`
  - With a ! to indicate complex logic
- Should we require a separator between the `base`/`this` and the `init`?
  - `:` separator
  - ',' separator
- Is `required` the right modifier? Other alternatives that have been suggested:
  - `req`
  - `require`
  - `mustinit`
  - `must`
  - `explicit`

Conclusion: We have removed the `init` constructor clause for now, and are proceeding with `required` as the property modifier.

Init clause restrictions

Should we allow access to `this` in the init clause? If we want the assignment in `init` to be a shorthand for assigning the member in the constructor itself, it seems like we should.

Additionally, does it create a new scope, like `base()` does, or does it share the same scope as the method body? This is particularly important for things like local functions, which the init clause may want to access, or for name shadowing, if an init expression introduces a variable via `out` parameter.

Conclusion: `init` clause has been removed.

Accessibility requirements and `init`

In versions of this proposal with the `init` clause, we talked about being able to have the following scenario:

```C#
public class Base
{
    protected required int _field;

    protected Base() {} // Contract required that _field is set
}
public class Derived : Base
{
    public Derived() : init(_field = 1) // Contract is fulfilled and _field
is removed from the required members list
    {
    }
}
```

However, we have removed the `init` clause from the proposal at this point, so we need to decide whether to allow this scenario in a limited fashion. The options we have are:

1. Disallow the scenario. This is the most conservative approach, and the rules in the Accessibility are currently written with this assumption in mind. The rule is that any member that is required must be at least as visible as its containing type.
2. Require that all constructors are either:
   a. No more visible than the least-visible required member.
   b. Have the `SetsRequiredMembersAttribute` applied to the constructor. These would ensure that anyone who can see a constructor can either set all the things it exports, or there is nothing to set. This could be useful for types that are only ever created via static `Create` methods or similar builders, but the utility seems overall limited.
3. Readd a way to remove specific parts of the contract to the proposal, as discussed in LDM previously.

Conclusion: Option 1, all required members must be at least as visible as their containing type.

Override rules

The current spec says that the `required` keyword needs to be copied over and that overrides can make a member *more* required, but not less. Is that what we want to do?

Allowing removal of requirements needs more contract modification abilities than we are currently proposing.

Conclusion: Adding `required` on override is allowed. If the overridden member is `required`, the overridding member must also be `required`.

Alternative metadata representation

We could also take a different approach to metadata representation, taking a page from extension methods. We could put a `RequiredMemberAttribute` on the type to indicate that the type contains required members, and then put a `RequiredMemberAttribute` on each member that is required. This would simplify the lookup sequence (no need to do member lookup, just look for members with the attribute).

Conclusion: Alternative approved.

Metadata Representation

The Metadata Representation needs to be approved. We additionally need to decide whether these attributes should be included in the BCL.

1. For `RequiredMemberAttribute`, this attribute is more akin to the general embedded attributes we use for nullable/nint/tuple member names, and will not be manually applied by the user in C#. It's possible that other languages might want to manually apply this attribute, however.
2. `SetsRequiredMembersAttribute`, on the other hand, is directly used by consumers, and thus should likely be in the BCL.

If we go with the alternative representation in the previous section, that might change the calculus on `RequiredMemberAttribute`: instead of being similar to the general embedded attributes for `nint`/nullable/tuple member names, it's closer to `System.Runtime.CompilerServices.ExtensionAttribute`, which has been in the framework since extension methods shipped.

Conclusion: We will put both attributes in the BCL.

Warning vs Error

Should not setting a required member be a warning or an error? It is certainly possible to trick the system, via `Activator.CreateInstance(typeof(C))` or similar, which means we may not be able to fully guarantee all properties are always set. We also allow

suppression of the diagnostics at the constructor-site by using the `!`, which we generally do not allow for errors. However, the feature is similar to readonly fields or init properties, in that we hard error if users attempt to set such a member after initialization, but they can be circumvented by reflection.

Conclusion: Errors.

AsyncMethodBuilder override

Article • 09/21/2021

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Allow per-method override of the async method builder to use. For some async methods we want to customize the invocation of `Builder.Create()` to use a different *builder type*.

C#

```
[AsyncMethodBuilderAttribute(typeof(PoolingAsyncValueTaskMethodBuilder<>))]
// new usage of AsyncMethodBuilderAttribute type
static async ValueTask<int> ExampleAsync() { ... }
```

Motivation

Today, async method builders are tied to a given type used as a return type of an async method. For example, any method that's declared as `async Task` uses `AsyncTaskMethodBuilder`, and any method that's declared as `async ValueTask<T>` uses `AsyncValueTaskMethodBuilder<T>`. This is due to the `[AsyncMethodBuilder(Type)]` attribute on the type used as a return type, e.g. `ValueTask<T>` is attributed as `[AsyncMethodBuilder(typeof(AsyncValueTaskMethodBuilder<>))]`. This addresses the majority common case, but it leaves a few notable holes for advanced scenarios.

In .NET 5, an experimental feature was shipped that provides two modes in which `AsyncValueTaskMethodBuilder` and `AsyncValueTaskMethodBuilder<T>` operate. The on-by-default mode is the same as has been there since the functionality was introduced: when

the state machine needs to be lifted to the heap, an object is allocated to store the state, and the async method returns a `ValueTask{<T>}` backed by a `Task{<T>}`. However, if an environment variable is set, all builders in the process switch to a mode where, instead, the `ValueTask{<T>}` instances are backed by reusable `IValueTaskSource{<T>}` implementations that are pooled. Each async method has its own pool with a fixed maximum number of instances allowed to be pooled, and as long as no more than that number are ever returned to the pool to be pooled at the same time, `async ValueTask<{T}>` methods effectively become free of any GC allocation overhead.

There are several problems with this experimental mode, however, which is both why a) it's off by default and b) we're likely to remove it in a future release unless very compelling new information emerges (https://github.com/dotnet/runtime/issues/13633).

- It introduces a behavioral difference for consumers of the returned `ValueTask{<T>}` if that `ValueTask` isn't being consumed according to spec. When it's backed by a `Task`, you can do with the `ValueTask` things you can do with a `Task`, like await it multiple times, await it concurrently, block waiting for it to complete, etc. But when it's backed by an arbitrary `IValueTaskSource`, such operations are prohibited, and automatically switching from the former to the latter can lead to bugs. With the switch at the process level and affecting all `async ValueTask` methods in the process, whether you control them or not, it's too big a hammer.
- It's not necessarily a performance win, and could represent a regression in some situations. The implementation is trading the cost of pooling (accessing a pool isn't free) with the cost of GC, and in various situations the GC can win. Again, applying the pooling to all `async ValueTask` methods in the process rather than being selective about the ones it would most benefit is too big a hammer.
- It adds to the IL size of a trimmed application, even if the flag isn't set, and then to the resulting asm size. It's possible that can be worked around with improvements to the implementation to teach it that for a given deployment the environment variable will always be false, but as it stands today, every `async ValueTask` method saw for example an ~2K binary footprint increase in aot images due to this option, and, again, that applies to all `async ValueTask` methods in the whole application closure.
- Different methods may benefit from differing levels of control, e.g. the size of the pool employed because of knowledge of the method and how it's used, but the same setting is applied to all uses of the builder. One could imagine working around that by having the builder code use reflection at runtime to look for some attribute, but that adds significant run-time expense, and likely on the startup path.

On top of all of these issues with the existing pooling, it's also the case that developers are prevented from writing their own customized builders for types they don't own. If, for example, a developer wants to implement their own pooling support, they also have to introduce a brand new task-like type, rather than just being able to use `{Value}Task{<T>}`, because the attribute specifying the builder is only specifiable on the type declaration of the return type.

We need a way to have an individual async method opt-in to a specific builder.

Detailed design

Using AsyncMethodBuilderAttribute on methods

In `dotnet/runtime`, add `AttributeTargets.Method` to the targets for `System.Runtime.CompilerServices.AsyncMethodBuilderAttribute`:

C#

```
namespace System.Runtime.CompilerServices
{
    /// <summary>
    /// Indicates the type of the async method builder that should be used
by a language compiler:
    /// - to build the return type of an async method that is attributed,
    /// - to build the attributed type when used as the return type of an
async method.
    /// </summary>
    [AttributeUsage(AttributeTargets.Method | AttributeTargets.Class |
AttributeTargets.Struct | AttributeTargets.Interface |
AttributeTargets.Delegate | AttributeTargets.Enum, Inherited = false,
AllowMultiple = false)]
    public sealed class AsyncMethodBuilderAttribute : Attribute
    {
        /// <summary>Initializes the <see
cref="AsyncMethodBuilderAttribute"/>.</summary>
        /// <param name="builderType">The <see cref="Type"/> of the
associated builder.</param>
        public AsyncMethodBuilderAttribute(Type builderType) => BuilderType
= builderType;

        /// <summary>Gets the <see cref="Type"/> of the associated builder.
</summary>
        public Type BuilderType { get; }
    }
}
```

This allows the attribute to be applied on methods or local functions or lambdas.

Example of usage on a method:

```C#
[AsyncMethodBuilder(typeof(PoolingAsyncValueTaskMethodBuilder<>))] // new usage, referring to some custom builder type
static async ValueTask<int> ExampleAsync() { ... }
```

It is an error to apply the attribute multiple times on a given method.
It is an error to apply the attribute to a lambda with an implicit return type.

A developer who wants to use a specific custom builder for all of their methods can do so by putting the relevant attribute on each method.

Determining the builder type for an async method

When compiling an async method, the builder type is determined by:

1. using the builder type from the `AsyncMethodBuilder` attribute if one is present,
2. otherwise, falling back to the builder type determined by previous approach. (see spec for task-like types).

If an `AsyncMethodBuilder` attribute is present, we take the builder type specified by the attribute and construct it if necessary.
If the override type is an open generic type, take the single type argument of the async method's return type and substitute it into the override type.
If the override type is a bound generic type, then we produce an error.
If the async method's return type does not have a single type argument, then we produce an error.

We verify that the builder type is compatible with the return type of the async method:

1. look for the public `Create` method with no type parameters and no parameters on the constructed builder type.
   It is an error if the method is not found. It is an error if the method returns a type other than the constructed builder type.
2. look for the public `Task` property.
   It is an error if the property is not found.
3. consider the type of that `Task` property (a task-like type):
   It is an error if the task-like type does not matches the return type of the async method.

Note that it is not necessary for the return type of the method to be a task-like type.

Execution

The builder type determined above is used as part of the existing async method design.

For example, today if a method is defined as:

```
C#

public async ValueTask<T> ExampleAsync() { ... }
```

the compiler will generate code akin to:

```
C#

[AsyncStateMachine(typeof(<ExampleAsync>d__29))]
[CompilerGenerated]
static ValueTask<int> ExampleAsync()
{
    <ExampleAsync>d__29 stateMachine;
    stateMachine.<>t__builder = AsyncValueTaskMethodBuilder<int>.Create();
    stateMachine.<>1__state = -1;
    stateMachine.<>t__builder.Start(ref stateMachine);
    return stateMachine.<>t__builder.Task;
}
```

With this change, if the developer wrote:

```
C#

[AsyncMethodBuilder(typeof(PoolingAsyncValueTaskMethodBuilder<>))] // new
usage, referring to some custom builder type
static async ValueTask<int> ExampleAsync() { ... }
```

it would instead be compiled to:

```
C#

[AsyncStateMachine(typeof(<ExampleAsync>d__29))]
[CompilerGenerated]
[AsyncMethodBuilder(typeof(PoolingAsyncValueTaskMethodBuilder<>))] //
retained but not necessary anymore
static ValueTask<int> ExampleAsync()
{
    <ExampleAsync>d__29 stateMachine;
    stateMachine.<>t__builder =
PoolingAsyncValueTaskMethodBuilder<int>.Create(); // <>t__builder now a
different type
    stateMachine.<>1__state = -1;
    stateMachine.<>t__builder.Start(ref stateMachine);
```

```
    return stateMachine.<>t__builder.Task;
}
```

Just those small additions enable:

- Anyone to write their own builder that can be applied to async methods that return `Task<T>` and `ValueTask<T>`
- As "anyone", the runtime to ship the experimental builder support as new public builder types that can be opted into on a method-by-method basis; the existing support would be removed from the existing builders. Methods (including some we care about in the core libraries) can then be attributed on a case-by-case basis to use the pooling support, without impacting any other unattributed methods.

and with minimal surface area changes or feature work in the compiler.

Note that we need the emitted code to allow a different type being returned from `Create` method:

```

AsyncPooledBuilder _builder = AsyncPooledBuilderWithSize4.Create();
```

Note that this mechanism to change the the builder type cannot be used when the synthesized entry-point for top-level statements is async. An explicit entry-point should be used instead.

Drawbacks

- The syntax for applying such an attribute to a method is verbose. The impact of this is lessened if a developer can apply it to multiple methods en mass, e.g. at the type or module level.

Alternatives

- Implement a different task-like type and expose that difference to consumers. `ValueTask` was made extensible via the `IValueTaskSource` interface to avoid that need, however.
- Address just the ValueTask pooling part of the issue by enabling the experiment as the on-by-default-and-only implementation. That doesn't address other aspects, such as configuring the pooling, or enabling someone else to provide their own builder.

- Earlier versions of this document allowed for scoped override of builder types.

Unresolved questions

1. **Replace or also create.** All of the examples in this proposal are about replacing a buildable task-like's builder. Should the feature be scoped to just that? Or should you be able to use this attribute on a method with a return type that doesn't already have a builder (e.g. some common interface)? That could impact overload resolution.
2. **Private Builders**. Should the compiler support non-public async method builders? This is not spec'd today, but experimentally we only support public ones. That makes some sense when the attribute is applied to a type to control what builder is used with that type, since anyone writing an async method with that type as the return type would need access to the builder. However, with this new feature, when that attribute is applied to a method, it only impacts the implementation of that method, and thus could reasonably reference a non-public builder. Likely we will want to support library authors who have non-public ones they want to use.

`params Collections`

Article • 05/16/2024

> **Note**
>
> This article is a feature specification. The specification serves as the design document for the feature. It includes proposed specification changes, along with information needed during the design and development of the feature. These articles are published until the proposed spec changes are finalized and incorporated in the current ECMA specification.
>
> There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent language design meeting (LDM) notes.
>
> You can learn more about the process for adopting feature speclets into the C# language standard in the article on the specifications.

Summary

In C# 12 language added support for creating instances of collection types beyond just arrays. See collection expressions. This proposal extends `params` support to all such collection types.

Motivation

A `params` array parameter provides a convenient way to call a method that takes an arbitrary length list of arguments. Today `params` parameter must be an array type. However, it might be beneficial for a developer to be able to have the same convenience when calling APIs that take other collection types. For example, an `ImmutableArray<T>`, `ReadOnlySpan<T>`, or plain `IEnumerable`. Especially in cases when compiler is able to avoid an implicit array allocation for the purpose of creating the collection (`ImmutableArray<T>`, `ReadOnlySpan<T>`, etc).

Today, in situations when an API takes a collection type, developers usually add `params` overload that takes an array, construct the target collection and call the original overload with that collection, thus consumers of the API have to trade an extra array allocation for convenience.

Another motivation is ability to add a params span overload and have it take precedence over the array version, just by recompiling existing source code.

Detailed design

Method parameters

The Method parameters section is adjusted as follows.

diff

```
formal_parameter_list
    : fixed_parameters
-    | fixed_parameters ',' parameter_array
+    | fixed_parameters ',' parameter_collection
-    | parameter_array
+    | parameter_collection
    ;

-parameter_array
+parameter_collection
-    : attributes? 'params' array_type identifier
+    : attributes? 'params' 'scoped'? type identifier
    ;
```

A *parameter_collection* consists of an optional set of *attributes*, a `params` modifier, an optional `scoped` modifier, a *type*, and an *identifier*. A parameter collection declares a single parameter of the given type with the given name. The *type* of a parameter collection shall be one of the following valid target types for a collection expression (see https://github.com/dotnet/csharplang/blob/main/proposals/csharp-12.0/collection-expressions.md#conversions):

- A single dimensional *array type* `T[]`, in which case the *element type* is `T`
- A *span type*
  - `System.Span<T>`
  - `System.ReadOnlySpan<T>`

    in which cases the *element type* is `T`
- A *type* with an appropriate *create method*, which is at least as accessible as the declaring member, and with a corresponding *element type* resulting from that determination
- A *struct* or *class type* that implements `System.Collections.IEnumerable` where:

    - The *type* has a constructor that can be invoked with no arguments, and the constructor is at least as accessible as the declaring member.
    - The *type* has an instance (not an extension) method `Add` where:
        - The method can be invoked with a single value argument.
        - If the method is generic, the type arguments can be inferred from the argument.
        - The method is at least as accessible as the declaring member.

    In which case the *element type* is the *iteration type* of the *type*.
- An *interface type*
    - `System.Collections.Generic.IEnumerable<T>`,
    - `System.Collections.Generic.IReadOnlyCollection<T>`,
    - `System.Collections.Generic.IReadOnlyList<T>`,
    - `System.Collections.Generic.ICollection<T>`,
    - `System.Collections.Generic.IList<T>`

    in which cases the *element type* is `T`

In a method invocation, a parameter collection permits either a single argument of the given parameter type to be specified, or it permits zero or more arguments of the collection's *element type* to be specified. Parameter collections are described further in *Parameter collections*.

A *parameter_collection* may occur after an optional parameter, but cannot have a default value – the omission of arguments for a *parameter_collection* would instead result in the creation of an empty collection.

Parameter collections

The Parameter arrays section is renamed and adjusted as follows.

A parameter declared with a `params` modifier is a parameter collection. If a formal parameter list includes a parameter collection, it shall be the last parameter in the list and it shall be of type specified in *Method parameters* section.

> *Note*: It is not possible to combine the `params` modifier with the modifiers `in`, `out`, or `ref`. *end note*

A parameter collection permits arguments to be specified in one of two ways in a method invocation:

- The argument given for a parameter collection can be a single expression that is implicitly convertible to the parameter collection type. In this case, the parameter collection acts precisely like a value parameter.
- Alternatively, the invocation can specify zero or more arguments for the parameter collection, where each argument is an expression that is implicitly convertible to the parameter collection's *element type*. In this case, the invocation creates an instance of the parameter collection type according to the rules specified in Collection expressions as though the arguments were used as expression elements in a collection expression in the same order, and uses the newly created collection instance as the actual argument. When constructing the collection instance, the original *unconverted* arguments are used.

Except for allowing a variable number of arguments in an invocation, a parameter collection is precisely equivalent to a value parameter of the same type.

When performing overload resolution, a method with a parameter collection might be applicable, either in its normal form or in its expanded form. The expanded form of a method is available only if the normal form of the method is not applicable and only if an applicable method with the same signature as the expanded form is not already declared in the same type.

A potential ambiguity arises between the normal form and the expanded form of the method with a single parameter collection argument when it can be used as the parameter collection itself and as the element of the parameter collection at the same time. The ambiguity presents no problem, however, since it can be resolved by inserting a cast or using a collection expression, if needed.

Signatures and overloading

All the rules around `params` modifier in Signatures and overloading remain as is.

Applicable function member

The Applicable function member section is adjusted as follows.

If a function member that includes a parameter collection is not applicable in its normal form, the function member might instead be applicable in its ***expanded form***:

- If parameter collection is not an array, an expanded form is not applicable for language versions C# 12 and below.
- The expanded form is constructed by replacing the parameter collection in the function member declaration with zero or more value parameters of the parameter

collection's *element type* such that the number of arguments in the argument list `A` matches the total number of parameters. If `A` has fewer arguments than the number of fixed parameters in the function member declaration, the expanded form of the function member cannot be constructed and is thus not applicable.

- Otherwise, the expanded form is applicable if for each argument in `A`, one of the following is true:
  - the parameter-passing mode of the argument is identical to the parameter-passing mode of the corresponding parameter, and
    - for a fixed value parameter or a value parameter created by the expansion, an implicit conversion exists from the argument expression to the type of the corresponding parameter, or
    - for an `in`, `out`, or `ref` parameter, the type of the argument expression is identical to the type of the corresponding parameter.
  - the parameter-passing mode of the argument is value, and the parameter-passing mode of the corresponding parameter is input, and an implicit conversion exists from the argument expression to the type of the corresponding parameter

Better function member

The Better function member section is adjusted as follows.

Given an argument list `A` with a set of argument expressions `{E₁, E₂, ..., Eᵥ}` and two applicable function members `Mᵥ` and `Mₓ` with parameter types `{P₁, P₂, ..., Pᵥ}` and `{Q₁, Q₂, ..., Qᵥ}`, `Mᵥ` is defined to be a ***better function member*** than `Mₓ` if

- for each argument, the implicit conversion from `Eᵥ` to `Qᵥ` is not better than the implicit conversion from `Eᵥ` to `Pᵥ`, and
- for at least one argument, the conversion from `Eᵥ` to `Pᵥ` is better than the conversion from `Eᵥ` to `Qᵥ`.

In case the parameter type sequences `{P₁, P₂, ..., Pᵥ}` and `{Q₁, Q₂, ..., Qᵥ}` are equivalent (i.e., each `Pᵢ` has an identity conversion to the corresponding `Qᵢ`), the following tie-breaking rules are applied, in order, to determine the better function member.

- If `Mᵢ` is a non-generic method and `Mₑ` is a generic method, then `Mᵢ` is better than `Mₑ`.
- Otherwise, if `Mᵢ` is applicable in its normal form and `Mₑ` has a params collection and is applicable only in its expanded form, then `Mᵢ` is better than `Mₑ`.

- Otherwise, if both methods have params collections and are applicable only in their expanded forms, and if the params collection of `Mᵢ` has fewer elements than the params collection of `Mₑ`, then `Mᵢ` is better than `Mₑ`.
- Otherwise, if `Mᵥ` has more specific parameter types than `Mₓ`, then `Mᵥ` is better than `Mₓ`. Let `{R1, R2, ..., Rn}` and `{S1, S2, ..., Sn}` represent the uninstantiated and unexpanded parameter types of `Mᵥ` and `Mₓ`. `Mᵥ`'s parameter types are more specific than `Mₓ`s if, for each parameter, `Rx` is not less specific than `Sx`, and, for at least one parameter, `Rx` is more specific than `Sx`:
  - A type parameter is less specific than a non-type parameter.
  - Recursively, a constructed type is more specific than another constructed type (with the same number of type arguments) if at least one type argument is more specific and no type argument is less specific than the corresponding type argument in the other.
  - An array type is more specific than another array type (with the same number of dimensions) if the element type of the first is more specific than the element type of the second.
- Otherwise if one member is a non-lifted operator and the other is a lifted operator, the non-lifted one is better.
- If neither function member was found to be better, and all parameters of `Mᵥ` have a corresponding argument whereas default arguments need to be substituted for at least one optional parameter in `Mₓ`, then `Mᵥ` is better than `Mₓ`.
- If for at least one parameter `Mᵥ` uses the ***better parameter-passing choice*** (§12.6.4.4) than the corresponding parameter in `Mₓ` and none of the parameters in `Mₓ` use the better parameter-passing choice than `Mᵥ`, `Mᵥ` is better than `Mₓ`.
- **Otherwise, if both methods have params collections and are applicable only in their expanded forms then `Mᵢ` is better than `Mₑ` if the same set of arguments corresponds to the collection elements for both methods, and one of the following holds (this corresponds to https://github.com/dotnet/csharplang/blob/main/proposals/csharp-12.0/collection-expressions.md#overload-resolution):**
  - **params collection of `Mᵢ` is `System.ReadOnlySpan<Eᵢ>`, and params collection of `Mₑ` is `System.Span<Eₑ>`, and an implicit conversion exists from `Eᵢ` to `Eₑ`**
  - **params collection of `Mᵢ` is `System.ReadOnlySpan<Eᵢ>` or `System.Span<Eᵢ>`, and params collection of `Mₑ` is an *array_or_array_interface_type* with *element type* `Eₑ`, and an implicit conversion exists from `Eᵢ` to `Eₑ`**
  - **both params collections are not *span_types*, and an implicit conversion exists from params collection of `Mᵢ` to params collection of `Mₑ`**
- Otherwise, no function member is better.

The reason why the new tie-breaking rule is placed at the end of the list is the last sub item

> - **both params collections are not *span_types*, and an implicit conversion exists from params collection of M_i to params collection of M_e**

it is applicable to arrays and, therefore, performing the tie-break earlier will introduce a behavior change for existing scenarios.

For example:

C#

```
class Program
{
    static void Main()
    {
        Test(1);
    }

    static void Test(in int x, params C2[] y) {} // There is an implicit
conversion from `C2[]` to `C1[]`
    static void Test(int x, params C1[] y) {} // Better candidate because of
"better parameter-passing choice"
}

class C1 {}
class C2 : C1 {}
```

If any of the previous tie-breaking rules apply (including the "better arguments conversions" rule), the overload resolution result can be different by comparison to the case when an explicit collection expression is used as an argument instead.

For example:

C#

```
class Program
{
    static void Test1()
    {
        M1(['1', '2', '3']); // Span overload is used
        M1('1', '2', '3');   // IEnumerable<char> overload is used because
`char` is an exact match
    }

    static void M1(params IEnumerable<char> value) {}
    static void M1(params System.ReadOnlySpan<MyChar> value) {}
```

```
    class MyChar
    {
        private readonly int _i;
        public MyChar(int i) { _i = i; }
        public static implicit operator MyChar(int i) => new MyChar(i);
        public static implicit operator char(MyChar c) => (char)c._i;
    }

    static void Test2()
    {
        M2([1]); // Span overload is used
        M2(1);   // Array overload is used, not generic
    }

    static void M2<T>(params System.Span<T> y){}
    static void M2(params int[] y){}

    static void Test3()
    {
        M3("3", ["4"]); // Span overload is used, better on the first
argument conversion, none is better on the second
        M3("3", "4");   // Ambiguity, better-ness of argument conversions
goes in opposite directions.
                        // Since parameter types are different ("object,
string" vs. "string, object"), tie-breaking rules do not apply
    }

    static void M3(object x, params string[] y) {}
    static void M3(string x, params Span<object> y) {}
}
```

However, our primary concern are scenarios where overloads differ only by params collection type, but the collection types have the same element type. The behavior should be consistent with explicit collection expressions for these cases.

The "**if the same set of arguments corresponds to the collection elements for both methods**" condition is important for scenarios like:

C#

```
class Program
{
    static void Main()
    {
        Test(x: 1, y: 2); // Ambiguous
    }

    static void Test(int x, params System.ReadOnlySpan<int> y) {}
    static void Test(int y, params System.Span<int> x) {}
}
```

It doesn't feel reasonable to "compare" collections that are built from different elements.

> This section was reviewed at LDM and was approved.

Dynamic vs. Static Binding

Expanded forms of candidates utilizing non-array params collections won't be considered as valid candidates by the current C# runtime binder.

Recap of the current rules

From Static and Dynamic Binding:

> When no dynamic expressions are involved, C# defaults to static binding, which means that the compile-time types of subexpressions are used in the selection process. However, when one of the subexpressions in the operations listed above is a dynamic expression, **the operation is instead dynamically bound**.
>
> It is a compile time error if a method invocation is dynamically bound and any of the parameters, including the receiver, has the `in` modifier.

However, an exception to this rule exists for local functions. From Local function declarations:

> If the type of the argument to a local function is `dynamic`, **the function to be called must be resolved at compile time, not runtime**.

From Compile-time checking of dynamic member invocation:

> Even though overload resolution of a dynamically bound operation takes place at run-time, it is sometimes possible at compile-time to know the list of function members from which an overload will be chosen:
>
> - For a delegate invocation (§12.8.9.4), the list is a single function member with the same parameter list as the *delegate_type* of the invocation
> - For a method invocation (§12.8.9.2) on a type, or on a value whose static type is not dynamic, the set of accessible methods in the method group is known at compile-time.
> - For an object creation expression (§12.8.16.2) the set of accessible constructors in the type is known at compile-time.

> - For an indexer access (§12.8.11.3) the set of accessible indexers in the receiver is known at compile-time.
>
> In these cases a limited compile-time check is performed on each member in the known set of function members, to see if it can be known for certain never to be invoked at run-time. For each function member `F` a modified parameter and argument list are constructed:
>
> - First, if `F` is a generic method and type arguments were provided, then those are substituted for the type parameters in the parameter list. However, if type arguments were not provided, no such substitution happens.
> - Then, any parameter whose type is open (i.e., contains a type parameter; see §8.4.3) is elided, along with its corresponding parameter(s).
>
> For `F` to pass the check, all of the following shall hold:
>
> - The modified parameter list for `F` is applicable to the modified argument list in terms of §12.6.4.2.
> - All constructed types in the modified parameter list satisfy their constraints (§8.4.5).
> - If the type parameters of `F` were substituted in the step above, their constraints are satisfied.
> - If `F` is a static method, the method group shall not have resulted from a *member_access* whose receiver is known at compile-time to be a variable or value.
> - If `F` is an instance method, the method group shall not have resulted from a *member_access* whose receiver is known at compile-time to be a type.
>
> If no candidate passes this test, a compile-time error occurs.

From Invocation expressions:

> An *invocation_expression* is dynamically bound (§12.3.3) if at least one of the following holds:
>
> - The *primary_expression* has compile-time type `dynamic`.
> - At least one argument of the optional *argument_list* has compile-time type `dynamic`.
>
> In this case, the compiler classifies the *invocation_expression* as a value of type `dynamic`. The rules below to determine the meaning of the *invocation_expression* are then applied at run-time, using the run-time type instead of the compile-time type

> of those of the *primary_expression* and arguments that have the compile-time type `dynamic`. If the *primary_expression* does not have compile-time type `dynamic`, then the method invocation undergoes a limited compile-time check as described in §12.6.5 Compile-time checking of dynamic member invocation.

Similar wording exists for Element access and Object creation expressions.

