


Programming in C#

Online module to accompany *Invitation to Computer Science, 7th Edition*, ISBN-10: 1305075773; ISBN-13: 9781305075771 (Cengage Learning, 2016).

1. Introduction to C#
 - 1.1 A Simple C# Program
 - 1.2 Creating and Running a C# Program
2. Virtual Data Storage
3. Statement Types
 - 3.1 Input/Output Statements
 - 3.2 The Assignment Statement
 - 3.3 Control Statements
4. Another Example
5. Managing Complexity
 - 5.1 Divide and Conquer
 - 5.2 Using Functions
 - 5.3 Writing Functions
6. Object-Oriented Programming
 - 6.1 What Is It?
 - 6.2 C# and OOP
 - 6.3 One More Example
 - 6.4 What Have We Gained?
7. Graphical Programming
 - 7.1 Graphics Hardware
 - 7.2 Graphics Software
8. Conclusion

EXERCISES

ANSWERS TO PRACTICE PROBLEMS



1 Introduction to C#

Hundreds of high-level programming languages have been developed; a fraction of these have become viable, commercially successful languages. There are a half-dozen or so languages that can illustrate some of the concepts of a high-level programming language, but this module uses C# for this purpose. The C# language was developed by Microsoft along with the .NET environment. C# is the programming language, and .NET is the runtime support environment.

Our intent here is not to make you an expert programmer—any more than our purpose in Chapter 4 was to make you an expert circuit designer. Indeed, there is much about the language that we will not even discuss. You will, however, get a sense of what programming in a high-level language is like, and perhaps you will see why some people think it is one of the most fascinating of human endeavors.

▶ 1.1 A Simple C# Program

Figure 1 shows a simple but complete C# program. Even if you know nothing about the C# language, it is not hard to get the general drift of what the program is doing.

FIGURE 1


A Simple C# Program

```
//Computes and outputs travel time
//for a given speed and distance
//Written by J. Q. Programmer, 6/15/16

using System;

namespace InvitationCSharp
{
    class TravelPlanner
    {
        static void Main(string[] args)
        {
            int speed;           //rate of travel
            double distance;    //miles to travel
            double time;        //time needed for this travel

            Console.WriteLine("Enter your speed in mph: ");
            speed = Convert.ToInt32(Console.ReadLine());
```

**FIGURE 1**

A Simple C# Program
(continued)

```

Console.WriteLine("Enter your distance in miles: ");
distance = Convert.ToDouble(Console.ReadLine());

time = distance / speed;

Console.WriteLine("At " + speed + " mph, ");
Console.WriteLine("it will take ");
Console.WriteLine(time + " hours to travel ");
Console.WriteLine(distance + " miles.");
    }
}
}

```

Someone running this program (the “user”) could have the following dialogue with the program, where boldface indicates what the user types:

```

Enter your speed in mph: 58
Enter your distance in miles: 657.5
At 58 mph, it will take
11.3362068965517 hours to travel 657.5 miles.
Press any key to continue . . .

```

The general form of a typical C# program is shown in Figure 2. To compare our simple example program with this form, we have reproduced the example program in Figure 3 with a number in front of each line. The numbers are there for reference purposes only; they are *not* part of the program.

Lines 1–3 in the program of Figure 3 are C# **comments**. Anything appearing on a line after the double slash symbol (//) is ignored by the compiler, just as anything following the double dash (--) is treated as a comment in the assembly language programs of Chapter 6. Although the computer ignores comments, they are important to include in a program because they

**FIGURE 2**

The Overall Form of a Typical
C# Program

```

prologue comment      [optional]
using directive
namespace
{
    class
    {
        functions      [optional]
        Main function
        {
            declarations [optional]
            Main function body
        }
    }
}
}

```

**FIGURE 3**

The Program of Figure 1
(line numbers added for
reference)

```

1. //Computes and outputs travel time
2. //for a given speed and distance
3. //Written by J. Q. Programmer, 6/15/16
4.
5. using System;
6.
7. namespace InvitationCSharp
8. {
9.     class TravelPlanner
10.    {
11.        static void Main(string[] args)
12.        {
13.            int speed;           //rate of travel
14.            double distance;    //miles to travel
15.            double time;       //time needed for this travel
16.
17.            Console.Write("Enter your speed in mph: ");
18.            speed = Convert.ToInt32(Console.ReadLine());
19.            Console.Write("Enter your distance in miles: ");
20.            distance = Convert.ToDouble(Console.ReadLine());
21.
22.            time = distance / speed;
23.
24.            Console.Write("At " + speed + " mph, ");
25.            Console.WriteLine("it will take ");
26.            Console.Write(time + " hours to travel ");
27.            Console.WriteLine(distance + " miles.");
28.        }
29.    }
30. }

```

give information to the human readers of the code. Every high-level language has some facility for including comments, because understanding code that someone else has written (or understanding your own code after a period of time has passed) is very difficult without the notes and explanations that comments provide. Comments are one way to *document* a computer program to make it more understandable. The comments in the program of Figure 3 describe what the program does plus tell who wrote the program and when. These three comment lines together make up the program's **prologue comment** (the introductory comment that comes first). According to the general form of Figure 2, the prologue comment is optional, but providing it is always a good idea. It's almost like the headline in a newspaper, giving the big picture up front.

Blank lines in C# programs are ignored and are used, like comments, to make the program more readable by human beings. In our example program, we've used blank lines (lines 4, 6, 16, 21, 23) to separate sections of the program, visually indicating groups of statements that are related.

Line 5 is a **using directive** to the compiler that refers to the *System* library. The eventual effect is that the linker includes object code from this library. The core C# language does not provide a way to get data into a

What Is C#?

In June 2000, Microsoft introduced the C# programming language (pronounced “C-sharp”). It was designed to improve upon the C++ language developed in the early 1980s at AT&T Labs. Because it is a totally new language, it has no backward compatibility issues, as C++ had with its predecessor, the C language. C# also shares many features with Java, a programming language released in 1996 by Sun Microsystems.

It is impossible to discuss C# without discussing the **Microsoft .NET Framework** that supports C# and many other programming languages. The .NET Framework is essentially

a giant collection of tools for software development. It was designed to allow traditional text-based applications, GUI applications, and Web-based programs to all be built with equal ease. Currently, over 90 languages are supported by the .NET Framework. That means applications written in any of these languages have access to the tools provided within the .NET Framework. Applications can even be written that mix and match modules in various languages.

