5

Methods and Parameters

FROM WHAT YOU HAVE LEARNED about C# programming so far, you should be able to write straightforward programs consisting of a list of statements, similar to the way programs were created in the 1970s. Programming has come a long way since the 1970s, however; as programs have become more complex, new paradigms have emerged to manage that complexity. *Procedural* or *structured* programming provides constructs by which statements are grouped together to form units. Furthermore, with structured programming, it is possible to pass data to a group of statements and then have data returned once the statements have executed.

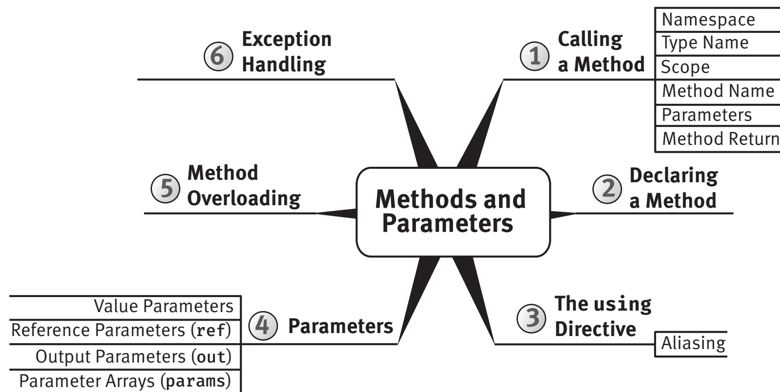

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Besides the basics of calling and defining methods, this chapter covers some slightly more advanced concepts—namely, recursion, method overloading, optional parameters, and named arguments. All method calls discussed so far and through the end of this chapter are static (a concept that Chapter 6 explores in detail).

Even as early as the `HelloWorld` program in Chapter 1, you learned how to define a method. In that example, you defined the `Main()` method. In this chapter, you will learn about method creation in more detail, including the special C# syntaxes (`ref` and `out`) for parameters that pass variables rather than values to methods. Lastly, we will touch on some rudimentary error handling.

Calling a Method

BEGINNER TOPIC

What Is a Method?

Up to this point, all of the statements in the programs you have written have appeared together in one grouping called a `Main()` method. When programs become any more complex than those we have seen thus far, a single method implementation quickly becomes difficult to maintain and complex to read through and understand.

A **method** is a means of grouping together a sequence of statements to perform a particular action or compute a particular result. This provides greater structure and organization for the statements that comprise a program. Consider, for example, a `Main()` method that counts the lines of source code in a directory. Instead of having one large `Main()` method, you can provide a shorter version that allows you to hone in on the details of each method implementation as necessary. Listing 5.1 shows an example.

LISTING 5.1: Grouping Statements into Methods

```
class LineCount
{
  static void Main()
  {
      int lineCount;
      string files;
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
            DisplayHelpText();
            files = GetFiles();
            lineCount = CountLines(files);
            DisplayLineCount(lineCount);
        }
        // ...
    }
```

Instead of placing all of the statements into `Main()`, the listing breaks them into groups called methods. The `System.Console.WriteLine()` statements that display the help text have been moved to the `DisplayHelpText()` method. All of the statements used to determine which files to count appear in the `GetFiles()` method. To actually count the files, the code calls the `CountLines()` method before displaying the results using the `DisplayLineCount()` method. With a quick glance, it is easy to review the code and gain an overview, because the method name describes the purpose of the method.

Guidelines

DO give methods names that are verbs or verb phrases.

A method is always associated with a type—usually a **class**—that provides a means of grouping related methods together.

Methods can receive data via **arguments** that are supplied for their **parameters**. Parameters are variables used for passing data from the **caller** (the code containing the method call) to the invoked method (`Write()`, `WriteLine()`, `GetFiles()`, `CountLines()`, and so on). In Listing 5.1, `files` and `lineCount` are examples of arguments passed to the `CountLines()` and `DisplayLineCount()` methods via their parameters. Methods can also return data to the caller via a **return value** (in Listing 5.1, the `GetFiles()` method call has a return value that is assigned to `files`).

To begin, we reexamine `System.Console.Write()`, `System.Console.WriteLine()`, and `System.Console.ReadLine()` from Chapter 1. This time we look at them as examples of method calls in general instead of looking at the specifics of printing and retrieving data from the console. Listing 5.2 shows each of the three methods in use.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Listing 5.2: A Simple Method Call

```
class HeyYou
{
  static void Main()
  {
      string firstName;
      string lastName;

      System.Console.WriteLine("Hey you!");

      System.Console.Write("Enter your first name: ");

      firstName = System.Console.ReadLine();
      System.Console.Write("Enter your last name: ");
      lastName = System.Console.ReadLine();
      System.Console.WriteLine(
          $"Your full name is { firstName } { lastName }.");
  }
}
```

The parts of the method call include the method name, argument list, and returned value. A fully qualified method name includes a namespace, type name, and method name; a period separates each part of a fully qualified method name. As we will see, methods are often called with only a part of their fully qualified name.

Namespaces

Namespaces are a categorization mechanism for grouping all types related to a particular area of functionality. Namespaces are hierarchical and can have arbitrarily many levels in the hierarchy, though namespaces with more than half a dozen levels are rare. Typically the hierarchy begins with a company name, and then a product name, and then the functional area. For example, in `Microsoft.Win32.Networking`, the outermost namespace is `Microsoft`, which contains an inner namespace `Win32`, which in turn contains an even more deeply nested `Networking` namespace.

Namespaces are primarily used to organize types by area of functionality so that they can be more easily found and understood. However, they can also be used to avoid type name collisions. For example, the compiler can distinguish between two types with the name `Button` as long as each type has a different namespace. Thus you can disambiguate types `System.Web.UI.WebControls.Button` and `System.Windows.Controls.Button`.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

In Listing 5.2, the `Console` type is found within the `System` namespace. The `System` namespace contains the types that enable the programmer to perform many fundamental programming activities. Almost all C# programs use types within the `System` namespace. Table 5.1 provides a listing of other common namespaces.

Table 5.1: Common Namespaces

Begin 4.0

Namespace	Description
`System`	Contains the fundamental types and types for conversion between types, mathematics, program invocation, and environment management.
`System.Collections.Generics`	Contains strongly typed collections that use generics.
`System.Data`	Contains types used for working with databases.
`System.Drawing`	Contains types for drawing to the display device and working with images.
`System.IO`	Contains types for working with directories and manipulating, loading, and saving files.
`System.Linq`	Contains classes and interfaces for querying data in collections using a Language Integrated Query.
`System.Text`	Contains types for working with strings and various text encodings, and for converting between those encodings.
`System.Text.RegularExpressions`	Contains types for working with regular expressions.
`System.Threading`	Contains types for multithreaded programming.
`System.Threading.Tasks`	Contains types for task-based asynchrony.
`System.Web`	Contains types that enable browser-to-server communication, generally over HTTP. The functionality within this namespace is used to support ASP.NET.

continues

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

TABLE 5.1: Common Namespaces (*continued*)

Namespace	Description
`System.Windows`	Contains types for creating rich user interfaces starting with .NET 3.0 using a UI technology called Windows Presentation Framework (WPF) that leverages Extensible Application Markup Language (XAML) for declarative design of the UI.
`System.Xml`	Contains standards-based support for XML processing.

End 4.0

It is not always necessary to provide the namespace when calling a method. For example, if the call expression appears in a type in the same namespace as the called method, the compiler can infer the namespace to be the namespace that contains the type. Later in this chapter, you will see how the `using` directive eliminates the need for a namespace qualifier as well.

Guidelines

DO use PascalCasing for namespace names.

CONSIDER organizing the directory hierarchy for source code files to match the namespace hierarchy.

Type Name

Calls to static methods require the type name qualifier as long as the target method is not within the same type.[1] (As discussed later in the chapter, a `using static` directive allows you to omit the type name.) For example, a call expression of `Console.WriteLine()` found in the method `HelloWorld.Main()` requires the type, `Console`, to be stated. However, just as with the namespace, C# allows the omission of the type name from a method call whenever the method is a member of the type containing the call expression. (Examples of method calls such as this appear in Listing 5.4.) The type name is unnecessary in such cases because the compiler

1. Or base class.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

infers the type from the location of the call. If the compiler can make no such inference, the name must be provided as part of the method call.

At their core, types are a means of grouping together methods and their associated data. For example, `Console` is the type that contains the `Write()`, `WriteLine()`, and `ReadLine()` methods (among others). All of these methods are in the same *group* because they belong to the `Console` type.

Scope

In the previous chapter, you learned that the *scope* of a program element is the region of text in which it can be referred to by its unqualified name. A call that appears inside a type declaration to a method declared in that type does not require the type qualifier because the method is in scope throughout its containing type. Similarly, a type is in scope throughout the namespace that declares it; therefore, a method call that appears in a type in a particular namespace need not specify that namespace in the method call name.

Method Name

Every method call contains a method name, which might or might not be qualified with a namespace and type name, as we have discussed. After the method name comes the argument list; the argument list is a parenthesized, comma-separated list of the values that correspond to the parameters of the method.

Parameters and Arguments

A method can take any number of parameters, and each parameter is of a specific data type. The values that the caller supplies for parameters are called the **arguments**; every argument must correspond to a particular parameter. For example, the following method call has three arguments:

```
System.IO.File.Copy(
    oldFileName, newFileName, false)
```

The method is found on the class `File`, which is located in the namespace `System.IO`. It is declared to have three parameters, with the first and second being of type `string` and the third being of type `bool`. In this example, we use variables (`oldFileName` and `newFileName`) of type `string` for the old and new filenames, and then specify `false` to indicate that the copy should fail if the new filename already exists.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Method Return Values

In contrast to `System.Console.WriteLine()`, the method call `System.Console.ReadLine()` in Listing 5.2 does not have any arguments because the method is declared to take no parameters. However, this method happens to have a **method return value**. The method return value is a means of transferring results from a called method back to the caller. Because `System.Console.ReadLine()` has a return value, it is possible to assign the return value to the variable `firstName`. In addition, it is possible to pass this method return value itself as an argument to another method call, as shown in Listing 5.3.