New rules

As mentioned above, there is already an exception around binding invocations of local functions involving dynamic arguments. A comment in the code implies that the exception is needed due to limitations of C# runtime binder. Apparently it cannot handle invocations of local functions due to the way compiler emits them.
Here are the precise rules for these cases:

- Invocations of local functions are bound statically
- If local function is generic and its type arguments are not specified explicitly, an error is reported
- If there is an ambiguity between normal and expanded forms of the function that cannot be resolved at compile time, an error is reported. Such ambiguity occurs when a single argument corresponds to params parameter, and the argument has type dynamic.

New rules generalize and expand this behavior to Invocation expressions, Element access, and Object creation expressions

If the *primary_expression* does not have compile-time type `dynamic`, then the method invocation undergoes a limited compile-time check as described in §12.6.5 Compile-time checking of dynamic member invocation.

If no candidate passes this test, a compile-time error occurs.

If only a single candidate passes the test, the invocation of the candidate is statically bound when all the following conditions are met:

- the candidate is either not generic, or its type arguments are explicitly specified;
- there is no ambiguity between normal and expanded forms of the candidate that cannot be resolved at compile time.

Otherwise, the *invocation_expression* is dynamically bound.

- If only a single candidate passed the test above:
  - if that candidate is a local function, a compile-time error occurs;

    - if that candidate has a non-array params parameter, a compile-time error occurs.
- Otherwise, if any candidate passing the test has non-array params parameter and it could possibly be applicable only in an expanded form, a compile-time warning occurs.

> The proposed new rules were reviewed at LDM and were approved.

We also should consider reverting/fixing spec violation that affects local functions today, see https://github.com/dotnet/roslyn/issues/71399 .

> LDM confirmed that we want to fix this spec violation.

Order of evaluation with non-array collections in non-trivial scenarios

> This section was reviewed at LDM and was approved. Despite the fact that array cases deviate from other collections, the official language specification doesn't have to specify different rules for arrays. The deviations could simply be treated as an implementation artifact. At the same time we do not intend to change the existing behavior around arrays.

Named arguments

A collection instance is created and populated after the lexically previous argument is evaluated, but before the lexically following argument is evaluated.

For example:

C#

```
class Program
{
    static void Main()
    {
        Test(b: GetB(), c: GetC(), a: GetA());
    }

    static void Test(int a, int b, params MyCollection c) {}

    static int GetA() => 0;
    static int GetB() => 0;
    static int GetC() => 0;
}
```

The order of evaluation is the following:

1. `GetB` is called
2. `MyCollection` is created and populated, `GetC` is called in the process
3. `GetA` is called
4. `Test` is called

Note, in params array case, the array is created right before the target methos is invoked, after all arguments are evaluated in their lexical order.

Compound assignment

A collection instance is created and populated after the lexically previous index is evaluated, but before the lexically following index is evaluated. The instance is used to invoke getter and setter of the target indexer.

For example:

C#

```
class Program
{
    static void Test(Program p)
    {
        p[GetA(), GetC()]++;
    }

    int this[int a, params MyCollection c] { get => 0; set {} }

    static int GetA() => 0;
    static int GetC() => 0;
}
```

The order of evaluation is the following:

1. `GetA` is called and cached
2. `MyCollection` is created, populated and cached, `GetC` is called in the process
3. Indexer's getter is invoked with cached values for indexes
4. Result is incremented
5. Indexer's setter is invoked with cached values for indexes and the result of the increment

An example with an empty collection:

C#

```
class Program
{
    static void Test(Program p)
    {
        p[GetA()]++;
    }

    int this[int a, params MyCollection c] { get => 0; set {} }

    static int GetA() => 0;
}
```

The order of evaluation is the following:

1. `GetA` is called and cached
2. An empty `MyCollection` is created and cached
3. Indexer's getter is invoked with cached values for indexes
4. Result is incremented
5. Indexer's setter is invoked with cached values for indexes and the result of the increment

Object Initializer

A collection instance is created and populated after the lexically previous index is evaluated, but before the lexically following index is evaluated. The instance is used to invoke indexer's getter as many times as necessary, if any.

For example:

C#

```
class C1
{
    public int F1;
    public int F2;
}

class Program
{
    static void Test()
    {
        _ = new Program() { [GetA(), GetC()] = { F1 = GetF1(), F2 = GetF2()
} };
    }

    C1 this[int a, params MyCollection c] => new C1();

    static int GetA() => 0;
```

```
    static int GetC() => 0;
    static int GetF1() => 0;
    static int GetF2() => 0;
}
```

The order of evaluation is the following:

1. `GetA` is called and cached
2. `MyCollection` is created, populated and cached, `GetC` is called in the process
3. Indexer's getter is invoked with cached values for indexes
4. `GetF1` is evaluated and assigned to `F1` field of `C1` retuned on the previous step
5. Indexer's getter is invoked with cached values for indexes
6. `GetF2` is evaluated and assigned to `F2` field of `C1` retuned on the previous step

Note, in params array case, its elements are evaluated and cached, but a new instance of an array (with the same values inside) is used for each invocation of indexer's getter instead. For the example above, the order of evaluation is the following:

1. `GetA` is called and cached
2. `GetC` is called and cached
3. Indexer's getter is invoked with cached `GetA` and new array populated with cached `GetC`
4. `GetF1` is evaluated and assigned to `F1` field of `C1` retuned on the previous step
5. Indexer's getter is invoked with cached `GetA` and new array populated with cached `GetC`
6. `GetF2` is evaluated and assigned to `F2` field of `C1` retuned on the previous step

An example with an empty collection:

C#

```
class C1
{
    public int F1;
    public int F2;
}

class Program
{
    static void Test()
    {
        _ = new Program() { [GetA()] = { F1 = GetF1(), F2 = GetF2() } };
    }

    C1 this[int a, params MyCollection c] => new C1();
```

```
    static int GetA() => 0;
    static int GetF1() => 0;
    static int GetF2() => 0;
}
```

The order of evaluation is the following:

1. `GetA` is called and cached
2. An empty `MyCollection` is created and cached
3. Indexer's getter is invoked with cached values for indexes
4. `GetF1` is evaluated and assigned to `F1` field of `C1` retuned on the previous step
5. Indexer's getter is invoked with cached values for indexes
6. `GetF2` is evaluated and assigned to `F2` field of `C1` retuned on the previous step

Ref safety

The collection expressions ref safety section is applicable to the construction of parameter collections when APIs are invoked in their expanded form.

Params parameters are implicitly `scoped` when their type is a ref struct. UnscopedRefAttribute can be used to override that.

Metadata

In metadata we could mark non-array `params` parameters with `System.ParamArrayAttribute`, as `params` arrays are marked today. However, it looks like we will be much safer to use a different attribute for non-array `params` parameters. For example, the current VB compiler will not be able to consume them decorated with `ParamArrayAttribute` neither in normal, nor in expanded form. Therefore, an addition of 'params' modifier is likely to break VB consumers, and very likely consumers from other languages or tools.

Given that, non-array `params` parameters are marked with a new `System.Runtime.CompilerServices.ParamCollectionAttribute`.

```
namespace System.Runtime.CompilerServices
{
    [AttributeUsage(AttributeTargets.Parameter, Inherited = true,
AllowMultiple = false)]
    public sealed class ParamCollectionAttribute : Attribute
    {
        public ParamCollectionAttribute() { }
```

```
    }
}
```

> This section was reviewed at LDM and was approved.

Open questions

Stack allocations

Here is a quote from https://github.com/dotnet/csharplang/blob/main/proposals/csharp-12.0/collection-expressions.md#unresolved-questions: "Stack allocations for huge collections might blow the stack. Should the compiler have a heuristic for placing this data on the heap? Should the language be unspecified to allow for this flexibility? We should follow the spec for params Span<T>." It sounds like we have to answer the questions in context of this proposal.

[Resolved] Implicitly `scoped` params

There was a suggestion that, when `params` modifies a `ref struct` parameter, it should be considered as declared `scoped`. The argument is made that number of cases where you want the parameter to be scoped is virtually 100% when looking through the BCL cases. In a few cases that need that, the default could be overwritten with `[UnscopedRef]`.

However, it might be undesirable to change the default simply based on presence of `params` modifier. Especially, that in overrides/implements scenarios `params` modifier doesn't have to match.

Resolution:

Params parameters are implicitly scoped - https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-11-15.md#params-improvements.

[Resolved] Consider enforcing `scoped` or `params` across overrides

We've previously stated that `params` parameters should be `scoped` by default. However, this introduces odd behavior in overridding, due to our existing rules around restating `params`:

```cs
class Base
{
    internal virtual Span<int> M1(scoped Span<int> s1, params Span<int> s2)
=> throw null!;
}

class Derived : Base
{
    internal override Span<int> M1(Span<int> s1, // Error, missing `scoped`
on override
                                   Span<int> s2  // No error: parameter is
implicitly params, and therefore implicitly scoped
                                  ) => throw null!;
}
```

We have a difference in behavior between carrying the `params` and carrying the `scoped` across overrides here: `params` is inherited implicitly, and with it `scoped`, while `scoped` by itself is *not* inherited implicitly and must be repeated at every level.

Proposal: We should enforce that overrides of `params` parameters must explicitly state `params` or `scoped` if the original definition is a `scoped` parameter. In other words, `s2` in `Derived` must have `params`, `scoped`, or both.

Resolution:

We will require explicitly stating `scoped` or `params` on override of a `params` parameter when a non-`params` parameter would be required to do so - https://github.com/dotnet/csharplang/blob/main/meetings/2024/LDM-2024-02-21.md#params-and-scoped-across-overrides .

[Resolved] Should presence of required members prevent declaration of `params` parameter?

Consider the following example:

```C#
using System.Collections;
using System.Collections.Generic;
```

```
public class MyCollection1 : IEnumerable<long>
{
    IEnumerator<long> IEnumerable<long>.GetEnumerator() => throw null;
    IEnumerator IEnumerable.GetEnumerator() => throw null;
    public void Add(long l) => throw null;

    public required int F; // Collection has required member and constructor
doesn't initialize it explicitly
}

class Program
{
    static void Main()
    {
        Test(2, 3); // error CS9035: Required member 'MyCollection1.F' must
be set in the object initializer or attribute constructor.
    }

    // Should an error be reported for the parameter indicating that the
constructor that is required
    // to be available doesn't initialize required members? In other words,
should one be able
    // to declare such a parameter under the specified conditions?
    static void Test(params MyCollection1 a)
    {
    }
}
```

Resolution:

We will validate `required` members against the constructor that is used to determine eligibility to be a `params` parameter at the declaration site - https://github.com/dotnet/csharplang/blob/main/meetings/2024/LDM-2024-02-21.md#required-members-and-params-parameters.

Alternatives

There is an alternative proposal that extends `params` only for `ReadOnlySpan<T>`.

Also, one might say, that with collection expressions now in the language, there is no need to extend `params` support at all. For any collection type. To consume an API with collection type, a developer simply needs to add two characters, `[` before the expanded list of arguments, and `]` after it. Given that, extending `params` support might be an overkill, especially that other languages are unlikely to support consumption of non-array `params` parameters any time soon.

Related proposals

- https://github.com/dotnet/csharplang/issues/1757
- https://github.com/dotnet/csharplang/blob/main/proposals/format.md#extending-params

Related design meetings

- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-11-15.md#params-improvements
- https://github.com/dotnet/csharplang/blob/main/meetings/2024/LDM-2024-01-08.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2024/LDM-2024-01-10.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2024/LDM-2024-01-29.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2024/LDM-2024-01-31.md#params-collections-evaluation-orders
- https://github.com/dotnet/csharplang/blob/main/meetings/2024/LDM-2024-02-21.md#params-collections

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# feature specifications feedback

C# feature specifications is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

Readonly Instance Members

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Championed Issue: https://github.com/dotnet/csharplang/issues/1710

Summary

Provide a way to specify individual instance members on a struct do not modify state, in the same way that `readonly struct` specifies no instance members modify state.

It is worth noting that `readonly instance member` != `pure instance member`. A `pure` instance member guarantees no state will be modified. A `readonly` instance member only guarantees that instance state will not be modified.

All instance members on a `readonly struct` could be considered implicitly `readonly instance members`. Explicit `readonly instance members` declared on non-readonly structs would behave in the same manner. For example, they would still create hidden copies if you called an instance member (on the current instance or on a field of the instance) which was itself not-readonly.

Motivation

Today, users have the ability to create `readonly struct` types which the compiler enforces that all fields are readonly (and by extension, that no instance members modify the state). However, there are some scenarios where you have an existing API that exposes accessible fields or that has a mix of mutating and non-mutating members. Under these circumstances, you cannot mark the type as `readonly` (it would be a breaking change).

This normally doesn't have much impact, except in the case of `in` parameters. With `in` parameters for non-readonly structs, the compiler will make a copy of the parameter for each instance member invocation, since it cannot guarantee that the invocation does not modify internal state. This can lead to a multitude of copies and worse overall performance than if you had just passed the struct directly by value. For an example, see this code on sharplab

Some other scenarios where hidden copies can occur include `static readonly fields` and `literals`. If they are supported in the future, `blittable constants` would end up in the same boat; that is they all currently necessitate a full copy (on instance member invocation) if the struct is not marked `readonly`.

Design

Allow a user to specify that an instance member is, itself, `readonly` and does not modify the state of the instance (with all the appropriate verification done by the compiler, of course). For example:

C#

```
public struct Vector2
{
    public float x;
    public float y;

    public readonly float GetLengthReadonly()
    {
        return MathF.Sqrt(LengthSquared);
    }

    public float GetLength()
    {
        return MathF.Sqrt(LengthSquared);
    }

    public readonly float GetLengthIllegal()
    {
        var tmp = MathF.Sqrt(LengthSquared);

        x = tmp;    // Compiler error, cannot write x
        y = tmp;    // Compiler error, cannot write y

        return tmp;
    }

    public readonly float LengthSquared
    {
        get
```

```
        {
            return (x * x) +
                   (y * y);
        }
    }
}

public static class MyClass
{
    public static float ExistingBehavior(in Vector2 vector)
    {
        // This code causes a hidden copy, the compiler effectively emits:
        //    var tmpVector = vector;
        //    return tmpVector.GetLength();
        //
        // This is done because the compiler doesn't know that `GetLength()`
        // won't mutate `vector`.

        return vector.GetLength();
    }

    public static float ReadonlyBehavior(in Vector2 vector)
    {
        // This code is emitted exactly as listed. There are no hidden
        // copies as the `readonly` modifier indicates that the method
        // won't mutate `vector`.

        return vector.GetLengthReadonly();
    }
}
```

Readonly can be applied to property accessors to indicate that `this` will not be mutated in the accessor. The following examples have readonly setters because those accessors modify the state of member field, but do not modify the value of that member field.

C#

```
public readonly int Prop1
{
    get
    {
        return this._store["Prop1"];
    }
    set
    {
        this._store["Prop1"] = value;
    }
}
```

When `readonly` is applied to the property syntax, it means that all accessors are `readonly`.

C#

```
public readonly int Prop2
{
    get
    {
        return this._store["Prop2"];
    }
    set
    {
        this._store["Prop2"] = value;
    }
}
```

Readonly can only be applied to accessors which do not mutate the containing type.

C#

```
public int Prop3
{
    readonly get
    {
        return this._prop3;
    }
    set
    {
        this._prop3 = value;
    }
}
```

Readonly can be applied to some auto-implemented properties, but it won't have a meaningful effect. The compiler will treat all auto-implemented getters as readonly whether or not the `readonly` keyword is present.

C#

```
// Allowed
public readonly int Prop4 { get; }
public int Prop5 { readonly get; }
public int Prop6 { readonly get; set; }

// Not allowed
public readonly int Prop7 { get; set; }
public int Prop8 { get; readonly set; }
```

Readonly can be applied to manually-implemented events, but not field-like events. Readonly cannot be applied to individual event accessors (add/remove).

C#

```
// Allowed
public readonly event Action<EventArgs> Event1
{
    add { }
    remove { }
}

// Not allowed
public readonly event Action<EventArgs> Event2;
public event Action<EventArgs> Event3
{
    readonly add { }
    readonly remove { }
}
public static readonly event Event4
{
    add { }
    remove { }
}
```

Some other syntax examples:

- Expression bodied members: `public readonly float ExpressionBodiedMember => (x * x) + (y * y);`
- Generic constraints: `public readonly void GenericMethod<T>(T value) where T : struct { }`

The compiler would emit the instance member, as usual, and would additionally emit a compiler recognized attribute indicating that the instance member does not modify state. This effectively causes the hidden `this` parameter to become `in T` instead of `ref T`.

This would allow the user to safely call said instance method without the compiler needing to make a copy.

The restrictions would include:

- The `readonly` modifier cannot be applied to static methods, constructors or destructors.
- The `readonly` modifier cannot be applied to delegates.
- The `readonly` modifier cannot be applied to members of class or interface.

Drawbacks

Same drawbacks as exist with `readonly struct` methods today. Certain code may still cause hidden copies.

Notes

Using an attribute or another keyword may also be possible.

This proposal is somewhat related to (but is more a subset of) `functional purity` and/or `constant expressions`, both of which have had some existing proposals.

Parameterless struct constructors

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Support parameterless constructors and instance field initializers for struct types.

Motivation

Explicit parameterless constructors would give more control over minimally constructed instances of the struct type. Instance field initializers would allow simplified initialization across multiple constructors. Together these would close an obvious gap between `struct` and `class` declarations.

Support for field initializers would also allow initialization of fields in `record struct` declarations without explicitly implementing the primary constructor.

C#

```
record struct Person(string Name)
{
    public object Id { get; init; } = GetNextId();
}
```

If struct field initializers are supported for constructors with parameters, it seems natural to extend that to parameterless constructors as well.

C#

```
record struct Person()
{
    public string Name { get; init; }
    public object Id { get; init; } = GetNextId();
}
```

Proposal

Instance field initializers

Instance field declarations for a struct may include initializers.

As with class field initializers §14.5.6.3:

> A variable initializer for an instance field cannot reference the instance being created.

An error is reported if a struct has field initializers and no declared instance constructors since the field initializers will not be run.

C#

```
struct S { int F = 42; } // error: 'struct' with field initializers must include an explicitly declared constructor
```

Constructors

A struct may declare a parameterless instance constructor.

A parameterless instance constructor is valid for all struct kinds including `struct`, `readonly struct`, `ref struct`, and `record struct`.

If no parameterless instance constructor is declared, the struct (see §15.4.9) ...

> implicitly has a parameterless instance constructor which always returns the value that results from setting all value type fields to their default value and all reference type fields to null.

Modifiers

A parameterless instance struct constructor must be declared `public`.

C#

```
struct S0 { }                   // ok
struct S1 { public S1() { } }   // ok
struct S2 { internal S2() { } } // error: parameterless constructor must be
'public'
```

Non-public constructors are ignored when importing types from metadata.

Constructors can be declared `extern` or `unsafe`. Constructors cannot be `partial`.

Executing field initializers

Instance variable initializers (§14.11.3) is **modified** as follows:

> When **a class** instance constructor has no constructor initializer, or it has a constructor initializer of the form `base(...)`, that constructor implicitly performs the initializations specified by the *variable_initializer*s of the instance fields declared in its class. This corresponds to a sequence of assignments that are executed immediately upon entry to the constructor and before the implicit invocation of the direct base class constructor.
>
> **When a struct instance constructor has no constructor initializer, that constructor implicitly performs the initializations specified by the *variable_initializer*s of the instance fields declared in its struct. This corresponds to a sequence of assignments that are executed immediately upon entry to the constructor.**
>
> **When a struct instance constructor has a `this()` constructor initializer that represents the *default parameterless constructor*, the declared constructor implicitly clears all instance fields and performs the initializations specified by the *variable_initializer*s of the instance fields declared in its struct. Immediately upon entry to the constructor, all value type fields are set to their default value and all reference type fields are set to `null`. Immediately after that, a sequence of assignments corresponding to the *variable_initializer*s are executed.**

Definite assignment

Instance fields (other than `fixed` fields) must be definitely assigned in struct instance constructors that do not have a `this()` initializer (see §15.4.9).

C#

```
struct S0 // ok
{
    int x;
    object y;
}

struct S1 // error: 'struct' with field initializers must include an
explicitly declared constructor
{
    int x = 1;
    object y;
}

struct S2
{
    int x = 1;
    object y;
    public S2() { } // error: field 'y' must be assigned
}

struct S3 // ok
{
    int x = 1;
    object y;
    public S3() { y = 2; }
}
```

No `base()` initializer

A `base()` initializer is disallowed in struct constructors.

The compiler will not emit a call to the base `System.ValueType` constructor from struct instance constructors.

`record struct`

An error is reported if a `record struct` has field initializers and does not contain a primary constructor nor any instance constructors since the field initializers will not be run.

C#

```
record struct R0;                   // ok
record struct R1 { int F = 42; }    // error: 'struct' with field
initializers must include an explicitly declared constructor
record struct R2() { int F = 42; } // ok
record struct R3(int F);            // ok
```

A `record struct` with an empty parameter list will have a parameterless primary constructor.

```
C#

record struct R3();                // primary .ctor: public R3() { }
record struct R4() { int F = 42; } // primary .ctor: public R4() { F = 42; }
```

An explicit parameterless constructor in a `record struct` must have a `this` initializer that calls the primary constructor or an explicitly declared constructor.

```
C#

record struct R5(int F)
{
    public R5() { }                  // error: must have 'this' initializer
that calls explicit .ctor
    public R5(object o) : this() { } // ok
    public int F =  F;
}
```

Fields

The implicitly-defined parameterless constructor will zero fields rather than calling any parameterless constructors for the field types. No warnings are reported that field constructors are ignored. *No change from C#9.*

```
C#

struct S0
{
    public S0() { }
}

struct S1
{
    S0 F; // S0 constructor ignored
}

struct S<T> where T : struct
{
    T F; // constructor (if any) ignored
}
```

`default` expression

`default` ignores the parameterless constructor and generates a zeroed instance. *No change from C#9.*

```
C#

// struct S { public S() { } }

_ = default(S); // constructor ignored, no warning
```

`new()`

Object creation invokes the parameterless constructor if public; otherwise the instance is zeroed. *No change from C#9.*

```
C#

// public struct PublicConstructor { public PublicConstructor() { } }
// public struct PrivateConstructor { private PrivateConstructor() { } }

_ = new PublicConstructor();  // call PublicConstructor::.ctor()
_ = new PrivateConstructor(); // initobj PrivateConstructor
```

A warning wave may report a warning for use of `new()` with a struct type that has constructors but no parameterless constructor. No warning will be reported when using substituting such a struct type for a type parameter with a `new()` or `struct` constraint.

```
C#

struct S { public S(int i) { } }
static T CreateNew<T>() where T : new() => new T();

_ = new S();        // warning: no constructor called
_ = CreateNew<S>(); // ok
```

Uninitialized values

A local or field of a struct type that is not explicitly initialized is zeroed. The compiler reports a definite assignment error for an uninitialized struct that is not empty. *No change from C#9.*

```
C#

NoConstructor s1;
PublicConstructor s2;
```

```
s1.ToString(); // error: use of unassigned local (unless type is empty)
s2.ToString(); // error: use of unassigned local (unless type is empty)
```

Array allocation

Array allocation ignores any parameterless constructor and generates zeroed elements. *No change from C#9.*

C#

```
// struct S { public S() { } }

var a = new S[1]; // constructor ignored, no warning
```

Parameter default value `new()`

A parameter default value of `new()` binds to the parameterless constructor if public (and reports an error that the value is not constant); otherwise the instance is zeroed. *No change from C#9.*

C#

```
// public struct PublicConstructor { public PublicConstructor() { } }
// public struct PrivateConstructor { private PrivateConstructor() { } }

static void F1(PublicConstructor s1 = new()) { }  // error: default value
must be constant
static void F2(PrivateConstructor s2 = new()) { } // ok: initobj
```

Type parameter constraints: `new()` and `struct`

The `new()` and `struct` type parameter constraints require the parameterless constructor to be `public` if defined (see Satisfying constraints - §8.4.5).

The compiler assumes all structs satisfy `new()` and `struct` constraints. *No change from C#9.*

C#

```
// public struct PublicConstructor { public PublicConstructor() { } }
// public struct InternalConstructor { internal InternalConstructor() { } }

static T CreateNew<T>() where T : new() => new T();
static T CreateStruct<T>() where T : struct => new T();
```

```
_ = CreateNew<PublicConstructor>();      // ok
_ = CreateStruct<PublicConstructor>();   // ok

_ = CreateNew<InternalConstructor>();    // compiles; may fail at runtime
_ = CreateStruct<InternalConstructor>(); // compiles; may fail at runtime
```

`new T()` is emitted as a call to `System.Activator.CreateInstance<T>()`, and the compiler assumes the implementation of `CreateInstance<T>()` invokes the `public` parameterless constructor if defined.

With .NET Framework, `Activator.CreateInstance<T>()` invokes the parameterless constructor if the constraint is `where T : new()` but appears to ignore the parameterless constructor if the constraint is `where T : struct`.

Optional parameters

Constructors with optional parameters are not considered parameterless constructors. *No change from C#9.*

C#

```
struct S1 { public S1(string s = "") { } }
struct S2 { public S2(params object[] args) { } }

_ = new S1(); // ok: ignores constructor
_ = new S2(); // ok: ignores constructor
```

Metadata

Explicit parameterless struct instance constructors will be emitted to metadata.

Public parameterless struct instance constructors will be imported from metadata; non-public struct instance constructors will be ignored. *No change from C#9.*

See also

- https://github.com/dotnet/roslyn/issues/1029

Design meetings

- https://github.com/dotnet/csharplang/blob/main/meetings/2021/LDM-2021-04-28.md#open-questions-in-record-and-parameterless-structs
- https://github.com/dotnet/csharplang/blob/main/meetings/2021/LDM-2021-03-10.md#parameterless-struct-constructors
- https://github.com/dotnet/csharplang/blob/main/meetings/2021/LDM-2021-01-27.md#field-initializers

Auto-default structs

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

https://github.com/dotnet/csharplang/issues/5737

Summary

This feature makes it so that in struct constructors, we identify fields which were not explicitly assigned by the user before returning or before use, and initialize them implicitly to `default` instead of giving definite assignment errors.

Motivation

This proposal is raised as a possible mitigation for usability issues found in dotnet/csharplang#5552 and dotnet/csharplang#5635, as well as addressing #5563 (all fields must be definitely assigned, but `field` is not accessible within the constructor).

Since C# 1.0, struct constructors have been required to definitely assign `this` as if it were an `out` parameter.

C#

```
public struct S
{
    public int x, y;
    public S() // error: Fields 'S.x' and 'S.y' must be fully assigned
before control is returned to the caller
    {
    }
}
```

This presents issues when setters are manually defined on semi-auto properties, since the compiler can't treat assignment of the property as equivalent to assignment of the backing field.

C#

```
public struct S
{
    public int X { get => field; set => field = value; }
    public S() // error: struct fields aren't fully assigned. But caller can
only assign 'this.field' by assigning 'this'.
    {
    }
}
```

We assume that introducing finer-grained restrictions for setters, such as a scheme where the setter doesn't take `ref this` but rather takes `out field` as a parameter, is going to be too niche and incomplete for some use cases.

One fundamental tension we are struggling with is that when struct properties have manually implemented setters, users often have to do some form of "repetition" of either repeatedly assigning or repeating their logic:

C#

```
struct S
{
    private int _x;
    public int X
    {
        get => _x;
        set => _x = value >= 0 ? value : throw new
ArgumentOutOfRangeException();
    }

    // Solution 1: assign some value in the constructor before "really"
assigning through the property setter.
    public S(int x)
    {
        _x = default;
        X = x;
    }

    // Solution 2: assign the field once in the constructor, repeating the
implementation of the setter.
    public S(int x)
    {
        _x = x >= 0 ? x : throw new ArgumentOutOfRangeException();
    }
}
```

Previous discussion

A small group has looked at this issue and considered a few possible solutions:

1. Require users to assign `this = default` when semi-auto properties have manually implemented setters. We agree this is the wrong solution since it blows away values set in field initializers.
2. Implicitly initialize all backing fields of auto/semi-auto properties.

    - This solves the "semi-auto property setters" problem, and it squarely places explicitly declared fields under different rules: "don't implicitly initialize my fields, but do implicitly initialize my auto-properties."

3. Provide a way to assign the backing field of a semi-auto property and require users to assign it.

    - This could be cumbersome compared to (2). An auto property is supposed to be "automatic", and perhaps that includes "automatic" initialization of the field. It could introduce confusion as to when the underlying field is being assigned by an assignment to the property, and when the property setter is being called.

We've also received feedback from users who want to, for example, include a few field initializers in structs without having to explicitly assign everything. We can solve this issue as well as the "semi-auto property with manually implemented setter" issue at the same time.

C#

```
struct MagnitudeVector3d
{
    double X, Y, Z;
    double Magnitude = 1;
    public MagnitudeVector3d() // error: must assign 'X', 'Y', 'Z' before
returning
    {
    }
}
```

Adjusting definite assignment

Instead of performing a definite assignment analysis to give errors for unassigned fields on `this`, we do it to determine *which fields need to be initialized implicitly*. Such initialization is inserted at the *beginning of the constructor*.

C#

```
struct S
{
    int x, y;

    // Example 1
    public S()
    {
        // ok. Compiler inserts an assignment of `this = default`.
    }

    // Example 2
    public S()
    {
        // ok. Compiler inserts an assignment of `y = default`.
        x = 1;
    }

    // Example 3
    public S()
    {
        // valid since C# 1.0. Compiler inserts no implicit assignments.
        x = 1;
        y = 2;
    }

    // Example 4
    public S(bool b)
    {
        // ok. Compiler inserts assignment of `this = default`.
        if (b)
            x = 1;
        else
            y = 2;
    }

    // Example 5
    void M() { }
    public S(bool b)
    {
        // ok. Compiler inserts assignment of `y = default`.
        x = 1;
        if (b)
            M();

        y = 2;
    }
}
```

In examples (4) and (5), the resulting codegen sometimes has "double assignments" of fields. This is generally fine, but for users who are concerned with such double assignments, we can emit what used to be definite assignment error diagnostics as *disabled-by-default* warning diagnostics.