In April 2003, only three years after the first release of C# and .NET, C# and the CLI (Common Language Infrastructure—a significant subset of the .NET tools) were adopted as ISO (International Organization for Standardization) standards. C# continues to grow in popularity.

program or for a program to display results. The *System* library (namespace) contains code for these purposes. Line 5 also tells the compiler to look in the *System* namespace for the definition of any names not specifically defined within the program. In this program, *ReadLine*, *Write*, and *Write-Line* (used to read input data and write output, respectively) get their meaning from the *System* namespace. In addition to *System*, C# has many other code libraries (namespaces), such as mathematical and graphics libraries, and therefore many other using directives are possible. Using directives are optional, but it would be a trivial program indeed that did not need input data or produce output results, so virtually every C# program has at least the using directive shown in our example.

Line 7 creates a namespace for the code. In C# (and .NET), namespaces help to organize libraries and hierarchies to prevent ambiguity when referencing objects. So, all the code for this program will be in the namespace *InvitationCSharp*. The name of this namespace, “InvitationCSharp”, was created by the author of the program. Microsoft chose the name “System” for the namespace mentioned earlier.

Line 9 is a **class header**, which announces that a **class** is about to be defined. The class is named *TravelPlanner*, and the curly braces at lines 10 and 29 mark the beginning and end of this class definition. All C# code (except for comments, using directives, and namespace statements) must be either a class header or inside a class definition. We will have much more to say about classes later. For now, just think of a class as a collection of sections of code called **functions** that are able to perform various related services. In the *TravelPlanner* class, there is only one function, the *Main* function. The service it performs is to compute and write out the time to travel a given distance at a given speed. Line 11:

```
static void Main(string[] args)
```

is the header for the *Main* function. It is not necessary to understand this somewhat obscure code; just remember that every C# program must have a

Main function, and that all Main functions start out exactly this way. The curly braces at lines 12 and 28 enclose the Main **function body**, which is the heart of the sample program. Lines 13–15 are declarations that name and describe the items of data that are used within the Main function. Descriptive names—*speed*, *distance*, and *time*—are used for these quantities to help document their purpose in the program, and comments provide further clarification. Line 13 describes an integer quantity (type *int*) called *speed*. Lines 14 and 15 declare *distance* and *time* as real number quantities (type *double*). A real number quantity is one containing a decimal point, such as 28.3, 102.0, or -17.5. Declarations are also optional in the sense that if a program does not use any data, no declarations are needed, but again, it would be unusual to find such a trivial program.

Statements with *Write* or *WriteLine* write messages to the user; statements with *ReadLine* get the strings the user entered for speed and distance. The prefix *Console* indicates that the user will enter input at the computer keyboard and will view the output on the computer's screen. The difference between *Write* and *WriteLine* is that *WriteLine* forces any subsequent output to begin on a new line, whereas after a *Write* statement, more can be written later on the same line. The *Convert* functions change the input strings, which are strings of characters, into numerical integer or double values, respectively, and store them in *speed* and *distance*. Line 22 computes the time required to travel this distance at this speed. Finally, lines 24–27 print the output to the user's screen. The values of *speed*, *time*, and *distance* are inserted in appropriate places among the strings of text shown in double quotes. When the *Write* and *WriteLine* functions encounter a numeric quantity, e.g., *speed*, they convert the numeric value to a string. That string is then concatenated to (joined to) the previous string that was output. The + sign is the C# **concatenation operator**—it joins two strings together.

You may have noticed that most of the statements in this program end with a semicolon. A semicolon must appear at the end of every executable C# instruction, which means everywhere except at the end of a comment, a namespace declaration, a class header, or a function header such as

```
static void Main (. . .)
```

The semicolon requirement is a bit of a pain in the neck, but the C# compiler generates one or more error messages if you omit the semicolon, so after the first few hundred times this happens, you tend to remember to put it in.

C#, along with every other programming language, has very specific rules of **syntax**—the correct form for each component of the language. Having a semicolon at the end of every executable statement is a C# syntax rule. Any violation of the syntax rules generates an error message from the compiler because the compiler does not recognize or know how to translate the offending code. In the case of a missing semicolon, the compiler cannot tell where the instruction ends. The syntax rules for a programming language are often defined by a formal grammar, much as correct English syntax is defined by rules of grammar.

C# is a **free-format language**, which means that it does not matter where things are placed on a line. For example, we could have written

```
time =
distance /
speed;
```


although this is clearly harder to read. The free-format characteristic explains why a semicolon is needed to mark the end of an instruction, which might be spread over several lines.

1.2 Creating and Running a C# Program

Creating and running a C# program is basically a three-step process. The first step is to type the program into a text editor. When you are finished, you save the file, giving it a name with the extension `.cs`. So the file for Figure 1 could be named

```
TravelPlanner.cs
```

As the second step, the program must be compiled using a C# compiler for your computer, and the resulting object code linked with any C# library object code. In our example, the program in the file *TravelPlanner.cs* would be compiled, resulting in a file called

```
TravelPlanner.exe
```

The third step loads and executes the program file, in this case *TravelPlanner.exe*. Depending on your system, you may have to type operating system commands for the last two steps.

Another approach is to do all of your work in an **Integrated Development Environment**, or **IDE**. The IDE lets the programmer perform a number of tasks within the shell of a single application program, rather than having to use a separate program for each task. A modern programming IDE provides a text editor, a file manager, a compiler, a linker and loader, and tools for debugging, all within this one piece of software. The IDE usually has a GUI (graphical user interface) with menu choices for the different tasks. This can significantly speed up program development.

C# COMPILERS

The C# examples in this module were written and executed in Microsoft Visual C# 2013, part of Microsoft Visual Studio 2013. This is an IDE with a GUI interface that supports many programming languages. Visual Studio Express 2013 for Windows Desktop, which supports Visual C# as well as other languages, is freely downloadable from Microsoft at

www.microsoft.com/express/product/default.aspx

and runs on Windows 7 or Windows 8 operating systems. Its use requires the Microsoft .NET Framework. If the .NET Framework is not already on your Windows system, you will be alerted at installation, and you can go to

www.microsoft.com/downloads

and look under Developer Tools to download it.

This C# exercise is just a beginning. In the rest of this module, we'll examine the features of the language that will enable you to write your own C# programs to carry out more sophisticated tasks.