Listing 5.3: Passing a Method Return Value as an Argument to Another Method Call

```
class Program
{
  static void Main()
  {
      System.Console.Write("Enter your first name: ");
      System.Console.WriteLine("Hello {0}!",
          System.Console.ReadLine());
  }
}
```

Instead of assigning the returned value to a variable and then using that variable as an argument to the call to `System.Console.WriteLine()`, Listing 5.3 calls the `System.Console.ReadLine()` method within the call to `System.Console.WriteLine()`. At execution time, the `System.Console.ReadLine()` method executes first, and its return value is passed directly into the `System.Console.WriteLine()` method, rather than into a variable.

Not all methods return data. Both versions of `System.Console.Write()` and `System.Console.WriteLine()` are examples of such methods. As you will see shortly, these methods specify a return type of `void`, just as the `HelloWorld` declaration of `Main` returned `void`.

Statement versus Method Call

Listing 5.3 provides a demonstration of the difference between a statement and a method call. Although `System.Console.WriteLine("Hello {0}!", System.Console.ReadLine());` is a single statement, it contains two method calls. A statement often contains one or more expressions, and in this example, two of those expressions are method calls. Therefore, method calls form parts of statements.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Although coding multiple method calls in a single statement often reduces the amount of code, it does not necessarily increase the readability and seldom offers a significant performance advantage. Developers should favor readability over brevity.

> **NOTE**
>
> In general, developers should favor readability over brevity. Readability is critical to writing code that is self-documenting and therefore more maintainable over time.

Declaring a Method

Begin 6.0

This section expands on the explanation of declaring a method to include parameters or a return type. Listing 5.4 contains examples of these concepts, and Output 5.1 shows the results.

Listing 5.4: Declaring a Method

```
class IntroducingMethods
{
  public static void Main()
  {
      string firstName;
      string lastName;
      string fullName;
      string initials;

      System.Console.WriteLine("Hey you!");

      firstName = GetUserInput("Enter your first name: ");
      lastName = GetUserInput("Enter your last name: ");

      fullName = GetFullName(firstName, lastName);
      initials = GetInitials(firstName, lastName);
      DisplayGreeting(fullName, initials);
  }

  static string GetUserInput(string prompt)
  {
      System.Console.Write(prompt);
      return System.Console.ReadLine();
  }

  static string GetFullName(  // C# 6.0 expression-bodied method
      string firstName, string lastName) =>
          $"{ firstName } { lastName }";
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
    static void DisplayGreeting(string fullName, string initials)
    {
        System.Console.WriteLine(
            $"Hello { fullName }! Your initials are { initials }");
        return;
    }

    static string GetInitials(string firstName, string lastName)
    {
        return $"{ firstName[0] }. { lastName[0] }.";
    }
}
```

Output 5.1

```
Hey you!
Enter your first name: Inigo
Enter your last name: Montoya
Your full name is Inigo Montoya.
```

End 6.0

Five methods are declared in Listing 5.4. From `Main()` the code calls `GetUserInput()`, followed by a call to `GetFullName()` and `GetInitials()`. All of the last three methods return a value and take arguments. In addition, the listing calls `DisplayGreeting()`, which doesn't return any data. No method in C# can exist outside the confines of an enclosing type; in this case, the enclosing type is the `IntroducingMethods` class. Even the `Main` method examined in Chapter 1 must be within a type.

Language Contrast: C++/Visual Basic—Global Methods

C# provides no global method support; everything must appear within a type declaration. This is why the `Main()` method was marked as `static`—the C# equivalent of a C++ global and Visual Basic "shared" method.

BEGINNER TOPIC

Refactoring into Methods

Moving a set of statements into a method instead of leaving them inline within a larger method is a form of **refactoring**. Refactoring reduces code duplication, because you can call the method from multiple places instead of duplicating the code. Refactoring also increases code readability. As part

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

of the coding process, it is a best practice to continually review your code and look for opportunities to refactor. This involves looking for blocks of code that are difficult to understand at a glance and moving them into a method with a name that clearly defines the code's behavior. This practice is often preferred over commenting a block of code, because the method name serves to describe what the implementation does.

For example, the `Main()` method that is shown in Listing 5.4 results in the same behavior as does the `Main()` method that is shown in Listing 1.16 in Chapter 1. Perhaps even more noteworthy is that although both listings are trivial to follow, Listing 5.4 is easier to grasp at a glance by just viewing the `Main()` method and not worrying about the details of each called method's implementation.

In earlier versions of Visual Studio, you can select a group of statements, right-click on it, and then select the Extract Method refactoring from the Refactoring section of the context menu to automatically move a group of statements to a new method. In Visual Studio 2015, the refactorings are available from the Quick Actions section of the context menu.

Formal Parameter Declaration

Consider the declarations of the `DisplayGreeting()`, `GetFullName()`, and the `GetInitials()` methods. The text that appears between the parentheses of a method declaration is the **formal parameter list**. (As we will see when we discuss generics, methods may also have a **type parameter list**. When it is clear from context which kind of parameters we are discussing, we simply refer to them as *parameters* in a *parameter list*.) Each parameter in the parameter list includes the type of the parameter along with the parameter name. A comma separates each parameter in the list.

Behaviorally, most parameters are virtually identical to local variables, and the naming convention of parameters follows accordingly. Therefore, parameter names use camelCase. Also, it is not possible to declare a local variable (a variable declared inside a method) with the same name as a parameter of the containing method, because this would create two *local variables* of the same name.

Guidelines

DO use camelCasing for parameter names.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Method Return Type Declaration

In addition to GetUserInput(), GetFullName(), and the GetInitials() methods requiring parameters to be specified, each of these methods also includes a **method return type**. You can tell that a method returns a value because a data type appears immediately before the method name in the method declaration. Each of these method examples specifies a string return type. Unlike with parameters, of which there can be any number, only one method return type is allowable.

As with GetUserInput() and GetInitials(), methods with a return type almost always contain one or more return statements that return control to the caller. A return statement consists of the return keyword followed by an expression that computes the value the method is returning. For example, the GetInitials() method's return statement is return $"{ firstName[0] }. { lastName[0] }.";. The expression (an interpolated string in this case) following the return keyword must be compatible with the stated return type of the method.

If a method has a return type, the block of statements that makes up the body of the method must have an *unreachable end point*. That is, there must be no way for control to "fall off the end" of a method without it returning a value. Often the easiest way to ensure that this condition is met is to make the last statement of the method a return statement. However, return statements can appear in locations other than at the end of a method implementation. For example, an if or switch statement in a method implementation could include a return statement within it; see Listing 5.5 for an example.

LISTING 5.5: A return Statement before the End of a Method

```
class Program
{
  static bool MyMethod()
  {
      string command = ObtainCommand();
      switch(command)
      {
          case "quit":
              return false;
          // ... omitted, other cases
          default:
              return true;
      }
  }
}
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

(Note that a `return` statement transfers control out of the `switch`, so no `break` statement is required to prevent illegal fall-through in a switch section that ends with a `return` statement.)

In Listing 5.5, the last statement in the method is not a `return` statement; it is a `switch` statement. However, the compiler can deduce that every possible code path through the method results in a `return`, so that the end point of the method is not reachable. Thus this method is legal even though it does not end with a `return` statement.

If particular code paths include unreachable statements following the `return`, the compiler will issue a warning that indicates the additional statements will never execute.

Though C# allows a method to have multiple `return` statements, code is generally more readable and easier to maintain if there is a single exit location rather than multiple returns sprinkled through various code paths of the method.

Specifying `void` as a return type indicates that there is no return value from the method. As a result, a call to the method may not be assigned to a variable or used as a parameter type at the call site. A `void` method call may be used only as a statement. Furthermore, within the body of the method the `return` statement becomes optional, and when it is specified, there must be no value following the `return` keyword. For example, the return of `Main()` in Listing 5.4 is `void`, and there is no `return` statement within the method. However, `DisplayGreeting()` includes an (optional) `return` statement that is not followed by any returned result.

Begin 7.0

Although, technically, a method can have only one return type, the return type could be a tuple. As a result, starting with C# 7.0, it is possible to return multiple values packaged as a tuple using C# tuple syntax. For example, you could declare a `GetName()` method, as shown in Listing 5.6.

Listing 5.6: Returning Multiple Values Using a Tuple

```
class Program
{
  static string GetUserInput(string prompt)
  {
      System.Console.Write(prompt);
      return System.Console.ReadLine();
  }
  static (string First, string Last) GetName()
  {
      string firstName, lastName;
      firstName = GetUserInput("Enter your first name: ");
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
            lastName = GetUserInput("Enter your last name: ");
            return (firstName, lastName);
        }
        static public void Main()
        {
            (string First, string Last) name = GetName();
            System.Console.WriteLine($"Hello { name.First } { name.Last }!");
        }
    }
```

End 7.0

Technically, of course, we are still returning only one data type, a `ValueTuple<string, string>`; however, effectively, you can return any (preferably reasonable) number you like.

Expression Bodied Methods

To support the simplest of method declarations without the formality of a method body, C# 6.0 introduced **expression bodied methods**, which are declared using an expression rather than a full method body. Listing 5.4's `GetFullName()` method provides an example of the expression bodied method:

```
static string GetFullName( string firstName, string lastName) =>
        $"{ firstName } { lastName }";
```

In place of the curly brackets typical of a method body, an expression bodied method uses the "goes to" operator (fully introduced in Chapter 13), for which the resulting data type must match the return type of the method. In other words, even though there is no explicit `return` statement in the expression bodied method implementation, it is still necessary that the return type from the expression match the method declaration's return type.

Expression bodied methods are syntactic shortcuts to the fuller method body declaration. As such, their use should be limited to the simplest of method implementations—generally expressible on a single line.

Language Contrast: C++—Header Files

Unlike in C++, C# classes never separate the implementation from the declaration. In C#, there is no header (`.h`) file or implementation (`.cpp`) file. Instead, declaration and implementation appear together in the same file. (C# does support an advanced feature called *partial methods*, in which the method's defining declaration is separate from its implementation, but for the purposes of this chapter, we consider only nonpartial methods.) The lack of separate declaration and implementation in C# removes the requirement to maintain redundant declaration information in two places found in languages that have separate header and implementation files, such as C++.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

BEGINNER TOPIC

Namespaces

As described earlier, **namespaces** are an organizational mechanism for categorizing and grouping together related types. Developers can discover related types by examining other types within the same namespace as a familiar type. Additionally, through namespaces, two or more types may have the same name as long as they are disambiguated by different namespaces.

The `using` Directive

Fully qualified namespace names can become quite long and unwieldy. It is possible, however, to import all the types from one or more namespaces into a file so that they can be used without full qualification. To achieve this, the C# programmer includes a `using` directive, generally at the top of the file. For example, in Listing 5.7, `Console` is not prefixed with `System`. The namespace may be omitted because of the `using System` directive that appears at the top of the listing.

Listing 5.7: using Directive Example

```
// The using directive imports all types from the
// specified namespace into the entire file
using System;

class HelloWorld
{
  static void Main()
  {
      // No need to qualify Console with System
      // because of the using directive above
      Console.WriteLine("Hello, my name is Inigo Montoya");
  }
}
```

The results of Listing 5.7 appear in Output 5.2.

Output 5.2

```
Hello, my name is Inigo Montoya
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

A `using` directive such as `using System` does not enable you to omit `System` from a type declared within a child namespace of `System`. For example, if your code accessed the `StringBuilder` type from the `System.Text` namespace, you would have to either include an additional `using System.Text;` directive or fully qualify the type as `System.Text.StringBuilder`, not just `Text.StringBuilder`. In short, a `using` directive does not import types from any **nested namespaces**. Nested namespaces, which are identified by the period in the namespace, always need to be imported explicitly.