C#

```
struct S
{
    int x;
    public S() // warning: 'S.x' is implicitly initialized to 'default'.
    {
    }
}
```

Users who set the severity of this diagnostic to "error" will opt in to the pre-C# 11 behavior. Such users are essentially "shut out" of semi-auto properties with manually implemented setters.

C#

```
struct S
{
    public int X
    {
        get => field;
        set => field = field < value ? value : field;
    }

    public S() // error: backing field of 'S.X' is implicitly initialized to
'default'.
    {
        X = 1;
    }
}
```

At first glance, this feels like a "hole" in the feature, but it's **actually the right thing to do**. By enabling the diagnostic, the user is telling us that they don't want the compiler to implicitly initialize their fields in the constructor. There's no way to avoid the implicit initialization here, so the solution for them is to use a different way of initializing the field than a manually implemented setter, such as manually declaring the field and assigning it, or by including a field initializer.

Currently, the JIT does not eliminate dead stores through refs, which means that these implicit initializations do have a real cost. But that might be fixable.
https://github.com/dotnet/runtime/issues/13727

It's worth noting that initializing individual fields instead of the entire instance is really just an optimization. The compiler should probably be free to implement whatever heuristic it wants, as long as it meets the invariant that fields which are not definitely assigned at all return points or before any non-field member access of `this` are implicitly initialized.

For example, if a struct has 100 fields, and just one of them is explicitly initialized, it might make more sense to do an `initobj` on the entire thing, than to implicitly emit `initobj` for the 99 other fields. However, an implementation which implicitly emits `initobj` for the 99 other fields would still be valid.

Changes to language specification

We adjust the following section of the standard:

https://github.com/dotnet/csharpstandard/blob/draft-v6/standard/expressions.md#11712-this-access

> If the constructor declaration has no constructor initializer, the `this` variable behaves exactly the same as an `out` parameter of the struct type. In particular, this means that the variable shall be definitely assigned in every execution path of the instance constructor.

We adjust this language to read:

If the constructor declaration has no constructor initializer, the `this` variable behaves similarly to an `out` parameter of the struct type, except that it is not an error when the definite assignment requirements (§9.4.1) are not met. Instead, we introduce the following behaviors:

1. When the `this` variable itself does not meet the requirements, then all unassigned instance variables within `this` at all points where requirements are violated are implicitly initialized to the default value (§9.3) in an *initialization* phase before any other code in the constructor runs.
2. When an instance variable *v* within `this` does not meet the requirements, or any instance variable at any level of nesting within *v* does not meet the requirements, then *v* is implicitly initialized to the default value in an *initialization* phase before any other code in the constructor runs.

Design meetings

https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-14.md#definite-assignment-in-structs

Low Level Struct Improvements

Article • 10/24/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

This proposal is an aggregation of several different proposals for `struct` performance improvements: `ref` fields and the ability to override lifetime defaults. The goal being a design which takes into account the various proposals to create a single overarching feature set for low level `struct` improvements.

> Note: Previous versions of this spec used the terms "ref-safe-to-escape" and "safe-to-escape", which were introduced in the Span safety feature specification. The ECMA standard committee changed the names to "ref-safe-context" and "safe-context", respectively. The values of the safe context have been refined to use "declaration-block", "function-member", and "caller-context" consistently. The speclets had used different phrasing for these terms, and also used "safe-to-return" as a synonym for "caller-context". This speclet has been updated to use the terms in the C# 7.3 standard.

Motivation

Earlier versions of C# added a number of low level performance features to the language: `ref` returns, `ref struct`, function pointers, etc. ... These enabled .NET developers to write highly performant code while continuing to leverage the C# language rules for type and memory safety. It also allowed the creation of fundamental performance types in the .NET libraries like `Span<T>`.

As these features have gained traction in the .NET ecosystem developers, both internal and external, have been providing us with information on remaining friction points in the ecosystem. Places where they still need to drop to `unsafe` code to get their work done, or require the runtime to special case types like `Span<T>`.

Today `Span<T>` is accomplished by using the `internal` type `ByReference<T>` which the runtime effectively treats as a `ref` field. This provides the benefit of `ref` fields but with the downside that the language provides no safety verification for it, as it does for other uses of `ref`. Further only dotnet/runtime can use this type as it's `internal`, so 3rd parties can not design their own primitives based on `ref` fields. Part of the motivation for this work is to remove `ByReference<T>` and use proper `ref` fields in all code bases.

This proposal plans to address these issues by building on top of our existing low level features. Specifically it aims to:

- Allow `ref struct` types to declare `ref` fields.
- Allow the runtime to fully define `Span<T>` using the C# type system and remove special case type like `ByReference<T>`
- Allow `struct` types to return `ref` to their fields.
- Allow runtime to remove `unsafe` uses caused by limitations of lifetime defaults
- Allow the declaration of safe `fixed` buffers for managed and unmanaged types in `struct`

Detailed Design

The rules for `ref struct` safety are defined in the span safety document using the previous terms. Those rules have been incorporated into the C# 7 standard in §9.7.2 and §16.4.12. This document will describe the required changes to this document as a result of this proposal. Once accepted as an approved feature these changes will be incorporated into that document.

Once this design is complete our `Span<T>` definition will be as follows:

```c#
readonly ref struct Span<T>
{
    readonly ref T _field;
    readonly int _length;

    // This constructor does not exist today but will be added as a part
    // of changing Span<T> to have ref fields. It is a convenient, and
    // safe, way to create a length one span over a stack value that today
```

```
    // requires unsafe code.
    public Span(ref T value)
    {
        _field = ref value;
        _length = 1;
    }
}
```

Provide ref fields and scoped

The language will allow developers to declare `ref` fields inside of a `ref struct`. This can be useful for example when encapsulating large mutable `struct` instances or defining high performance types like `Span<T>` in libraries besides the runtime.

C#

```
ref struct S
{
    public ref int Value;
}
```

A `ref` field will be emitted into metadata using the `ELEMENT_TYPE_BYREF` signature. This is no different than how we emit `ref` locals or `ref` arguments. For example `ref int _field` will be emitted as `ELEMENT_TYPE_BYREF ELEMENT_TYPE_I4`. This will require us to update ECMA335 to allow this entry but this should be rather straight forward.

Developers can continue to initialize a `ref struct` with a `ref` field using the `default` expression in which case all declared `ref` fields will have the value `null`. Any attempt to use such fields will result in a `NullReferenceException` being thrown.

c#

```
ref struct S
{
    public ref int Value;
}

S local = default;
local.Value.ToString(); // throws NullReferenceException
```

While the C# language pretends that a `ref` cannot be `null` this is legal at the runtime level and has well defined semantics. Developers who introduce `ref` fields into their types need to be aware of this possibility and should be **strongly** discouraged from leaking this detail into consuming code. Instead `ref` fields should be validated as non-

null using the runtime helpers and throwing when an uninitialized `struct` is used incorrectly.

```c#
ref struct S1
{
    private ref int Value;

    public int GetValue()
    {
        if (System.Runtime.CompilerServices.Unsafe.IsNullRef(ref Value))
        {
            throw new InvalidOperationException(...);
        }

        return Value;
    }
}
```

A `ref` field can be combined with `readonly` modifiers in the following ways:

- `readonly ref`: this is a field that cannot be ref reassigned outside a constructor or `init` methods. It can be value assigned though outside those contexts
- `ref readonly`: this is a field that can be ref reassigned but cannot be value assigned at any point. This how an `in` parameter could be ref reassigned to a `ref` field.
- `readonly ref readonly`: a combination of `ref readonly` and `readonly ref`.

```c#
ref struct ReadOnlyExample
{
    ref readonly int Field1;
    readonly ref int Field2;
    readonly ref readonly int Field3;

    void Uses(int[] array)
    {
        Field1 = ref array[0];  // Okay
        Field1 = array[0];      // Error: can't assign ref readonly value
(value is readonly)
        Field2 = ref array[0];  // Error: can't repoint readonly ref
        Field2 = array[0];      // Okay
        Field3 = ref array[0];  // Error: can't repoint readonly ref
        Field3 = array[0];      // Error: can't assign ref readonly value
(value is readonly)
```

```
    }
}
```

A `readonly ref struct` will require that `ref` fields are declared `readonly ref`. There is no requirement that they are declared `readonly ref readonly`. This does allow a `readonly struct` to have indirect mutations via such a field but that is no different than a `readonly` field that pointed to a reference type today (more details)

A `readonly ref` will be emitted to metadata using the `initonly` flag, same as any other field. A `ref readonly` field will be attributed with `System.Runtime.CompilerServices.IsReadOnlyAttribute`. A `readonly ref readonly` will be emitted with both items.

This feature requires runtime support and changes to the ECMA spec. As such these will only be enabled when the corresponding feature flag is set in corelib. The issue tracking the exact API is tracked here https://github.com/dotnet/runtime/issues/64165

The set of changes to our safe context rules necessary to allow `ref` fields is small and targeted. The rules already account for `ref` fields existing and being consumed from APIs. The changes need to focus on only two aspects: how they are created and how they are ref reassigned.

First the rules establishing *ref-safe-context* values for fields need to be updated for `ref` fields as follows:

> An expression in the form `ref e.F` *ref-safe-context* as follows:
>
> 1. If `F` is a `ref` field its *ref-safe-context* is the *safe-context* of `e`.
> 2. Else if `e` is of a reference type, it has *ref-safe-context* of *caller-context*
> 3. Else its *ref-safe-context* is taken from the *ref-safe-context* of `e`.

This does not represent a rule change though as the rules have always accounted for `ref` state to exist inside a `ref struct`. This is in fact how the `ref` state in `Span<T>` has always worked and the consumption rules correctly account for this. The change here is just accounting for developers to be able to access `ref` fields directly and ensure they do so by the existing rules implicitly applied to `Span<T>`.

This does mean though that `ref` fields can be returned as `ref` from a `ref struct` but normal fields cannot.

```
c#
```

```
ref struct RS
{
    ref int _refField;
    int _field;

    // Okay: this falls into bullet one above.
    public ref int Prop1 => ref _refField;

    // Error: This is bullet four above and the ref-safe-context of `this`
    // in a `struct` is function-member.
    public ref int Prop2 => ref _field;
}
```

This may seem like an error at first glance but this is a deliberate design point. Again though, this is not a new rule being created by this proposal, it is instead acknowledging the existing rules `Span<T>` behaved by now that developers can declare their own `ref` state.

Next the rules for ref reassignment need to be adjusted for the presence of `ref` fields. The primary scenario for ref reassignment is `ref struct` constructors storing `ref` parameters into `ref` fields. The support will be more general but this is the core scenario. To support this the rules for ref reassignment will be adjusted to account for `ref` fields as follows:

Ref reassignment rules

The left operand of the `= ref` operator must be an expression that binds to a ref local variable, a ref parameter (other than `this`), an out parameter, **or a ref field**.

> For a ref reassignment in the form `e1 = ref e2` both of the following must be true:
>
> 1. `e2` must have *ref-safe-context* at least as large as the *ref-safe-context* of `e1`
> 2. `e1` must have the same *safe-context* as `e2` Note

That means the desired `Span<T>` constructor works without any extra annotation:

```c#
readonly ref struct Span<T>
{
    readonly ref T _field;
    readonly int _length;

    public Span(ref T value)
    {
```

```
        // Falls into the `x.e1 = ref e2` case, where `x` is the implicit
`this`. The
        // safe-context of `this` is *return-only* and ref-safe-context of
`value` is
        // *caller-context* hence this is legal.
        _field = ref value;
        _length = 1;
    }
}
```

The change to ref reassignment rules means `ref` parameters can now escape from a method as a `ref` field in a `ref struct` value. As discussed in the compat considerations section this can change the rules for existing APIs that never intended for `ref` parameters to escape as a `ref` field. The lifetime rules for parameters are based solely on their declaration not on their usage. All `ref` and `in` parameters have *ref-safe-context* of *caller-context* and hence can now be returned by `ref` or a `ref` field. In order to support APIs having `ref` parameters that can be escaping or non-escaping, and thus restore C# 10 call site semantics, the language will introduce limited lifetime annotations.

`scoped` modifier

The keyword `scoped` will be used to restrict the lifetime of a value. It can be applied to a `ref` or a value that is a `ref struct` and has the impact of restricting the *ref-safe-context* or *safe-context* lifetime, respectively, to the *function-member*. For example:

Parameter or Local	ref-safe-context	safe-context
`Span<int> s`	*function-member*	*caller-context*
`scoped Span<int> s`	*function-member*	*function-member*
`ref Span<int> s`	*caller-context*	*caller-context*
`scoped ref Span<int> s`	*function-member*	*caller-context*

In this relationship the *ref-safe-context* of a value can never be wider the *safe-context*.

This allows for APIs in C# 11 to be annotated such that they have the same rules as C# 10:

```
c#

Span<int> CreateSpan(scoped ref int parameter)
{
```

```
    // Just as with C# 10, the implementation of this method isn't relevant
to callers.
}

Span<int> BadUseExamples(int parameter)
{
    // Legal in C# 10 and legal in C# 11 due to scoped ref
    return CreateSpan(ref parameter);

    // Legal in C# 10 and legal in C# 11 due to scoped ref
    int local = 42;
    return CreateSpan(ref local);

    // Legal in C# 10 and legal in C# 11 due to scoped ref
    Span<int> span = stackalloc int[42];
    return CreateSpan(ref span[0]);
}
```

The `scoped` annotation also means that the `this` parameter of a `struct` can now be defined as `scoped ref T`. Previously it had to be special cased in the rules as `ref` parameter that had different *ref-safe-context* rules than other `ref` parameters (see all the references to including or excluding the receiver in the safe context rules). Now it can be expressed as a general concept throughout the rules which further simplifies them.

The `scoped` annotation can also be applied to the following locations:

- locals: This annotation sets the lifetime as *safe-context*, or *ref-safe-context* in case of a `ref` local, to of *function-member* irrespective of the initializer lifetime.

c#

```
Span<int> ScopedLocalExamples()
{
    // Error: `span` has a safe-context of *function-member*. That is true
even though the
    // initializer has a safe-context of *caller-context*. The annotation
overrides the
    // initializer
    scoped Span<int> span = default;
    return span;

    // Okay: the initializer has safe-context of *caller-context* hence so
does `span2`
    // and the return is legal.
    Span<int> span2 = default;
    return span2;

    // The declarations of `span3` and `span4` are functionally identical
because the
```

```
    // initializer has a safe-context of *function-member* meaning the
`scoped` annotation
    // is effectively implied on `span3`
    Span<int> span3 = stackalloc int[42];
    scoped Span<int> span4 = stackalloc int[42];
}
```

Other uses for `scoped` on locals are discussed below.

The `scoped` annotation cannot be applied to any other location including returns, fields, array elements, etc ... Further while `scoped` has impact when applied to any `ref`, `in` or `out` it only has impact when applied to values which are `ref struct`. Having declarations like `scoped int` has no impact because a non `ref struct` is always safe to return. The compiler will create a diagnostic for such cases to avoid developer confusion.

Change the behavior of `out` parameters

To further limit the impact of the compat change of making `ref` and `in` parameters returnable as `ref` fields, the language will change the default *ref-safe-context* value for `out` parameters to be *function-member*. Effectively `out` parameters are implicitly `scoped out` going forward. From a compat perspective this means they cannot be returned by `ref`:

```c#
ref int Sneaky(out int i)
{
    i = 42;

    // Error: ref-safe-context of out is now function-member
    return ref i;
}
```

This will increase the flexibility of APIs that return `ref struct` values and have `out` parameters because it does not have to consider the parameter being captured by reference anymore. This is important because it's a common pattern in reader style APIs:

```c#
Span<byte> Read(Span<byte> buffer, out int read)
{
    // ..
}
```

```
Span<byte> Use()
{
    var buffer = new byte[256];

    // If we keep current `out` ref-safe-context this is an error. The
language must consider
    // the `read` parameter as returnable as a `ref` field
    //
    // If we change `out` ref-safe-context this is legal. The language does
not consider the
    // `read` parameter to be returnable hence this is safe
    int read;
    return Read(buffer, out read);
}
```

The language will also no longer consider arguments passed to an `out` parameter to be returnable. Treating the input to an `out` parameter as returnable was extremely confusing to developers. It essentially subverts the intent of `out` by forcing developers to consider the value passed by the caller which is never used except in languages that don't respect `out`. Going forward languages that support `ref struct` must ensure the original value passed to an `out` parameter is never read.

C# achieves this via it's definite assignment rules. That both achieves our ref safe context rules as well as allowing for existing code which assigns and then returns `out` parameters values.

```
c#

Span<int> StrangeButLegal(out Span<int> span)
{
    span = default;
    return span;
}
```

Together these changes mean the argument to an `out` parameter does not contribute *safe-context* or *ref-safe-context* values to method invocations. This significantly reduces the overall compat impact of `ref` fields as well as simplifies how developers think about `out`. An argument to an `out` parameter does not contribute to the return, it is simply an output.

Infer *safe-context* of declaration expressions

The *safe-context* of a declaration variable from an `out` argument (`M(x, out var y)`) or deconstruction (`(var x, var y) = M()`) is the *narrowest* of the following:

- caller-context
- if out variable is marked `scoped`, then *declaration-block* (i.e. function-member or narrower).
- if out variable's type is ref struct, consider all arguments to the containing invocation, including the receiver:
  - *safe-context* of any argument where its corresponding parameter is not `out` and has *safe-context* of *return-only* or wider
  - *ref-safe-context* of any argument where its corresponding parameter has *ref-safe-context* of *return-only* or wider

See also Examples of inferred *safe-context* of declaration expressions.

Implicitly `scoped` parameters

Overall there are two `ref` location which are implicitly declared as `scoped`:

- `this` on a `struct` instance method
- `out` parameters

The ref safe context rules will be written in terms of `scoped ref` and `ref`. For ref safe context purposes an `in` parameter is equivalent to `ref` and `out` is equivalent to `scoped ref`. Both `in` and `out` will only be specifically called out when it is important to the semantic of the rule. Otherwise they are just considered `ref` and `scoped ref` respectively.

When discussing the *ref-safe-context* of arguments that correspond to `in` parameters they will be generalized as `ref` arguments in the spec. In the case the argument is an lvalue then the *ref-safe-context* is that of the lvalue, otherwise it is *function-member*. Again `in` will only be called out here when it is important to the semantic of the current rule.

Return-only safe context

The design also requires that the introduction of a new safe-context: *return-only*. This is similar to *caller-context* in that it can be returned but it can **only** be returned through a `return` statement.

The details of *return-only* is that it's a context which is greater than *function-member* but smaller than *caller-context*. An expression provided to a `return` statement must be at least *return-only*. As such most existing rules fall out. For example assignment into a `ref` parameter from an expression with a *safe-context* of *return-only* will fail because it's

smaller than the `ref` parameter's *safe-context* which is *caller-context*. The need for this new escape context will be discussed below.

There are three locations which default to *return-only*:

- A `ref` or `in` parameter with have a *ref-safe-context* of *return-only*. This is done in part for `ref struct` to prevent silly cyclic assignment issues. It is done uniformly though to simplify the model as well as minimize compat changes.
- A `out` parameter for a `ref struct` will have *safe-context* of *return-only*. This allows for return and `out` to be equally expressive. This does not have the silly cyclic assignment problem because `out` is implicitly `scoped` so the *ref-safe-context* is still smaller than the *safe-context*.
- A `this` parameter for a `struct` constructor will have a *safe-context* of *return-only*. This falls out due to being modeled as `out` parameters.

Any expression or statement which explicitly returns a value from a method or lambda must have a *safe-context*, and if applicable a *ref-safe-context*, of at least *return-only*. That includes `return` statements, expression bodied members and lambda expressions.

Likewise any assignment to an `out` must have a *safe-context* of at least *return-only*. This is not a special case though, this just follows from the existing assignment rules.

Note: An expression whose type is not a `ref struct` type always has a *safe-context* of *caller-context*.

Rules for method invocation

The ref safe context rules for method invocation will be updated in several ways. The first is by recognizing the impact that `scoped` has on arguments. For a given argument `expr` that is passed to parameter `p`:

> 1. If `p` is `scoped ref` then `expr` does not contribute *ref-safe-context* when considering arguments.
> 2. If `p` is `scoped` then `expr` does not contribute *safe-context* when considering arguments.
> 3. If `p` is `out` then `expr` does not contribute *ref-safe-context* or *safe-context* more details

The language "does not contribute" means the arguments are simply not considered when calculating the *ref-safe-context* or *safe-context* value of the method return

respectively. That is because the values can't contribute to that lifetime as the `scoped` annotation prevents it.

The method invocation rules can now be simplified. The receiver no longer needs to be special cased, in the case of `struct` it is now simply a `scoped ref T`. The value rules need to change to account for `ref` field returns:

> A value resulting from a method invocation `e1.M(e2, ...)`, where `M()` does not return ref-to-ref-struct, has a *safe-context* taken from the narrowest of the following:
>
> 1. The *caller-context*
> 2. When the return is a `ref struct` the *safe-context* contributed by all argument expressions
> 3. When the return is a `ref struct` the *ref-safe-context* contributed by all `ref` arguments
>
> If `M()` does return ref-to-ref-struct, the *safe-context* is the same as the *safe-context* of all arguments which are ref-to-ref-struct. It is an error if there are multiple arguments with different *safe-context* because of method arguments must match.

The `ref` calling rules can be simplified to:

> A value resulting from a method invocation `ref e1.M(e2, ...)`, where `M()` does not return ref-to-ref-struct, is *ref-safe-context* the narrowest of the following contexts:
>
> 1. The *caller-context*
> 2. The *safe-context* contributed by all argument expressions
> 3. The *ref-safe-context* contributed by all `ref` arguments
>
> If `M()` does return ref-to-ref-struct, the *ref-safe-context* is the narrowest *ref-safe-context* contributed by all arguments which are ref-to-ref-struct.

This rule now lets us define the two variants of desired methods:

```c#
Span<int> CreateWithoutCapture(scoped ref int value)
{
    // Error: value Rule 3 specifies that the safe-context be limited to the ref-safe-context
    // of the ref argument. That is the *function-member* for value hence this is not allowed.
    return new Span<int>(ref value);
}
```

```
Span<int> CreateAndCapture(ref int value)
{
    // Okay: value Rule 3 specifies that the safe-context be limited to the
ref-safe-context
    // of the ref argument. That is the *caller-context* for value hence
this is not allowed.
    return new Span<int>(ref value);
}

Span<int> ComplexScopedRefExample(scoped ref Span<int> span)
{
    // Okay: the safe-context of `span` is *caller-context* hence this is
legal.
    return span;

    // Okay: the local `refLocal` has a ref-safe-context of *function-
member* and a
    // safe-context of *caller-context*. In the call below it is passed to a
    // parameter that is `scoped ref` which means it does not contribute
    // ref-safe-context. It only contributes its safe-context hence the
returned
    // rvalue ends up as safe-context of *caller-context*
    Span<int> local = default;
    ref Span<int> refLocal = ref local;
    return ComplexScopedRefExample(ref refLocal);

    // Error: similar analysis as above but the safe-context of `stackLocal`
is
    // *function-member* hence this is illegal
    Span<int> stackLocal = stackalloc int[42];
    return ComplexScopedRefExample(ref stackLocal);
}
```

Method arguments must match

The presence of `ref` fields means the rules around method arguments must match need to be updated as a `ref` parameter can now be stored as a field in a `ref struct` argument to the method. Previously the rule only had to consider another `ref struct` being stored as a field. The impact of this is discussed in the compat considerations. The new rule is ...

> For any method invocation `e.M(a1, a2, ... aN)`
>
> 1. Calculate the narrowest *safe-context* from:
>
>   - *caller-context*
>   - The *safe-context* of all arguments

    - The *ref-safe-context* of all ref arguments whose corresponding parameters have a *ref-safe-context* of *caller-context*
2. All `ref` arguments of `ref struct` types must be assignable by a value with that *safe-cpmtext*. This is a case where `ref` does **not** generalize to include `in` and `out`

> For any method invocation `e.M(a1, a2, ... aN)`
>
> 1. Calculate the narrowest *safe-context* from:
>
>   - *caller-context*
>   - The *safe-context* of all arguments
>   - The *ref-safe-context* of all ref arguments whose corresponding parameters are not `scoped`
>
> 2. All `out` arguments of `ref struct` types must be assignable by a value with that *safe-context*.

The presence of `scoped` allows developers to reduce the friction this rule creates by marking parameters which are not returned as `scoped`. This removes their arguments from (1) in both cases above and provides greater flexibility to callers.

Impact of this change is discussed more deeply below. Overall this will allow developers to make call sites more flexible by annotating non-escaping ref-like values with `scoped`.

Parameter scope variance

The `scoped` modifier and `[UnscopedRef]` attribute (see below) on parameters also impacts our object overriding, interface implementation and `delegate` conversion rules. The signature for an override, interface implementation or `delegate` conversion can:

- Add `scoped` to a `ref` or `in` parameter
- Add `scoped` to a `ref struct` parameter
- Remove `[UnscopedRef]` from an `out` parameter
- Remove `[UnscopedRef]` from a `ref` parameter of a `ref struct` type

Any other difference with respect to `scoped` or `[UnscopedRef]` is considered a mismatch.

The compiler will report a diagnostic for *unsafe scoped mismatches* across overrides, interface implementations, and delegate conversions when:

- The method returns a `ref struct` or returns a `ref` or `ref readonly`, or the method has a `ref` or `out` parameter of `ref struct` type, and
- The method has at least one additional `ref`, `in`, or `out` parameter, or a parameter of `ref struct` type.

The rules above ignore `this` parameters because `ref struct` instance methods cannot be used for overrides, interface implementations, or delegate conversions.

The diagnostic is reported as an *error* if the mismatched signatures are both using C#11 ref safe context rules; otherwise, the diagnostic is a *warning*.

The scoped mismatch warning may be reported on a module compiled with C#7.2 ref safe context rules where `scoped` is not available. In some such cases, it may be necessary to suppress the warning if the other mismatched signature cannot be modified.

The `scoped` modifier and `[UnscopedRef]` attribute also have the following effects on method signatures:

- The `scoped` modifier and `[UnscopedRef]` attribute do not affect hiding
- Overloads cannot differ only on `scoped` or `[UnscopedRef]`

The section on `ref` field and `scoped` is long so wanted to close with a brief summary of the proposed breaking changes:

- A value that has *ref-safe-context* to the *caller-context* is returnable by `ref` or `ref` field.
- A `out` parameter would have a *safe-context* of *function-member*.

Detailed Notes:

- A `ref` field can only be declared inside of a `ref struct`
- A `ref` field cannot be declared `static`, `volatile` or `const`
- A `ref` field cannot have a type that is `ref struct`
- The reference assembly generation process must preserve the presence of a `ref` field inside a `ref struct`
- A `readonly ref struct` must declare its `ref` fields as `readonly ref`
- For by-ref values the `scoped` modifier must appear before `in`, `out`, or `ref`
- The span safety rules document will be updated as outlined in this document
- The new ref safe context rules will be in effect when either
  - The core library contains the feature flag indicating support for `ref` fields
  - The `langversion` value is 11 or higher

Syntax

12.6.2 Local variable declarations: added `'scoped'?`.

```antlr
local_variable_declaration
    : 'scoped'? local_variable_mode_modifier? local_variable_type
local_variable_declarators
    ;

local_variable_mode_modifier
    : 'ref' 'readonly'?
    ;
```

12.9.4 The for statement: added `'scoped'?` *indirectly* from `local_variable_declaration`.

12.9.5 The foreach statement: added `'scoped'?`.

```antlr
foreach_statement
    : 'foreach' '(' 'scoped'? local_variable_type identifier 'in' expression
')'
      embedded_statement
    ;
```

11.6.2 Argument lists: added `'scoped'?` for `out` declaration variable.

```antlr
argument_value
    : expression
    | 'in' variable_reference
    | 'ref' variable_reference
    | 'out' ('scoped'? local_variable_type)? identifier
    ;
```

--.-.- Deconstruction expressions:

```antlr
[TBD]
```

14.6.2 Method parameters: added `'scoped'?` to `parameter_modifier`.

antlr

```
fixed_parameter
    : attributes? parameter_modifier? type identifier default_argument?
    ;

parameter_modifier
    | 'this' 'scoped'? parameter_mode_modifier?
    | 'scoped' parameter_mode_modifier?
    | parameter_mode_modifier
    ;

parameter_mode_modifier
    : 'in'
    | 'ref'
    | 'out'
    ;
```

19.2 Delegate declarations: added `'scoped'?` *indirectly* from `fixed_parameter`.

11.16 Anonymous function expressions: added `'scoped'?`.

antlr

```
explicit_anonymous_function_parameter
    : 'scoped'? anonymous_function_parameter_modifier? type identifier
    ;

anonymous_function_parameter_modifier
    : 'in'
    | 'ref'
    | 'out'
    ;
```

Sunset restricted types

The compiler has a concept of a set of "restricted types" which is largely undocumented. These types were given a special status because in C# 1.0 there was no general purpose way to express their behavior. Most notably the fact that the types can contain references to the execution stack. Instead the compiler had special knowledge of them and restricted their use to ways that would always be safe: disallowed returns, cannot use as array elements, cannot use in generics, etc ...

Once `ref` fields are available and extended to support `ref struct` these types can be correctly defined in C# using a combination of `ref struct` and `ref` fields. Therefore

when the compiler detects that a runtime supports `ref` fields it will no longer have a notion of restricted types. It will instead use the types as they are defined in the code.