2 Virtual Data Storage

One of the improvements we seek in a high-level language is freedom from having to manage data movement within memory. Assembly language does not require us to give the actual memory address of the storage location to be used for each item, as in machine language. However, we still have to move values, one by one, back and forth between memory and the arithmetic logic unit (ALU) as simple modifications are made, such as setting the value of A to the sum of the values of B and C. We want the computer to let us use data values by name in any appropriate computation without thinking about where they are stored or what is currently in some register in the ALU. In fact, we do not even want to know that there *is* such a thing as an ALU, where data are moved to be operated on; instead, we want the virtual machine to manage the details when we request that a computation be performed. A high-level language allows this, and it also allows the names for data items to be more meaningful than in assembly language.

Names in a programming language are called **identifiers**. Each language has its own specific rules for what a legal identifier can look like. In C# an identifier can be any combination of letters, digits, the underscore symbol (`_`), and the “at” sign (`@`), as long as it does not begin with a digit. However, identifiers beginning with underscore characters should be avoided; they are generally used for special purposes. An additional restriction is that an identifier cannot be one of the few **reserved words**, such as *void*, *int*, *double*, and so forth, that have a special meaning in C# and that you would not be likely to use anyway. The three integers *B*, *C*, and *A* in our assembly language

CAPITALIZATION OF IDENTIFIERS

There are two standard capitalization patterns for identifiers, particularly “multiple word” identifiers:

camel case: First word begins with a lowercase letter, additional words begin with uppercase letters (*finalTotal*)

Pascal case: All words begin with an uppercase letter (*FinalTotal*)

The code in this chapter uses the following convention for creating identifiers (examples included):

Simple variables – camel case: *speed*, *time*, *finalTotal*

Function names – camel case: *myFunction*, *getInput*

Class names – Pascal case: *MyClass*

Object names – camel case: *myObject*

The underscore character is not used. Occasionally, however, we'll use single capital letters for identifiers in quick code fragments.

program can therefore have more descriptive names, such as *subTotal*, *tax*, and *finalTotal*. The use of descriptive identifiers is one of the greatest aids to human understanding of a program. Identifiers can be almost arbitrarily long, so be sure to use a meaningful identifier such as *finalTotal* instead of something like *A*; the improved readability is well worth the extra typing time. C# is a **case-sensitive** language, which means that uppercase letters are distinguished from lowercase letters. Thus, *FinalTotal*, *Finaltotal*, and *finalTotal* are three different identifiers.

Data that a program uses can come in two varieties. Some quantities are fixed throughout the duration of the program, and their values are known ahead of time. These quantities are called **constants**. An example of a constant is the integer value 2. Another is an approximation to π , say 3.1416. The integer 2 is a constant that we don't have to name by an identifier, nor do we have to build the value 2 in memory manually by the equivalent of a `.DATA` pseudo-op. We can just use the symbol "2" in any program statement. When "2" is first encountered in a program statement, the binary representation of the integer 2 is automatically generated and stored in a memory location. Likewise, we can use "3.1416" for the real number value 3.1416, but if we are really using this number as an approximation to π , it is more informative to use the identifier *pi*.

Some quantities used in a program have values that change as the program executes, or values that are not known ahead of time but must be obtained from the computer user (or from a data file previously prepared by the user) as the program runs. These quantities are called **variables**. For example, in a program doing computations with circles (where we might use the constant *pi*), we might need to obtain from the user or a data file the radius of the circle. This variable can be given the identifier *radius*.

Identifiers for variables serve the same purpose in program statements as pronouns do in ordinary English statements. The English statement "He will be home today" has specific meaning only when we plug in the value for which "He" stands. Similarly, a program statement such as

```
time = distance/speed;
```

becomes an actual computation only when numeric values have been stored in the memory locations referenced by the *distance* and *speed* identifiers.

We know that all data are represented internally in binary form. In Chapter 4 we noted that any one sequence of binary digits can be interpreted as a whole number, a negative number, a real number (one containing a decimal point, such as -17.5 or 28.342), or as a letter of the alphabet. C# requires the following information about each variable in the program:

- What identifier we want to use for it (its name)
- What **data type** it represents (e.g., an integer or a letter of the alphabet)

The data type determines how many bytes will be needed to store the variable—that is, how many memory cells are to be considered as one **memory location** referenced by one identifier—and also how the string of bits in that memory location is to be interpreted. C# provides several

“primitive” data types that represent a single unit of information, as shown in Figure 4. The way to give the necessary information within a C# program is to declare each variable. A **variable declaration** consists of a data type followed by a list of one or more identifiers of that type. Our sample program used three declaration statements:

```
int speed;           //rate of travel
double distance;    //miles to travel
double time;        //time needed for this travel
```

but these could have been combined into two:

```
int speed;           //rate of travel
double distance, time; //miles to travel and time
                        //needed for this travel
```

Where do the variable declarations go? Although the only requirement is that a variable must be declared before it can be used, all variable declarations are usually collected together at the top of the Main function, as in our sample program. This gives the reader of the code quick information about the data that the program will be using.

What about the constant π ? We want to assign the fixed value 3.1416 to the π identifier. Constant declarations are just like variable declarations, with the addition of the keyword **const** and the assignment of the fixed value to the constant identifier.

```
const double pi = 3.1416;
```

Some programmers use all uppercase letters to denote constant identifiers, but the compiler identifies a constant quantity only by the presence of “const” in the declaration. Once a quantity has been declared as a constant, any attempt later in the program to change its value generates an error message from the compiler.

In addition to variables of a primitive data type that hold only one unit of information, it is possible to declare a whole collection of related variables at one time. This allows storage to be set aside as needed to contain each of the values in this collection. For example, suppose we want to record the number of “hits” on a Web site for each month of the year. The value for each month is a single integer. We want a collection of 12 such integers, ordered in a particular way. An **array** groups together a collection of


FIGURE 4

Some of the C# Primitive Data Types

int	An integer quantity
double	A real number
char	A character (a single keyboard character, such as 'a')

memory locations, all storing data of the same type. The following statement declares an array:

```
int[] hits = new int[12];
```


The left side of the equals sign says that *hits* is an array of integers; the right side of the equals sign actually generates (new) memory locations for 12 integer quantities. The 12 individual array elements are numbered from *hits[0]* to *hits[11]*. (Notice that a C# array counts from 0 up to 11, instead of from 1 up to 12.) Thus, we use *hits[0]* to refer to the first entry in *hits*, which represents the number of visits to the Web site during the first month of the year, January. Continuing this numbering scheme, *hits[2]* refers to the number of visits during March, and *hits[11]* to the number of visits during December. In this way we use one declaration to set up 12 separate (but related) integer storage locations. Figure 5 illustrates this array.