Language Contrast: Java—Wildcards in `import` Directive

Java allows for importing namespaces using a wildcard such as the following:

```
import javax.swing.*;
```

In contrast, C# does not support a wildcard `using` directive but instead requires each namespace to be imported explicitly.

Language Contrast: Visual Basic .NET—Project Scope `Imports` Directive

Unlike C#, Visual Basic .NET supports the ability to specify the `using` directive equivalent, `Imports`, for an entire project rather than for just a specific file. In other words, Visual Basic .NET provides a command-line means of the `using` directive that will span an entire compilation.

Frequent use of types within a particular namespace implies that the addition of a `using` directive for that namespace is a good idea, instead of fully qualifying all types within the namespace. Accordingly, almost all C# files include the `using System` directive at the top. Throughout the remainder of this book, code listings often omit the `using System` directive. Other namespace directives are included explicitly, however.

One interesting effect of the `using System` directive is that the string data type can be identified with varying case: `String` or `string`. The former version relies on the `using System` directive and the latter uses the `string` keyword. Both are valid C# references to the `System.String` data

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

type, and the resultant Common Intermediate Language (CIL) code is unaffected by which version is chosen.[2]

ADVANCED TOPIC

Nested using Directives

Not only can you have `using` directives at the top of a file, but you also can include them at the top of a namespace declaration. For example, if a new namespace, `EssentialCSharp`, were declared, it would be possible to add a `using` declarative at the top of the namespace declaration (see Listing 5.8).

Listing 5.8: Specifying the using Directive inside a Namespace Declaration

```
namespace EssentialCSharp
{
  using System;

  class HelloWorld
  {
    static void Main()
    {
        // No need to qualify Console with System
        // because of the using directive above
        Console.WriteLine("Hello, my name is Inigo Montoya");
    }
  }
}
```

The results of Listing 5.8 appear in Output 5.3.

Output 5.3

```
Hello, my name is Inigo Montoya
```

The difference between placing the `using` directive at the top of a file and placing it at the top of a namespace declaration is that the directive is active only within the namespace declaration. If the code includes a new

2. I prefer the `string` keyword, but whichever representation a programmer selects, the code within a project ideally should be consistent.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

namespace declaration above or below the `EssentialCSharp` declaration, the `using System` directive within a different namespace would not be active. Code seldom is written this way, especially given the standard practice of providing a single type declaration per file.

Begin 6.0

`using static` Directive

The `using` directive allows you to abbreviate a type name by omitting the namespace portion of the name—such that just the type name can be specified for any type within the stated namespace. In contrast, the `using static` directive allows you to omit both the namespace and the type name from any member of the stated type. A `using static System.Console` directive, for example, allows you to specify `WriteLine()` rather than the fully qualified method name of `System.Console.WriteLine()`. Continuing with this example, we can update Listing 5.2 to leverage the `using static System.Console` directive to create Listing 5.9.

LISTING 5.9: `using static` Directive

```
using static System.Console;

class HeyYou
{
  static void Main()
  {
      string firstName;
      string lastName;

      WriteLine("Hey you!");

      Write("Enter your first name: ");

      firstName = ReadLine();
      Write("Enter your last name: ");
      lastName = ReadLine();
      WriteLine(
          $"Your full name is { firstName } { lastName }.");
  }
}
```

In this case, there is no loss of readability of the code: `WriteLine()`, `Write()`, and `ReadLine()` all clearly relate to a console directive. In fact, one could argue that the resulting code is simpler and therefore clearer than before.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

However, sometimes this is not the case. For example, if your code uses classes that have overlapping behavior names, such as an `Exists()` method on a file and an `Exists()` method on a directory, then perhaps a `using static` directive would reduce clarity when you invoke `Exists()`. Similarly, if the class you were writing had its own members with overlapping behavior names—for example, `Display()` and `Write()`—then perhaps clarity would be lost to the reader.

This ambiguity would not be allowed by the compiler. If two members with the same signature were available (through either `using static` directives or separately declared members), any invocation of them that was ambiguous would result in a compile error.

End 6.0

Aliasing

The `using` directive also allows **aliasing** a namespace or type. An alias is an alternative name that you can use within the text to which the `using` directive applies. The two most common reasons for aliasing are to disambiguate two types that have the same name and to abbreviate a long name. In Listing 5.10, for example, the `CountDownTimer` alias is declared as a means of referring to the type `System.Timers.Timer`. Simply adding a `using System.Timers` directive will not sufficiently enable the code to avoid fully qualifying the `Timer` type. The reason is that `System.Threading` also includes a type called `Timer`; therefore, using just `Timer` within the code will be ambiguous.

Listing 5.10: Declaring a Type Alias

```
using System;
using System.Threading;
using CountDownTimer = System.Timers.Timer;

class HelloWorld
{
  static void Main()
  {
      CountDownTimer timer;

      // ...
  }
}
```

Listing 5.10 uses an entirely new name, `CountDownTimer`, as the alias. It is possible, however, to specify the alias as `Timer`, as shown in Listing 5.11.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Listing 5.11: Declaring a Type Alias with the Same Name

```
using System;
using System.Threading;

// Declare alias Timer to refer to System.Timers.Timer to
// avoid code ambiguity with System.Threading.Timer
using Timer = System.Timers.Timer;

class HelloWorld
{
  static void Main()
  {
      Timer timer;

      // ...
  }
}
```

Because of the alias directive, "Timer" is not an ambiguous reference. Furthermore, to refer to the `System.Threading.Timer` type, you will have to either qualify the type or define a different alias.

Returns and Parameters on `Main()`

So far, declaration of an executable's `Main()` method has been the simplest declaration possible. You have not included any parameters or non-`void` return type in your `Main()` method declarations. However, C# supports the ability to retrieve the command-line arguments when executing a program, and it is possible to return a status indicator from the `Main()` method.

The runtime passes the command-line arguments to `Main()` using a single `string` array parameter. All you need to do to retrieve the parameters is to access the array, as demonstrated in Listing 5.12. The purpose of this program is to download a file whose location is given by a URL. The first command-line argument identifies the URL, and the optional second argument is the filename to which to save the file. The listing begins with a `switch` statement that evaluates the number of parameters (`args.Length`) as follows:

1. If there are not two parameters, display an error indicating that it is necessary to provide the URL and filename.
2. The presence of two arguments indicates the user has provided both the URL of the resource and the download target filename.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Listing 5.12: Passing Command-Line Arguments to Main

```
using System;
using System.Net;

class Program
{
  static int Main(string[] args)
  {
      int result;
      string targetFileName;
      string url;
      switch (args.Length)
      {
          default:
              // Exactly two arguments must be specified; give an error
              Console.WriteLine(
                  "ERROR:  You must specify the "
                  + "URL and the file name");
              targetFileName = null;
              url = null;
              break;
          case 2:
              url = args[0];
              targetFileName = args[1];
              break;
      }

      if (targetFileName != null && url != null)
      {
          WebClient webClient = new WebClient();
          webClient.DownloadFile(url, targetFileName);
          result = 0;
      }
      else
      {
          Console.WriteLine(
              "Usage: Downloader.exe <URL> <TargetFileName>");
          result = 1;
      }
      return result;
  }


}
```

The results of Listing 5.12 appear in Output 5.4.

Output 5.4

```
>Downloader.exe
ERROR:  You must specify the URL to be downloaded
Downloader.exe <URL> <TargetFileName>
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

If you were successful in calculating the target filename, you would use it to save the downloaded file. Otherwise, you would display the help text. The `Main()` method also returns an `int` rather than a `void`. This is optional for a `Main()` declaration, but if it is used, the program can return a status code to a caller (such as a script or a batch file). By convention, a return other than zero indicates an error.

Although all command-line arguments can be passed to `Main()` via an array of strings, sometimes it is convenient to access the arguments from inside a method other than `Main()`. The `System.Environment.GetCommandLineArgs()` method returns the command-line arguments array in the same form that `Main(string[] args)` passes the arguments into `Main()`.

ADVANCED TOPIC

Disambiguate Multiple `Main()` Methods

If a program includes two classes with `Main()` methods, it is possible to specify on the command line which class to use for the `Main()` declaration. `csc.exe` includes an `/m` option to specify the fully qualified class name of `Main()`.

BEGINNER TOPIC

Call Stack and Call Site

As code executes, methods call more methods, which in turn call additional methods, and so on. In the simple case of Listing 5.4, `Main()` calls `GetUserInput()`, which in turn calls `System.Console.ReadLine()`, which in turn calls even more methods internally. Every time a new method is invoked, the runtime creates an *activation frame* that contains information about the arguments passed to the new call, the local variables of the new call, and information about where control should resume when the new method returns. The set of calls within calls within calls, and so on, produces a series of activation frames that is termed the **call stack**.[3]

3. Except for async or iterator methods, which move their activator records onto the heap.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

As program complexity increases, the call stack generally gets larger and larger as each method calls another method. As calls complete, however, the call stack shrinks until another method is invoked. The process of removing activation frames from the call stack is termed **stack unwinding**. Stack unwinding always occurs in the reverse order of the method calls. When the method completes, execution returns to the **call site**—that is, the location from which the method was invoked.

Advanced Method Parameters

So far this chapter's examples have returned data via the method return value. This section demonstrates how methods can return data via their method parameters and how a method may take a variable number of arguments.

Value Parameters

Arguments to method calls are usually **passed by value**, which means the value of the argument expression is copied into the target parameter. For example, in Listing 5.13, the value of each variable that `Main()` uses when calling `Combine()` will be copied into the parameters of the `Combine()` method. Output 5.5 shows the results of this listing.

Listing 5.13: Passing Variables by Value