To support this our ref safe context rules will be updated as follows:

- `__makeref` will be treated as a method with the signature `static TypedReference __makeref<T>(ref T value)`
- `__refvalue` will be treated as a method with the signature `static ref T __refvalue<T>(TypedReference tr)`. The expression `__refvalue(tr, int)` will effectively use the second argument as the type parameter.
- `__arglist` as a parameter will have a *ref-safe-context* and *safe-context* of *function-member*.
- `__arglist(...)` as an expression will have a *ref-safe-context* and *safe-context* of *function-member*.

Conforming runtimes will ensure that `TypedReference`, `RuntimeArgumentHandle` and `ArgIterator` are defined as `ref struct`. Further `TypedReference` must be viewed as having a `ref` field to a `ref struct` for any possible type (it can store any value). That combined with the above rules will ensure references to the stack do not escape beyond their lifetime.

Note: strictly speaking this is a compiler implementation detail vs. part of the language. But given the relationship with `ref` fields it is being included in the language proposal for simplicity.

Provide unscoped

One of the most notable friction points is the inability to return fields by `ref` in instance members of a `struct`. This means developers can't create `ref` returning methods / properties and have to resort to exposing fields directly. This reduces the usefulness of `ref` returns in `struct` where it is often the most desired.

```c#
struct S
{
    int _field;

    // Error: this, and hence _field, can't return by ref
    public ref int Prop => ref _field;
}
```

The rationale for this default is reasonable but there is nothing inherently wrong with a `struct` escaping `this` by reference, it is simply the default chosen by the ref safe context rules.

To fix this the language will provide the opposite of the `scoped` lifetime annotation by supporting an `UnscopedRefAttribute`. This can be applied to any `ref` and it will change the *ref-safe-context* to be one level wider than its default. For example:

- if applied to a `struct` instance method it will become *return only* where previously it was *containing method*.
- if applied to a `ref` parameter it will become *caller-context* where previously it was *return only*

When applying `[UnscopedRef]` to an instance method of a `struct` it has the impact of modifying the implicit `this` parameter. This means `this` acts as an unannotated `ref` of the same type.

```c#
struct S
{
    int field;

    // Error: `field` has the ref-safe-context of `this` which is *function-member* because
    // it is a `scoped ref`
    ref int Prop1 => ref field;

    // Okay: `field` has the ref-safe-context of `this` which is *caller-context* because
    // it is a `ref`
    [UnscopedRef] ref int Prop1 => ref field;
}
```

The annotation can also be placed on `out` parameters to restore them to C# 10 behavior.

```c#
ref int SneakyOut([UnscopedRef] out int i)
{
    i = 42;
    return ref i;
}
```

For the purposes of ref safe cotnext rules, such an `[UnscopedRef] out` is considered simply a `ref`. Similar to how `in` is considered `ref` for lifetime purposes.

The `[UnscopedRef]` annotation will be disallowed on `init` members and constructors inside `struct`. Those members are already special with respect to `ref` semantics as they view `readonly` members as mutable. This means taking `ref` to those members appears as a simple `ref`, not `ref readonly`. This is allowed within the boundary of constructors and `init`. Allowing `[UnscopedRef]` would permit such a `ref` to incorrectly escape outside the constructor and permit mutation after `readonly` semantics had taken place.

The attribute type will have the following definition:

```c#
namespace System.Diagnostics.CodeAnalysis
{
    [AttributeUsage(
        AttributeTargets.Method | AttributeTargets.Property |
AttributeTargets.Parameter,
        AllowMultiple = false,
        Inherited = false)]
    public sealed class UnscopedRefAttribute : Attribute
    {
    }
}
```

Detailed Notes:

- An instance method or property annotated with `[UnscopedRef]` has *ref-safe-context* of `this` set to the *caller-context*.
- A member annotated with `[UnscopedRef]` cannot implement an interface.
- It is an error to use `[UnscopedRef]` on
  - A member that is not declared on a `struct`
  - A `static` member, `init` member or constructor on a `struct`
  - A parameter marked `scoped`
  - A parameter passed by value
  - A parameter passed by reference that is not implicitly scoped

ScopedRefAttribute

The `scoped` annotations will be emitted into metadata via the type `System.Runtime.CompilerServices.ScopedRefAttribute` attribute. The attribute will be

matched by namespace-qualified name so the definition does not need to appear in any specific assembly.

The `ScopedRefAttribute` type is for compiler use only - it is not permitted in source. The type declaration is synthesized by the compiler if not already included in the compilation.

The type will have the following definition:

```c#
namespace System.Runtime.CompilerServices
{
    [AttributeUsage(AttributeTargets.Parameter, AllowMultiple = false,
Inherited = false)]
    internal sealed class ScopedRefAttribute : Attribute
    {
    }
}
```

The compiler will emit this attribute on the parameter with `scoped` syntax. This will only be emitted when the syntax causes the value to differ from its default state. For example `scoped out` will cause no attribute to be emitted.

RefSafetyRulesAttribute

There are several differences in the *ref safe context* rules between C#7.2 and C#11. Any of these differences could result in breaking changes when recompiling with C#11 against references compiled with C#10 or earlier.

1. unscoped `ref`/`in`/`out` parameters may escape a method invocation as a `ref` field of a `ref struct` in C#11, not in C#7.2
2. `out` parameters are implicitly scoped in C#11, and unscoped in C#7.2
3. `ref`/`in` parameters to `ref struct` types are implicitly scoped in C#11, and unscoped in C#7.2

To reduce the chance of breaking changes when recompiling with C#11, we will update the C#11 compiler to use the ref safe context rules *for method invocation* that *match the rules that were used to analyze the method declaration*. Essentially, when analyzing a call to a method compiled with an older compiler, the C#11 compiler will use C#7.2 ref safe context rules.

To enable this, the compiler will emit a new `[module: RefSafetyRules(11)]` attribute when the module is compiled with `-langversion:11` or higher or compiled with a corlib

containing the feature flag for `ref` fields.

The argument to the attribute indicates the language version of the *ref safe context* rules used when the module was compiled. The version is currently fixed at `11` regardless of the actual language version passed to the compiler.

The expectation is that future versions of the compiler will update the ref safe context rules and emit attributes with distinct versions.

If the compiler loads a module that includes a `[module: RefSafetyRules(version)]` *with a* `version` *other than* `11`, the compiler will report a warning for the unrecognized version if there are any calls to methods declared in that module.

When the C#11 compiler *analyzes a method call*:

- If the module containing the method declaration includes `[module: RefSafetyRules(version)]`, regardless of `version`, the method call is analyzed with C#11 rules.
- If the module containing the method declaration is from source, and compiled with `-langversion:11` or with a corlib containing the feature flag for `ref` fields, the method call is analyzed with C#11 rules.
- *If the module containing the method declaration references* `System.Runtime { ver: 7.0 }`*, the method call is analyzed with C#11 rules. This rule is a temporary mitigation for modules compiled with earlier previews of C#11 / .NET 7 and will be removed later.*
- Otherwise, the method call is analyzed with C#7.2 rules.

A pre-C#11 compiler will ignore any `RefSafetyRulesAttribute` and analyze method calls with C#7.2 rules only.

The `RefSafetyRulesAttribute` will be matched by namespace-qualified name so the definition does not need to appear in any specific assembly.

The `RefSafetyRulesAttribute` type is for compiler use only - it is not permitted in source. The type declaration is synthesized by the compiler if not already included in the compilation.

C#

```
namespace System.Runtime.CompilerServices
{
    [AttributeUsage(AttributeTargets.Module, AllowMultiple = false,
Inherited = false)]
    internal sealed class RefSafetyRulesAttribute : Attribute
    {
```

```
        public RefSafetyRulesAttribute(int version) { Version = version; }
        public readonly int Version;
    }
}
```

Safe fixed size buffers

Safe fixed size buffers was not delivered in C# 11. This feature may be implemented in a future version of C#.

The language will relax the restrictions on fixed sized arrays such that they can be declared in safe code and the element type can be managed or unmanaged. This will make types like the following legal:

```
c#

internal struct CharBuffer
{
    internal char Data[128];
}
```

These declarations, much like their `unsafe` counter parts, will define a sequence of `N` elements in the containing type. These members can be accessed with an indexer and can also be converted to `Span<T>` and `ReadOnlySpan<T>` instances.

When indexing into a `fixed` buffer of type `T` the `readonly` state of the container must be taken into account. If the container is `readonly` then the indexer returns `ref readonly T` else it returns `ref T`.

Accessing a `fixed` buffer without an indexer has no natural type however it is convertible to `Span<T>` types. In the case the container is `readonly` the buffer is implicitly convertible to `ReadOnlySpan<T>`, else it can implicitly convert to `Span<T>` or `ReadOnlySpan<T>` (the `Span<T>` conversion is considered *better*).

The resulting `Span<T>` instance will have a length equal to the size declared on the `fixed` buffer. The *safe-context* of the returned value will be equal to the *safe-context* of the container, just as it would if the backing data was accessed as a field.

For each `fixed` declaration in a type where the element type is `T` the language will generate a corresponding `get` only indexer method whose return type is `ref T`. The indexer will be annotated with the `[UnscopedRef]` attribute as the implementation will be returning fields of the declaring type. The accessibility of the member will match the accessibility on the `fixed` field.

For example, the signature of the indexer for `CharBuffer.Data` will be the following:

```c#
[UnscopedRef] internal ref char DataIndexer(int index) => ...;
```

If the provided index is outside the declared bounds of the `fixed` array then an `IndexOutOfRangeException` will be thrown. In the case a constant value is provided then it will be replaced with a direct reference to the appropriate element. Unless the constant is outside the declared bounds in which case a compile time error would occur.

There will also be a named accessor generated for each `fixed` buffer that provides by value `get` and `set` operations. Having this means that `fixed` buffers will more closely resemble existing array semantics by having a `ref` accessor as well as byval `get` and `set` operations. This means compilers will have the same flexibility when emitting code consuming `fixed` buffers as they do when consuming arrays. This should make operations like `await` over `fixed` buffers easier to emit.

This also has the added benefit that it will make `fixed` buffers easier to consume from other languages. Named indexers is a feature that has existed since the 1.0 release of .NET. Even languages which cannot directly emit a named indexer can generally consume them (C# is actually a good example of this).

The backing storage for the buffer will be generated using the `[InlineArray]` attribute. This is a mechanism discussed in issue 12320 which allows specifically for the case of efficiently declaring sequence of fields of the same type. This particular issue is still under active discussion and the expectation is that the implementation of this feature will follow however that discussion goes.

Initializers with `ref` values in `new` and `with` expressions

In section 11.7.15.3 Object initializers, we update the grammar to:

```antlr
initializer_value
    : 'ref' expression // added
    | expression
    | object_or_collection_initializer
    ;
```

In the section for with expression, we update the grammar to:

antlr

```
member_initializer
    : identifier '=' 'ref' expression // added
    | identifier '=' expression
    ;
```

The left operand of the assignment must be an expression that binds to a ref field.
The right operand must be an expression that yields an lvalue designating a value of the same type as the left operand.

We add a similar rule to ref local reassignment:
If the left operand is a writeable ref (i.e. it designates anything other than a `ref readonly` field), then the right operand must be a writeable lvalue.

The escape rules for constructor invocations remain:

> A `new` expression that invokes a constructor obeys the same rules as a method invocation that is considered to return the type being constructed.

Namely the rules of method invocation updated above:

> An rvalue resulting from a method invocation `e1.M(e2, ...)` has *safe-context* from the smallest of the following contexts:
>
> 1. The *caller-context*
> 2. The *safe-context* contributed by all argument expressions
> 3. When the return is a `ref struct` then *ref-safe-context* contributed by all `ref` arguments

For a `new` expression with initializers, the initializer expressions count as arguments (they contribute their *safe-context*) and the `ref` initializer expressions count as `ref` arguments (they contribute their *ref-safe-context*), recursively.

Changes in unsafe context

Pointer types (section 22.3) are extended to allow managed types as referent type. Such pointer types are written as a managed type followed by a `*` token. They produce a warning.

The address-of operator (section 22.6.5) is relaxed to accept a variable with a managed type as its operand.

The `fixed` statement (section 22.7) is relaxed to accept *fixed_pointer_initializer* that is the address of a variable of managed type `T` or that is an expression of an *array_type* with elements of a managed type `T`.

The stack allocation initializer (section 22.9) is similarly relaxed.

Considerations

There are considerations other parts of the development stack should consider when evaluating this feature.

Compat Considerations

The challenge in this proposal is the compatibility implications this design has to our existing span safety rules, or §9.7.2. While those rules fully support the concept of a `ref struct` having `ref` fields they do not allow for APIs, other than `stackalloc`, to capture `ref` state that refers to the stack. The ref safe cpmtext rules have a hard assumption, or §16.4.12.8 that a constructor of the form `Span(ref T value)` does not exist. That means the safety rules do not account for a `ref` parameter being able to escape as a `ref` field hence it allows for code like the following.

```c#
Span<int> CreateSpanOfInt()
{
    // This is legal according to the 7.2 span rules because they do not
account
    // for a constructor in the form Span(ref T value) existing.
    int local = 42;
    return new Span<int>(ref local);
}
```

Effectively there are three ways for a `ref` parameter to escape from a method invocation:

1. By value return
2. By `ref` return
3. By `ref` field in `ref struct` that is returned or passed as `ref` / `out` parameter

The existing rules only account for (1) and (2). They do not account for (3) hence gaps like returning locals as `ref` fields are not accounted for. This design must change the

rules to account for (3). This will have a small impact to compatibility for existing APIs. Specifically it will impact APIs that have the following properties.

- Have a `ref struct` in the signature
  - Where the `ref struct` is a return type, `ref` or `out` parameter
  - Has an additional `in` or `ref` parameter excluding the receiver

In C# 10 callers of such APIs never had to consider that `ref` state input to the API could be captured as a `ref` field. That allowed for several patterns to exist, safely in C# 10, that will be unsafe in C# 11 due to the ability for `ref` state to escape as a `ref` field. For example:

```c#
Span<int> CreateSpan(ref int parameter)
{
    // The implementation of this method is irrelevant when considering the lifetime of the
    // returned Span<T>. The ref safe context rules only look at the method signature, not the
    // implementation. In C# 10 ref fields didn't exist hence there was no way for `parameter`
    // to escape by ref in this method
}

Span<int> BadUseExamples(int parameter)
{
    // Legal in C# 10 but would be illegal with ref fields
    return CreateSpan(ref parameter);

    // Legal in C# 10 but would be illegal with ref fields
    int local = 42;
    return CreateSpan(ref local);

    // Legal in C# 10 but would be illegal with ref fields
    Span<int> span = stackalloc int[42];
    return CreateSpan(ref span[0]);
}
```

The impact of this compatibility break is expected to be very small. The impacted API shape made little sense in the absence of `ref` fields hence it is unlikely customers created many of these. Experiments running tools to spot this API shape over existing repositories back up that assertion. The only repository with any significant counts of this shape is dotnet/runtime and that is because that repo can create `ref` fields via the `ByReference<T>` intrinsic type.

Even so the design must account for such APIs existing because it expresses a valid pattern, just not a common one. Hence the design must give developers the tools to restore the existing lifetime rules when upgrading to C# 10. Specifically it must provide mechanisms that allow developers to annotate `ref` parameters as unable to escape by `ref` or `ref` field. That allows customers to define APIs in C# 11 that have the same C# 10 callsite rules.

Reference Assemblies

A reference assembly for a compilation using features described in this proposal must maintain the elements that convey ref safe context information. That means all lifetime annotation attributes must be preserved in their original position. Any attempt to replace or omit them can lead to invalid reference assemblies.

Representing `ref` fields is more nuanced. Ideally a `ref` field would appear in a reference assembly as would any other field. However a `ref` field represents a change to the metadata format and that can cause issues with tool chains that are not updated to understand this metadata change. A concrete example is C++/CLI which will likely error if it consumes a `ref` field. Hence it's advantageous if `ref` fields can be omitted from reference assemblies in our core libraries.

A `ref` field by itself has no impact on ref safe context rules. As a concrete example consider that flipping the existing `Span<T>` definition to use a `ref` field has no impact on consumption. Hence the `ref` itself can be omitted safely. However a `ref` field does have other impacts to consumption that must be preserved:

- A `ref struct` which has a `ref` field is never considered `unmanaged`
- The type of the `ref` field impacts infinite generic expansion rules. Hence if the type of a `ref` field contains a type parameter that must be preserved

Given those rules here is a valid reference assembly transformation for a `ref struct`:

```c#
// Impl assembly
ref struct S<T>
{
    ref T _field;
}

// Ref assembly
ref struct S<T>
{
    object _o; // force managed
```

```
    T _f; // maintain generic expansion protections
}
```

Annotations

Lifetimes are most naturally expressed using types. A given program's lifetimes are safe when the lifetime types type check. While the syntax of C# implicitly adds lifetimes to values, there is an underlying type system that describes the fundamental rules here. It's often easier to discuss the implication of changes to the design in terms of these rules so they are included here for discussion sake.

Note that this is not meant to be a 100% complete documentation. Documenting every single behavior isn't a goal here. Instead it's meant to establish a general understanding and common verbiage by which the model, and potential changes to it, can be discussed.

Usually it's not necessary to directly talk about lifetime types. The exceptions are places where lifetimes can vary based on particular "instantiation" sites. This is a kind of polymorphism and we call these varying lifetimes "generic lifetimes", represented as generic parameters. C# does not provide syntax for expressing lifetime generics, so we define an implicit "translation" from C# to an expanded lowered language that contains explicit generic parameters.

The below examples make use of named lifetimes. The syntax `$a` refers to a lifetime named `a`. It is a lifetime that has no meaning by itself but can be given a relationship to other lifetimes via the `where $a : $b` syntax. This establishes that `$a` is convertible to `$b`. It may help to think of this as establishing that `$a` is a lifetime at least as long as `$b`.

There are a few predefined lifetimes for convenience and brevity below:

- `$heap`: this is the lifetime of any value that exists on the heap. It is available in all contexts and method signatures.
- `$local`: this is the lifetime of any value that exists on the method stack. It's effectively a name place holder for *function-member*. It is implicitly defined in methods and can appear in method signatures except for any output position.
- `$ro`: name place holder for *return only*
- `$cm`: name place holder for *caller-context*

There are a few predefined relationships between lifetimes:

- `where $heap : $a` for all lifetimes `$a`
- `where $cm : $ro`

- `where $x : $local` for all predefined lifetimes. User defined lifetimes have no relationship to local unless explicitly defined.

Lifetime variables when defined on types can be invariant or covariant. These are expressed using the same syntax as generic parameters:

```
C#

// $this is covariant
// $a is invariant
ref struct S<out $this, $a>
```

The lifetime parameter `$this` on type definitions is *not* predefined but it does have a few rules associated with it when it is defined:

- It must be the first lifetime parameter.
- It must be covariant: `out $this`.
- The lifetime of `ref` fields must be convertible to `$this`
- The `$this` lifetime of all non-ref fields must be `$heap` or `$this`.

The lifetime of a ref is expressed by providing a lifetime argument to the ref. For example a `ref` that refers to the heap is expressed as `ref<$heap>`.

When defining a constructor in the model the name `new` will be used for the method. It is necessary to have a parameter list for the returned value as well as the constructor arguments. This is necessary to express the relationship between constructor inputs and the constructed value. Rather than having `Span<$a><$ro>` the model will use `Span<$a> new<$ro>` instead. The type of `this` in the constructor, including lifetimes, will be the defined return value.

The basic rules for the lifetime are defined as:

- All lifetimes are expressed syntactically as generic arguments, coming before type arguments. This is true for predefined lifetimes except `$heap` and `$local`.
- All types `T` that are not a `ref struct` implicitly have lifetime of `T<$heap>`. This is implicit, there is no need to write `int<$heap>` in every sample.
- For a `ref` field defined as `ref<$l0> T<$l1, $l2, ... $ln>`:
  - All lifetimes `$l1` through `$ln` must be invariant.
  - The lifetime of `$l0` must be convertible to `$this`
- For a `ref` defined as `ref<$a> T<$b, ...>`, `$b` must be convertible to `$a`
- The `ref` of a variable has a lifetime defined by:
  - For a `ref` local, parameter, field or return of type `ref<$a> T` the lifetime is `$a`

    - `$heap` for all reference types and fields of reference types
    - `$local` for everything else
- An assignment or return is legal when the underlying type conversion is legal
- Lifetimes of expressions can be made explicit by using cast annotations:
    - `(T<$a> expr)` the value lifetime is explicitly `$a` for `T<...>`
    - `ref<$a> (T<$b>)expr` the value lifetime is `$b` for `T<...>` and the ref lifetime is `$a`.

For the purpose of lifetime rules a `ref` is considered part of the type of the expression for purposes of conversions. It is logically represented by converting `ref<$a> T<...>` to `ref<$a, T<...>>` where `$a` is covariant and `T` is invariant.

Next let's define the rules that allow us to map C# syntax to the underlying model.

For brevity sake a type which has no explicit lifetime parameters treated as if there is `out $this` defined and applied to all fields of the type. A type with a `ref` field must define explicit lifetime parameters.

These rules exists to support our existing invariant that `T` can be assigned to `scoped T` for all types. That maps down to `T<$a, ...>` being assignable to `T<$local, ...>` for all lifetimes known to be convertible to `$local`. Further this supports other items like being able to assign `Span<T>` from the heap to those on the stack. This does exclude types where fields have differing lifetimes for non-ref values but that is the reality of C# today. Changing that would require a significant change of C# rules that would need to be mapped out.

The type of `this` for a type `S<out $this, ...>` inside an instance method is implicitly defined as the following:

- For normal instance method: `ref<$local> S<$cm, ...>`
- For instance method annotated with `[UnscopedRef]`: `ref<$ro> S<$cm, ...>`

The lack of an explicit `this` parameter forces the implicit rules here. For complex samples and discussions consider writing as a `static` method and making `this` an explicit parameter.

C#

```
ref struct S<out $this>
{
    // Implicit this can make discussion confusing
    void M<$ro, $cm>(ref<$ro> S<$cm> s) {  }

    // Rewrite as explicit this to simplify discussion
```

```
    static void M<$ro, $cm>(ref<$local> S<$cm> this, ref<$ro> S<$cm> s) { }
}
```

The C# method syntax maps to the model in the following ways:

- `ref` parameters have a ref lifetime of `$ro`
- parameters of type `ref struct` have a this lifetime of `$cm`
- ref returns have a ref lifetime of `$ro`
- returns of type `ref struct` have a value lifetime of `$ro`
- `scoped` on a parameter or `ref` changes the ref lifetime to be `$local`

Given that let's explore a simple example that demonstrates the model here:

C#

```
ref int M1(ref int i) => ...

// Maps to the following.

ref<$ro> int Identity<$ro>(ref<$ro> int i)
{
    // okay: has ref lifetime $ro which is equal to $ro
    return ref i;

    // okay: has ref lifetime $heap which convertible $ro
    int[] array = new int[42];
    return ref array[0];

    // error: has ref lifetime $local which has no conversion to $a hence
    // it's illegal
    int local = 42;
    return ref local;
}
```

Now let's explore the same example using a `ref struct`:

C#

```
ref struct S
{
    ref int Field;

    S(ref int f)
    {
        Field = ref f;
    }
}

S M2(ref int i, S span1, scoped S span2) => ...
```

```
// Maps to

ref struct S<out $this>
{
    // Implicitly
    ref<$this> int Field;

    S<$ro> new<$ro>(ref<$ro> int f)
    {
        Field = ref f;
    }
}

S<$ro> M2<$ro>(
    ref<$ro> int i,
    S<$ro> span1)
    S<$local> span2)
{
    // okay: types match exactly
    return span1;

    // error: has lifetime $local which has no conversion to $ro
    return span2;

    // okay: type S<$heap> has a conversion to S<$ro> because $heap has a
    // conversion to $ro and the first lifetime parameter of S<> is
covariant
    return default(S<$heap>)

    // okay: the ref lifetime of ref $i is $ro so this is just an
    // identity conversion
    S<$ro> local = new S<$ro>(ref $i);
    return local;

    int[] array = new int[42];
    // okay: S<$heap> is convertible to S<$ro>
    return new S<$heap>(ref<$heap> array[0]);

    // okay: the parameter of the ctor is $ro ref int and the argument is
$heap ref int. These
    // are convertible.
    return new S<$ro>(ref<$heap> array[0]);

    // error: has ref lifetime $local which has no conversion to $a hence
    // it's illegal
    int local = 42;
    return ref local;
}
```

Next let's see how this helps with the cyclic self assignment problem:

```
C#
```

```
ref struct S
{
    int field;
    ref int refField;

    static void SelfAssign(ref S s)
    {
        s.refField = ref s.field;
    }
}

// Maps to

ref struct S<out $this>
{
    int field;
    ref<$this> int refField;

    static void SelfAssign<$ro, $cm>(ref<$ro> S<$cm> s)
    {
        // error: the types work out here to ref<$cm> int = ref<$ro> int and
that is
        // illegal as $ro has no conversion to $cm (the relationship is the
other direction)
        s.refField = ref<$ro> s.field;
    }
}
```

Next let's see how this helps with the silly capture parameter problem:

C#

```
ref struct S
{
    ref int refField;

    void Use(ref int parameter)
    {
        // error: this needs to be an error else every call to this.Use(ref
local) would fail
        // because compiler would assume the `ref` was captured by ref.
        this.refField = ref parameter;
    }
}

// Maps to

ref struct S<out $this>
{
    ref<$this> int refField;

    // Using static form of this method signature so the type of this is
```

```
explicit.
    static void Use<$ro, $cm>(ref<$local> S<$cm> @this, ref<$ro> int
parameter)
    {
        // error: the types here are:
        //  - refField is ref<$cm> int
        //  - ref parameter is ref<$ro> int
        // That means the RHS is not convertible to the LHS ($ro is not
covertible to $cm) and
        // hence this reassignment is illegal
        @this.refField = ref<$ro> parameter;
    }
}
```

Open Issues

Change the design to avoid compat breaks

This design proposes several compatibility breaks with our existing ref-safe-context rules. Even though the breaks are believed to be minimally impactful significant consideration was given to a design which had no breaking changes.

The compat preserving design though was significantly more complex than this one. In order to preserve compat `ref` fields need distinct lifetimes for the ability to return by `ref` and return by `ref` field. Essentially it requires us to provide *ref-field-safe-context* tracking for all parameters to a method. This needs to be calculated for all expressions and tracked in all values virtually everywhere that *ref-safe-context* is tracked today.

Further this value has relationships with *ref-safe-context*. For example it's non-sensical to have a value can be returned as a `ref` field but not directly as `ref`. That is because `ref` fields can be trivially returned by `ref` already (`ref` state in a `ref struct` can be returned by `ref` even when the containing value cannot). Hence the rules further need constant adjustment to ensure these values are sensible with respect to each other.

Also it means the language needs syntax to represent `ref` parameters that can be returned in three different ways: by `ref` field, by `ref` and by value. The default being returnable by `ref`. Going forward though the more natural return, particularly when `ref struct` are involved is expected to be by `ref` field or `ref`. That means new APIs require an extra syntax annotation to be correct by default. This is undesirable.

These compat changes though will impact methods that have the following properties:

- Have a `Span<T>` or `ref struct`

- Where the `ref struct` is a return type, `ref` or `out` parameter
- Has an additional `in` or `ref` parameter (excluding the receiver)

To understand the impact it's helpful to break APIs into categories:

1. Want consumers to account for `ref` being captured as a `ref` field. Prime example is the `Span(ref T value)` constructors
2. Do not want consumers to account for `ref` being captured as a `ref` field. These though break into two categories
   a. Unsafe APIs. These are APIS inside the `Unsafe` and `MemoryMarshal` types, of which `MemoryMarshal.CreateSpan` is the most prominent. These APIs do capture the `ref` unsafely but they are also known to be unsafe APIs.
   b. Safe APIs. These are APIs which take `ref` parameters for efficiency but it is not actually captured anywhere. Examples are small but one is `AsnDecoder.ReadEnumeratedBytes`

This change primarily benefits (1) above. These are expected to make up the majority of APIs that take a `ref` and return a `ref struct` going forward. The changes negatively impact (2.1) and (2.2) as it breaks the existing calling semantics because the lifetime rules change.

The APIs in category (2.1) though are largely authored by Microsoft or by developers who stand the most to benefit from `ref` fields (the Tanner's of the world). It is reasonable to assume this class of developers would be amenable to a compatibility tax on upgrade to C# 11 in the form of a few annotations to retain the existing semantics if `ref` fields were provided in return.

The APIs in category (2.2) are the biggest issue. It is unknown how many such APIs exist and it's unclear if these would be more / less frequent in 3rd party code. The expectation is there is a very small number of them, particularly if we take the compat break on `out`. Searches so far have revealed a very small number of these existing in `public` surface area. This is a hard pattern to search for though as it requires semantic analysis. Before taking this change a tool based approach would be needed to verify the assumptions around this impacting a small number of known cases.

For both cases in category (2) though the fix is straight forward. The `ref` parameters that do not want to be considered capturable must add `scoped` to the `ref`. In (2.1) this will likely also force the developer to use `Unsafe` or `MemoryMarshal` but that is expected for unsafe style APIs.

Ideally the language could reduce the impact of silent breaking changes by issuing a warning when an API silently falls into the troublesome behavior. That would be a

method that both takes a `ref`, returns `ref struct` but does not actually capture the `ref` in the `ref struct`. The compiler could issue a diagnostic in that case informing developers such `ref` should be annotated as `scoped ref` instead.

Decision This design can be achieved but the resulting feature is more difficult to use to the point the decision was made to take the compat break.

Decision The compiler will provide a warning when a method meets the criteria but does not capture the `ref` parameter as a `ref` field. This should suitably warn customers on upgrade about the potential issues they are creating

Keywords vs. attributes

This design calls for using attributes to annotate the new lifetime rules. This also could've been done just as easily with contextual keywords. For instance `[DoesNotEscape]` could map to `scoped`. However keywords, even the contextual ones, generally must meet a very high bar for inclusion. They take up valuable language real estate and are more prominent parts of the language. This feature, while valuable, is going to serve a minority of C# developers.