Here is an example of the power of a high-level language. In assembly language we can name only individual memory locations—that is, individual items of data—but in C# we can also assign a name to an entire collection of related data items. An array thus enables us to talk about an entire table of values, or the individual elements making up that table. If we are writing C# programs to implement the data cleanup algorithms of Chapter 3, we can use an array of integers to store the 10 data items.



FIGURE 5

A 12-Element Array *hits*



PRACTICE PROBLEMS

- Which of the following are legitimate C# identifiers?

martinBradley C3P_OH Amy3 3Right const

- Write a declaration statement for a C# program that uses one integer quantity called *number*.
- Write a C# statement that declares a type *double* constant called *taxRate* that has the value 5.5.
- Using the *hits* array of Figure 5, how do you reference the number of hits on the Web page for August?

3 Statement Types

Now that we can reserve memory for data items by simply naming what we want to store and describing its data type, we will examine additional kinds of programming instructions (statements) that C# provides. These statements enable us to manipulate the data items and do something useful with them. The instructions in C#, or indeed in any high-level language, are designed as components for algorithmic problem solving, rather than as one-to-one translations of the underlying machine language instruction set of the computer. Thus they allow the programmer to work at a higher level of abstraction. In this section we examine three types of high-level programming language statements. They are consistent with the pseudocode operations we described in Chapter 2 (see Figure 2.9).

Input/output statements make up one type of statement. An **input statement** collects a value from the user for a variable within the program. In our TravelPlanner program, we need input statements to get the specific values of the speed and distance that are to be used in the computation. An **output statement** writes a message or the value of a program variable to the user's screen. Once the TravelPlanner program computes the time required to travel the given distance at the given speed, the output statement displays that value on the screen, along with other information about what that value means.

Another type of statement is the **assignment statement**, which assigns a value to a program variable. This is similar to what an input statement does, except that the value is not collected directly from the user, but is computed by the program. In pseudocode we called this a "computation operation."

Control statements, the third type of statement, affect the order in which instructions are executed. A program executes one instruction or program statement at a time. Without directions to the contrary, instructions are executed sequentially, from first to last in the program. (In Chapter 2 we called this a straight-line algorithm.) Imagine beside each program statement a light bulb that lights up while that statement is being executed; you would see a ripple of lights from the top to the bottom of the program. Sometimes, however, we want to interrupt this sequential progression and jump around in the program (which is accomplished by the instructions JUMP, JUMPGT, and so on, in assembly language). The progression of lights, which may no longer be sequential, illustrates the **flow of control** in the program—that is, the path through the program that is traced by following the currently executing statement. Control statements direct this flow of control.

▶ 3.1 Input/Output Statements

Remember that the job of an input statement is to collect from the user specific values for variables in the program. In pseudocode, to get the value for *speed* in the TravelPlanner program, we would say something like

Get value for *speed*

C# can do this task using a function named *ReadLine*. The input statement is

```
speed = Convert.ToInt32(Console.ReadLine());
```

Because all variables must be declared before they can be used, the declaration statement that says *speed* is to be a variable (of data type *int*) precedes this input statement.

Let's say that we have written the entire *TravelPlanner* program and it is now executing. When the preceding input statement is encountered, the program stops and waits for the user to enter a value for *speed* (by typing it at the keyboard, followed by pressing the ENTER key). For example, the user could type

58 <ENTER>

The *ReadLine* function captures the string consisting of a 5 followed by an 8; this is just a two-character string, similar to the string "ab" consisting of an *a* followed by a *b*. In other words, the two-length string of characters "58" is not the same as the integer numeric value of 58, and we could not do any numerical computations with it. It is necessary to convert the string of numeric characters into an integer. That task is performed by the *Convert.ToInt32* function (it is also part of the *System* namespace). If the user enters a decimal number as the input value for *speed*, e.g., 48.7, *ReadLine* will gather the character string and pass it to *Convert.ToInt32*, which will fail and report an error that the input is not in the correct format for an integer.¹

The value of *distance* is input using the statement

```
distance = Convert.ToDouble(Console.ReadLine());
```

Note that this time the conversion of the string of characters gathered by *ReadLine* is to type *double*. Here it would be acceptable to enter an integer value, say 657, instead of 657.0. The conversion process knows that it can make a *double* value from a string of numeric characters that does not contain a decimal point.

After the two input statements, the value of the time can be computed and stored in the memory location referenced by *time*. A pseudocode operation for producing output would be something like

Print the value of *time*

This could be done by the following statement:

```
Console.Write(time);
```

Output in C# is handled as the opposite of input. A value stored in memory—in this case the value of the variable *time*—is converted into a string and copied to the Console (the screen). But we don't want the program to simply print a number with no explanation; we want some words to make the output meaningful.

The general form of the output statement is

```
Console.Write(string); or Console.WriteLine(string);
```

¹ Another form of this input statement is

```
speed = int.Parse(Console.ReadLine());
```

Parse converts the string supplied by *ReadLine* into its integer form.