```
class Program
{
  static void Main()
  {
      // ...
      string fullName;
      string driveLetter = "C:";
      string folderPath  = "Data";
      string fileName    = "index.html";

      fullName = Combine(driveLetter, folderPath, fileName);

      Console.WriteLine(fullName);
      // ...
  }

  static string Combine(
      string driveLetter, string folderPath, string fileName)
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
        {
            string path;
            path = string.Format("{1}{0}{2}{0}{3}",
                System.IO.Path.DirectorySeparatorChar,
                driveLetter, folderPath, fileName);
            return path;
        }
    }
```

Output 5.5

```
C:\Data\index.html
```

Even if the `Combine()` method assigns `null` to `driveLetter`, `folderPath`, and `fileName` before returning, the corresponding variables within `Main()` will maintain their original values because the variables are copied when calling a method. When the call stack unwinds at the end of a call, the copied data is thrown away.

BEGINNER TOPIC

Matching Caller Variables with Parameter Names

In Listing 5.13, the variable names in the caller exactly matched the parameter names in the called method. This matching is provided simply for readability purposes; whether names match is entirely irrelevant to the behavior of the method call. The parameters of the called method and the local variables of the calling method are found in different declaration spaces and have nothing to do with each other.

ADVANCED TOPIC

Reference Types versus Value Types

For the purposes of this section, it is inconsequential whether the parameter passed is a value type or a reference type. Rather, the important issue is whether the called method can write a value into the caller's original variable. Since a copy of the caller variable's value is made, the caller's variable cannot be reassigned. Nevertheless, it is helpful to understand the difference between a variable that contains a value type and a variable that contains a reference type.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

The value of a reference type variable is, as the name implies, a reference to the location where the data associated with the object is stored. How the runtime chooses to represent the value of a reference type variable is an implementation detail of the runtime; typically it is represented as the address of the memory location in which the object's data is stored, but it need not be.

If a reference type variable is passed by value, the reference itself is copied from the caller to the method parameter. As a result, the target method cannot update the caller variable's value but it may update the data referred to by the reference.

Alternatively, if the method parameter is a value type, the value itself is copied into the parameter, and changing the parameter in the called method will not affect the original caller's variable.

Reference Parameters (`ref`)

Consider Listing 5.14, which calls a function to swap two values, and Output 5.6, which shows the results.

Listing 5.14: Passing Variables by Reference

```
class Program
{
  static void Main()
  {
      // ...
      string first = "hello";
      string second = "goodbye";
      Swap(ref first, ref second);

      Console.WriteLine(
          $@"first = ""{ first }"", second = ""{ second }""" );
      // ...
  }

  static void Swap(ref string x, ref string y)
  {
      string temp = x;
      x = y;
      y = temp;
  }
}
```

Output 5.6

```
first = "goodbye", second = "hello"
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

The values assigned to first and second are successfully switched. To do this, the variables are **passed by reference**. The obvious difference between the call to Swap() and Listing 5.13's call to Combine() is the inclusion of the keyword ref in front of the parameter's data type. This keyword changes the call such that the variables used as arguments are passed by reference, so the called method can update the original caller's variables with new values.

When the called method specifies a parameter as ref, the caller is required to supply a variable, not a value, as an argument and to place ref in front of the variables passed. In so doing, the caller explicitly recognizes that the target method could reassign the values of the variables associated with any ref parameters it receives. Furthermore, it is necessary to initialize any local variables passed as ref because target methods could read data from ref parameters without first assigning them. In Listing 5.14, for example, temp is assigned the value of first, assuming that the variable passed in first was initialized by the caller. Effectively, a ref parameter is an alias for the variable passed. In other words, it is essentially giving a parameter name to an existing variable, rather than creating a new variable and copying the value of the argument into it.

Begin 7.0

Output Parameters (out)

As mentioned earlier, a variable used as a ref parameter must be assigned before it is passed to the called method, because the called method might read from the variable. The "swap" example given previously must read and write from both variables passed to it. However, it is often the case that a method that takes a reference to a variable intends to write to the variable but not to read from it. In such cases, clearly it could be safe to pass an uninitialized local variable by reference.

To achieve this, code needs to decorate parameter types with the keyword out. This is demonstrated in the TryGetPhoneButton() method in Listing 5.15, which returns the phone button corresponding to a character.

Listing 5.15: Passing Variables Out Only

```
class ConvertToPhoneNumber
{
  static int Main(string[] args)
  {
      if(args.Length == 0)
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

7.0

```
        {
            Console.WriteLine(
                "ConvertToPhoneNumber.exe <phrase>");
            Console.WriteLine(
                "'_' indicates no standard phone button");
            return 1;
        }
        foreach(string word in args)
        {
            foreach(char character in word)
            {
                if(TryGetPhoneButton(character, out char button))
                {
                    Console.Write(button);
                }
                else
                {
                    Console.Write('_');
                }
            }
        }
        Console.WriteLine();
        return 0;
    }

    static bool TryGetPhoneButton(char character, out char button)
    {
        bool success = true;
        switch( char.ToLower(character) )
        {
            case '1':
                button = '1';
                break;
            case '2': case 'a': case 'b': case 'c':
                button = '2';
                break;

            // ...

            case '-':
                button = '-';
                break;
            default:
                  // Set the button to indicate an invalid value
                  button = '_';
                success = false;
                break;
        }
        return success;
    }
}
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Output 5.7 shows the results of Listing 5.15.

7.0

Output 5.7

```
>ConvertToPhoneNumber.exe CSharpIsGood
274277474663
```

In this example, the `TryGetPhoneButton()` method returns `true` if it can successfully determine the `character`'s corresponding phone button. The function also returns the corresponding button by using the `button` parameter, which is decorated with `out`.

An `out` parameter is functionally identical to a `ref` parameter; the only difference is which requirements the language enforces regarding how the aliased variable is read from and written to. Whenever a parameter is marked with `out`, the compiler checks that the parameter is set for all code paths within the method that return normally (i.e., the code paths that do not throw an exception). If, for example, the code does not assign `button` a value in some code path, the compiler will issue an error indicating that the code didn't initialize `button`. Listing 5.15 assigns `button` to the underscore character because even though it cannot determine the correct phone button, it is still necessary to assign a value.

A common coding error when working with `out` parameters is to forget to declare the `out` variable before you use it. Starting with C# 7.0, it is possible to declare the `out` variable inline when invoking the function. Listing 5.15 uses this feature with the statement `TryGetPhoneButton(character, out char button)` without ever declaring the `button` variable beforehand. Prior to C# 7.0, it would be necessary to first declare the `button` variable and then invoke the function with `TryGetPhoneButton(character, out button)`.

Another C# 7.0 feature is the ability to discard an `out` parameter entirely. If, for example, you simply wanted to know whether a character was a valid phone button but not actually return the numeric value, you could discard the `button` parameter using an underscore: `TryGetPhoneButton(character, out _)`.

Prior to C# 7.0's tuple syntax, a developer of a method might declare one or more `out` parameters to get around the restriction that a method may have only one return type; a method that needs to return two values can do so by returning one value normally, as the return value of the method,

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

and a second value by writing it into an aliased variable passed as an `out` parameter. Although this pattern is both common and legal, there are usually better ways to achieve that aim. For example, if you are considering returning two or more values from a method and C# 7.0 is available, it is likely preferable to use C# 7.0 tuple syntax. Prior to that, consider writing two methods, one for each value, or still using the `System.ValueTuple` type (which would require referencing the `System.ValueTuple NuGet` package) but without C# 7.0 syntax.

7.0

> **NOTE**
>
> Each and every normal code path must result in the assignment of all `out` parameters.

Read-Only Pass by Reference (in)

Begin 7.2

In C# 7.2, support was added for passing a value type by reference that was read only. Rather than passing the value type to a function so that it could be changed, read-only pass by reference was added so that the value type could be passed by reference so that not only copy of the value type occurred but, in addition, the invoked method could not change the value type. In other words, the purpose of the feature is to reduce the memory copied when passing a value while still identifying it as read only, thus improving the performance. This syntax is to add an `in` modifier to the parameter. For example:

```
int Method(in int number) { ... }
```

With the `in` modifier, any attempts to reassign number (`number++`, for example) will result in a compile error indicating that number is read only.

End 7.2

Return by Reference

Another C# 7.0 addition is support for returning a reference to a variable. Consider, for example, a function that returns the first pixel in an image that is associated with red-eye, as shown in Listing 5.16.

Listing 5.16: ref Return and ref Local Declaration

```
// Returning a reference
public static ref byte FindFirstRedEyePixel(byte[] image)
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

7.0

```
    {
      // Do fancy image detection perhaps with machine learning
      for (int counter = 0; counter < image.Length; counter++)
      {
        if(image[counter] == (byte)ConsoleColor.Red)
        {
          return ref image[counter];
        }
      }
      throw new InvalidOperationException("No pixels are red.");
    }
    public static void Main()
    {
      byte[] image = new byte[254];
      // Load image
      int index = new Random().Next(0, image.Length - 1);
      image[index] =
          (byte)ConsoleColor.Red;
      System.Console.WriteLine(
          $"image[{index}]={(ConsoleColor)image[index]}");
      // ...