On the surface that would seem to favor not using keywords but there are two important points to consider:

1. The annotations will effect program semantics. Having attributes impact program semantics is a line C# is reluctant to cross and it's unclear if this is the feature that should justify the language taking that step.
2. The developers most likely to use this feature intersect strongly with the set of developers that use function pointers. That feature, while also used by a minority of developers, did warrant a new syntax and that decision is still seen as sound.

Taken together this means syntax should be considered.

A rough sketch of the syntax would be:

- `[RefDoesNotEscape]` maps to `scoped ref`
- `[DoesNotEscape]` maps to `scoped`
- `[RefDoesEscape]` maps to `unscoped`

Decision Use syntax for `scoped` and `scoped ref`; use attribute for `unscoped`.

Allow fixed buffer locals

This design allows for safe `fixed` buffers that can support any type. One possible extension here is allowing such `fixed` buffers to be declared as local variables. This would allow a number of existing `stackalloc` operations to be replaced with a `fixed` buffer. It would also expand the set of scenarios we could have stack style allocations as `stackalloc` is limited to unmanaged element types while `fixed` buffers are not.

```c#
class FixedBufferLocals
{
    void Example()
    {
        Span<int> span = stackalloc int[42];
        int buffer[42];
    }
}
```

This holds together but does require us to extend the syntax for locals a bit. Unclear if this is or isn't worth the extra complexity. Possible we could decide no for now and bring back later if sufficient need is demonstrated.

Example of where this would be beneficial:
https://github.com/dotnet/runtime/pull/34149

Decision hold off on this for now

To use modreqs or not

A decision needs to be made if methods marked with new lifetime attributes should or should not translate to `modreq` in emit. There would be effectively a 1:1 mapping between annotations and `modreq` if this approach was taken.

The rationale for adding a `modreq` is the attributes change the semantics of ref safe context rules. Only languages which understand these semantics should be calling the methods in question. Further when applied to OHI scenarios, the lifetimes become a contract that all derived methods must implement. Having the annotations exist without `modreq` can lead to situations where `virtual` method chains with conflicting lifetime annotations are loaded (can happen if only one part of `virtual` chain is compiled and other is not).

The initial ref safe context work did not use `modreq` but instead relied on languages and the framework to understand. At the same time though all of the elements that contribute to the ref safe context rules are a strong part of the method signature: `ref`,

`in`, `ref struct`, etc ... Hence any change to the existing rules of a method already results in a binary change to the signature. To give the new lifetime annotations the same impact they will need `modreq` enforcement.

The concern is whether or not this is overkill. It does have the negative impact that making signatures more flexible, by say adding `[DoesNotEscape]` to a parameter, will result in a binary compat change. That trade off means that over time frameworks like BCL likely won't be able to relax such signatures. It could be mitigated to a degree by taking some approach the language does with `in` parameters and only apply `modreq` in virtual positions.

Decision Do not use `modreq` in metadata. The difference between `out` and `ref` is not `modreq` but they now have different ref safe context values. There is no real benefit to only half enforcing the rules with `modreq` here.

Allow multi-dimensional fixed buffers

Should the design for `fixed` buffers be extended to include multi-dimensional style arrays? Essentially allowing for declarations like the following:

```c#
struct Dimensions
{
    int array[42, 13];
}
```

Decision Do not allow for now

Violating scoped

The runtime repository has several non-public APIs that capture `ref` parameters as `ref` fields. These are unsafe because the lifetime of the resulting value is not tracked. For example the `Span<T>(ref T value, int length)` constructor.

The majority of these APIs will likely choose to have proper lifetime tracking on the return which will be achieved simply by updating to C# 11. A few though will want to keep their current semantics of not tracking the return value because their entire intent is to be unsafe. The most notable examples are `MemoryMarshal.CreateSpan` and `MemoryMarshal.CreateReadOnlySpan`. This will be achieved by marking the parameters as `scoped`.

That means the runtime needs an established pattern for unsafely removing `scoped` from a parameter:

1. `Unsafe.AsRef<T>(in T value)` could expand its existing purpose by changing to `scoped in T value`. This would allow it to both remove `in` and `scoped` from parameters. It then becomes the universal "remove ref safety" method
2. Introduce a new method whose entire purpose is to remove `scoped`: `ref T Unsafe.AsUnscoped<T>(scoped in T value)`. This removes `in` as well because if it did not then callers still need a combination of method calls to "remove ref safety" at which point the existing solution is likely sufficient.

Unscoped this by default?

The design only has two locations which are `scoped` by default:

- `this` is `scoped ref`
- `out` is `scoped ref`

The decision on `out` is to significantly reduce the compat burden of `ref` fields and at the same time is a more natural default. It lets developers actually think of `out` as data flowing outward only where as if it's `ref` then the rules must consider data flowing in both directions. This leads to significant developer confusion.

The decision on `this` is undesirable because it means a `struct` cannot return a field by `ref`. This is an important scenario to high perf developers and the `[UnscopedRef]` attribute was added essentially for this one scenario.

Keywords have a high bar and adding it for a single scenario is suspect. As such thought was given to whether we could avoid this keyword at all by making `this` simply `ref` by default and not `scoped ref`. All members that need `this` to be `scoped ref` could do so by marking the method `scoped` (as a method can be marked `readonly` to create a `readonly ref` today).

On a normal `struct` this is mostly a positive change as it only introduces compat issues when a member has a `ref` return. There are **very** few of these methods and a tool could spot these and convert them to `scoped` members quickly.

On a `ref struct` this change introduces significantly bigger compat issues. Consider the following:

```
c#
```

```
ref struct Sneaky
{
    int Field;
    ref int RefField;

    public void SelfAssign()
    {
        // This pattern of ref reassign to fields on this inside instance
methods is now
        // completely legal.
        RefField = ref Field;
    }

    static Sneaky UseExample()
    {
        Sneaky local = default;

        // Error: this is illegal, and must be illegal, by our existing
rules as the
        // ref-safe-context of local is now an input into method arguments
must match.
        local.SelfAssign();

        // This would be dangerous as local now has a dangerous `ref` but
the above
        // prevents us from getting here.
        return local;
    }
}
```

Essentially it would mean all instance method invocations on *mutable* `ref struct` locals would be illegal unless the local was further marked as `scoped`. The rules have to consider the case where fields were ref reassigned to other fields in `this`. A `readonly ref struct` doesn't have this problem because the `readonly` nature prevents ref reassignment. Still this would be a significant back compat breaking change as it would impact virtually every existing mutable `ref struct`.

A `readonly ref struct` though is still problematic once we expand to having `ref` fields to `ref struct`. It allows for the same basic problem by just moving the capture into the value of the `ref` field:

```
c#

readonly ref struct ReadOnlySneaky
{
    readonly int Field;
    readonly ref ReadOnlySpan<int> Span;

    public void SelfAssign()
```

```
    {
        // Instance method captures a ref to itself
        Span = new ReadOnlySpan<int>(ref Field, 1);
    }
}
```

Some thought was given to the idea of having `this` have different defaults based on the type of `struct` or member. For example:

- `this` as `ref`: `struct`, `readonly ref struct` or `readonly member`
- `this` as `scoped ref`: `ref struct` or `readonly ref struct` with `ref` field to `ref struct`

This minimizes compat breaks and maximizes flexibility but at the cost of complicating the story for customers. It also doesn't fully solve the problem because future features, like safe `fixed` buffers, require that a mutable `ref struct` have `ref` returns for fields which don't work by this design alone as it would fall into the `scoped ref` category.

Decision Keep `this` as `scoped ref`

ref fields to ref struct

This feature opens up a new set of ref safe context rules because it allows for a `ref` field to refer to a `ref struct`. This generic nature of `ByReference<T>` meant that up until now the runtime could not have such a construct. As a result all of our rules are written under the assumption this is not possible. The `ref` field feature is largely not about making new rules but codifying the existing rules in our system. Allowing `ref` fields to `ref struct` requires us to codify new rules because there are several new scenarios to consider.

The first is that a `readonly ref` is now capable of storing `ref` state. For example:

```
c#

readonly ref struct Container
{
    readonly ref Span<int> Span;

    void Store(Span<int> span)
    {
        Span = span;
    }
}
```

This means when thinking about method arguments must match rules we must consider `readonly ref T` is potential method output when `T` potentially has a `ref` field to a `ref struct`.

The second issue is language must consider a new type of safe context: *ref-field-safe-context*. All `ref struct` which transitively contain a `ref` field have another escape scope representing the value(s) in the `ref` field(s). In the case of multiple `ref` fields they can be collectively tracked as a single value. The default value for this for parameters is *caller-context*.

```c#
ref struct Nested
{
    ref Span<int> Span;
}

Span<int> M(ref Nested nested) => nested.Span;
```

This value is not related to the *safe-context* of the container; that is as the container context gets smaller it has no impact on the *ref-field-safe-context* of the `ref` field values. Further the *ref-field-safe-context* can never be smaller than the *safe-context* of the container.

```c#
ref struct Nested
{
    ref Span<int> Span;
}

void M(ref Nested nested)
{
    scoped ref Nested refLocal = ref nested;

    // the ref-field-safe-context of local is still *caller-context* which
means the following
    // is illegal
    refLocal.Span = stackalloc int[42];

    scoped Nested valLocal = nested;

    // the ref-field-safe-context of local is still *caller-context* which
means the following
    // is still illegal
    valLocal.Span = stackalloc int[42];
}
```

This *ref-field-safe-context* has essentially always existed. Up until now `ref` fields could only point to normal `struct` hence it was trivially collapsed to *caller-context*. To support `ref` fields to `ref struct` our existing rules need to be updated to take into account this new *ref-safe-context*.

Third the rules for ref reassignment need to be updated to ensure that we don't violate *ref-field-context* for the values. Essentially for `x.e1 = ref e2` where the type of `e1` is a `ref struct` the *ref-field-safe-context* must be equal.

These problems are very solvable. The compiler team has sketched out a few versions of these rules and they largely fall out from our existing analysis. The problem is there is no consuming code for such rules that helps prove out there correctness and usability. This makes us very hesitant to add support because of the fear we'll pick wrong defaults and back the runtime into usability corner when it does take advantage of this. This concern is particularly strong because .NET 8 likely pushes us in this direction with `allow T: ref struct` and `Span<Span<T>>`. The rules would be better written if it's done in conjunction with consumption code.

Decision Delay allowing `ref` field to `ref struct` until .NET 8 where we have scenarios that will help drive the rules around these scenarios.

What will make C# 11.0?

The features outlined in this document don't need to be implemented in a single pass. Instead they can be implemented in phases across several language releases in the following buckets:

1. `ref` fields and `scoped`
2. `[UnscopedRef]`
3. `ref` fields to `ref struct`
4. Sunset restricted types
5. fixed sized buffers

What gets implemented in which release is merely a scoping exercise.

Decision Only (1) and (2) made C# 11.0. The rest will be considered in future versions of C#.

Future Considerations

Advanced lifetime annotations

The lifetime annotations in this proposal are limited in that they allow developers to change the default escape / don't escape behavior of values. This does add powerful flexibility to our model but it does not radically change the set of relationships that can be expressed. At the core the C# model is still effectively binary: can a value be returned or not?

That allows limited lifetime relationships to be understood. For example a value that can't be returned from a method has a smaller lifetime than one that can be returned from a method. There is no way to describe the lifetime relationship between values that can be returned from a method though. Specifically there is no way to say that one value has a larger lifetime than the other once it's established both can be returned from a method. The next step in our lifetime evolution would be allowing such relationships to be described.

Other methods such as Rust allow this type of relationship to be expressed and hence can implement handle more complex `scoped` style operations. Our language could similarly benefit if such a feature were included. At the moment there is no motivating pressure to do this but if there is in the future our `scoped` model could be expanded to included it in a fairly straight forward fashion.

Every `scoped` could be assigned a named lifetime by adding a generic style argument to the syntax. For example `scoped<'a>` is a value that has lifetime `'a`. Constraints like `where` could then be used to describe the relationships between these lifetimes.

```c#
void M(scoped<'a> ref MyStruct s, scoped<'b> Span<int> span)
  where 'b >= 'a
{
    s.Span = span;
}
```

This method defines two lifetimes `'a` and `'b` and there relationship, specifically that `'b` is greater than `'a`. This allows for the callsite to have more granular rules for how values can be safely passed into methods vs. the more coarse grained rules present today.

Related Information

Issues

The following issues are all related to this proposal:

- https://github.com/dotnet/csharplang/issues/1130
- https://github.com/dotnet/csharplang/issues/1147
- https://github.com/dotnet/csharplang/issues/992
- https://github.com/dotnet/csharplang/issues/1314
- https://github.com/dotnet/csharplang/issues/2208
- https://github.com/dotnet/runtime/issues/32060
- https://github.com/dotnet/runtime/issues/61135
- https://github.com/dotnet/csharplang/discussions/78

Proposals

The following proposals are related to this proposal:

- https://github.com/dotnet/csharplang/blob/725763343ad44a9251b03814e6897d87fe553769/proposals/fixed-sized-buffers.md

Existing samples

Utf8JsonReader

This particular snippet requires unsafe because it runs into issues with passing around a `Span<T>` which can be stack allocated to an instance method on a `ref struct`. Even though this parameter is not captured the language must assume it is and hence needlessly causes friction here.

Utf8JsonWriter

This snippet wants to mutate a parameter by escaping elements of the data. The escaped data can be stack allocated for efficiency. Even though the parameter is not escaped the compiler assigns it a *safe-context* of outside the enclosing method because it is a parameter. This means in order to use stack allocation the implementation must use `unsafe` in order to assign back to the parameter after escaping the data.

Fun Samples

ReadOnlySpan<T>

c#

```
public readonly ref struct ReadOnlySpan<T>
{
    readonly ref readonly T _value;
```

```
    readonly int _length;

    public ReadOnlySpan(in T value)
    {
        _value = ref value;
        _length = 1;
    }
}
```

Frugal list

```c#
struct FrugalList<T>
{
    private T _item0;
    private T _item1;
    private T _item2;

    public int Count = 3;

    public FrugalList(){}

    public ref T this[int index]
    {
        [UnscopedRef] get
        {
            switch (index)
            {
                case 0: return ref _item0;
                case 1: return ref _item1;
                case 2: return ref _item2;
                default: throw null;
            }
        }
    }
}
```

Examples and Notes

Below are a set of examples demonstrating how and why the rules work the way they do. Included are several examples showing dangerous behaviors and how the rules prevent them from happening. It's important to keep these in mind when making adjustments to the proposal.

Ref reassignment and call sites

Demonstrating how ref reassignment and method invocation work together.

c#

```
ref struct RS
{
    ref int _refField;

    public ref int Prop => ref _refField;

    public RS(int[] array)
    {
        _refField = ref array[0];
    }

    public RS(ref int i)
    {
        _refField = ref i;
    }

    public RS CreateRS() => ...;

    public ref int M1(RS rs)
    {
        // The call site arguments for Prop contribute here:
        //   - `rs` contributes no ref-safe-context as the corresponding
parameter,
        //      which is `this`, is `scoped ref`
        //   - `rs` contribute safe-context of *caller-context*
        //
        // This is an lvalue invocation and the arguments contribute only
safe-context
        // values of *caller-context*. That means `local1` has ref-safe-
context of
        // *caller-context*
        ref int local1 = ref rs.Prop;

        // Okay: this is legal because `local` has ref-safe-context of
*caller-context*
        return ref local1;

        // The arguments contribute here:
        //   - `this` contributes no ref-safe-context as the corresponding
parameter
        //     is `scoped ref`
        //   - `this` contributes safe-context of *caller-context*
        //
        // This is an rvalue invocation and following those rules the safe-
context of
        // `local2` will be *caller-context*
        RS local2 = CreateRS();

        // Okay: this follows the same analysis as `ref rs.Prop` above
        return ref local2.Prop;
```

```
        // The arguments contribute here:
        //   - `local3` contributes ref-safe-context of *function-member*
        //   - `local3` contributes safe-context of *caller-context*
        //
        // This is an rvalue invocation which returns a `ref struct` and
following those
        // rules the safe-context of `local4` will be *function-member*
        int local3 = 42;
        var local4 = new RS(ref local3);

        // Error:
        // The arguments contribute here:
        //   - `local4` contributes no ref-safe-context as the corresponding
parameter
        //     is `scoped ref`
        //   - `local4` contributes safe-context of *function-member*
        //
        // This is an lvalue invocation and following those rules the ref-
safe-context
        // of the return is *function-member*
        return ref local4.Prop;
    }
}
```

Ref reassignment and unsafe escapes

The reason for the following line in the ref reassignment rules may not be obvious at first glance:

> `e1` must have the same *safe-context* as `e2`

This is because the lifetime of the values pointed to by `ref` locations are invariant. The indirection prevents us from allowing any kind of variance here, even to narrower lifetimes. If narrowing is allowed then it opens up the following unsafe code:

C#

```
void Example(ref Span<int> p)
{
    Span<int> local = stackalloc int[42];
    ref Span<int> refLocal = ref local;

    // The safe-context of refLocal is narrower than p. For a non-ref
reassignment
    // this would be allowed as its safe to assign wider lifetimes to
narrower ones.
    // In the case of ref reassignment though this rule prevents it as the
    // safe-context values are different.
```

```
    refLocal = ref p;

    // If it were allowed this would be legal as the safe-context of
refLocal
    // is *caller-context* and that is satisfied by stackalloc. At the same
time
    // it would be assigning through p and escaping the stackalloc to the
calling
    // method
    //
    // This is equivalent of saying p = stackalloc int[13]!!!
    refLocal = stackalloc int[13];
}
```

For a `ref` to non `ref struct` this rule is trivially satisfied as the values all have the same *safe-context*. This rule really only comes into play when the value is a `ref struct`.

This behavior of `ref` will also be important in a future where we allow `ref` fields to `ref struct`.

scoped locals

The use of `scoped` on locals will be particularly helpful to code patterns which conditionally assign values with different *safe-context* to locals. It means code no longer needs to rely on initialization tricks like `= stackalloc byte[0]` to define a local *safe-context* but now can simply use `scoped`.

c#

```
// Old way
// Span<byte> span = stackalloc byte[0];
// New way
scoped Span<byte> span;
int len = ...;
if (len < MaxStackLen)
{
    span = stackalloc byte[len];
}
else
{
    span = new byte[len];
}
```

This pattern comes up frequently in low level code. When the `ref struct` involved is `Span<T>` the above trick can be used. It is not applicable to other `ref struct` types though and can result in low level code needing to resort to `unsafe` to work around the inability to properly specify the lifetime.

scoped parameter values

One source of repeated friction in low level code is the default escape for parameters is permissive. They are *safe-context* to the *caller-context*. This is a sensible default because it lines up with the coding patterns of .NET as a whole. In low level code though there is a larger use of `ref struct` and this default can cause friction with other parts of the ref safe context rules.

The main friction point occurs because of the method arguments must match rule. This rule most commonly comes into play with instance methods on `ref struct` where at least one parameter is also a `ref struct`. This is a common pattern in low level code where `ref struct` types commonly leverage `Span<T>` parameters in their methods. For example it will occur on any writer style `ref struct` that uses `Span<T>` to pass around buffers.

This rule exists to prevent scenarios like the following:

```c#
ref struct RS
{
    Span<int> _field;
    void Set(Span<int> p)
    {
        _field = p;
    }

    static void DangerousCode(ref RS p)
    {
        Span<int> span = stackalloc int[] { 42 };

        // Error: if allowed this would let the method return a reference to
        // the stack
        p.Set(span);
    }
}
```

Essentially this rule exists because the language must assume that all inputs to a method escape to their maximum allowed *safe-context*. When there are `ref` or `out` parameters, including the receivers, it's possible for the inputs to escape as fields of those `ref` values (as happens in `RS.Set` above).

In practice though there are many such methods which pass `ref struct` as parameters that never intend to capture them in output. It is just a value that is used within the current method. For example:

c#

```
ref struct JsonReader
{
    Span<char> _buffer;
    int _position;

    internal bool TextEquals(ReadOnlySpan<char> text)
    {
        var current = _buffer.Slice(_position, text.Length);
        return current == text;
    }
}

class C
{
    static void M(ref JsonReader reader)
    {
        Span<char> span = stackalloc char[4];
        span[0] = 'd';
        span[1] = 'o';
        span[2] = 'g';

        // Error: The safe-context of `span` is function-member
        // while `reader` is outside function-member hence this fails
        // by the above rule.
        if (reader.TextEquals(span))
        {
            ...
        }
    }
}
```

In order to work around this low level code will resort to `unsafe` tricks to lie to the compiler about the lifetime of their `ref struct`. This significantly reduces the value proposition of `ref struct` as they are meant to be a means to avoid `unsafe` while continuing to write high performance code.

This is where `scoped` is an effective tool on `ref struct` parameters because it removes them from consideration as being returned from the method according to the updated method arguments must match rule. A `ref struct` parameter which is consumed, but never returned, can be labeled as `scoped` to make call sites more flexible.

c#

```
ref struct JsonReader
{
    Span<char> _buffer;
    int _position;
```

```
    internal bool TextEquals(scoped ReadOnlySpan<char> text)
    {
        var current = _buffer.Slice(_position, text.Length);
        return current == text;
    }
}

class C
{
    static void M(ref JsonReader reader)
    {
        Span<char> span = stackalloc char[4];
        span[0] = 'd';
        span[1] = 'o';
        span[2] = 'g';

        // Okay: the compiler never considers `span` as capturable here
hence it doesn't
        // contribute to the method arguments must match rule
        if (reader.TextEquals(span))
        {
            ...
        }
    }
}
```

Preventing tricky ref assignment from readonly mutation

When a `ref` is taken to a `readonly` field in a constructor or `init` member the type is `ref` not `ref readonly`. This is a long standing behavior that allows for code like the following:

c#

```
struct S
{
    readonly int i;

    public S(string s)
    {
        M(ref i);
    }

    static void M(ref int i) { }
}
```

That does pose a potential problem though if such a `ref` were able to be stored into a `ref` field on the same type. It would allow for direct mutation of a `readonly struct` from an instance member:

```c#
readonly ref struct S
{
    readonly int i;
    readonly ref int r;
    public S()
    {
        i = 0;
        r = ref i;
    }

    public void Oops()
    {
        r++;
    }
```

The proposal prevents this though because it violates the ref safe context rules. Consider the following:

- The *ref-safe-context* of `this` is *function-member* and *safe-context* is *caller-context*. These are both standard for `this` in a `struct` member.
- The *ref-safe-context* of `i` is *function-member*. This falls out from the field lifetimes rules. Specifically rule 4.

At that point the line `r = ref i` is illegal by ref reassignment rules.

These rules were not intended to prevent this behavior but do so as a side effect. It's important to keep this in mind for any future rule update to evaluate the impact to scenarios like this.

Silly cyclic assignment

One aspect this design struggled with is how freely a `ref` can be returned from a method. Allowing all `ref` to be returned as freely as normal values is likely what most developers intuitively expect. However it allows for pathological scenarios that the compiler must consider when calculating ref safety. Consider the following:

```c#
ref struct S
{
    int field;
    ref int refField;

    static void SelfAssign(ref S s)
    {
```

```
        s.refField = ref s.field;
    }
}
```

This is not a code pattern that we expect any developers to use. Yet when a `ref` can be returned with the same lifetime as a value it is legal under the rules. The compiler must consider all legal cases when evaluating a method call and this leads to such APIs being effectively unusable.

c#

```
void M(ref S s)
{
    ...
}

void Usage()
{
    // safe-context to caller-context
    S local = default;

    // Error: compiler is forced to assume the worst and concludes a self assignment
    // is possible here and must issue an error.
    M(ref local);
}
```

To make these APIs usable the compiler ensures that the `ref` lifetime for a `ref` parameter is smaller than lifetime of any references in the associated parameter value. This is the rationale for having *ref-safe-context* for `ref` to `ref struct` be *return-only* and `out` be *caller-context*. That prevents cyclic assignment because of the difference in lifetimes.

It is also why `[UnscopedRef]` only promotes the *ref-safe-context* of any `ref` to `ref struct` values to *return-only* and not *caller-context*. Consider that using *caller-context* allows for cyclic assignment and would force a viral use of `[UnscopedRef]` for a `ref struct`:

c#

```
ref struct S
{
    byte Field;

    [UnscopedRef]
    public Span<byte> Data => new Span<byte>(ref Field, 1);
}
```

```
void M(ref S s)
{
    // Error: passing a scoped ref to [UnscopedRef] ref
    Span<byte> span = s.Data;
}
```

This is correctly illegal in that case because the compiler has to consider the pathological case that `S.Data` could cyclic assign via `this`. That forces methods all methods that call `S.Data` to further mark their `ref` parameters as `[UnscopedRef]`. This is viral until the method which creates the value as a local. This is why *return-only* exists as a *safe-context*. It does complicate the spec / implementation but it serves to make the feature significantly more usable.

Note: this cyclic assignment problem does continue to exist for `[UnscopedRef] out` to `ref struct` because that causes the *safe-context* and *ref-safe-context* to be equivalent.

c#

```
ref struct RS
{
    int field;
    ref int refField;
}

void M1(out RS p)
{
    // Error: from method arguments must match:
    // Step 1 would calculate the narrowest escape as *caller-context*
    // Step 2 would fail the assignment check because p has safe-context of
*return-only*
    M2(out p);
}

void M2([UnscopedRef] out RS p)
{
    // The lifetimes of LHS and RHS are equivalent here and hence this is
legal
    p.refField = ref p.Field;
}
```

In terms of advanced annotations the `[UnscopedRef]` design creates the following:

```
ref struct S { }

// C# code
S Create1(ref S p)
```

```
S Create2([UnscopedRef] ref S p)

// Annotation equivalent
scoped<'b> S Create1(scoped<'a> ref scoped<'b> S)
scoped<'a> S Create2(scoped<'a> ref scoped<'b> S)
  where 'b >= 'a
```

readonly cannot be deep through ref fields

Consider the below code sample:

```c#
ref struct S
{
    ref int Field;

    readonly void Method()
    {
        // Legal or illegal?
        Field = 42;
    }
}
```

When designing the rules for `ref` fields on `readonly` instances in a vacuum the rules can be validly designed such that the above is legal or illegal. Essentially `readonly` can validly be deep through a `ref` field or it can apply only to the `ref`. Applying only to the `ref` prevents ref reassignment but allows normal assignment which changes the referred to value.

This design does not exist in a vacuum though, it is designing rules for types that already effectively have `ref` fields. The most prominent of which, `Span<T>`, already has a strong dependency on `readonly` not being deep here. Its primary scenario is the ability to assign to the `ref` field through a `readonly` instance.

```c#
readonly ref struct SpanOfOne
{
    readonly ref int Field;

    public ref int this[int index]
    {
        get
        {
            if (index != 1)
                throw new Exception();
```

```
            return ref Field;
        }
    }
}
```

This means we must choose the shallow interpretation of `readonly`.

Modeling constructors

One subtle design question is: How are constructors bodies modeled for ref safety? Essentially how is the following constructor analyzed?

```c#
ref struct S
{
    int field;

    public S(ref int f)
    {
        field = ref f;
    }
}
```

There are roughly two approaches:

1. Model as a `static` method where `this` is a local where its *safe-context* is *caller-context*
2. Model as a `static` method where `this` is an `out` parameter.

Further a constructor must meet the following invariants:

1. Ensure that `ref` parameters can be captured as `ref` fields.
2. Ensure that `ref` to fields of `this` cannot be escaped through `ref` parameters. That would violate tricky ref assignment.

The intent is to pick the form that satisfies our invariants without introduction of any special rules for constructors. Given that the best model for constructors is viewing `this` as an `out` parameter. The *return only* nature of the `out` allows us to satisfy all the invariants above without any special casing:

```c#
public static void ctor(out S @this, ref int f)
{
    // The ref-safe-context of `ref f` is *return-only* which is also the
```

```
    // safe-context of `this.field` hence this assignment is allowed
    @this.field = ref f;
}
```

Method arguments must match

The method arguments must match rule is a common source of confusion for developers. It's a rule which has a number of special cases that are hard to understand unless you are familiar with the reasoning behind the rule. For the sake of better understanding the reasons for the rule we will simplify *ref-safe-context* and *safe-context* to simply *context*.

Methods can pretty liberally return state passed to them as parameters. Essentially any reachable state which is unscoped can be returned (including returning by `ref`). This can be returned directly through a `return` statement or indirectly by assigning into a `ref` value.

Direct returns don't pose much problems for ref safety. The compiler simply needs to look at all the returnable inputs to a method and then it effectively restricts the return value to be the minimum *context* of the input. That return value then goes through normal processing.

Indirect returns pose a significant problem because all `ref` are both an input and output to the method. These outputs already have a known *context*. The compiler can't infer new ones, it has to consider them at their current level. That means the compiler has to look at every single `ref` which is assignable in the called method, evaluate it's *context*, and then verify no returnable input to the method has a smaller *context* than that `ref`. If any such case exists then the method call must be illegal because it could violate `ref` safety.

Method arguments must match is the process by which the compiler asserts this safety check.

A different way to evaluate this which is often easier for developers to consider is to do the following exercise:

1. Look at the method definition identify all places where state can be indirectly returned: a. Mutable `ref` parameters pointing to `ref struct` b. Mutable `ref` parameters with ref assignable `ref` fields c. Assignable `ref` params or `ref` fields pointing to `ref struct` (consider recursively)
2. Look at the call site a. Identify the contexts that line up with the locations identified above b. Identify the contexts of all inputs to the method that are returnable (don't

line up with `scoped` parameters)

If any value in 2.b is smaller than 2.a then the method call must be illegal. Let's look at a few examples to illustrate the rules:

```c#
ref struct R { }

class Program
{
    static void F0(ref R a, scoped ref R b) => throw null;

    static void F1(ref R x, scoped R y)
    {
        F0(ref x, ref y);
    }
}
```

Looking at the call to `F0` lets go through (1) and (2). The parameters with potential for indirect return are `a` and `b` as both can be directly assigned. The arguments which line up to those parameters are:

- `a` which maps to `x` that has *context* of *caller-context*
- `b` which maps to `y` that has with *context* of *function-member*

The set of returnable input to the method are

- `x` with *escape-scope* of *caller-context*
- `ref x` with *escape-scope* of *caller-context*
- `y` with *escape-scope* of *function-member*

The value `ref y` is not returnable since it maps to a `scoped ref` hence it is not considered an input. But given that there is at least one input with a smaller *escape scope* (`y` argument) than one of the outputs (`x` argument) the method call is illegal.