The string could be empty, as follows:

```
Console.WriteLine();
```

This just prints a blank line, which is useful for formatting the output to make it easier to read. The string can also be a **literal string** (enclosed in double quotes). Literal strings are printed out exactly as is. For example,

```
Console.WriteLine("Here's your answer.");
```

prints

```
Here's your answer.
```

A string can also be composed of items joined by the concatenation operator `+`. The items can be literal strings, numbers, or variables. Items that are not themselves literal strings are converted to strings for the purposes of writing them out. In the `TravelPlanner` program we used

```
Console.Write(time + " hours to travel ");
```

to write out the value of `time`. Notice the spaces at the beginning and end of the literal string, within the quotation marks so that they are part of the text. Without these spaces, running the `TravelPlanner` program with our original data of 58 mph and 657.5 miles, the output would be printed as

```
11.3362068965517hours to travel657.5 miles.
```

Spacing aside, the value for `time` of

```
11.3362068965517
```

is fairly ridiculous. It does not make sense to display the result to 13 decimal digits. The appearance of numerical output can be controlled, rather than leaving it up to the system to decide, by including a **format specifier** in the output statement. If only two digits to the right of the decimal point are to be displayed for `time`, the output statement would take the following form:

```
Console.Write(time.ToString("0.00")
+ " hours to travel ");
```

Here the conversion to string is forced in the variable identifier list by using the `ToString` function, rather than waiting for the `Write` function to do it by default. The value `"0.00"` is the format specifier. It forces two digits to the right of the decimal point (rounding if the value has more than two decimal digits), and one or more digits to the left of the decimal point (even though there is only one zero in the format specifier to the left of the decimal point). Using this statement, the output would be rounded to two decimal places:

```
11.34 hours to travel 657.5 miles.
```

An alternative statement

```
Console.WriteLine(time.ToString("#.##")
    + " hours to travel ");
```

would produce the same result. The difference between the 0 and the # is that the # specifier will suppress leading and trailing zeros, while the 0 specifier will fill the indicated number of columns, using leading and trailing zeros if necessary. Both will use more columns to the left of the decimal point if the integer part is too large to fit in the indicated column space. For example:

<i>Format Specifier</i>	<i>Value</i>	<i>Printed Result</i>
"0.00"	275.687	275.69 rounded to 2 decimal digits
"00000.00"	2.6	00002.60 leading and trailing 0's
"#####.##"	2.6	2.6 no leading or trailing 0's
"#####.00"	2.6	2.60 no leading 0's, 2 decimal digits
"0000"	53	0053 leading 0's

The ability to specify the number of decimal digits in the output is particularly handy for dealing with dollar-and-cent values, where we always expect to see two digits behind the decimal point.

A real number written in the form 11.34 is said to be in **fixed-point format**, but with the appropriate format specifier, the output can be written in **scientific notation**. The fixed-point format

```
11.34
```

written in scientific notation is

```
1.134E + 01
```

which means 1.134×10^1 . (The "E" means "times 10 to the power of. . .".) The output statement to produce this would be

```
Console.WriteLine(time.ToString("0.000E+00")
    + " hours to travel ");
```

where the E in the format specifier indicates scientific notation, and the 00 after the plus sign forces two digits in the exponent.

Let's back up a bit and note that we also need to print some text information before the input statement, to alert the user that the program expects some input. A statement such as

```
Console.WriteLine("Enter your speed in mph: ");
```

acts as a user **prompt**. Without a prompt, the user may be unaware that the program is waiting for some input; instead, it may simply seem to the user that the program is "hung up."

Assembling all of these bits and pieces, we can see that

```
Console.Write("Enter your speed in mph: ");
speed = Convert.ToInt32(Console.ReadLine());
Console.Write("Enter your distance in miles: ");
distance = Convert.ToDouble(Console.ReadLine());
```

is a series of prompt, input, prompt, input statements to get the data, and then

```
Console.Write("At " + speed + " mph, ");
Console.WriteLine("it will take ");
Console.Write(time + " hours to travel ");
Console.WriteLine(distance + " miles.");
```

writes out the computed value of *time* along with the associated input values in an informative message. In the middle, we need a program statement to compute the value of *time*. We can do this with a single assignment statement; the assignment statement is explained in the next section.

Finally, whereas a single C# statement can be spread over multiple lines, a line break cannot occur in the middle of a literal string. The solution is to make two smaller substrings and join them together by concatenation, as in

```
Console.WriteLine("Oh for a sturdy ship to sail, "
    + "and a star to steer her by.");
```

which will produce a single line of output.

PRACTICE PROBLEMS

1. Write two statements that prompt the user to enter an integer value and store that value in a (previously declared) variable called *quantity*.
2. A program has computed a value of 37 for the variable *height*. Write an output statement that prints this variable using six columns, with successive output to appear on the next line.
3. What appears on the screen after execution of the following statements?

```
Console.Write("This is");
Console.WriteLine("goodbye");
```


What's in a Name?

Why is C# called C#? According to Anders Hejlsberg, leader of Microsoft's C# development team, the language was known as Cool when it was under development. But "Cool" presented trademark issues when it was time to market the language, so another name had to be found. C#, as we have noted, is a successor to C++, which in turn

is built on C. So C#, in a musical sense, is a "half-note" above C. But there's an even better explanation. C++ ("plus plus") is an improvement over C. Since the sharp symbol

#

contains 4 "+" marks, C# can be thought of as C++ plus plus, i.e., an improvement over C++.

3.2 The Assignment Statement

As we said earlier, an assignment statement assigns a value to a program variable. This is accomplished by evaluating some expression and then writing the resulting value in the memory location referenced by the program variable. The general pseudocode operation

Set the value of "variable" to "arithmetic expression"

has as its C# equivalent

```
variable = expression;
```

The expression on the right is evaluated, and the result is then written into the memory location named on the left. For example, suppose that *A*, *B*, and *C* have all been declared as integer variables in some program. The assignment statements

```
B = 2;
C = 5;
```

result in *B* taking on the value 2 and *C* taking on the value 5. After execution of

```
A = B + C;
```

A has the value that is the sum of the current values of *B* and *C*. Assignment is a destructive operation, so whatever *A*'s previous value was, it is gone. Note that this one assignment statement says to add the values of *B* and *C* and assign the result to *A*. This one high-level language statement is equivalent to three assembly language statements needed to do this same task (LOAD *B*, ADD *C*, STORE *A*). A high-level language program thus packs more power per line than an assembly language program. To state it another way, whereas a single assembly language instruction is equivalent to a single machine language instruction, a single C# instruction is usually equivalent to many assembly language instructions or machine language instructions, which allows us to think at a higher level of problem solving.

In the assignment statement, the expression on the right is evaluated first. Only then is the value of the variable on the left changed. This means that an assignment statement like

```
A = A + 1;
```

makes sense. If *A* has the value 7 before this statement is executed, then the expression evaluates to

```
7 + 1, or 8
```

and 8 then becomes the new value of *A*. (Here it becomes obvious that the assignment instruction symbol = is not the same as the mathematical equals sign =, because $A = A + 1$ does not make sense mathematically.)