      // Obtain a reference to the first red pixel
      ref byte redPixel = ref FindFirstRedEyePixel(image);
      // Update it to be Black
      redPixel = (byte)ConsoleColor.Black;
      System.Console.WriteLine(
          $"image[{index}]={(ConsoleColor)image[redPixel]}");
    }
```

By returning a reference to the variable, the caller is then able to update the pixel to a different color, as shown in the highlighted line of Listing 5.16. Checking for the update via the array shows that the value is now black.

There are two important restrictions on return by reference—both due to object lifetime: Object references shouldn't be garbage collected while they're still referenced, and they shouldn't consume memory when they no longer have any references. To enforce these restrictions, you can only return the following from a reference-returning function:

- References to fields or array elements
- Other reference-returning properties or functions
- References that were passed in as parameters to the by-reference-returning function

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

For example, `FindFirstRedEyePixel()` returns a reference to an item in the image array, which was a parameter to the function. Similarly, if the image was stored as a field within the class, you could return the field by reference:

```
byte[] _Image;
public ref byte[] Image { get {  return ref _Image; } }
```

Second, `ref` locals are initialized to refer to a particular variable and can't be modified to refer to a different variable.

There are several return-by-reference characteristics of which to be cognizant:

- If you're returning a reference, you obviously must return it. This means, therefore, that in the example in Listing 5.16, even if no red-eye pixel exists, you still need to return a reference byte. The only workaround would be to throw an exception. In contrast, the by-reference parameter approach allows you to leave the parameter unchanged and return a `bool` indicating success. In many cases, this might be preferable.
- When declaring a reference local variable, initialization is required. This involves assigning it a `ref` return from a function or a reference to a variable:

  ```
  ref string text;  // Error
  ```

- Although it's possible in C# 7.0 to declare a reference local variable, declaring a field of type `ref` isn't allowed:

  ```
  class Thing { ref string _Text;  /* Error */ }
  ```

- You can't declare a by-reference type for an auto-implemented property:

  ```
  class Thing { ref string Text { get;set; }  /* Error */ }
  ```

- Properties that return a reference are allowed:

  ```
  class Thing { string _Text = "Inigo Montoya";
  ref string Text { get { return ref _Text; } } }
  ```

- A reference local variable can't be initialized with a value (such as null or a constant). It must be assigned from a by-reference-returning member or a local variable, field, or array element:

  ```
  ref int number = null; ref int number = 42;  // ERROR
  ```

End 7.0

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Parameter Arrays (params)

In the examples so far, the number of arguments that must be passed has been fixed by the number of parameters declared in the target method declaration. However, sometimes it is convenient if the number of arguments may vary. Consider the `Combine()` method from Listing 5.13. In that method, you passed the drive letter, folder path, and filename. What if the path had more than one folder, and the caller wanted the method to join additional folders to form the full path? Perhaps the best option would be to pass an array of strings for the folders. However, this would make the calling code a little more complex, because it would be necessary to construct an array to pass as an argument.

To make it easier on the callers of such a method, C# provides a keyword that enables the number of arguments to vary in the calling code instead of being set by the target method. Before we discuss the method declaration, observe the calling code declared within `Main()`, as shown in Listing 5.17.

Listing 5.17: Passing a Variable Parameter List

```
using System;
using System.IO;
class PathEx
{
  static void Main()
  {
      string fullName;

      // ...

      // Call Combine() with four arguments
      fullName = Combine(
          Directory.GetCurrentDirectory(),
          "bin", "config", "index.html");
      Console.WriteLine(fullName);

      // ...

      // Call Combine() with only three arguments
      fullName = Combine(
          Environment.SystemDirectory,
          "Temp", "index.html");
      Console.WriteLine(fullName);

      // ...
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
        // Call Combine() with an array
        fullName = Combine(
            new string[] {
                "C:\\", "Data",
                "HomeDir", "index.html"} );
        Console.WriteLine(fullName);
        // ...
    }

    static string Combine(params string[] paths)
    {
        string result = string.Empty;
        foreach (string path in paths)
        {
            result = Path.Combine(result, path);
        }
        return result;
    }
}
```

Output 5.8 shows the results of Listing 5.17.

OUTPUT 5.8

```
C:\Data\mark\bin\config\index.html
C:\WINDOWS\system32\Temp\index.html
C:\Data\HomeDir\index.html
```

In the first call to `Combine()`, four arguments are specified. The second call contains only three arguments. In the final call, a single argument is passed using an array. In other words, the `Combine()` method takes a variable number of arguments—presented either as any number of string arguments separated by commas or as a single array of strings. The former syntax is called the *expanded* form of the method call, and the latter form is called the *normal* form.

To allow invocation using either form, the `Combine()` method does the following:

1. Places `params` immediately before the last parameter in the method declaration
2. Declares the last parameter as an array

With a **parameter array** declaration, it is possible to access each corresponding argument as a member of the `params` array. In the `Combine()`

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

method implementation, you iterate over the elements of the `paths` array and call `System.IO.Path.Combine()`. This method automatically combines the parts of the path, appropriately using the platform-specific directory-separator character. Note that `PathEx.Combine()` is identical to `Path.Combine()` except that `PathEx.Combine()` handles a variable number of parameters rather than simply two.

There are a few notable characteristics of the parameter array:

- The parameter array is not necessarily the only parameter on a method.
- The parameter array must be the last parameter in the method declaration. Since only the last parameter may be a parameter array, a method cannot have more than one parameter array.
- The caller can specify zero arguments that correspond to the parameter array parameter, which will result in an array of zero items being passed as the parameter array.
- Parameter arrays are type-safe: The arguments given must be compatible with the element type of the parameter array.
- The caller can use an explicit array rather than a comma-separated list of parameters. The resulting CIL code is identical.
- If the target method implementation requires a minimum number of parameters, those parameters should appear explicitly within the method declaration, forcing a compile error instead of relying on runtime error handling if required parameters are missing. For example, if you have a method that requires one or more integer arguments, declare the method as `int Max(int first, params int[] operands)` rather than as `int Max(params int[] operands)` so that at least one value is passed to `Max()`.

Using a parameter array, you can pass a variable number of arguments of the same type into a method. The section "Method Overloading," which appears later in this chapter, discusses a means of supporting a variable number of arguments that are not necessarily of the same type.

Guidelines

DO use parameter arrays when a method can handle any number—including zero—of additional arguments.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Recursion

Calling a method **recursively** or implementing the method using **recursion** refers to use of a method that calls itself. Recursion is sometimes the simplest way to implement a particular algorithm. Listing 5.18 counts the lines of all the C# source files (`*.cs`) in a directory and its subdirectory.

Listing 5.18: Counting the Lines within *.cs Files, Given a Directory

```
using System.IO;

public static class LineCounter
{
  // Use the first argument as the directory
  // to search, or default to the current directory
  public static void Main(string[] args)
  {
      int totalLineCount = 0;
      string directory;
      if (args.Length > 0)
      {
          directory = args[0];
      }
      else
      {
          directory = Directory.GetCurrentDirectory();
      }
      totalLineCount = DirectoryCountLines(directory);
      System.Console.WriteLine(totalLineCount);
    }

  static int DirectoryCountLines(string directory)
  {
      int lineCount = 0;
      foreach (string file in
          Directory.GetFiles(directory, "*.cs"))
      {
            lineCount += CountLines(file);
        }

      foreach (string subdirectory in
          Directory.GetDirectories(directory))
      {
          lineCount += DirectoryCountLines(subdirectory);
      }

      return lineCount;
  }
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
    private static int CountLines(string file)
    {
        string line;
        int lineCount = 0;
        FileStream stream =
            new FileStream(file, FileMode.Open);⁴
        StreamReader reader = new StreamReader(stream);
        line = reader.ReadLine();

        while(line != null)
        {
            if (line.Trim() != "")
            {
                lineCount++;
            }
            line = reader.ReadLine();
        }

        reader.Close();  // Automatically closes the stream
        return lineCount;
    }
}
```

Output 5.9 shows the results of Listing 5.18.

Output 5.9

```
104
```

The program begins by passing the first command-line argument to `DirectoryCountLines()` or by using the current directory if no argument is provided. This method first iterates through all the files in the current directory and totals the source code lines for each file. After processing each file in the directory, the code processes each subdirectory by passing the subdirectory back into the `DirectoryCountLines()` method, rerunning the method using the subdirectory. The same process is repeated recursively through each subdirectory until no more directories remain to process.

Readers unfamiliar with recursion may find it confusing at first. Regardless, it is often the simplest pattern to code, especially with hierarchical type data such as the filesystem. However, although it may be the most readable approach, it is generally not the fastest implementation. If

4. This code could be improved with a `using` statement, a construct that we have avoided because it has not yet been introduced.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

performance becomes an issue, developers should seek an alternative solution to a recursive implementation. The choice generally hinges on balancing readability with performance.

BEGINNER TOPIC

Infinite Recursion Error

A common programming error in recursive method implementations appears in the form of a stack overflow during program execution. This usually happens because of **infinite recursion**, in which the method continually calls back on itself, never reaching a point that triggers the end of the recursion. It is a good practice for programmers to review any method that uses recursion and to verify that the recursion calls are finite.

A common pattern for recursion using pseudocode is as follows:

```
M(x)
{
  if x is trivial
    return the result
  else
    a. Do some work to make the problem smaller
    b. Recursively call M to solve the smaller problem
    c. Compute the result based on a. and b.
    return the result
}
```

Things go wrong when this pattern is not followed. For example, if you don't make the problem smaller or if you don't handle all possible "smallest" cases, the recursion never terminates.

Method Overloading

Listing 5.18 called `DirectoryCountLines()`, which counted the lines of `*.cs` files. However, if you want to count code in `*.h`/`*.cpp` files or in `*.vb` files, `DirectoryCountLines()` will not work. Instead, you need a method that takes the file extension but still keeps the existing method definition so that it handles `*.cs` files by default.

All methods within a class must have a unique signature, and C# defines uniqueness by variation in the method name, parameter data types, or number of parameters. This does not include method return data types; defining two methods that differ only in their return data types

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

will cause a compile error. This is true even if the return type is two different tuples. **Method overloading** occurs when a class has two or more methods with the same name and the parameter count and/or data types vary between the overloaded methods.

> **NOTE**
>
> A method is considered unique as long as there is variation in the method name, parameter data types, or number of parameters.

Method overloading is a type of **operational polymorphism**. Polymorphism occurs when the same logical operation takes on many ("poly") forms ("morphs") because the data varies. Calling `WriteLine()` and passing a format string along with some parameters is implemented differently than calling `WriteLine()` and specifying an integer. However, logically, to the caller, the method takes care of writing the data, and it is somewhat irrelevant how the internal implementation occurs. Listing 5.19 provides an example, and Output 5.10 shows the results.

LISTING 5.19: Counting the Lines within *.cs Files Using Overloading