A different variation is the following:

```c#
ref struct R { }

class Program
{
    static void F0(ref R a, ref int b) => throw null;

    static void F1(ref R x)
```

```
    {
        int y = 42;
        F0(ref x, ref y);
    }
}
```

Again the parameters with potential for indirect return are `a` and `b` as both can be directly assigned. But `b` can be excluded because it does not point to a `ref struct` hence cannot be used to store `ref` state. Thus we have:

- `a` which maps to `x` that has *context* of *caller-context*

The set of returnable input to the method are:

- `x` with *context* of *caller-context*
- `ref x` with *context* of *caller-context*
- `ref y` with *context* of *function-member*

Given that there is at least one input with a smaller *escape scope* (`ref y` argument) than one of the outputs (`x` argument) the method call is illegal.

This is the logic that the method arguments must match rule is trying to encompass. It goes further as it considers both `scoped` as a way to remove inputs from consideration and `readonly` as a way to remove `ref` as an output (can't assign into a `readonly ref` so it can't be a source of output). These special cases do add complexity to the rules but it's done so for the benefit of the developer. The compiler seeks to remove all inputs and outputs it knows can't contribute to the result to give developers maximum flexibility when calling a member. Much like overload resolution it's worth the effort to make our rules more complex when it creates more flexibility for consumers.

Examples of inferred *safe-context* of declaration expressions

Related to Infer *safe-context* of declaration expressions.

C#

```
ref struct RS
{
    public RS(ref int x) { } // assumed to be able to capture 'x'

    static void M0(RS input, out RS output) => output = input;

    static void M1()
    {
        var i = 0;
```

```
        var rs1 = new RS(ref i); // safe-context of 'rs1' is function-member
        M0(rs1, out var rs2); // safe-context of 'rs2' is function-member
    }

    static void M2(RS rs1)
    {
        M0(rs1, out var rs2); // safe-context of 'rs2' is function-member
    }

    static void M3(RS rs1)
    {
        M0(rs1, out scoped var rs2); // 'scoped' modifier forces safe-
context of 'rs2' to the current local context (function-member or narrower).
    }
}
```

Note that the local context which results from the `scoped` modifier is the narrowest which could possibly be used for the variable--to be any narrower would mean the expression refers to variables which are only declared in a narrower context than the expression.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# feature specification feedback

The C# feature specifications are open source. Provide feedback here.

Open a documentation issue

Provide product feedback

Inline Arrays

Article • 08/16/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Provide a general-purpose and safe mechanism for consuming struct types utilizing InlineArrayAttribute feature. Provide a general-purpose and safe mechanism for declaring inline arrays within C# classes, structs, and interfaces.

> Note: Previous versions of this spec used the terms "ref-safe-to-escape" and "safe-to-escape", which were introduced in the Span safety feature specification. The ECMA standard committee changed the names to "ref-safe-context" and "safe-context", respectively. The values of the safe context have been refined to use "declaration-block", "function-member", and "caller-context" consistently. The speclets had used different phrasing for these terms, and also used "safe-to-return" as a synonym for "caller-context". This speclet has been updated to use the terms in the C# 7.3 standard.

Motivation

This proposal plans to address the many limitations of https://github.com/dotnet/csharpstandard/blob/draft-v8/standard/unsafe-code.md#228-fixed-size-buffers. Specifically it aims to allow:

- accessing elements of struct types utilizing InlineArrayAttribute feature;
- the declaration of inline arrays for managed and unmanaged types in a `struct`, `class`, or `interface`.

And provide language safety verification for them.

Detailed Design

Recently runtime added InlineArrayAttribute feature. In short, a user can declare a structure type like the following:

```C#
[System.Runtime.CompilerServices.InlineArray(10)]
public struct Buffer
{
    private object _element0;
}
```

Runtime provides a special type layout for the `Buffer` type:

- The size of the type is extended to fit 10 (the number comes from the InlineArray attribute) elements of `object` type (the type comes from the type of the only instance field in the struct, `_element0` in this example).
- The first element is aligned with the instance field and with the beginning of the struct
- The elements are laid out sequentially in memory as though they are elements of an array.

Runtime provides regular GC tracking for all elements in the struct.

This proposal will refer to types like this as "inline array types".

Elements of an inline array type can be accessed through pointers or through span instances returned by System.Runtime.InteropServices.MemoryMarshal.CreateSpan/System.Runtime.InteropServices.MemoryMarshal.CreateReadOnlySpan APIs. However, neither the pointer approach, nor the APIs provide type and bounds checking out of the box.

Language will provide a type-safe/ref-safe way for accessing elements of inline array types. The access will be span based. This limits support to inline array types with element types that can be used as a type argument. For example, a pointer type cannot be used as an element type. Other examples the span types.

Obtaining instances of span types for an inline array type

Since there is a guarantee that the first element in an inline array type is aligned at the beginning of the type (no gap), compiler will use the following code to get a `Span` value:

```C#
```

```
MemoryMarshal.CreateSpan(ref Unsafe.As<TBuffer, TElement>(ref buffer), size)
```

And the following code to get a `ReadOnlySpan` value:

C#

```
MemoryMarshal.CreateReadOnlySpan(ref Unsafe.As<TBuffer, TElement>(ref
Unsafe.AsRef(in buffer)), size)
```

In order to reduce IL size at use sites compiler should be able to add two generic reusable helpers into private implementation detail type and use them across all use sites in the same program.

C#

```
public static System.Span<TElement> InlineArrayAsSpan<TBuffer, TElement>(ref
TBuffer buffer, int size) where TBuffer : struct
{
    return MemoryMarshal.CreateSpan(ref Unsafe.As<TBuffer, TElement>(ref
buffer), size);
}

public static System.ReadOnlySpan<TElement>
InlineArrayAsReadOnlySpan<TBuffer, TElement>(in TBuffer buffer, int size)
where TBuffer : struct
{
    return MemoryMarshal.CreateReadOnlySpan(ref Unsafe.As<TBuffer, TElement>
(ref Unsafe.AsRef(in buffer)), size);
}
```

Element access

The Element access will be extended to support inline array element access.

An *element_access* consists of a *primary_no_array_creation_expression*, followed by a "`[`" token, followed by an *argument_list*, followed by a "`]`" token. The *argument_list* consists of one or more *arguments*, separated by commas.

ANTLR

```
element_access
    : primary_no_array_creation_expression '[' argument_list ']'
    ;
```

The *argument_list* of an *element_access* is not allowed to contain `ref` or `out` arguments.

An *element_access* is dynamically bound (§11.3.3) if at least one of the following holds:

- The *primary_no_array_creation_expression* has compile-time type `dynamic`.
- At least one expression of the *argument_list* has compile-time type `dynamic` and the *primary_no_array_creation_expression* does not have an array type, **and the *primary_no_array_creation_expression* does not have an inline array type or there is more than one item in the argument list**.

In this case, the compiler classifies the *element_access* as a value of type `dynamic`. The rules below to determine the meaning of the *element_access* are then applied at run-time, using the run-time type instead of the compile-time type of those of the *primary_no_array_creation_expression* and *argument_list* expressions which have the compile-time type `dynamic`. If the *primary_no_array_creation_expression* does not have compile-time type `dynamic`, then the element access undergoes a limited compile-time check as described in §11.6.5.

If the *primary_no_array_creation_expression* of an *element_access* is a value of an *array_type*, the *element_access* is an array access (§11.7.10.2). **If the *primary_no_array_creation_expression* of an *element_access* is a variable or value of an inline array type and the *argument_list* consists of a single argument, the *element_access* is an inline array element access.** Otherwise, the *primary_no_array_creation_expression* shall be a variable or value of a class, struct, or interface type that has one or more indexer members, in which case the *element_access* is an indexer access (§11.7.10.3).

Inline array element access

For an inline array element access, the *primary_no_array_creation_expression* of the *element_access* must be a variable or value of an inline array type. Furthermore, the *argument_list* of an inline array element access is not allowed to contain named arguments. The *argument_list* must contain a single expression, and the expression must be

- of type `int`, or
- implicitly convertible to `int`, or
- implicitly convertible to `System.Index`, or
- implicitly convertible to `System.Range`.

When the expression type is int

If *primary_no_array_creation_expression* is a writable variable, the result of evaluating an inline array element access is a writable variable equivalent to invoking public ref T this[int index] { get; } with that integer value on an instance of `System.Span<T>` returned by `System.Span<T> InlineArrayAsSpan` method on *primary_no_array_creation_expression*. For the purpose of ref-safety analysis the *ref-safe-context/safe-context* of the access are equivalent to the same for an invocation of a method with the signature `static ref T GetItem(ref InlineArrayType array)`. The resulting variable is considered movable if and only if *primary_no_array_creation_expression* is movable.

If *primary_no_array_creation_expression* is a readonly variable, the result of evaluating an inline array element access is a readonly variable equivalent to invoking public ref readonly T this[int index] { get; } with that integer value on an instance of `System.ReadOnlySpan<T>` returned by `System.ReadOnlySpan<T> InlineArrayAsReadOnlySpan` method on *primary_no_array_creation_expression*. For the purpose of ref-safety analysis the *ref-safe-context/safe-context* of the access are equivalent to the same for an invocation of a method with the signature `static ref readonly T GetItem(in InlineArrayType array)`. The resulting variable is considered movable if and only if *primary_no_array_creation_expression* is movable.

If *primary_no_array_creation_expression* is a value, the result of evaluating an inline array element access is a value equivalent to invoking public ref readonly T this[int index] { get; } with that integer value on an instance of `System.ReadOnlySpan<T>` returned by `System.ReadOnlySpan<T> InlineArrayAsReadOnlySpan` method on *primary_no_array_creation_expression*. For the purpose of ref-safety analysis the *ref-safe-context/safe-context* of the access are equivalent to the same for an invocation of a method with the signature `static T GetItem(InlineArrayType array)`.

For example:

C#

```
[System.Runtime.CompilerServices.InlineArray(10)]
public struct Buffer10<T>
{
    private T _element0;
}

void M1(Buffer10<int> x)
{
    ref int a = ref x[0]; // Ok, equivalent to `ref int a = ref
InlineArrayAsSpan<Buffer10<int>, int>(ref x, 10)[0]`
}

void M2(in Buffer10<int> x)
{
```

```
    ref readonly int a = ref x[0]; // Ok, equivalent to `ref readonly int a
= ref InlineArrayAsReadOnlySpan<Buffer10<int>, int>(in x, 10)[0]`
    ref int b = ref x[0]; // An error, `x` is a readonly variable => `x[0]`
is a readonly variable
}

Buffer10<int> GetBuffer() => default;

void M3()
{
    int a = GetBuffer()[0]; // Ok, equivalent to `int a =
InlineArrayAsReadOnlySpan<Buffer10<int>, int>(GetBuffer(), 10)[0]`
    ref readonly int b = ref GetBuffer()[0]; // An error, `GetBuffer()[0]`
is a value
    ref int c = ref GetBuffer()[0]; // An error, `GetBuffer()[0]` is a value
}
```

Indexing into an inline array with a constant expression outside of the declared inline array bounds is a compile time error.

When the expression is implicitly convertible to `int`

The expression is converted to int and then the element access is interpreted as described in **When the expression type is int** section.

When the expression implicitly convertible to `System.Index`

The expression is converted to `System.Index`, which is then transformed to an int-based index value as described at https://github.com/dotnet/csharplang/blob/main/proposals/csharp-8.0/ranges.md#implicit-index-support, assuming that the length of the collection is known at compile time and is equal to the amount of elements in the inline array type of the *primary_no_array_creation_expression*. Then the element access is interpreted as described in **When the expression type is int** section.

When the expression implicitly convertible to `System.Range`

If *primary_no_array_creation_expression* is a writable variable, the result of evaluating an inline array element access is a value equivalent to invoking public Span<T> Slice (int start, int length) on an instance of `System.Span<T>` returned by `System.Span<T> InlineArrayAsSpan` method on *primary_no_array_creation_expression*. For the purpose of ref-safety analysis the *ref-safe-context*/*safe-context* of the access are equivalent to the same for an invocation of a method with the signature `static System.Span<T> GetSlice(ref InlineArrayType array)`.

If *primary_no_array_creation_expression* is a readonly variable, the result of evaluating an inline array element access is a value equivalent to invoking public ReadOnlySpan<T> Slice (int start, int length) on an instance of `System.ReadOnlySpan<T>` returned by `System.ReadOnlySpan<T> InlineArrayAsReadOnlySpan` method on *primary_no_array_creation_expression*. For the purpose of ref-safety analysis the *ref-safe-context/safe-context* of the access are equivalent to the same for an invocation of a method with the signature `static System.ReadOnlySpan<T> GetSlice(in InlineArrayType array)`.

If *primary_no_array_creation_expression* is a value, an error is reported.

The arguments for the `Slice` method invocation are calculated from the index expression converted to `System.Range` as described at https://github.com/dotnet/csharplang/blob/main/proposals/csharp-8.0/ranges.md#implicit-range-support, assuming that the length of the collection is known at compile time and is equal to the amount of elements in the inline array type of the *primary_no_array_creation_expression*.

Compiler can omit the `Slice` call if it is known at compile time that `start` is 0 and `length` is less or equal to the amount of elements in the inline array type. Compiler can also report an error if it is known at compile time that slicing goes out of inline array bounds.

For example:

C#

```
void M1(Buffer10<int> x)
{
    System.Span<int> a = x[..]; // Ok, equivalent to `System.Span<int> a = InlineArrayAsSpan<Buffer10<int>, int>(ref x, 10).Slice(0, 10)`
}

void M2(in Buffer10<int> x)
{
    System.ReadOnlySpan<int> a = x[..]; // Ok, equivalent to `System.ReadOnlySpan<int> a = InlineArrayAsReadOnlySpan<Buffer10<int>, int>(in x, 10).Slice(0, 10)`
    System.Span<int> b = x[..]; // An error, System.ReadOnlySpan<int> cannot be converted to System.Span<int>
}

Buffer10<int> GetBuffer() => default;

void M3()
{
```

```
    _ = GetBuffer()[..]; // An error, `GetBuffer()` is a value
}
```

Conversions

A new conversion, an inline array conversion, from expression will be added. The inline array conversion is a standard conversion.

There is an implicit conversion from expression of an inline array type to the following types:

- `System.Span<T>`
- `System.ReadOnlySpan<T>`

However, converting a readonly variable to `System.Span<T>` or converting a value to either type is an error.

For example:

C#

```
void M1(Buffer10<int> x)
{
    System.ReadOnlySpan<int> a = x; // Ok, equivalent to
`System.ReadOnlySpan<int> a = InlineArrayAsReadOnlySpan<Buffer10<int>, int>
(in x, 10)`
    System.Span<int> b = x; // Ok, equivalent to `System.Span<int> b =
InlineArrayAsSpan<Buffer10<int>, int>(ref x, 10)`
}

void M2(in Buffer10<int> x)
{
    System.ReadOnlySpan<int> a = x; // Ok, equivalent to
`System.ReadOnlySpan<int> a = InlineArrayAsReadOnlySpan<Buffer10<int>, int>
(in x, 10)`
    System.Span<int> b = x; // An error, readonly mismatch
}

Buffer10<int> GetBuffer() => default;

void M3()
{
    System.ReadOnlySpan<int> a = GetBuffer(); // An error, ref-safety
    System.Span<int> b = GetBuffer(); // An error, ref-safety
}
```

For the purpose of ref-safety analysis the *safe-context* of the conversion is equivalent to *safe-context* for an invocation of a method with the signature `static System.Span<T>`

`Convert(ref InlineArrayType array)`, or `static System.ReadOnlySpan<T> Convert(in InlineArrayType array)`.

List patterns

List patterns will not be supported for instances of inline array types.

Definite assignment checking

Regular definite assignment rules are applicable to variables that have an inline array type.

Collection literals

An inline array type is a valid *constructible collection* target type for a collection literal.

For example:

```
C#

Buffer10<int> b = [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]; // initializes user-defined inline array
```

The length of the collection literal must match the length of the target inline array type. If the length of the literal is known at compile time and it doesn't match the target length, an error is reported. Otherwise, an exception is going to be thrown at runtime once the mismatch is encountered. The exact exception type is TBD. Some candidates are: System.NotSupportedException, System.InvalidOperationException.

An instance of an inline array type is a valid expression in a *spread_element*.

Validation of the InlineArrayAttribute applications

Compiler will validate the following aspects of the InlineArrayAttribute applications:

- The target type is a struct
- The target type has only one field
- Specified length > 0
- The target struct doesn't have an explicit layout specified

Inline Array elements in an object initializer

By default, element initialization will not be supported via *initializer_target* of form `'['` `argument_list ']'` (see https://github.com/dotnet/csharpstandard/blob/draft-v8/standard/expressions.md#128163-object-initializers):

C#

```
static C M2() => new C() { F = {[0] = 111} }; // error CS1913: Member '[0]'
cannot be initialized. It is not a field or property.

class C
{
    public Buffer10<int> F;
}
```

However, if the inline array type explicitly defines suitable indexer, object initializer will use it:

C#

```
static C M2() => new C() { F = {[0] = 111} }; // Ok, indexer is invoked

class C
{
    public Buffer10<int> F;
}

[System.Runtime.CompilerServices.InlineArray(10)]
public struct Buffer10<T>
{
    private T _element0;

    public T this[int i]
    {
        get => this[i];
        set => this[i] = value;
    }
}
```

The foreach statement

The foreach statement will be adjusted to allow usage of an inline array type as a collection in a foreach statement.

For example:

C#

```
foreach (var a in getBufferAsValue())
{
    WriteLine(a);
}

foreach (var b in getBufferAsWritableVariable())
{
    WriteLine(b);
}

foreach (var c in getBufferAsReadonlyVariable())
{
    WriteLine(c);
}

Buffer10<int> getBufferAsValue() => default;
ref Buffer10<int> getBufferAsWritableVariable() => default;
ref readonly Buffer10<int> getBufferAsReadonlyVariable() => default;
```

is equivalent to:

C#

```
Buffer10<int> temp = getBufferAsValue();
foreach (var a in (System.ReadOnlySpan<int>)temp)
{
    WriteLine(a);
}

foreach (var b in (System.Span<int>)getBufferAsWritableVariable())
{
    WriteLine(b);
}

foreach (var c in (System.ReadOnlySpan<int>)getBufferAsReadonlyVariable())
{
    WriteLine(c);
}
```

We will support `foreach` over inline arrays, even if it starts as restricted in `async` methods due to involvement of the span types into the translation.

Open design questions

Alternatives

Inline array type syntax

The grammar at https://github.com/dotnet/csharpstandard/blob/draft-v8/standard/types.md#821-general will be adjusted as follows:

diff

```
array_type
    : non_array_type rank_specifier+
    ;

rank_specifier
    : '[' ','* ']'
+   | '[' constant_expression ']'
    ;
```

The type of the *constant_expression* must be implicitly convertible to type `int`, and the value must be a non-zero positive integer.

The relevant part of the https://github.com/dotnet/csharpstandard/blob/draft-v8/standard/arrays.md#1621-general section will be adjusted as follows.

The grammar productions for array types are provided in https://github.com/dotnet/csharpstandard/blob/draft-v8/standard/types.md#821-general.

An array type is written as a *non_array_type* followed by one or more *rank_specifier*s.

A *non_array_type* is any *type* that is not itself an *array_type*.

The rank of an array type is given by the leftmost *rank_specifier* in the *array_type*: A *rank_specifier* indicates that the array is an array with a rank of one plus the number of "`,`" tokens in the *rank_specifier*.

The element type of an array type is the type that results from deleting the leftmost *rank_specifier*:

- An array type of the form `T[ constant_expression ]` is an anonymous inline array type with length denoted by *constant_expression* and a non-array element type `T`.
- An array type of the form `T[ constant_expression ][R₁]...[Rₓ]` is an anonymous inline array type with length denoted by *constant_expression* and an element type `T[R₁]...[Rₓ]`.
- An array type of the form `T[R]` (where R is not a *constant_expression*) is a regular array type with rank `R` and a non-array element type `T`.

- An array type of the form `T[R][R₁]...[Rₓ]` (where R is not a *constant_expression*) is a regular array type with rank `R` and an element type `T[R₁]...[Rₓ]`.

In effect, the *rank_specifier*s are read from left to right *before* the final non-array element type.

> *Example*: The type in `int[][,,][,]` is a single-dimensional array of three-dimensional arrays of two-dimensional arrays of `int`. *end example*

At run-time, a value of a regular array type can be `null` or a reference to an instance of that array type.

> *Note*: Following the rules of https://github.com/dotnet/csharpstandard/blob/draft-v8/standard/arrays.md#166-array-covariance, the value may also be a reference to a covariant array type. *end note*

An anonymous inline array type is a compiler synthesized inline array type with internal accessibility. The element type must be a type that can be used as a type argument. Unlike an explicitly declared inline array type, an anonymous inline array type cannot be referenced by name, it can be referenced only by *array_type* syntax. In context of the same program, any two *array_type*s denoting inline array types of the same element type and of the same length, refer to the same anonymous inline array type.

Besides internal accessibility, compiler will prevent consumption of APIs utilizing anonymous inline array types across assembly boundaries by using a required custom modifier (exact type TBD) applied to an anonymous inline array type reference in the signature.

Array creation expressions

Array creation expressions

ANTLR

```
array_creation_expression
    : 'new' non_array_type '[' expression_list ']' rank_specifier*
      array_initializer?
    | 'new' array_type array_initializer
    | 'new' rank_specifier array_initializer
    ;
```

Given the current grammar, use of a *constant_expression* in place of the *expression_list* already has meaning of allocating a regular single-dimensional array type of the specified length. Therefore, *array_creation_expression* will continue to represent an allocation of a regular array.

However, the new form of the *rank_specifier* could be used to incorporate an anonymous inline array type into the element type of the allocated array.

For example, the following expressions create a regular array of length 2 with an element type of an anonymous inline array type with element type int and length 5:

C#

```
new int[2][5];
new int[][5] {default, default};
new [] {default(int[5]), default(int[5])};
```

Array initializers

The Array initializers section will be adjusted to allow use of *array_initializer* to initialize inline array types (no changes to the grammar necessary).

ANTLR

```
array_initializer
    : '{' variable_initializer_list? '}'
    | '{' variable_initializer_list ',' '}'
    ;

variable_initializer_list
    : variable_initializer (',' variable_initializer)*
    ;

variable_initializer
    : expression
    | array_initializer
    ;
```

The length of the inline array must be explicitly provided by the target type.

For example:

C#

```
int[5] a = {1, 2, 3, 4, 5}; // initializes anonymous inline array of length
5
Buffer10<int> b = {1, 2, 3, 4, 5, 6, 7, 8, 9, 10}; // initializes user-
```

```
defined inline array
var c = new int[][] {{11, 12}, {21, 22}, {31, 32}}; // An error for the
nested array initializer
var d = new int[][2] {{11, 12}, {21, 22}, {31, 32}}; // An error for the
nested array initializer
```

Detailed Design (Option 2)

Note, that for the purpose of this proposal a term "fixed-size buffer" refers to a the proposed "safe fixed-size buffer" feature rather than to a buffer described at https://github.com/dotnet/csharpstandard/blob/draft-v8/standard/unsafe-code.md#228-fixed-size-buffers.

In this design, fixed-size buffer types do not get general special treatment by the language. There is a special syntax to declare members that represent fixed-size buffers and new rules around consuming those members. They are not fields from the language point of view.

The grammar for *variable_declarator* in https://github.com/dotnet/csharpstandard/blob/draft-v8/standard/classes.md#145-fields will be extended to allow specifying the size of the buffer:

diff

```
field_declaration
    : attributes? field_modifier* type variable_declarators ';'
    ;

field_modifier
    : 'new'
    | 'public'
    | 'protected'
    | 'internal'
    | 'private'
    | 'static'
    | 'readonly'
    | 'volatile'
    | unsafe_modifier   // unsafe code support
    ;

variable_declarators
    : variable_declarator (',' variable_declarator)*
    ;

variable_declarator
    : identifier ('=' variable_initializer)?
+   | fixed_size_buffer_declarator
    ;
```

```
fixed_size_buffer_declarator
    : identifier '[' constant_expression ']'
    ;
```

A *fixed_size_buffer_declarator* introduces a fixed-size buffer of a given element type.

The buffer element type is the *type* specified in `field_declaration`. A fixed-size buffer declarator introduces a new member and consists of an identifier that names the member, followed by a constant expression enclosed in `[` and `]` tokens. The constant expression denotes the number of elements in the member introduced by that fixed-size buffer declarator. The type of the constant expression must be implicitly convertible to type `int`, and the value must be a non-zero positive integer.

The elements of a fixed-size buffer shall be laid out sequentially in memory as though they are elements of an array.

A *field_declaration* with a *fixed_size_buffer_declarator* in an interface must have `static` modifier.

Depending on the situation (details are specified below), an access to a fixed-size buffer member is classified as a value (never a variable) of either `System.ReadOnlySpan<S>` or `System.Span<S>`, where S is the element type of the fixed-size buffer. Both types provide indexers returning a reference to a specific element with appropriate "readonly-ness", which prevents direct assignment to the elements when language rules don't permit that.

This limits the set of types that can be used as a fixed-size buffer element type to types that can be used as type arguments. For example, a pointer type cannot be used as an element type.

The resulting span instance will have a length equal to the size declared on the fixed-size buffer. Indexing into the span with a constant expression outside of the declared fixed-size buffer bounds is a compile time error.

The *safe-context* of the value will be equal to the *safe-context* of the container, just as it would if the backing data was accessed as a field.

Fixed-size buffers in expressions

Member lookup of a fixed-size buffer member proceeds exactly like member lookup of a field.

A fixed-size buffer can be referenced in an expression using a *simple_name* or a *member_access* .

When an instance fixed-size buffer member is referenced as a simple name, the effect is the same as a member access of the form `this.I`, where `I` is the fixed-size buffer member. When a static fixed-size buffer member is referenced as a simple name, the effect is the same as a member access of the form `E.I`, where `I` is the fixed-size buffer member and `E` is the declaring type.

Non-readonly fixed-size buffers

In a member access of the form `E.I`, if `E` is of a struct type and a member lookup of `I` in that struct type identifies a non-readonly instance fixed-size member, then `E.I` is evaluated and classified as follows:

- If `E` is classified as a value, then `E.I` can be used only as a *primary_no_array_creation_expression* of an element access with index of `System.Index` type, or of a type implicitly convertible to int. Result of the element access is a fixed-size member's element at the specified position, classified as a value.
- Otherwise, if `E` is classified as a readonly variable and the result of the expression is classified as a value of type `System.ReadOnlySpan<S>`, where S is the element type of `I`. The value can be used to access member's elements.
- Otherwise, `E` is classified as a writable variable and the result of the expression is classified as a value of type `System.Span<S>`, where S is the element type of `I`. The value can be used to access member's elements.

In a member access of the form `E.I`, if `E` is of a class type and a member lookup of `I` in that class type identifies a non-readonly instance fixed-size member, then `E.I` is evaluated and classified as a value of type `System.Span<S>`, where S is the element type of `I`.

In a member access of the form `E.I`, if member lookup of `I` identifies a non-readonly static fixed-size member, then `E.I` is evaluated and classified as a value of type `System.Span<S>`, where S is the element type of `I`.

Readonly fixed-size buffers

When a *field_declaration* includes a `readonly` modifier, the member introduced by the fixed_size_buffer_declarator is a ***readonly fixed-size buffer***. Direct assignments to

elements of a readonly fixed-size buffer can only occur in an instance constructor, init member or static constructor in the same type. Specifically, direct assignments to an element of readonly fixed-size buffer are permitted only in the following contexts:

- For an instance member, in the instance constructors or init member of the type that contains the member declaration; for a static member, in the static constructor of the type that contains the member declaration. These are also the only contexts in which it is valid to pass an element of readonly fixed-size buffer as an `out` or `ref` parameter.

Attempting to assign to an element of a readonly fixed-size buffer or pass it as an `out` or `ref` parameter in any other context is a compile-time error. This is achieved by the following.

A member access for a readonly fixed-size buffer is evaluated and classified as follows:

- In a member access of the form `E.I`, if `E` is of a struct type and `E` is classified as a value, then `E.I` can be used only as a *primary_no_array_creation_expression* of an element access with index of `System.Index` type, or of a type implicitly convertible to int. Result of the element access is a fixed-size member's element at the specified position, classified as a value.
- If access occurs in a context where direct assignments to an element of readonly fixed-size buffer are permitted, the result of the expression is classified as a value of type `System.Span<S>`, where S is the element type of the fixed-size buffer. The value can be used to access member's elements.
- Otherwise, the expression is classified as a value of type `System.ReadOnlySpan<S>`, where S is the element type of the fixed-size buffer. The value can be used to access member's elements.

Definite assignment checking

Fixed-size buffers are not subject to definite assignment-checking, and fixed-size buffer members are ignored for purposes of definite-assignment checking of struct type variables.

When a fixed-size buffer member is static or the outermost containing struct variable of a fixed-size buffer member is a static variable, an instance variable of a class instance, or an array element, the elements of the fixed-size buffer are automatically initialized to their default values. In all other cases, the initial content of a fixed-size buffer is undefined.

Metadata

Metadata emit and code generation

For metadata encoding compiler will rely on recently added System.Runtime.CompilerServices.InlineArrayAttribute.