All four basic arithmetic operations can be done in C#, where they are denoted by

```
+ Addition
- Subtraction
* Multiplication
/ Division
```

For the most part, this is standard mathematical notation, rather than the somewhat verbose assembly language op code mnemonics such as SUBTRACT. The reason a special symbol is used for multiplication is that \times would be confused with *x*, an identifier, \cdot (a multiplication dot) doesn't appear on the keyboard, and juxtaposition—writing *AB* for $A*B$ —would look like a single identifier named *AB*.

We do have to pay some attention to data types. In particular, division has one peculiarity. If at least one of the two values being divided is a real number, then division behaves as we expect. Thus,

```
7.0/2 7/2.0 7.0/2.0
```

all result in the value 3.5. However, if the two values being divided are both integers, the result is an integer value; if the division doesn't "come out even," the integer value is obtained by truncating the answer to an integer quotient. Thus,

```
7/2
```

results in the value 3. Think of grade-school long division of integers:

```
  3
2)7
 6
 1
```

Here the quotient is 3 and the remainder is 1. C# also provides an operation, with the symbol %, to obtain the integer remainder. Using this operation,

```
7 % 2
```

results in the value 1. If the values are stored in type *int* variables, the same thing happens. For example,

```
int numerator;
int denominator;
numerator = 7;
```

```

denominator = 2;
Console.Write("The result of " + numerator);
Console.Write("/" + denominator);
Console.WriteLine(" is " + numerator / denominator);

```

produces the output

```
The result of 7/2 is 3
```

As soon as an arithmetic operation involves one or more real (decimal) numbers, any integers are converted to their real number equivalent, and the calculations are done with real numbers.

Data types also play a role in assignment statements. Suppose the expression in an assignment statement evaluates to a real number and is then assigned to an identifier that has been declared as an integer. The real number is truncated, and the digits behind the decimal point are lost. We mentioned that this same situation occurs if you input a decimal value for an integer variable, but there, due to the conversion from string to integer, an error occurs. Unlike the input situation, the C# compiler can see what you are doing with the assignment statement and will usually give you a warning that says something about “possible loss of data.” But assigning an integer value to a type *double* identifier merely changes the integer to its real number equivalent. C# does this **type casting** (changing of data type) automatically. This type cast would merely change the integer 3, for example, to its real number equivalent 3.0.

This explains why we declared *distance* to be type *double* in the *TravelPlanner* program. The user can enter an integer value for distance, and C# uses `Convert.ToDouble` to type cast it to a real number. But if we had declared both *speed* and *distance* to be integers, then the division to compute *time* would only produce integer answers.

You should assign only an expression that has a character value to a variable that has been declared to be type *char*. Suppose that *letter* is a variable of type *char*. Then

```
letter = 'm';
```

is a legitimate assignment statement, giving *letter* the value of the character 'm'. Note that single quotation marks are used here, as opposed to the double quotation marks that enclose a literal string.

PRACTICE PROBLEMS

1. *newNumber* and *next* are integer variables in a C# program. Write a statement to assign the value of *newNumber* to *next*.
2. What is the value of *average* after the following statements are executed? (*total* and *number* are type *int*, and *average* is type *double*.)

```

total = 277;
number = 5;
average = total/number;

```


The assignment

```
letter = '4';
```

is also acceptable; the single quotes around the 4 mean that it is being treated as just another character on the keyboard, not as the integer 4.

▶ 3.3 Control Statements


We mentioned earlier that sequential flow of control is the default; that is, a program executes instructions sequentially from first to last. The flowchart in Figure 6 illustrates this, where S_1, S_2, \dots, S_k are program instructions (i.e., program statements).

As stated in Chapter 2, no matter how complicated the task to be done, only three types of control mechanisms are needed:

1. **Sequential:** Instructions are executed in order.
2. **Conditional:** Which instruction executes next depends on some condition.
3. **Looping:** A group of instructions may be executed many times.

Sequential flow of control, the default, is what occurs if the program does not contain any instances of the other two control structures. In the TravelPlanner program, for example, instructions are executed sequentially, beginning with the input statements, next the computation, and finally the output statements.

FIGURE 6
Sequential Flow of Control



In Chapter 2 we introduced pseudocode notation for conditional operations and looping. In Chapter 6 we learned how to write somewhat laborious assembly language code to implement conditional operations and looping. Now we'll see how C# provides instructions that directly carry out these control structure mechanisms—more evidence of the power of high-level language instructions. We can think in a pseudocode algorithm design mode, as we did in Chapter 2, and then translate that pseudocode directly into C# code.

Conditional flow of control begins with the evaluation of a **Boolean condition**, also called a **Boolean expression**, which can be either true or false. We discussed these “true/false conditions” in Chapter 2, and we also encountered Boolean expressions in Chapter 4, where they were used to design circuits. A Boolean condition often involves comparing the values of two expressions and determining whether they are equal, whether the first is greater than the second, and so on. Again assuming that *A*, *B*, and *C* are integer variables in a program, the following are legitimate Boolean conditions:

<code>A == 0</code>	(Does <i>A</i> currently have the value 0?)
<code>B < (A + C)</code>	(Is the current value of <i>B</i> less than the sum of the current values of <i>A</i> and <i>C</i> ?)
<code>A != B</code>	(Does <i>A</i> currently have a different value than <i>B</i> ?)

If the current values of *A*, *B*, and *C* are 2, 5, and 7, respectively, then the first condition is false (*A* does not have the value zero), the second condition is true (5 is less than 2 plus 7), and the third condition is true (*A* and *B* do not have equal values). Comparisons need not be numeric. They can also be done between variables of type *char*, where the “ordering” is the usual alphabetic ordering. If *initial* is a value of type *char* with a current value of 'D', then

```
initial == 'F'
```

is false because *initial* does not have the value 'F', and

```
initial < 'P'
```

is true because 'D' precedes 'P' in the alphabet (or, more precisely, because the binary code for 'D' is numerically less than the binary code for 'P'). Note that the comparisons are case sensitive, so 'F' is not equal to 'f', but 'F' is less than 'f'.

Figure 7 shows the comparison operations available in C#. Note the use of the two equality signs to test whether two expressions have the same value. The single equality sign is used in an assignment statement, the double equality sign in a comparison.

Boolean conditions can be built up using the Boolean operators AND, OR, and NOT. Truth tables for these operators were given in Chapter 4 (Figures 4.12–4.14). The only new thing is the symbols that C# uses for these operators, shown in Figure 8.