```
using System.IO;

public static class LineCounter
{
  public static void Main(string[] args)
  {
      int totalLineCount;

      if (args.Length > 1)
      {
          totalLineCount =
              DirectoryCountLines(args[0], args[1]);
      }
      if (args.Length > 0)
      {
          totalLineCount = DirectoryCountLines(args[0]);
      }
      else
      {
          totalLineCount  = DirectoryCountLines();
      }

      System.Console.WriteLine(totalLineCount);
  }
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
    static int DirectoryCountLines()
    {
        return DirectoryCountLines(
            Directory.GetCurrentDirectory());
    }

    static int DirectoryCountLines(string directory)
    {
        return DirectoryCountLines(directory, "*.cs");
    }

    static int DirectoryCountLines(
        string directory, string extension)
    {
        int lineCount = 0;
        foreach (string file in
            Directory.GetFiles(directory, extension))
        {
            lineCount += CountLines(file);
        }

        foreach (string subdirectory in
            Directory.GetDirectories(directory))
        {
            lineCount += DirectoryCountLines(subdirectory);
        }

        return lineCount;
    }

    private static int CountLines(string file)
    {
        int lineCount = 0;
        string line;
        FileStream stream =
            new FileStream(file, FileMode.Open);5
        StreamReader reader = new StreamReader(stream);
        line = reader.ReadLine();
        while(line != null)
        {
            if (line.Trim() != "")
            {
                lineCount++;
            }
            line = reader.ReadLine();
        }
```

5. This code could be improved with a `using` statement, a construct that we have avoided because it has not yet been introduced.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
        reader.Close();  // Automatically closes the stream
        return lineCount;
    }
}
```

Output 5.10

```
>LineCounter.exe .\ *.cs
28
```

The effect of method overloading is to provide optional ways to call the method. As demonstrated inside `Main()`, you can call the `DirectoryCountLines()` method with or without passing the directory to search and the file extension.

Notice that the parameterless implementation of `DirectoryCountLines()` was changed to call the single-parameter version (`int DirectoryCountLines (string directory)`). This is a common pattern when implementing overloaded methods. The idea is that developers implement only the core logic in one method, and all the other overloaded methods will call that single method. If the core implementation changes, it needs to be modified in only one location rather than within each implementation. This pattern is especially prevalent when using method overloading to enable optional parameters that do not have values determined at compile time, so they cannot be specified using optional parameters.

> **NOTE**
>
> Placing the core functionality into a single method that all other overloading methods invoke means that you can make changes in implementation in just the core method, which the other methods will automatically take advantage of.

Begin 4.0

Optional Parameters

Starting with C# 4.0, the language designers added support for **optional parameters**. By allowing the association of a parameter with a constant value as part of the method declaration, it is possible to call a method without passing an argument for every parameter of the method (see Listing 5.20).

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Listing 5.20: Methods with Optional Parameters

```
using System.IO;

public static class LineCounter
{
  public static void Main(string[] args)
  {
      int totalLineCount;

      if (args.Length > 1)
      {
          totalLineCount =
              DirectoryCountLines(args[0], args[1]);
      }
      if (args.Length > 0)
      {
          totalLineCount = DirectoryCountLines(args[0]);
      }
      else
      {
          totalLineCount  = DirectoryCountLines();
      }

      System.Console.WriteLine(totalLineCount);
  }

  static int DirectoryCountLines()
  {
      // ...
  }

/*
  static int DirectoryCountLines(string directory)
  { ... }
*/

  static int DirectoryCountLines(
      string directory, string extension = "*.cs")
  {
      int lineCount = 0;
      foreach (string file in
          Directory.GetFiles(directory, extension))
      {
          lineCount += CountLines(file);
      }

      foreach (string subdirectory in
          Directory.GetDirectories(directory))
      {
          lineCount += DirectoryCountLines(subdirectory);
      }

      return lineCount;
  }
```

4.0

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
        private static int CountLines(string file)
        {
            // ...
        }
    }
```

In Listing 5.20, the `DirectoryCountLines()` method declaration with a single parameter has been removed (commented out), but the call from `Main()` (specifying one parameter) remains. When no `extension` parameter is specified in the call, the value assigned to `extension` within the declaration (`*.cs` in this case) is used. This allows the calling code to not specify a value if desired, and it eliminates the additional overload that would be required in C# 3.0 and earlier. Note that optional parameters must appear after all required parameters (those that don't have default values). Also, the fact that the default value needs to be a constant, compile-time–resolved value is fairly restrictive. You cannot, for example, declare a method like

```
DirectoryCountLines(
  string directory = Environment.CurrentDirectory,
  string extension = "*.cs")
```

because `Environment.CurrentDirectory` is not a constant. In contrast, because `"*.cs"` is a constant, C# does allow it for the default value of an optional parameter.

4.0

Guidelines

DO provide good defaults for all parameters where possible.

DO provide simple method overloads that have a small number of required parameters.

CONSIDER organizing overloads from the simplest to the most complex.

A second method call feature made available in C# 4.0 is the use of **named arguments**. With named arguments, it is possible for the caller to explicitly identify the name of the parameter to be assigned a value rather than relying solely on parameter and argument order to correlate them (see Listing 5.21).

Listing 5.21: Specifying Parameters by Name

```
class Program
{
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
    static void Main()
    {
        DisplayGreeting(
            firstName: "Inigo", lastName: "Montoya");
    }

    public static void DisplayGreeting(
        string firstName,
        string middleName = default(string),
        string lastName = default(string))
    {

        // ...

    }
}
```

In Listing 5.21, the call to `DisplayGreeting()` from within `Main()` assigns a value to a parameter by name. Of the two optional parameters (`middleName` and `lastName`), only `lastName` is given as an argument. For cases where a method has lots of parameters and many of them are optional (a common occurrence when accessing Microsoft COM libraries), using the named argument syntax is certainly a convenience. However, along with the convenience comes an impact on the flexibility of the method interface. In the past, parameter names could be changed without causing C# code that invokes the method to no longer compile. With the addition of named parameters, the parameter name becomes part of the interface because changing the name would cause code that uses the named parameter to no longer compile.

4.0

Guidelines

DO treat parameter names as part of the API, and avoid changing the names if version compatibility between APIs is important.

For many experienced C# developers, this is a surprising restriction. However, the restriction has been imposed as part of the Common Language Specification ever since .NET 1.0. Moreover, Visual Basic has always supported calling methods with named arguments. Therefore, library developers should already be following the practice of not changing parameter names to successfully interoperate with other .NET languages

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

from version to version. In essence, C# 4.0 now imposes the same restriction on changing parameter names that many other .NET languages already require.

Given the combination of method overloading, optional parameters, and named parameters, resolving which method to call becomes less obvious. A call is **applicable** (compatible) with a method if all parameters have exactly one corresponding argument (either by name or by position) that is type compatible, unless the parameter is optional (or is a parameter array). Although this restricts the possible number of methods that will be called, it doesn't identify a unique method. To further distinguish which specific method will be called, the compiler uses only explicitly identified parameters in the caller, ignoring all optional parameters that were not specified at the caller. Therefore, if two methods are applicable because one of them has an optional parameter, the compiler will resolve to the method without the optional parameter.

End 4.0

ADVANCED TOPIC

Method Resolution

When the compiler must choose which of several applicable methods is the best one for a particular call, the one with the *most specific* parameter types is chosen. Assuming there are two applicable methods, each requiring an implicit conversion from an argument to a parameter type, the method whose parameter type is the more derived type will be used.

For example, a method that takes a `double` parameter will be chosen over a method that takes an `object` parameter if the caller passes an argument of type `int`. This is because `double` is more specific than `object`. There are objects that are not doubles, but there are no doubles that are not objects, so `double` must be more specific.

If more than one method is applicable and no unique best method can be determined, the compiler will issue an error indicating that the call is ambiguous.

For example, given the following methods:

```
static void Method(object thing){}
static void Method(double thing){}
static void Method(long thing){}
static void Method(int thing){}
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

a call of the form `Method(42)` will resolve as `Method(int thing)` because that is an exact match from the argument type to the parameter type. Were that method to be removed, overload resolution would choose the `long` version, because `long` is more specific than either `double` or `object`.

The C# specification includes additional rules governing implicit conversion between `byte`, `ushort`, `uint`, `ulong`, and the other numeric types. In general, though, it is better to use a cast to make the intended target method more recognizable.

Basic Error Handling with Exceptions

This section examines how to handle error reporting via a mechanism known as **exception handling**.

With exception handling, a method is able to pass information about an error to a calling method without using a return value or explicitly providing any parameters to do so. Listing 5.22 contains a slight modification to Listing 1.16, the `HeyYou` program from Chapter 1. Instead of requesting the last name of the user, it prompts for the user's age.