Fixed-size buffers like:

C#

```
public partial class C
{
    public int buffer1[10];
    public readonly int buffer2[10];
}
```

will be emitted as fields of a specially decorated struct type.

Equivalent C# code will be:

C#

```
public partial class C
{
    public Buffer10<int> buffer1;
    public readonly Buffer10<int> buffer2;
}

[System.Runtime.CompilerServices.InlineArray(10)]
public struct Buffer10<T>
{
    private T _element0;

    [UnscopedRef]
    public System.Span<T> AsSpan()
    {
        return System.Runtime.InteropServices.MemoryMarshal.CreateSpan(ref
_element0, 10);
    }

    [UnscopedRef]
    public readonly System.ReadOnlySpan<T> AsReadOnlySpan()
    {
        return
System.Runtime.InteropServices.MemoryMarshal.CreateReadOnlySpan(
                    ref System.Runtime.CompilerServices.Unsafe.AsRef(in
_element0), 10);
```

```
    }
}
```

The actual naming conventions for the type and its members are TBD. The framework will likely include a set of predefined "buffer" types that cover a limited set of buffer sizes. When a predefined type doesn't exist, compiler will synthesize it in the module being built. Names of the generated types will be "speakable" in order to support consumption from other languages.

A code generated for an access like:

C#

```
public partial class C
{
    void M1(int val)
    {
        buffer1[1] = val;
    }

    int M2()
    {
        return buffer2[1];
    }
}
```

will be equivalent to:

C#

```
public partial class C
{
    void M1(int val)
    {
        buffer.AsSpan()[1] = val;
    }

    int M2()
    {
        return buffer2.AsReadOnlySpan()[1];
    }
}
```

Metadata import

When compiler imports a field declaration of type *T* and the following conditions are all met:

- *T* is a struct type decorated with the `InlineArray` attribute, and
- The first instance field declared within *T* has type *F*, and
- There is a `public System.Span<F> AsSpan()` within *T*, and
- There is a `public readonly System.ReadOnlySpan<T> AsReadOnlySpan()` or `public System.ReadOnlySpan<T> AsReadOnlySpan()` within *T*.

the field will be treated as C# fixed-size buffer with element type *F*. Otherwise, the field will be treated as a regular field of type *T*.

Method or property group like approach in the language

One thought is to treat these members more like method groups, in that they aren't automatically a value in and of themselves, but can be made into one if necessary. Here's how that would work:

- Safe fixed-size buffer accesses have their own classification (just like e.g. method groups and lambdas)
- They can be indexed directly as a language operation (not via span types) to produce a variable (which is readonly if the buffer is in a readonly context, just the same as fields of a struct)
- They have implicit conversions-from-expression to `Span<T>` and `ReadOnlySpan<T>`, but use of the former is an error if they are in a readonly context
- Their natural type is `ReadOnlySpan<T>`, so that's what they contribute if they participate in type inference (e.g., var, best-common-type or generic)

C/C++ fixed-size buffers

C/C++ has a different notion of fixed-size buffers. For example, there is a notion of "zero-length fixed sized buffers", which is often used as a way to indicate that the data is "variable length". It is not a goal of this proposal to be able to interop with that.

LDM meetings

- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-04-03.md#fixed-size-buffers
- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-04-10.md#fixed-size-buffers
- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-05-01.md#fixed-size-buffers

- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-05-03.md#inline-arrays
- https://github.com/dotnet/csharplang/blob/main/meetings/2023/LDM-2023-05-17.md#inline-arrays

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# feature specification feedback

C# feature specification is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

default interface methods

Article • 08/12/2022

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

Add support for *virtual extension methods* - methods in interfaces with concrete implementations. A class or struct that implements such an interface is required to have a single *most specific* implementation for the interface method, either implemented by the class or struct, or inherited from its base classes or interfaces. Virtual extension methods enable an API author to add methods to an interface in future versions without breaking source or binary compatibility with existing implementations of that interface.

These are similar to Java's "Default Methods".

(Based on the likely implementation technique) this feature requires corresponding support in the CLI/CLR. Programs that take advantage of this feature cannot run on earlier versions of the platform.

Motivation

The principal motivations for this feature are

- Default interface methods enable an API author to add methods to an interface in future versions without breaking source or binary compatibility with existing implementations of that interface.
- The feature enables C# to interoperate with APIs targeting Android (Java) and iOS (Swift), which support similar features.
- As it turns out, adding default interface implementations provides the elements of the "traits" language feature

(https://en.wikipedia.org/wiki/Trait_(computer_programming)). Traits have proven to be a powerful programming technique (http://scg.unibe.ch/archive/papers/Scha03aTraits.pdf).

Detailed design

The syntax for an interface is extended to permit

- member declarations that declare constants, operators, static constructors, and nested types;
- a *body* for a method or indexer, property, or event accessor (that is, a "default" implementation);
- member declarations that declare static fields, methods, properties, indexers, and events;
- member declarations using the explicit interface implementation syntax; and
- Explicit access modifiers (the default access is `public`).

Members with bodies permit the interface to provide a "default" implementation for the method in classes and structs that do not provide their own implementation.

Interfaces may not contain instance state. While static fields are now permitted, instance fields are not permitted in interfaces. Instance auto-properties are not supported in interfaces, as they would implicitly declare a hidden field.

Static and private methods permit useful refactoring and organization of code used to implement the interface's public API.

A method override in an interface must use the explicit interface implementation syntax.

It is an error to declare a class type, struct type, or enum type within the scope of a type parameter that was declared with a *variance_annotation*. For example, the declaration of `C` below is an error.

C#

```
interface IOuter<out T>
{
    class C { } // error: class declaration within the scope of variant type
parameter 'T'
}
```

Concrete methods in interfaces

The simplest form of this feature is the ability to declare a *concrete method* in an interface, which is a method with a body.

```
C#

interface IA
{
    void M() { WriteLine("IA.M"); }
}
```

A class that implements this interface need not implement its concrete method.

```
C#

class C : IA { } // OK

IA i = new C();
i.M(); // prints "IA.M"
```

The final override for `IA.M` in class `C` is the concrete method `M` declared in `IA`. Note that a class does not inherit members from its interfaces; that is not changed by this feature:

```
C#

new C().M(); // error: class 'C' does not contain a member 'M'
```

Within an instance member of an interface, `this` has the type of the enclosing interface.

Modifiers in interfaces

The syntax for an interface is relaxed to permit modifiers on its members. The following are permitted: `private`, `protected`, `internal`, `public`, `virtual`, `abstract`, `sealed`, `static`, `extern`, and `partial`.

> ***TODO***: check what other modifiers exist.

An interface member whose declaration includes a body is a `virtual` member unless the `sealed` or `private` modifier is used. The `virtual` modifier may be used on a function member that would otherwise be implicitly `virtual`. Similarly, although `abstract` is the default on interface members without bodies, that modifier may be given explicitly. A non-virtual member may be declared using the `sealed` keyword.

It is an error for a `private` or `sealed` function member of an interface to have no body. A `private` function member may not have the modifier `sealed`.

Access modifiers may be used on interface members of all kinds of members that are permitted. The access level `public` is the default but it may be given explicitly.

> ***Open Issue:*** We need to specify the precise meaning of the access modifiers such as `protected` and `internal`, and which declarations do and do not override them (in a derived interface) or implement them (in a class that implements the interface).

Interfaces may declare `static` members, including nested types, methods, indexers, properties, events, and static constructors. The default access level for all interface members is `public`.

Interfaces may not declare instance constructors, destructors, or fields.

> ***Closed Issue:*** Should operator declarations be permitted in an interface? Probably not conversion operators, but what about others? ***Decision***: Operators are permitted *except* for conversion, equality, and inequality operators.

> ***Closed Issue:*** Should `new` be permitted on interface member declarations that hide members from base interfaces? ***Decision***: Yes.

> ***Closed Issue:*** We do not currently permit `partial` on an interface or its members. That would require a separate proposal. ***Decision***: Yes. https://github.com/dotnet/csharplang/blob/master/meetings/2018/LDM-2018-10-17.md#permit-partial-in-interface

Explicit implementation in interfaces

Explicit implementations allow the programmer to provide a most specific implementation of a virtual member in an interface where the compiler or runtime would not otherwise find one. An implementation declaration is permitted to *explicitly* implement a particular base interface method by qualifying the declaration with the interface name (no access modifier is permitted in this case). Implicit implementations are not permitted.

```
C#

interface IA
{
```

```
    void M() { WriteLine("IA.M"); }
}
interface IB : IA
{
    void IA.M() { WriteLine("IB.M"); } // Explicit implementation
}
interface IC : IA
{
    void M() { WriteLine("IC.M"); } // Creates a new M, unrelated to `IA.M`.
Warning
}
```

Explicit implementations in interfaces may not be declared `sealed`.

Public `virtual` function members in an interface may only be implemented in a derived interface explicitly (by qualifying the name in the declaration with the interface type that originally declared the method, and omitting an access modifier). The member must be *accessible* where it is implemented.

Reabstraction

A virtual (concrete) method declared in an interface may be reabstracted in a derived interface

C#

```
interface IA
{
    void M() { WriteLine("IA.M"); }
}
interface IB : IA
{
    abstract void IA.M();
}
class C : IB { } // error: class 'C' does not implement 'IA.M'.
```

The `abstract` modifier is required in the declaration of `IB.M`, to indicate that `IA.M` is being reabstracted.

This is useful in derived interfaces where the default implementation of a method is inappropriate and a more appropriate implementation should be provided by implementing classes.

> ***Open Issue:*** Should reabstraction be permitted?

The most specific implementation rule

We require that every interface and class have a *most specific implementation* for every virtual member among the implementations appearing in the type or its direct and indirect interfaces. The *most specific implementation* is a unique implementation that is more specific than every other implementation. If there is no implementation, the member itself is considered the most specific implementation.

One implementation `M1` is considered *more specific* than another implementation `M2` if `M1` is declared on type `T1`, `M2` is declared on type `T2`, and either

1. `T1` contains `T2` among its direct or indirect interfaces, or
2. `T2` is an interface type but `T1` is not an interface type.

For example:

C#

```
interface IA
{
    void M() { WriteLine("IA.M"); }
}
interface IB : IA
{
    void IA.M() { WriteLine("IB.M"); }
}
interface IC : IA
{
    void IA.M() { WriteLine("IC.M"); }
}
interface ID : IB, IC { } // error: no most specific implementation for
'IA.M'
abstract class C : IB, IC { } // error: no most specific implementation for
'IA.M'
abstract class D : IA, IB, IC // ok
{
    public abstract void M();
}
```

The most specific implementation rule ensures that a conflict (i.e. an ambiguity arising from diamond inheritance) is resolved explicitly by the programmer at the point where the conflict arises.

Because we support explicit reabstractions in interfaces, we could do so in classes as well

C#

```
abstract class E : IA, IB, IC // ok
{
    abstract void IA.M();
}
```

> ***Open issue***: should we support explicit interface abstract implementations in classes?

In addition, it is an error if in a class declaration the most specific implementation of some interface method is an abstract implementation that was declared in an interface. This is an existing rule restated using the new terminology.

C#

```
interface IF
{
    void M();
}
abstract class F : IF { } // error: 'F' does not implement 'IF.M'
```

It is possible for a virtual property declared in an interface to have a most specific implementation for its `get` accessor in one interface and a most specific implementation for its `set` accessor in a different interface. This is considered a violation of the *most specific implementation* rule.

`static` and `private` methods

Because interfaces may now contain executable code, it is useful to abstract common code into private and static methods. We now permit these in interfaces.

> ***Closed issue***: Should we support private methods? Should we support static methods? **Decision: YES**

> ***Open issue***: should we permit interface methods to be `protected` or `internal` or other access? If so, what are the semantics? Are they `virtual` by default? If so, is there a way to make them non-virtual?

> ***Open issue***: If we support static methods, should we support (static) operators?

Base interface invocations

Code in a type that derives from an interface with a default method can explicitly invoke that interface's "base" implementation.

C#

```
interface I0
{
   void M() { Console.WriteLine("I0"); }
}
interface I1 : I0
{
   override void M() { Console.WriteLine("I1"); }
}
interface I2 : I0
{
   override void M() { Console.WriteLine("I2"); }
}
interface I3 : I1, I2
{
   // an explicit override that invoke's a base interface's default method
   void I0.M() { I2.base.M(); }
}
```

An instance (nonstatic) method is permitted to invoke the implementation of an accessible instance method in a direct base interface nonvirtually by naming it using the syntax `base(Type).M`. This is useful when an override that is required to be provided due to diamond inheritance is resolved by delegating to one particular base implementation.

C#

```
interface IA
{
    void M() { WriteLine("IA.M"); }
}
interface IB : IA
{
    override void IA.M() { WriteLine("IB.M"); }
}
interface IC : IA
{
    override void IA.M() { WriteLine("IC.M"); }
}

class D : IA, IB, IC
{
    void IA.M() { base(IB).M(); }
}
```

When a `virtual` or `abstract` member is accessed using the syntax `base(Type).M`, it is required that `Type` contains a unique *most specific override* for `M`.

Binding base clauses

Interfaces now contain types. These types may be used in the base clause as base interfaces. When binding a base clause, we may need to know the set of base interfaces to bind those types (e.g. to lookup in them and to resolve protected access). The meaning of an interface's base clause is thus circularly defined. To break the cycle, we add a new language rules corresponding to a similar rule already in place for classes.

While determining the meaning of the *interface_base* of an interface, the base interfaces are temporarily assumed to be empty. Intuitively this ensures that the meaning of a base clause cannot recursively depend on itself.

We used to have the following rules:

"When a class B derives from a class A, it is a compile-time error for A to depend on B. A class **directly depends on** its direct base class (if any) and **directly depends on** the ~~**class**~~ within which it is immediately nested (if any). Given this definition, the complete set of ~~**classes**~~ upon which a class depends is the reflexive and transitive closure of the **directly depends on** relationship."

It is a compile-time error for an interface to directly or indirectly inherit from itself. The **base interfaces** of an interface are the explicit base interfaces and their base interfaces. In other words, the set of base interfaces is the complete transitive closure of the explicit base interfaces, their explicit base interfaces, and so on.

We are adjusting them as follows:

When a class B derives from a class A, it is a compile-time error for A to depend on B. A class **directly depends on** its direct base class (if any) and **directly depends on** the ***type*** within which it is immediately nested (if any).

When an interface IB extends an interface IA, it is a compile-time error for IA to depend on IB. An interface **directly depends on** its direct base interfaces (if any) and **directly depends on** the type within which it is immediately nested (if any).

Given these definitions, the complete set of **types** upon which a type depends is the reflexive and transitive closure of the **directly depends on** relationship.

Effect on existing programs

The rules presented here are intended to have no effect on the meaning of existing programs.

Example 1:

C#

```
interface IA
{
    void M();
}
class C: IA // Error: IA.M has no concrete most specific override in C
{
    public static void M() { } // method unrelated to 'IA.M' because static
}
```

Example 2:

C#

```
interface IA
{
    void M();
}
class Base: IA
{
    void IA.M() { }
}
class Derived: Base, IA // OK, all interface members have a concrete most
specific override
{
    private void M() { } // method unrelated to 'IA.M' because private
}
```

The same rules give similar results to the analogous situation involving default interface methods:

C#

```
interface IA
{
    void M() { }
}
class Derived: IA // OK, all interface members have a concrete most specific
override
{
    private void M() { } // method unrelated to 'IA.M' because private
}
```

> ***Closed issue***: confirm that this is an intended consequence of the specification.
> **Decision: YES**

Runtime method resolution

> ***Closed Issue:*** The spec should describe the runtime method resolution algorithm in the face of interface default methods. We need to ensure that the semantics are consistent with the language semantics, e.g. which declared methods do and do not override or implement an `internal` method.

CLR support API

In order for compilers to detect when they are compiling for a runtime that supports this feature, libraries for such runtimes are modified to advertise that fact through the API discussed in https://github.com/dotnet/corefx/issues/17116 . We add

```C#
namespace System.Runtime.CompilerServices
{
    public static class RuntimeFeature
    {
        // Presence of the field indicates runtime support
        public const string DefaultInterfaceImplementation =
nameof(DefaultInterfaceImplementation);
    }
}
```

> ***Open issue***: Is that the best name for the *CLR* feature? The CLR feature does much more than just that (e.g. relaxes protection constraints, supports overrides in interfaces, etc). Perhaps it should be called something like "concrete methods in interfaces", or "traits"?

Further areas to be specified

- [ ] It would be useful to catalog the kinds of source and binary compatibility effects caused by adding default interface methods and overrides to existing interfaces.

Drawbacks

This proposal requires a coordinated update to the CLR specification (to support concrete methods in interfaces and method resolution). It is therefore fairly "expensive"

and it may be worth doing in combination with other features that we also anticipate would require CLR changes.

Alternatives

None.

Unresolved questions

- Open questions are called out throughout the proposal, above.
- See also https://github.com/dotnet/csharplang/issues/406 for a list of open questions.
- The detailed specification must describe the resolution mechanism used at runtime to select the precise method to be invoked.
- The interaction of metadata produced by new compilers and consumed by older compilers needs to be worked out in detail. For example, we need to ensure that the metadata representation that we use does not cause the addition of a default implementation in an interface to break an existing class that implements that interface when compiled by an older compiler. This may affect the metadata representation that we can use.
- The design must consider interoperation with other languages and existing compilers for other languages.

Resolved Questions

Abstract Override

The earlier draft spec contained the ability to "reabstract" an inherited method:

C#

```
interface IA
{
    void M();
}
interface IB : IA
{
    override void M() { }
}
interface IC : IB
{
    override void M(); // make it abstract again
}
```

My notes for 2017-03-20 showed that we decided not to allow this. However, there are at least two use cases for it:

1. The Java APIs, with which some users of this feature hope to interoperate, depend on this facility.
2. Programming with *traits* benefits from this. Reabstraction is one of the elements of the "traits" language feature (https://en.wikipedia.org/wiki/Trait_(computer_programming)). The following is permitted with classes:

```C#
public abstract class Base
{
    public abstract void M();
}
public abstract class A : Base
{
    public override void M() { }
}
public abstract class B : A
{
    public override abstract void M(); // reabstract Base.M
}
```

Unfortunately this code cannot be refactored as a set of interfaces (traits) unless this is permitted. By the *Jared principle of greed*, it should be permitted.

> ***Closed issue:*** Should reabstraction be permitted? [YES] My notes were wrong. The LDM notes say that reabstraction is permitted in an interface. Not in a class.

Virtual Modifier vs Sealed Modifier

From Aleksey Tsingauz:

> We decided to allow modifiers explicitly stated on interface members, unless there is a reason to disallow some of them. This brings an interesting question around virtual modifier. Should it be required on members with default implementation?
>
> We could say that:
>
> - if there is no implementation and neither virtual, nor sealed are specified, we assume the member is abstract.

> - if there is an implementation and neither abstract, nor sealed are specified, we assume the member is virtual.
> - sealed modifier is required to make a method neither virtual, nor abstract.
>
> Alternatively, we could say that virtual modifier is required for a virtual member. I.e, if there is a member with implementation not explicitly marked with virtual modifier, it is neither virtual, nor abstract. This approach might provide better experience when a method is moved from a class to an interface:
>
> - an abstract method stays abstract.
> - a virtual method stays virtual.
> - a method without any modifier stays neither virtual, nor abstract.
> - sealed modifier cannot be applied to a method that is not an override.
>
> What do you think?

> ***Closed Issue:*** Should a concrete method (with implementation) be implicitly `virtual`? [YES]

Decisions: Made in the LDM 2017-04-05:

1. non-virtual should be explicitly expressed through `sealed` or `private`.
2. `sealed` is the keyword to make interface instance members with bodies non-virtual
3. We want to allow all modifiers in interfaces
4. Default accessibility for interface members is public, including nested types
5. private function members in interfaces are implicitly sealed, and `sealed` is not permitted on them.
6. Private classes (in interfaces) are permitted and can be sealed, and that means sealed in the class sense of sealed.
7. Absent a good proposal, partial is still not allowed on interfaces or their members.

Binary Compatibility 1

When a library provides a default implementation

C#

```
interface I1
{
    void M() { Impl1 }
}
interface I2 : I1
{
```

```
}
class C : I2
{
}
```

We understand that the implementation of `I1.M` in `C` is `I1.M`. What if the assembly containing `I2` is changed as follows and recompiled

C#

```
interface I2 : I1
{
    override void M() { Impl2 }
}
```

but `C` is not recompiled. What happens when the program is run? An invocation of `(C as I1).M()`

1. Runs `I1.M`
2. Runs `I2.M`
3. Throws some kind of runtime error

Decision: Made 2017-04-11: Runs `I2.M`, which is the unambiguously most specific override at runtime.

Event accessors (closed)

> ***Closed Issue:*** Can an event be overridden "piecewise"?

Consider this case:

C#

```
public interface I1
{
    event T e1;
}
public interface I2 : I1
{
    override event T
    {
        add { }
        // error: "remove" accessor missing
    }
}
```

This "partial" implementation of the event is not permitted because, as in a class, the syntax for an event declaration does not permit only one accessor; both (or neither) must be provided. You could accomplish the same thing by permitting the abstract remove accessor in the syntax to be implicitly abstract by the absence of a body:

C#

```
public interface I1
{
    event T e1;
}
public interface I2 : I1
{
    override event T
    {
        add { }
        remove; // implicitly abstract
    }
}
```

Note that *this is a new (proposed) syntax*. In the current grammar, event accessors have a mandatory body.

> ***Closed Issue:*** Can an event accessor be (implicitly) abstract by the omission of a body, similarly to the way that methods in interfaces and property accessors are (implicitly) abstract by the omission of a body?

Decision: (2017-04-18) No, event declarations require both concrete accessors (or neither).

Reabstraction in a Class (closed)

Closed Issue: We should confirm that this is permitted (otherwise adding a default implementation would be a breaking change):

C#

```
interface I1
{
    void M() { }
}
abstract class C : I1
{
    public abstract void M(); // implement I1.M with an abstract method in C
}
```

Decision: (2017-04-18) Yes, adding a body to an interface member declaration shouldn't break C.

Sealed Override (closed)

The previous question implicitly assumes that the `sealed` modifier can be applied to an `override` in an interface. This contradicts the draft specification. Do we want to permit sealing an override? Source and binary compatibility effects of sealing should be considered.

> ***Closed Issue:*** Should we permit sealing an override?

Decision: (2017-04-18) Let's not allowed `sealed` on overrides in interfaces. The only use of `sealed` on interface members is to make them non-virtual in their initial declaration.

Diamond inheritance and classes (closed)

The draft of the proposal prefers class overrides to interface overrides in diamond inheritance scenarios:

> We require that every interface and class have a *most specific override* for every interface method among the overrides appearing in the type or its direct and indirect interfaces. The *most specific override* is a unique override that is more specific than every other override. If there is no override, the method itself is considered the most specific override.
>
> One override `M1` is considered *more specific* than another override `M2` if `M1` is declared on type `T1`, `M2` is declared on type `T2`, and either
>
> 1. `T1` contains `T2` among its direct or indirect interfaces, or
> 2. `T2` is an interface type but `T1` is not an interface type.

The scenario is this

```C#
interface IA
{
    void M();
}
interface IB : IA
{
    override void M() { WriteLine("IB"); }
```

```
}
class Base : IA
{
    void IA.M() { WriteLine("Base"); }
}
class Derived : Base, IB // allowed?
{
    static void Main()
    {
        IA a = new Derived();
        a.M();           // what does it do?
    }
}
```

We should confirm this behavior (or decide otherwise)

> ***Closed Issue:*** Confirm the draft spec, above, for *most specific override* as it applies to mixed classes and interfaces (a class takes priority over an interface). See https://github.com/dotnet/csharplang/blob/master/meetings/2017/LDM-2017-04-19.md#diamonds-with-classes .

Interface methods vs structs (closed)

There are some unfortunate interactions between default interface methods and structs.

C#

```
interface IA
{
    public void M() { }
}
struct S : IA
{
}
```

Note that interface members are not inherited:

C#

```
var s = default(S);
s.M(); // error: 'S' does not contain a member 'M'
```

Consequently, the client must box the struct to invoke interface methods

C#

```
IA s = default(S); // an S, boxed
s.M(); // ok
```

Boxing in this way defeats the principal benefits of a `struct` type. Moreover, any mutation methods will have no apparent effect, because they are operating on a *boxed copy* of the struct:

C#

```
interface IB
{
    public void Increment() { P += 1; }
    public int P { get; set; }
}
struct T : IB
{
    public int P { get; set; } // auto-property
}

T t = default(T);
Console.WriteLine(t.P); // prints 0
(t as IB).Increment();
Console.WriteLine(t.P); // prints 0
```

> ***Closed Issue:*** What can we do about this:
>
> 1. Forbid a `struct` from inheriting a default implementation. All interface methods would be treated as abstract in a `struct`. Then we may take time later to decide how to make it work better.
> 2. Come up with some kind of code generation strategy that avoids boxing. Inside a method like `IB.Increment`, the type of `this` would perhaps be akin to a type parameter constrained to `IB`. In conjunction with that, to avoid boxing in the caller, non-abstract methods would be inherited from interfaces. This may increase compiler and CLR implementation work substantially.
> 3. Not worry about it and just leave it as a wart.
> 4. Other ideas?

Decision: Not worry about it and just leave it as a wart. See https://github.com/dotnet/csharplang/blob/master/meetings/2017/LDM-2017-04-19.md#structs-and-default-implementations .

Base interface invocations (closed)

The draft spec suggests a syntax for base interface invocations inspired by Java: `Interface.base.M()`. We need to select a syntax, at least for the initial prototype. My favorite is `base<Interface>.M()`.

> ***Closed Issue:*** What is the syntax for a base member invocation?

Decision: The syntax is `base(Interface).M()`. See https://github.com/dotnet/csharplang/blob/master/meetings/2017/LDM-2017-04-19.md#base-invocation. The interface so named must be a base interface, but does not need to be a direct base interface.

> ***Open Issue:*** Should base interface invocations be permitted in class members?

Decision: Yes. https://github.com/dotnet/csharplang/blob/master/meetings/2017/LDM-2017-04-19.md#base-invocation

Overriding non-public interface members (closed)

In an interface, non-public members from base interfaces are overridden using the `override` modifier. If it is an "explicit" override that names the interface containing the member, the access modifier is omitted.

> ***Closed Issue:*** If it is an "implicit" override that does not name the interface, does the access modifier have to match?

Decision: Only public members may be implicitly overridden, and the access must match. See https://github.com/dotnet/csharplang/blob/master/meetings/2017/LDM-2017-04-18.md#dim-implementing-a-non-public-interface-member-not-in-list.

> ***Open Issue:*** Is the access modifier required, optional, or omitted on an explicit override such as `override void IB.M() {}`?

> ***Open Issue:*** Is `override` required, optional, or omitted on an explicit override such as `void IB.M() {}`?

How does one implement a non-public interface member in a class? Perhaps it must be done explicitly?

```
C#
```

```
interface IA
{
    internal void MI();
    protected void MP();
}
class C : IA
{
    // are these implementations?
    internal void MI() {}
    protected void MP() {}
}
```

> ***Closed Issue:*** How does one implement a non-public interface member in a class?

Decision: You can only implement non-public interface members explicitly. See https://github.com/dotnet/csharplang/blob/master/meetings/2017/LDM-2017-04-18.md#dim-implementing-a-non-public-interface-member-not-in-list .

Decision: No `override` keyword permitted on interface members. https://github.com/dotnet/csharplang/blob/master/meetings/2018/LDM-2018-10-17.md#does-an-override-in-an-interface-introduce-a-new-member

Binary Compatibility 2 (closed)

Consider the following code in which each type is in a separate assembly

C#

```
interface I1
{
    void M() { Impl1 }
}
interface I2 : I1
{
    override void M() { Impl2 }
}
interface I3 : I1
{
}
class C : I2, I3
{
}
```

We understand that the implementation of `I1.M` in `C` is `I2.M`. What if the assembly containing `I3` is changed as follows and recompiled

```C#
interface I3 : I1
{
    override void M() { Impl3 }
}
```

but `C` is not recompiled. What happens when the program is run? An invocation of `(C as I1).M()`

1. Runs `I1.M`
2. Runs `I2.M`
3. Runs `I3.M`
4. Either 2 or 3, deterministically
5. Throws some kind of runtime exception

Decision: Throw an exception (5). See https://github.com/dotnet/csharplang/blob/master/meetings/2018/LDM-2018-10-17.md#issues-in-default-interface-methods.

Permit `partial` in interface? (closed)

Given that interfaces may be used in ways analogous to the way abstract classes are used, it may be useful to declare them `partial`. This would be particularly useful in the face of generators.

> ***Proposal:*** Remove the language restriction that interfaces and members of interfaces may not be declared `partial`.

Decision: Yes. See https://github.com/dotnet/csharplang/blob/master/meetings/2018/LDM-2018-10-17.md#permit-partial-in-interface.

`Main` in an interface? (closed)

> ***Open Issue:*** Is a `static Main` method in an interface a candidate to be the program's entry point?

Decision: Yes. See https://github.com/dotnet/csharplang/blob/master/meetings/2018/LDM-2018-10-17.md#main-in-an-interface.

Confirm intent to support public non-virtual methods (closed)

Can we please confirm (or reverse) our decision to permit non-virtual public methods in an interface?

C#

```
interface IA
{
    public sealed void M() { }
}
```

> ***Semi-Closed Issue:*** (2017-04-18) We think it is going to be useful, but will come back to it. This is a mental model tripping block.

Decision: Yes. https://github.com/dotnet/csharplang/blob/master/meetings/2018/LDM-2018-10-17.md#confirm-that-we-support-public-non-virtual-methods .

Does an `override` in an interface introduce a new member? (closed)

There are a few ways to observe whether an override declaration introduces a new member or not.