A conditional statement relies on the value of a Boolean condition (true or false) to decide which programming statement to execute next. If the condition is true, one statement is executed next, but if the condition is false, a different statement is executed next. Control is therefore no longer in a straight-line (sequential) flow, but hops to one place or to another. Figure 9 illustrates this situation. If the condition is true, the statement S1 is executed (and statement S2 is not); if the condition is false, the statement S2 is

FIGURE 7

C# Comparison Operators

COMPARISON	SYMBOL	EXAMPLE	EXAMPLE RESULT
the same value as	==	2 == 5	false
less than	<	2 < 5	true
less than or equal to	<=	5 <= 5	true
greater than	>	2 > 5	false
greater than or equal to	>=	2 >= 5	false
not the same value as	!=	2 != 5	true

FIGURE 8

C# Boolean Operators

OPERATOR	SYMBOL	EXAMPLE	EXAMPLE RESULT
AND	&&	(2 < 5) && (2 > 7)	false
OR		(2 < 5) (2 > 7)	true
NOT	!	!(2 == 5)	true

executed (and statement S1 is not). In either case, the flow of control then continues on to statement S3. We saw this same scenario when we discussed pseudocode conditional statements in Chapter 2 (Figure 2.4).


The C# instruction that carries out conditional flow of control is called an **if-else** statement. It has the form shown below (note that the words *if* and *else* are lowercase and that the Boolean condition must be in parentheses).

```

if (Boolean condition)
    S1;
else
    S2;

```

FIGURE 9

Conditional Flow of Control
(if-else)

Below is a simple if-else statement, where we assume that A , B , and C are integer variables.

```
if (B < (A + C))
    A = 2*A;
else
    A = 3*A;
```

Suppose that when this statement is reached, the values of A , B , and C are 2, 5, and 7, respectively. As we noted before, the condition $B < (A + C)$ is then true, so the statement

```
A = 2*A;
```

is executed, and the value of A is changed to 4. However, suppose that when this statement is reached, the values of A , B , and C are 2, 10, and 7, respectively. Then the condition $B < (A + C)$ is false, the statement


```
A = 3*A;
```

is executed, and the value of A is changed to 6.

A variation on the if-else statement is to allow an “empty else” case. Here we want to do something if the condition is true, but if the condition is false, we want to do nothing. Figure 10 illustrates the empty else case. If the condition is true, statement $S1$ is executed, and after that the flow of control continues on to statement $S3$, but if the condition is false, nothing happens except to move the flow of control directly on to statement $S3$.

FIGURE 10

If-else with Empty else



This *if* variation on the if-else statement can be accomplished by omitting the word *else*. This form of the instruction therefore looks like

```
if (Boolean condition)
    S1;
```

We could write

```
if (B < (A + C))
    A = 2*A;
```

This has the effect of doubling the value of A if the condition is true and of doing nothing if the condition is false.

It is possible to combine statements into a group by putting them within the curly braces { and }. The group is then treated as a single statement, called a **compound statement**. A compound statement can be used anywhere a single statement is allowed. For example,

```
{
    Console.WriteLine("This is the first statement.");
    Console.WriteLine("This is the second statement.");
    Console.WriteLine("This is the third statement.");
}
```

is treated as a single statement. The implication is that in Figure 9, S1 or S2 might be compound statements. This makes the if-else statement potentially much more powerful and similar to the pseudocode conditional statement in Figure 2.9.

Let's expand on our TravelPlanner program and give the user of the program a choice of computing the time either as a decimal number (3.75 hours) or as hours and minutes (3 hours, 45 minutes). This situation is ideal for a conditional statement. Depending on what the user wants to do, the program does one of two tasks. For either task, the program still needs information about the speed and distance. The program must also collect information to indicate which task the user wishes to perform. We need an additional variable in the program to store this information. Let's use a variable called *choice* of type *char* to collect the user's choice of which task to perform. We also need two new integer variables to store the values of hours and minutes.

Figure 11 shows the new program, with the three additional declared variables. The condition evaluated at the beginning of the if-else statement tests whether *choice* has the value 'D'. If so, then the condition is true, and the first group of statements is executed—that is, the time is output in decimal format as we have been doing all along. If *choice* does not have the value 'D', then the condition is false. In this event, the second group of statements is executed. Note that because of the way the condition is written, if *choice* does not have the value 'D', it is assumed that the user wants to compute the time in hours and minutes, even though *choice* may have any other non-D value (including 'd') that the user may have typed in response to the prompt.

**FIGURE 11**

*The TravelPlanner Program with
a Conditional Statement*

```
//Computes and outputs travel time
//for a given speed and distance
//Written by J. Q. Programmer, 6/10/16

using System;

namespace InvitationCSharp
{
    class TravelPlanner
    {
        static void Main(string[] args)
        {
            int speed;           //rate of travel
            double distance;     //miles to travel
            double time;         //time needed for this travel
            int hours;           //time for travel in hours
            int minutes;         //leftover time in minutes
            char choice;         //choice of output as
                                //decimal hours
                                //or hours and minutes

            Console.Write("Enter your speed in mph: ");
            speed = Convert.ToInt32(Console.ReadLine());
            Console.Write("Enter your distance in miles: ");
            distance = Convert.ToDouble(Console.ReadLine());
            Console.WriteLine("Enter your choice of format" +
                " for time, ");
            Console.Write("decimal hours (D) " +
                "or hours and minutes (M): ");
            choice = Convert.ToChar(Console.ReadLine());

            if (choice == 'D')
            {
                time = distance / speed;
                Console.Write("At " + speed + " mph, ");
                Console.WriteLine("it will take ");
                Console.Write(time + " hours to travel ");
                Console.WriteLine(distance + " miles.");
            }
            else
            {
                time = distance / speed;
                hours = (int)time;
                minutes = (int)((time - hours) * 60);
                Console.Write("At " + speed + " mph, ");
                Console.WriteLine("it will take ");
                Console.WriteLine(hours + " hours and " +
                    minutes + " minutes to travel " +
                    distance + " miles.");
            }
        }
    }
}
```

To compute hours and minutes (the *else* clause of the if-else statement), time is computed in the usual way, which results in a decimal value. The whole number part of that decimal is the number of hours needed for the trip. We can get this number by type casting the decimal number to an integer. This is accomplished by

```
hours = (int)time;
```

which drops all digits behind the decimal point and stores the resulting integer value in *hours*. Based on earlier discussions of converting to integer, the following line of code would seem appropriate:

```
hours = Convert.ToInt32(time);
```

This line of code does not give the desired result because the *Convert.ToInt32* function actually rounds the type *double* value. (For the example below, it takes in 9.5 and produces 10.)