Listing 5.22: Converting a string to an int

```
using System;

class ExceptionHandling
{
  static void Main()
  {
      string firstName;
      string ageText;
      int age;

      Console.WriteLine("Hey you!");

      Console.Write("Enter your first name: ");
      firstName = System.Console.ReadLine();

      Console.Write("Enter your age: ");
      ageText = Console.ReadLine();
      age = int.Parse(ageText);

      Console.WriteLine(
          $"Hi { firstName }!  You are { age*12 } months old.");
  }
}
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Output 5.11 shows the results of Listing 5.22.

Output 5.11

```
Hey you!
Enter your first name: Inigo
Enter your age: 42
Hi Inigo!  You are 504 months old.
```

The return value from `System.Console.ReadLine()` is stored in a variable called `ageText` and is then passed to a method with the `int` data type, called `Parse()`. This method is responsible for taking a string value that represents a number and converting it to an `int` type.

BEGINNER TOPIC

42 as a String versus 42 as an Integer

C# requires that every non-null value have a well-defined type associated with it. Therefore, not only the data value but also the type associated with the data is important. A string value of 42, therefore, is distinctly different from an integer value of 42. The string is composed of the two characters 4 and 2, whereas the `int` is the number 42.

Given the converted string, the final `System.Console.WriteLine()` statement will print the age in months by multiplying the age value by 12.

But what happens if the user does not enter a valid integer string? For example, what happens if the user enters "forty-two"? The `Parse()` method cannot handle such a conversion. It expects the user to enter a string that contains only digits. If the `Parse()` method is sent an invalid value, it needs some way to report this fact back to the caller.

Trapping Errors

To indicate to the calling method that the parameter is invalid, `int.Parse()` will **throw an exception**. Throwing an exception halts further execution in the current control flow and jumps into the first code block within the call stack that handles the exception.

Since you have not yet provided any such handling, the program reports the exception to the user as an **unhandled exception**. Assuming there is no registered debugger on the system, the error will appear on the console with a message such as that shown in Output 5.12.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Output 5.12

```
Hey you!
Enter your first name: Inigo
Enter your age: forty-two

Unhandled Exception: System.FormatException: Input string was
        not in a correct format.
    at System.Number.ParseInt32(String s, NumberStyles style,
        NumberFormatInfo info)
    at ExceptionHandling.Main()
```

Obviously, such an error is not particularly helpful. To fix this, it is necessary to provide a mechanism that handles the error, perhaps reporting a more meaningful error message back to the user.

This process is known as **catching an exception**. The syntax is demonstrated in Listing 5.23, and the output appears in Output 5.13.

Listing 5.23: Catching an Exception

```
using System;

class ExceptionHandling
{
  static int Main()
  {
      string firstName;
      string ageText;
      int age;
      int result = 0;

      Console.Write("Enter your first name: ");
      firstName = Console.ReadLine();

      Console.Write("Enter your age: ");
      ageText = Console.ReadLine();

      try
      {
          age = int.Parse(ageText);
          Console.WriteLine(
              $"Hi { firstName }! You are { age*12 } months old.");
      }
      catch (FormatException )
      {
          Console.WriteLine(
              $"The age entered, { ageText }, is not valid.");
          result = 1;
      }
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
        catch(Exception exception)
        {
            Console.WriteLine(
                $"Unexpected error:  { exception.Message }");
            result = 1;
        }
        finally
        {
            Console.WriteLine($"Goodbye { firstName }");
        }

        return result;
    }
}
```

Output 5.13

```
Enter your first name: Inigo
Enter your age: forty-two
The age entered, forty-two, is not valid.
Goodbye Inigo
```

To begin, surround the code that could potentially throw an exception (`age = int.Parse()`) with a **try block**. This block begins with the `try` keyword. It indicates to the compiler that the developer is aware of the possibility that the code within the block might throw an exception, and if it does, one of the **catch blocks** will attempt to handle the exception.

One or more catch blocks (or the finally block) must appear immediately following a try block. The catch block header (see the Advanced Topic titled "General Catch" later in this chapter) optionally allows you to specify the data type of the exception, and as long as the data type matches the exception type, the catch block will execute. If, however, there is no appropriate catch block, the exception will fall through and go unhandled as though there were no exception handling. The resultant control flow appears in Figure 5.1.

For example, assume the user enters "forty-two" for the age in the previous example. In this case, `int.Parse()` will throw an exception of type `System.FormatException`, and control will jump to the set of catch blocks. (`System.FormatException` indicates that the string was not of the correct format to be parsed appropriately.) Since the first catch block matches the type of exception that `int.Parse()` threw, the code inside this block will execute. If a statement within the try block threw a different exception,

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

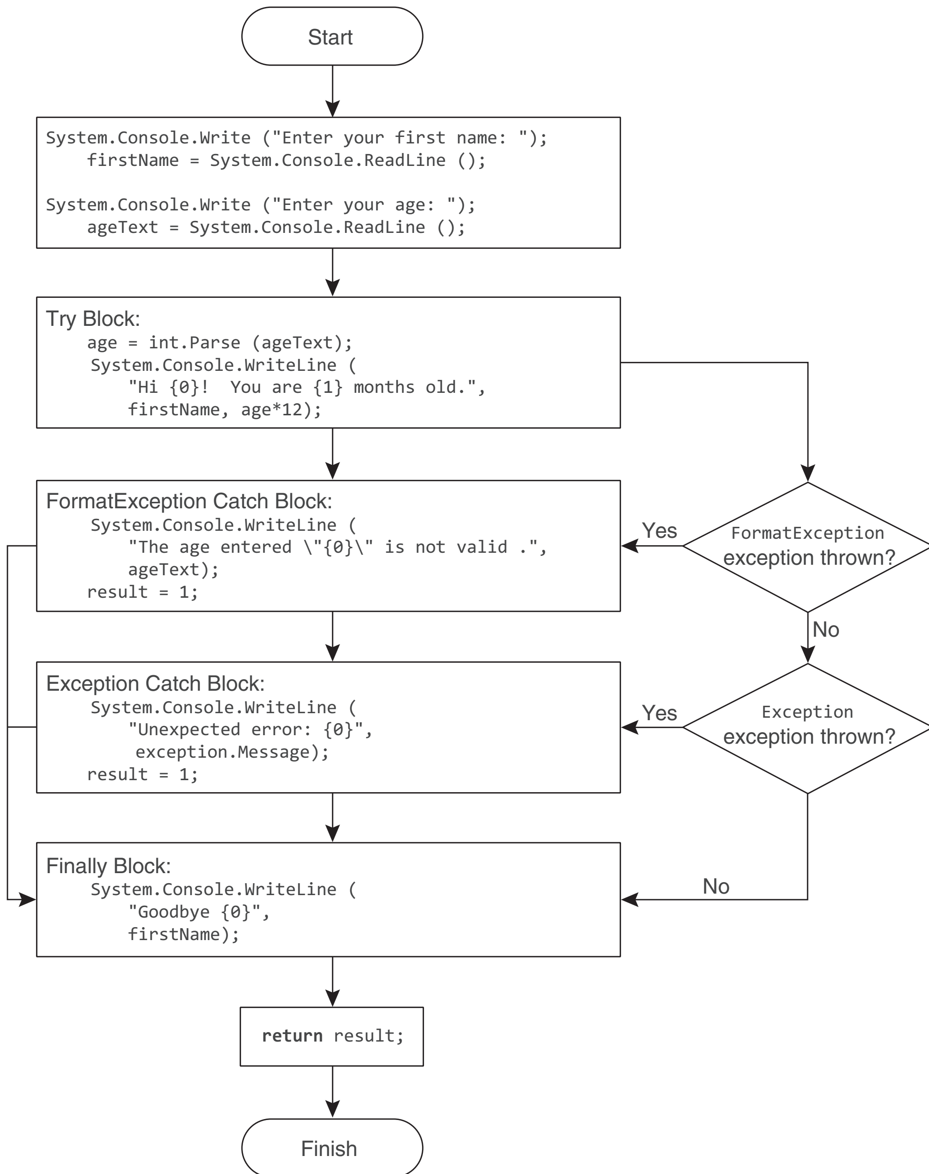

Figure 5.1: Exception-Handling Control Flow

the second catch block would execute because all exceptions are of type `System.Exception`.

If there were no `System.FormatException` catch block, the `System.Exception` catch block would execute even though `int.Parse` throws a

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

System.FormatException. This is because a System.FormatException is also of type System.Exception. (System.FormatException is a more specific implementation of the generic exception, System.Exception.)

The order in which you handle exceptions is significant. Catch blocks must appear from most specific to least specific. The System.Exception data type is least specific, so it appears last. System.FormatException appears first because it is the most specific exception that Listing 5.23 handles.

Regardless of whether control leaves the try block normally or because the code in the try block throws an exception, the **finally block** of code will execute after control leaves the try-protected region. The purpose of the finally block is to provide a location to place code that will execute regardless of how the try/catch blocks exit—with or without an exception. Finally blocks are useful for cleaning up resources regardless of whether an exception is thrown. In fact, it is possible to have a try block with a finally block and no catch block. The finally block executes regardless of whether the try block throws an exception or whether a catch block is even written to handle the exception. Listing 5.24 demonstrates the try/finally block, and Output 5.14 shows the results.

Listing 5.24: Finally Block without a Catch Block

```
using System;

class ExceptionHandling
{
  static int Main()
  {
      string firstName;
      string ageText;
      int age;
      int result = 0;

      Console.Write("Enter your first name: ");
      firstName = Console.ReadLine();

      Console.Write("Enter your age: ");
      ageText = Console.ReadLine();

      try
      {
          age = int.Parse(ageText);
          Console.WriteLine(
              $"Hi { firstName }! You are { age*12 } months old.");
      }
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
        finally
        {
            Console.WriteLine($"Goodbye { firstName }");
        }