C#

```
interface IA
{
    void M(int x) { }
}
interface IB : IA
{
    override void M(int y) { }
}
interface IC : IB
{
    static void M2()
    {
        M(y: 3); // permitted?
    }
    override void IB.M(int z) { } // permitted? What does it override?
}
```

> ***Open Issue:*** Does an override declaration in an interface introduce a new member? (closed)

In a class, an overriding method is "visible" in some senses. For example, the names of its parameters take precedence over the names of parameters in the overridden method. It may be possible to duplicate that behavior in interfaces, as there is always a most specific override. But do we want to duplicate that behavior?

Also, it is possible to "override" an override method? [Moot]

Decision: No `override` keyword permitted on interface members. https://github.com/dotnet/csharplang/blob/master/meetings/2018/LDM-2018-10-17.md#does-an-override-in-an-interface-introduce-a-new-member .

Properties with a private accessor (closed)

We say that private members are not virtual, and the combination of virtual and private is disallowed. But what about a property with a private accessor?

C#

```
interface IA
{
    public virtual int P
    {
        get => 3;
        private set => { }
    }
}
```

Is this allowed? Is the `set` accessor here `virtual` or not? Can it be overridden where it is accessible? Does the following implicitly implement only the `get` accessor?

C#

```
class C : IA
{
    public int P
    {
        get => 4;
        set { }
    }
}
```

Is the following presumably an error because IA.P.set isn't virtual and also because it isn't accessible?

```
C#
```

```
class C : IA
{
    int IA.P
    {
        get => 4;
        set { }
    }
}
```

Decision: The first example looks valid, while the last does not. This is resolved analogously to how it already works in C#. https://github.com/dotnet/csharplang/blob/master/meetings/2018/LDM-2018-10-17.md#properties-with-a-private-accessor

Base Interface Invocations, round 2 (closed)

Our previous "resolution" to how to handle base invocations doesn't actually provide sufficient expressiveness. It turns out that in C# and the CLR, unlike Java, you need to specify both the interface containing the method declaration and the location of the implementation you want to invoke.

I propose the following syntax for base calls in interfaces. I'm not in love with it, but it illustrates what any syntax must be able to express:

```
C#
```

```
interface I1 { void M(); }
interface I2 { void M(); }
interface I3 : I1, I2 { void I1.M() { } void I2.M() { } }
interface I4 : I1, I2 { void I1.M() { } void I2.M() { } }
interface I5 : I3, I4
{
    void I1.M()
    {
        base<I3>(I1).M(); // calls I3's implementation of I1.M
        base<I4>(I1).M(); // calls I4's implementation of I1.M
    }
    void I2.M()
    {
        base<I3>(I2).M(); // calls I3's implementation of I2.M
        base<I4>(I2).M(); // calls I4's implementation of I2.M
    }
}
```

If there is no ambiguity, you can write it more simply

C#

```
interface I1 { void M(); }
interface I3 : I1 { void I1.M() { } }
interface I4 : I1 { void I1.M() { } }
interface I5 : I3, I4
{
    void I1.M()
    {
        base<I3>.M(); // calls I3's implementation of I1.M
        base<I4>.M(); // calls I4's implementation of I1.M
    }
}
```

Or

C#

```
interface I1 { void M(); }
interface I2 { void M(); }
interface I3 : I1, I2 { void I1.M() { } void I2.M() { } }
interface I5 : I3
{
    void I1.M()
    {
        base(I1).M(); // calls I3's implementation of I1.M
    }
    void I2.M()
    {
        base(I2).M(); // calls I3's implementation of I2.M
    }
}
```

Or

C#

```
interface I1 { void M(); }
interface I3 : I1 { void I1.M() { } }
interface I5 : I3
{
    void I1.M()
    {
        base.M(); // calls I3's implementation of I1.M
    }
}
```

Decision: Decided on `base(N.I1<T>).M(s)`, conceding that if we have an invocation binding there may be problem here later on.
https://github.com/dotnet/csharplang/blob/master/meetings/2018/LDM-2018-11-14.md#default-interface-implementations

Warning for struct not implementing default method? (closed)

@vancem asserts that we should seriously consider producing a warning if a value type declaration fails to override some interface method, even if it would inherit an implementation of that method from an interface. Because it causes boxing and undermines constrained calls.

Decision: This seems like something more suited for an analyzer. It also seems like this warning could be noisy, since it would fire even if the default interface method is never called and no boxing will ever occur.
https://github.com/dotnet/csharplang/blob/master/meetings/2018/LDM-2018-10-17.md#warning-for-struct-not-implementing-default-method

Interface static constructors (closed)

When are interface static constructors run? The current CLI draft proposes that it occurs when the first static method or field is accessed. If there are neither of those then it might never be run??

[2018-10-09 The CLR team proposes "Going to mirror what we do for valuetypes (cctor check on access to each instance method)"]

Decision: Static constructors are also run on entry to instance methods, if the static constructor was not `beforefieldinit`, in which case static constructors are run before access to the first static field.
https://github.com/dotnet/csharplang/blob/master/meetings/2018/LDM-2018-10-17.md#when-are-interface-static-constructors-run

Design meetings

2017-03-08 LDM Meeting Notes 2017-03-21 LDM Meeting Notes 2017-03-23 meeting "CLR Behavior for Default Interface Methods" 2017-04-05 LDM Meeting Notes 2017-04-11 LDM Meeting Notes 2017-04-18 LDM Meeting Notes 2017-04-19 LDM Meeting Notes 2017-05-17 LDM Meeting Notes 2017-05-31 LDM

Meeting Notes 2017-06-14 LDM Meeting Notes 2018-10-17 LDM Meeting Notes 2018-11-14 LDM Meeting Notes

Static abstract members in interfaces

Article • 03/28/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Summary

An interface is allowed to specify abstract static members that implementing classes and structs are then required to provide an explicit or implicit implementation of. The members can be accessed off of type parameters that are constrained by the interface.

Motivation

There is currently no way to abstract over static members and write generalized code that applies across types that define those static members. This is particularly problematic for member kinds that *only* exist in a static form, notably operators.

This feature allows generic algorithms over numeric types, represented by interface constraints that specify the presence of given operators. The algorithms can therefore be expressed in terms of such operators:

c#

```
// Interface specifies static properties and operators
interface IAddable<T> where T : IAddable<T>
{
    static abstract T Zero { get; }
    static abstract T operator +(T t1, T t2);
}

// Classes and structs (including built-ins) can implement interface
struct Int32 : …, IAddable<Int32>
{
    static Int32 I.operator +(Int32 x, Int32 y) => x + y; // Explicit
```

```
    public static int Zero => 0;                          // Implicit
}

// Generic algorithms can use static members on T
public static T AddAll<T>(T[] ts) where T : IAddable<T>
{
    T result = T.Zero;                   // Call static operator
    foreach (T t in ts) { result += t; } // Use `+`
    return result;
}

// Generic method can be applied to built-in and user-defined types
int sixtyThree = AddAll(new [] { 1, 2, 4, 8, 16, 32 });
```

Syntax

Interface members

The feature would allow static interface members to be declared virtual.

Today's rules

Today, instance members in interfaces are implicitly abstract (or virtual if they have a default implementation), but can optionally have an `abstract` (or `virtual`) modifier. Non-virtual instance members must be explicitly marked as `sealed`.

Static interface members today are implicitly non-virtual, and do not allow `abstract`, `virtual` or `sealed` modifiers.

Proposal

Abstract static members

Static interface members other than fields are allowed to also have the `abstract` modifier. Abstract static members are not allowed to have a body (or in the case of properties, the accessors are not allowed to have a body).

c#

```
interface I<T> where T : I<T>
{
    static abstract void M();
    static abstract T P { get; set; }
    static abstract event Action E;
```

```
    static abstract T operator +(T l, T r);
    static abstract bool operator ==(T l, T r);
    static abstract bool operator !=(T l, T r);
    static abstract implicit operator T(string s);
    static abstract explicit operator string(T t);
}
```

Virtual static members

Static interface members other than fields are allowed to also have the `virtual` modifier. Virtual static members are required to have a body.

```c#
interface I<T> where T : I<T>
{
    static virtual void M() {}
    static virtual T P { get; set; }
    static virtual event Action E;
    static virtual T operator +(T l, T r) { throw new
NotImplementedException(); }
}
```

Explicitly non-virtual static members

For symmetry with non-virtual instance members, static members should be allowed an optional `sealed` modifier, even though they are non-virtual by default:

```c#
interface I0
{
    static sealed void M() => Console.WriteLine("Default behavior");

    static sealed int f = 0;

    static sealed int P1 { get; set; }
    static sealed int P2 { get => f; set => f = value; }

    static sealed event Action E1;
    static sealed event Action E2 { add => E1 += value; remove => E1 -=
value; }

    static sealed I0 operator +(I0 l, I0 r) => l;
}
```

Implementation of interface members

Today's rules

Classes and structs can implement abstract instance members of interfaces either implicitly or explicitly. An implicitly implemented interface member is a normal (virtual or non-virtual) member declaration of the class or struct that just "happens" to also implement the interface member. The member can even be inherited from a base class and thus not even be present in the class declaration.

An explicitly implemented interface member uses a qualified name to identify the interface member in question. The implementation is not directly accessible as a member on the class or struct, but only through the interface.

Proposal

No new syntax is needed in classes and structs to facilitate implicit implementation of static abstract interface members. Existing static member declarations serve that purpose.

Explicit implementations of static abstract interface members use a qualified name along with the `static` modifier.

c#

```
class C : I<C>
{
    string _s;
    public C(string s) => _s = s;
    static void I<C>.M() => Console.WriteLine("Implementation");
    static C I<C>.P { get; set; }
    static event Action I<C>.E;
    static C I<C>.operator +(C l, C r) => new C($"{l._s} {r._s}");
    static bool I<C>.operator ==(C l, C r) => l._s == r._s;
    static bool I<C>.operator !=(C l, C r) => l._s != r._s;
    static implicit I<C>.operator C(string s) => new C(s);
    static explicit I<C>.operator string(C c) => c._s;
}
```

Semantics

Operator restrictions

Today all unary and binary operator declarations have some requirement involving at least one of their operands to be of type `T` or `T?`, where `T` is the instance type of the enclosing type.

These requirements need to be relaxed so that a restricted operand is allowed to be of a type parameter that counts as "the instance type of the enclosing type".

In order for a type parameter `T` to count as " the instance type of the enclosing type", it must meet the following requirements:

- `T` is a direct type parameter on the interface in which the operator declaration occurs, and
- `T` is *directly* constrained by what the spec calls the "instance type" - i.e. the surrounding interface with its own type parameters used as type arguments.

Equality operators and conversions

Abstract/virtual declarations of `==` and `!=` operators, as well as abstract/virtual declarations of implicit and explicit conversion operators will be allowed in interfaces. Derived interfaces will be allowed to implement them too.

For `==` and `!=` operators, at least one parameter type must be a type parameter that counts as "the instance type of the enclosing type", as defined in the previous section.

Implementing static abstract members

The rules for when a static member declaration in a class or struct is considered to implement a static abstract interface member, and for what requirements apply when it does, are the same as for instance members.

TBD: *There may be additional or different rules necessary here that we haven't yet thought of.*

Interfaces as type arguments

We discussed the issue raised by https://github.com/dotnet/csharplang/issues/5955 and decided to add a restriction around usage of an interface as a type argument (https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-03-28.md#type-hole-in-static-abstracts). Here is the restriction as it was proposed by https://github.com/dotnet/csharplang/issues/5955 and approved by the LDM.

An interface containing or inheriting a static abstract/virtual member that does not have most specific implementation in the interface cannot be used as a type argument. If all static abstract/virtual members have most specific implementation, the interface can be used as a type argument.

Accessing static abstract interface members

A static abstract interface member `M` may be accessed on a type parameter `T` using the expression `T.M` when `T` is constrained by an interface `I` and `M` is an accessible static abstract member of `I`.

```c#
T M<T>() where T : I<T>
{
    T.M();
    T t = T.P;
    T.E += () => { };
    return t + T.P;
}
```

At runtime, the actual member implementation used is the one that exists on the actual type provided as a type argument.

```c#
C c = M<C>(); // The static members of C get called
```

Since query expressions are spec'ed as a syntactic rewrite, C# actually lets you use a *type* as the query source, as long as it has static members for the query operators you use! In other words, if the *syntax* fits, we allow it! We think this behavior was not intentional or important in the original LINQ, and we don't want to do the work to support it on type parameters. If there are scenarios out there we will hear about them, and can choose to embrace this later.

Variance safety §17.2.3.2

Variance safety rules should apply to signatures of static abstract members. The addition proposed in https://github.com/dotnet/csharplang/blob/main/proposals/variance-safety-for-static-interface-members.md#variance-safety should be adjusted from

These restrictions do not apply to occurrences of types within declarations of static members.

to

*These restrictions do not apply to occurrences of types within declarations of **non-virtual, non-abstract** static members.*

§10.5.4 User defined implicit conversions

The following bullet points

- Determine the types `S`, S_o and T_o.
  - If `E` has a type, let `S` be that type.
  - If `S` or `T` are nullable value types, let S_i and T_i be their underlying types, otherwise let S_i and T_i be `S` and `T`, respectively.
  - If S_i or T_i are type parameters, let S_o and T_o be their effective base classes, otherwise let S_o and T_o be S_x and T_i, respectively.
- Find the set of types, `D`, from which user-defined conversion operators will be considered. This set consists of `S0` (if `S0` is a class or struct), the base classes of `S0` (if `S0` is a class), and `T0` (if `T0` is a class or struct).
- Find the set of applicable user-defined and lifted conversion operators, `U`. This set consists of the user-defined and lifted implicit conversion operators declared by the classes or structs in `D` that convert from a type encompassing `S` to a type encompassed by `T`. If `U` is empty, the conversion is undefined and a compile-time error occurs.

are adjusted as follows:

- Determine the types `S`, S_o and T_o.
  - If `E` has a type, let `S` be that type.
  - If `S` or `T` are nullable value types, let S_i and T_i be their underlying types, otherwise let S_i and T_i be `S` and `T`, respectively.
  - If S_i or T_i are type parameters, let S_o and T_o be their effective base classes, otherwise let S_o and T_o be S_x and T_i, respectively.
- Find the set of applicable user-defined and lifted conversion operators, `U`.
  - Find the set of types, `D1`, from which user-defined conversion operators will be considered. This set consists of `S0` (if `S0` is a class or struct), the base classes of `S0` (if `S0` is a class), and `T0` (if `T0` is a class or struct).
  - Find the set of applicable user-defined and lifted conversion operators, `U1`. This set consists of the user-defined and lifted implicit conversion operators declared by the classes or structs in `D1` that convert from a type encompassing `S` to a type encompassed by `T`.

    - If `U1` is not empty, then `U` is `U1`. Otherwise,
        - Find the set of types, `D2`, from which user-defined conversion operators will be considered. This set consists of S_i *effective interface set* and their base interfaces (if S_i is a type parameter), and T_i *effective interface set* (if T_i is a type parameter).
        - Find the set of applicable user-defined and lifted conversion operators, `U2`. This set consists of the user-defined and lifted implicit conversion operators declared by the interfaces in `D2` that convert from a type encompassing `S` to a type encompassed by `T`.
        - If `U2` is not empty, then `U` is `U2`
- If `U` is empty, the conversion is undefined and a compile-time error occurs.

§10.5.5 User-defined explicit conversions

The following bullet points

- Determine the types `S`, S_o and T_o.
    - If `E` has a type, let `S` be that type.
    - If `S` or `T` are nullable value types, let S_i and T_i be their underlying types, otherwise let S_i and T_i be `S` and `T`, respectively.
    - If S_i or T_i are type parameters, let S_o and T_o be their effective base classes, otherwise let S_o and T_o be S_i and T_i, respectively.
- Find the set of types, `D`, from which user-defined conversion operators will be considered. This set consists of `S0` (if `S0` is a class or struct), the base classes of `S0` (if `S0` is a class), `T0` (if `T0` is a class or struct), and the base classes of `T0` (if `T0` is a class).
- Find the set of applicable user-defined and lifted conversion operators, `U`. This set consists of the user-defined and lifted implicit or explicit conversion operators declared by the classes or structs in `D` that convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`. If `U` is empty, the conversion is undefined and a compile-time error occurs.

are adjusted as follows:

- Determine the types `S`, S_o and T_o.
    - If `E` has a type, let `S` be that type.
    - If `S` or `T` are nullable value types, let S_i and T_i be their underlying types, otherwise let S_i and T_i be `S` and `T`, respectively.

    - If S_i or T_i are type parameters, let S_o and T_o be their effective base classes, otherwise let S_o and T_o be S_i and T_i, respectively.
- Find the set of applicable user-defined and lifted conversion operators, `U`.
    - Find the set of types, `D1`, from which user-defined conversion operators will be considered. This set consists of `S0` (if `S0` is a class or struct), the base classes of `S0` (if `S0` is a class), `T0` (if `T0` is a class or struct), and the base classes of `T0` (if `T0` is a class).
    - Find the set of applicable user-defined and lifted conversion operators, `U1`. This set consists of the user-defined and lifted implicit or explicit conversion operators declared by the classes or structs in `D1` that convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`.
    - If `U1` is not empty, then `U` is `U1`. Otherwise,
        - Find the set of types, `D2`, from which user-defined conversion operators will be considered. This set consists of S_i *effective interface set* and their base interfaces (if S_i is a type parameter), and T_i *effective interface set* and their base interfaces (if T_i is a type parameter).
        - Find the set of applicable user-defined and lifted conversion operators, `U2`. This set consists of the user-defined and lifted implicit or explicit conversion operators declared by the interfaces in `D2` that convert from a type encompassing or encompassed by `S` to a type encompassing or encompassed by `T`.
        - If `U2` is not empty, then `U` is `U2`
- If `U` is empty, the conversion is undefined and a compile-time error occurs.

Default implementations

An *additional* feature to this proposal is to allow static virtual members in interfaces to have default implementations, just as instance virtual/abstract members do.

One complication here is that default implementations would want to call other static virtual members "virtually". Allowing static virtual members to be called directly on the interface would require flowing a hidden type parameter representing the "self" type that the current static method really got invoked on. This seems complicated, expensive and potentially confusing.

We discussed a simpler version which maintains the limitations of the current proposal that static virtual members can *only* be invoked on type parameters. Since interfaces with static virtual members will often have an explicit type parameter representing a

"self" type, this wouldn't be a big loss: other static virtual members could just be called on that self type. This version is a lot simpler, and seems quite doable.

At https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-01-24.md#default-implementations-of-abstract-statics we decided to support Default Implementations of static members following/expanding the rules established in https://github.com/dotnet/csharplang/blob/main/proposals/csharp-8.0/default-interface-methods.md accordingly.

Pattern matching

Given the following code, a user might reasonably expect it to print True (as it would if the constant pattern was written inline):

```C#
M(1.0);

static void M<T>(T t) where T : INumberBase<T>
{
    Console.WriteLine(t is 1);
}
```

However, because the input type of the pattern is not `double`, the constant `1` pattern will first type check the incoming `T` against `int`. This is unintuitive, so we will block it until a future C# version adds better handling for numeric matching against types derived from `INumberBase<T>`. To do so, we will say that, we will explicitly recognize `INumberBase<T>` as the type that all "numbers" will derive from, and block the pattern if we're trying to match a numeric constant pattern against a number type that we can't represent the pattern in (ie, a type parameter constrained to `INumberBase<T>`, or a user-defined number type that inherits from `INumberBase<T>`).

Formally, we add an exception to the definition of *pattern-compatible* for constant patterns:

> A constant pattern tests the value of an expression against a constant value. The constant may be any constant expression, such as a literal, the name of a declared `const` variable, or an enumeration constant. When the input value is not an open type, the constant expression is implicitly converted to the type of the matched expression; if the type of the input value is not *pattern-compatible* with the type of the constant expression, the pattern-matching operation is an error. **If the constant expression being matched against is a numeric value, the input value is a type**

> **that inherits from `System.Numerics.INumberBase<T>`, and there is no constant conversion from the constant expression to the type of the input value, the pattern-matching operation is an error.**

We also add a similar exception for relational patterns:

> When the input is a type for which a suitable built-in binary relational operator is defined that is applicable with the input as its left operand and the given constant as its right operand, the evaluation of that operator is taken as the meaning of the relational pattern. Otherwise we convert the input to the type of the expression using an explicit nullable or unboxing conversion. It is a compile-time error if no such conversion exists. **It is a compile-time error if the input type is a type parameter constrained to or a type inheriting from `System.Numerics.INumberBase<T>` and the input type has no suitable built-in binary relational operator defined.** The pattern is considered not to match if the conversion fails. If the conversion succeeds then the result of the pattern-matching operation is the result of evaluating the expression e OP v where e is the converted input, OP is the relational operator, and v is the constant expression.

Drawbacks

- "static abstract" is a new concept and will meaningfully add to the conceptual load of C#.
- It's not a cheap feature to build. We should make sure it's worth it.

Alternatives

Structural constraints

An alternative approach would be to have "structural constraints" directly and explicitly requiring the presence of specific operators on a type parameter. The drawbacks of that are: - This would have to be written out every time. Having a named constraint seems better. - This is a whole new kind of constraint, whereas the proposed feature utilizes the existing concept of interface constraints. - It would only work for operators, not (easily) other kinds of static members.

Unresolved questions

Static abstract interfaces and static classes

See https://github.com/dotnet/csharplang/issues/5783 and https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-16.md#static-abstract-interfaces-and-static-classes for more information.

Design meetings

- https://github.com/dotnet/csharplang/blob/master/meetings/2021/LDM-2021-02-08.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2021/LDM-2021-04-05.md
- https://github.com/dotnet/csharplang/blob/master/meetings/2020/LDM-2020-06-29.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-01-24.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-02-16.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-03-28.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-04-06.md
- https://github.com/dotnet/csharplang/blob/main/meetings/2022/LDM-2022-06-06.md

Generic Attributes

Article • 06/23/2023

> **Note**
>
> This article is a feature specification. The specification serves as the design document for the feature. It includes proposed specification changes, along with information needed during the design and development of the feature. These articles are published until the proposed spec changes are finalized and incorporated in the current ECMA specification.
>
> There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent language design meeting (LDM) notes.
>
> You can learn more about the process for adopting feature speclets into the C# language standard in the article on the specifications.

Summary

When generics were introduced in C# 2.0, attribute classes were not allowed to participate. We can make the language more composable by removing (rather, loosening) this restriction. The .NET Core runtime has added support for generic attributes. Now, all that's missing is support for generic attributes in the compiler.

Motivation

Currently attribute authors can take a `System.Type` as a parameter and have users pass a `typeof` expression to provide the attribute with types that it needs. However, outside of analyzers, there's no way for an attribute author to constrain what types are allowed to be passed to an attribute via `typeof`. If attributes could be generic, then attribute authors could use the existing system of type parameter constraints to express the requirements for the types they take as input.

Detailed design

The following section is amended: §14.2.4.2

> The direct base class of a class type must not be any of the following types: System.Array, System.Delegate, System.MulticastDelegate, System.Enum, or System.ValueType. ~~Furthermore, a generic class declaration cannot use System.Attribute as a direct or indirect base class.~~

One important note is that the following section of the spec is *unaffected* when referencing the point of usage of an attribute, i.e. within an attribute list: Type parameters - §8.5.

> A type parameter cannot be used anywhere within an attribute.

This means that when a generic attribute is used, its construction needs to be fully "closed", i.e. not containing any type parameters, which means the following is still disallowed:

```cs
using System;
using System.Collections.Generic;

public class Attr<T1> : Attribute { }

public class Program<T2>
{
    [Attr<T2>] // error
    [Attr<List<T2>>] // error
    void M() { }
}
```

When a generic attribute is used in an attribute list, its type arguments have the same restrictions that `typeof` has on its argument. For example, `[Attr<dynamic>]` is an error. This is because "attribute-dependent" types like `dynamic`, `List<string?>`, `nint`, and so on can't be fully represented in the final IL for an attribute type argument, because there isn't a symbol to "attach" the `DynamicAttribute` or other well-known attribute to.

Drawbacks

Removing the restriction, reasoning out the implications, and adding the appropriate tests is work.

Alternatives

Attribute authors who want users to be able to discover the requirements for the types they provide to attributes need to write analyzers and guide their users to use those analyzers in their builds.

Unresolved questions

- [x] What does `AllowMultiple = false` mean on a generic attribute? If we have `[Attr<string>]` and `[Attr<object>]` both used on a symbol, does that mean "multiple" of the attribute are in use?
  - For now we are inclined to take the more restrictive route here and consider the attribute class's original definition when deciding whether multiple of it have been applied. In other words, `[Attr<string>]` and `[Attr<object>]` applied together is incompatible with `AllowMultiple = false`.

Design meetings

- https://github.com/dotnet/csharplang/blob/main/meetings/2017/LDM-2017-02-21.md#generic-attributes
  - At the time there was a concern that we would have to gate the feature on whether the target runtime supports it. (However, we now only support C# 10 on .NET 6. It would be nice for the implementation to be aware of what minimum target framework supports the feature, but seems less essential today.)

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

.NET

C# feature specifications feedback

C# feature specifications is an open source project. Select a link to provide feedback:

Open a documentation issue

Provide product feedback

ExperimentalAttribute

Article • 10/31/2023

> **Note**
>
> This article is a feature specification. The specification represents the *proposed* feature specification. There may be some discrepancies between the feature specification and the completed implementation. Those differences are captured in the pertinent **language design meeting (LDM) notes**. Links to pertinent meetings are included at the bottom of the spec. You can learn more about the process for merging feature speclets into the C# language standard in the article on the **specifications**.

Report warnings for references to types and members marked with `System.Diagnostics.CodeAnalysis.ExperimentalAttribute`.

C#

```
namespace System.Diagnostics.CodeAnalysis
{
    [AttributeUsage(AttributeTargets.Assembly |
                    AttributeTargets.Module |
                    AttributeTargets.Class |
                    AttributeTargets.Struct |
                    AttributeTargets.Enum |
                    AttributeTargets.Constructor |
                    AttributeTargets.Method |
                    AttributeTargets.Property |
                    AttributeTargets.Field |
                    AttributeTargets.Event |
                    AttributeTargets.Interface |
                    AttributeTargets.Delegate, Inherited = false)]
    public sealed class ExperimentalAttribute : Attribute
    {
        public ExperimentalAttribute(string diagnosticId)
        {
            DiagnosticId = diagnosticId;
        }

        public string DiagnosticId { get; }
        public string? UrlFormat { get; set; }
    }
}
```

Reported diagnostic

Although the diagnostic is technically a warning (so that the compiler allows suppressing it), it is treated as an error for purpose of reporting. This causes the build to fail if the diagnostic is not suppressed.

The diagnostic is reported for any reference to a type or member that is either:

- marked with the attribute,
- in an assembly or module marked with the attribute,

except when the reference occurs within `[Experimental]` members (automatic suppression).

It is also possible to suppress the diagnostic by usual means, such as an explicit compiler option or `#pragma`.
For example, if the API is marked with `[Experimental("DiagID")]` or `[Experimental("DiagID", UrlFormat = "https://example.org/{0}")]`, the diagnostic can be suppressed with `#pragma warning disable DiagID`.

An error is produced if the diagnostic ID given to the experimental attribute is not a valid C# identifier.

The diagnostic message is a specific message, where `'{0}'` is the fully-qualified type or member name.

```
'{0}' is for evaluation purposes only and is subject to change or removal in
future updates.
```

The attribute is not inherited from base types or overridden members.

ObsoleteAttribute and DeprecatedAttribute

Warnings for `[Experimental]` are reported within `[Obsolete]` or `[Deprecated]` members.
Warnings and errors for `[Obsolete]` and `[Deprecated]` are reported inside `[Experimental]` members.
But warnings and errors for `[Obsolete]` and `[Deprecated]` are reported instead of `[Experimental]` if there are multiple attributes.

Collaborate with us on GitHub

The source for this content can be found on GitHub, where you can also create and review issues and pull requests. For more information, see our contributor guide.

C# feature specification feedback

The C# feature specifications are open source. Provide feedback here.

Open a documentation issue

Provide product feedback

C# documentation

Learn how to write any application using the C# programming language on the .NET platform.

Learn to program

GET STARTED

Learn C# | Tutorials, courses, videos, and more

VIDEO

C# beginner video series

TUTORIAL

Self-guided tutorials

In-browser tutorial

REFERENCE

C# on Q&A

Languages on .NET tech community forums

C# on Stack Overflow

C# on Discord

Fundamentals

OVERVIEW

A tour of C#

Inside a C# program

C# highlights video series

CONCEPT

Type system

Object oriented programming

Functional techniques

Exceptions

Coding style

TUTORIAL

Display command-line

Intro to classes

Object oriented C#

Converting types

Pattern matching

Use LINQ to query data

What's new

WHAT'S NEW

What's new in C# 13

What's new in C# 12

What's new in C# 11

What's new in C# 10

TUTORIAL

Explore record types

Explore top level statements

Explore new patterns

Write a custom string interpolation handler

REFERENCE

Breaking changes in the C# compiler

Version compatibility

Key concepts

OVERVIEW

C# language strategy

Programming concepts

CONCEPT

Language Integrated Query (LINQ)

Asynchronous programming

Advanced concepts

REFERENCE

Reflection and attributes

Expression trees

Native interoperability

Performance engineering

.NET Compiler Platform SDK

Stay in touch

REFERENCE

.NET developer community

YouTube

Twitter

C# language reference

The language reference provides an informal reference to C# syntax and idioms for beginners and experienced C# and .NET developers.

C# language reference

OVERVIEW

C# language strategy

REFERENCE

C# keywords

C# operators and expressions

Configure language version

C# language specification - C# 8 draft in progress

What's new

WHAT'S NEW

What's new in C# 12

What's new in C# 11

What's new in C# 10

REFERENCE

Breaking changes in the C# compiler

Version compatibility

Stay in touch

REFERENCE

.NET developer community

YouTube

Twitter