        return result;
    }
}
```

Output 5.14

```
Enter your first name: Inigo
Enter your age: forty-two

Unhandled Exception: System.FormatException: Input string was not in a
correct format.
   at System.Number.StringToNumber(String str, NumberStyles options,
NumberBuffer& number, NumberFormatInfo info, Boolean parseDecimal)
   at System.Number.ParseInt32(String s, NumberStyles style,
NumberFormatInfo info)
   at ExceptionHandling.Main()
Goodbye Inigo
```

The attentive reader will have noticed something interesting here: The runtime first reported the unhandled exception and then ran the finally block. What explains this unusual behavior?

First, the behavior is legal because when an exception is unhandled, the behavior of the runtime is implementation defined; any behavior is legal! The runtime chooses this particular behavior because it knows before it chooses to run the finally block that the exception will be unhandled; the runtime has already examined all of the activation frames on the call stack and determined that none of them is associated with a catch block that matches the thrown exception.

As soon as the runtime determines that the exception will be unhandled, it checks whether a debugger is installed on the machine, because you might be the software developer who is analyzing this failure. If a debugger is present, it offers the user the chance to attach the debugger to the process *before* the finally block runs. If there is no debugger installed or if the user declines to debug the problem, the default behavior is to print the unhandled exception to the console and then see if there are any finally blocks that could run. Due to the “implementation-defined” nature of the situation, the runtime is not required to run finally blocks in this situation; an implementation may choose to do so or not.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Guidelines

AVOID explicitly throwing exceptions from finally blocks. (Implicitly thrown exceptions resulting from method calls are acceptable.)

DO favor try/finally and avoid using try/catch for cleanup code.

DO throw exceptions that describe which exceptional circumstance occurred, and if possible, how to prevent it.

ADVANCED TOPIC

Exception Class Inheritance

Starting in C# 2.0, all objects thrown as exceptions derive from `System.Exception`. (Objects thrown from other languages that do not derive from `System.Exception` are automatically "wrapped" by an object that does.) Therefore, they can be handled by the `catch(System.Exception exception)` block. It is preferable, however, to include a catch block that is specific to the most derived type (e.g., `System.FormatException`), because then it is possible to get the most information about an exception and handle it less generically. In so doing, the `catch` statement that uses the most derived type is able to handle the exception type specifically, accessing data related to the exception thrown and avoiding conditional logic to determine what type of exception occurred.

This is why C# enforces the rule that catch blocks appear from most derived to least derived. For example, a `catch` statement that catches `System.Exception` cannot appear before a statement that catches `System.FormatException` because `System.FormatException` derives from `System.Exception`.

A method could throw many exception types. Table 5.2 lists some of the more common ones within the framework.

TABLE 5.2: Common Exception Types

Exception Type	Description
`System.Exception`	The "base" exception from which all other exceptions derive.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Table 5.2: Common Exception Types (*continued*)

Exception Type	Description
`System.ArgumentException`	Indicates that one of the arguments passed into the method is invalid.
`System.ArgumentNullException`	Indicates that a particular argument is `null` and that this is not a valid value for that parameter.
`System.ApplicationException`	To be avoided. The original idea was that you might want to have one kind of handling for system exceptions and another for application exceptions, which, although plausible, doesn't actually work well in the real world.
`System.FormatException`	Indicates that the string format is not valid for conversion.
`System.IndexOutOfRangeException`	Indicates that an attempt was made to access an array or other collection element that does not exist.
`System.InvalidCastException`	Indicates that an attempt to convert from one data type to another was not a valid conversion.
`System.InvalidOperationException`	Indicates that an unexpected scenario has occurred such that the application is no longer in a valid state of operation.
`System.NotImplementedException`	Indicates that although the method signature exists, it has not been fully implemented.
`System.NullReferenceException`	Thrown when code tries to find the object referred to by a reference that is null.
`System.ArithmeticException`	Indicates an invalid math operation, not including divide by zero.
`System.ArrayTypeMismatchException`	Occurs when attempting to store an element of the wrong type into an array.
`System.StackOverflowException`	Indicates an unexpectedly deep recursion.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

ADVANCED TOPIC

General Catch

It is possible to specify a catch block that takes no parameters, as shown in Listing 5.25.

LISTING 5.25: General Catch Blocks

```
...
try
{
    age = int.Parse(ageText);
    System.Console.WriteLine(
        $"Hi { firstName }!  You are { age*12 } months old.");
}
catch (System.FormatException exception)
{
    System.Console.WriteLine(
        $"The age entered ,{ ageText }, is not valid.");
    result = 1;
}
catch(System.Exception exception)
{
    System.Console.WriteLine(
        $"Unexpected error:  { exception.Message }");
    result = 1;
}
catch
{
    System.Console.WriteLine("Unexpected error!");
    result = 1;
}
finally
{
    System.Console.WriteLine($"Goodbye { firstName }");
}
...
```

A catch block with no data type, called a **general catch block**, is equivalent to specifying a catch block that takes an `object` data type—for instance, `catch(object exception){...}`. Because all classes ultimately derive from `object`, a catch block with no data type must appear last.

General catch blocks are rarely used because there is no way to capture any information about the exception. In addition, C# doesn't support the ability to throw an exception of type `object`. (Only libraries written in languages such as C++ allow exceptions of any type.)

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Begin 2.0

The behavior starting in C# 2.0 varies slightly from the earlier C# behavior. In C# 2.0, if a language allows throwing non-System.Exceptions, the object of the thrown exception will be wrapped in a System.Runtime.CompilerServices.RuntimeWrappedException that does derive from System.Exception. Therefore, all exceptions, whether derived from System.Exception or not, will propagate into C# assemblies as if they were derived from System.Exception.

The result is that System.Exception catch blocks will catch all exceptions not caught by earlier blocks, and a general catch block, following a System.Exception catch block, will never be invoked. Consequently, following a System.Exception catch block with a general catch block in C# 2.0 or later will result in a compiler warning indicating that the general catch block will never execute.

End 2.0

Guidelines

AVOID general catch blocks and replace them with a catch of System.Exception.

AVOID catching exceptions for which the appropriate action is unknown. It is better to let an exception go unhandled than to handle it incorrectly.

AVOID catching and logging an exception before rethrowing it. Instead, allow the exception to escape until it can be handled appropriately.

Reporting Errors Using a throw Statement

C# allows developers to throw exceptions from their code, as demonstrated in Listing 5.26 and Output 5.15.

Listing 5.26: Throwing an Exception

```
using System;
public class ThrowingExceptions
{
  public static void Main()
  {
      try
      {
          Console.WriteLine("Begin executing");
```

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

```
            Console.WriteLine("Throw exception");
               throw new Exception("Arbitrary exception");
            Console.WriteLine("End executing");
        }
        catch(FormatException exception)
        {
            Console.WriteLine(
                "A FormateException was thrown");
        }
        catch(Exception exception)
        {
            Console.WriteLine(
                $"Unexpected error: { exception.Message }");
        }
        catch
        {
            Console.WriteLine("Unexpected error!");
        }

        Console.WriteLine(
            "Shutting down...");
    }
}
```

Output 5.15

```
Begin executing
Throw exception...
Unexpected error:  Arbitrary exception
Shutting down...
```

As the arrows in Listing 5.26 depict, throwing an exception causes execution to jump from where the exception is thrown into the first catch block within the stack that is compatible with the thrown exception type.[6] In this case, the second catch block handles the exception and writes out an error message. In Listing 5.26, there is no finally block, so execution falls through to the `System.Console.WriteLine()` statement following the try/catch block.

To throw an exception, it is necessary to have an instance of an exception. Listing 5.26 creates an instance using the keyword `new` followed by the type of the exception. Most exception types allow a message to be generated as part of throwing the exception, so that when the exception occurs, the message can be retrieved.

6. Technically it could be caught by a compatible catch filter as well.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Sometimes a catch block will trap an exception but be unable to handle it appropriately or fully. In these circumstances, a catch block can rethrow the exception using the `throw` statement without specifying any exception, as shown in Listing 5.27.

Listing 5.27: Rethrowing an Exception

```
...
      catch(Exception exception)
      {
          Console.WriteLine(
              $@"Rethrowing unexpected error:  {
                  exception.Message }");
          throw;
      }
...
```

In Listing 5.27, the `throw` statement is "empty" rather than specifying that the exception referred to by the `exception` variable is to be thrown. This illustrates a subtle difference: `throw;` preserves the *call stack* information in the exception, whereas `throw exception;` replaces that information with the current call stack information. For debugging purposes, it is usually better to know the original call stack.

Guidelines

DO prefer using an empty throw when catching and rethrowing an exception so as to preserve the call stack.

DO report execution failures by throwing exceptions rather than returning error codes.

DO NOT have public members that return exceptions as return values or an `out` parameter. Throw exceptions to indicate errors; do not use them as return values to indicate errors.

Avoid Using Exception Handling to Deal with Expected Situations

Developers should make an effort to avoid throwing exceptions for expected conditions or normal control flow. For example, developers should not expect users to enter valid text when specifying their age.[7]

7. In general, developers should expect their users to perform unexpected actions; in turn, they should code defensively to handle "stupid user tricks."

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Therefore, instead of relying on an exception to validate data entered by the user, developers should provide a means of checking the data before attempting the conversion. (Better yet, they should prevent the user from entering invalid data in the first place.) Exceptions are designed specifically for tracking exceptional, unexpected, and potentially fatal situations. Using them for an unintended purpose such as expected situations will cause your code to be hard to read, understand, and maintain.

Additionally, like most languages, C# incurs a slight performance hit when throwing an exception—taking microseconds compared to the nanoseconds most operations take. This delay is generally not noticeable in human time—except when the exception goes unhandled. For example, when Listing 5.22 is executed and the user enters an invalid age, the exception is unhandled and there is a noticeable delay while the runtime searches the environment to see whether there is a debugger to load. Fortunately, slow performance when a program is shutting down isn't generally a factor to be concerned with.

Guidelines

DO NOT use exceptions for handling normal, expected conditions; use them for exceptional, unexpected conditions.

Begin 2.0

ADVANCED TOPIC

Numeric Conversion with `TryParse()`

One of the problems with the `Parse()` method is that the only way to determine whether the conversion will be successful is to attempt the cast and then catch the exception if it doesn't work. Because throwing an exception is a relatively expensive operation, it is better to attempt the conversion without exception handling. In the first release of C#, the only data type that enabled this behavior was a `double` method called `double.TryParse()`. However, this method is included with all numeric primitive types starting with the Microsoft .NET Framework 2.0. It requires the use of the `out` keyword because the return from the `TryParse()` function is a `bool` rather than the converted value. Listing 5.28 is a code snippet that demonstrates the conversion using `int.TryParse()`.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.

Listing 5.28: Conversion Using int.TryParse()

```
if (int.TryParse(ageText, out int age))
{
    Console.WriteLine(
        $"Hi { firstName }!  "
        + $"You are { age*12 } months old.");
}
else
{
    Console.WriteLine(
        $"The age entered, { ageText }, is not valid.");
}
```

With the Microsoft .NET Framework 4, a `TryParse()` method was also added to enum types.

With the `TryParse()` method, it is no longer necessary to include a try/catch block simply for the purpose of handling the string-to-numeric conversion.

End 2.0

SUMMARY

This chapter discussed the details of declaring and calling methods, including the use of the keywords `out` and `ref` to pass and return variables rather than their values. In addition to method declaration, this chapter introduced exception handling.

A method is a fundamental construct that is a key to writing readable code. Instead of writing large methods with lots of statements, you should use methods to create "paragraphs" of roughly 10 or fewer statements within your code. The process of breaking large functions into smaller pieces is one of the ways you can refactor your code to make it more readable and maintainable.

The next chapter considers the class construct and describes how it encapsulates methods (behavior) and fields (data) into a single unit.

From Essential C# 7.0 by Mark Michaelis. Copyright 2018 Pearson Education, Inc. ISBN: 978-15-0930358-8.