To find the fractional part of the hour that we dropped, we subtract *hours* from *time*. We multiply this by 60 to turn it into some number of minutes, but this is still a decimal number. We do another type cast to truncate this to an integer value for *minutes*:

```
minutes = (int)((time - hours) * 60);
```

For example, if the user enters data of 50 mph and 475 miles and requests output in hours and minutes, the table below shows the computed values.

<i>Quantity</i>	<i>Value</i>
speed	50
distance	475
time = distance/speed	9.5
hours = int(time)	9
time - hours	0.5
(time - hours) * 60	30.0
minutes = int((time - hours)*60)	30

Here is the actual program output for this case:

```
Enter your speed in mph: 50
Enter your distance in miles: 475
Enter your choice of format for time,
decimal hours (D) or hours and minutes (M): M
At 50 mph, it will take
9 hours and 30 minutes to travel 475 miles.
```

The two statement groups in an if-else statement are identified by the enclosing curly braces, but in Figure 11 we also indented them to make them easier

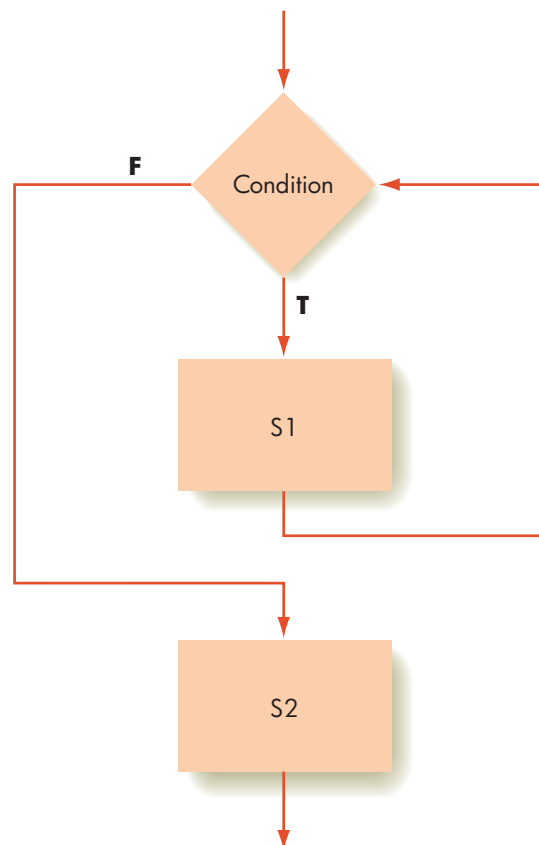
to pick out when looking at the program. Like comments, indentation is ignored by the computer but is valuable in helping people to more readily understand a program.

Now let's look at the third variation on flow of control, namely looping (iteration). We want to execute the same group of statements (called the **loop body**) repeatedly, depending on the result of a Boolean condition. As long as (while) the condition remains true, the loop body is executed. The condition is tested before each execution of the loop body. When the condition becomes false, the loop body is not executed again, which is usually expressed by saying that the algorithm *exits* the loop. To ensure that the algorithm ultimately exits the loop, the condition must be such that its truth value can be affected by what happens when the loop body is executed. Figure 12 illustrates the while loop. The loop body is statement S1 (which can be a compound statement), and S1 is executed *while* the condition is true. Once the condition is false, the flow of control moves on to statement S2. If the condition is false when it is first evaluated, then the body of the loop is never executed at all. We saw this same scenario when we discussed pseudocode looping statements in Chapter 2 (Figure 2.6).

C# uses a while statement to implement this type of looping. The form of the statement is

```
while (Boolean condition)  
    S1;
```

FIGURE 12
while Loop



For example, suppose we want to write a program to add up a number of nonnegative integers that the user supplies and write out the total. We need a variable to hold the total; we'll call this variable *sum*, and make its data type *int*. To handle the numbers to be added, we could declare a bunch of integer variables such as *n1*, *n2*, *n3*, . . . and do a series of input-and-add statements of the form

```
n1 = Convert.ToInt32(Console.ReadLine());
sum = sum + n1;
n2 = Convert.ToInt32(Console.ReadLine());
sum = sum + n2;
```

and so on. There are two problems with this approach. The first is that we may not know ahead of time how many numbers the user wants to add. If we declare variables *n1*, *n2*, . . . , *n25*, and the user wants to add 26 numbers, the program won't do the job. The second problem is that this approach requires too much effort. Suppose that we know the user wants to add 2000 numbers. We could declare 2000 variables (*n1*, . . . , *n2000*), and we could write the above input-and-add statements 2000 times, but it wouldn't be fun. Nor is it necessary—we are doing a very repetitive task here, and we should be able to use a loop mechanism to simplify the job. (We faced a similar situation in the first pass at a sequential search algorithm, Figure 2.11; our solution there was also to use iteration.)

Even if we use a loop mechanism, we are still adding a succession of values to *sum*. Unless we are sure that the value of *sum* is zero to begin with, we cannot be sure that the answer isn't nonsense. Remember that the identifier *sum* is simply an indirect way to designate a memory location in the computer. That memory location contains a pattern of bits, perhaps left over from whatever was stored there when some previous program was run. We cannot assume that just because this program hasn't used *sum*, its value is zero. (In contrast, the assembly language statement `SUM: .DATA 0` reserves a memory location, assigns it the identifier `SUM`, and fills it with the value zero.) If we want the beginning value of *sum* to be zero, we must use an assignment statement. Using assignment statements to set the values of certain variables before they are used by the program is called **initialization of variables**.

Now on to the loop mechanism. First, let's note that once a number has been read in and added to *sum*, the program doesn't need to know the value of the number any longer. We can declare just one integer variable called *number* and use it repeatedly to hold the first numerical value, then the second, and so on. The general idea is

```
sum = 0; //initialize sum
while (there are more numbers to add)
{
    number = Convert.ToInt32(Console.ReadLine());
    sum = sum + number;
}
Console.WriteLine("The total is " + sum);
```

Now we have to figure out what the condition "there are more numbers to add" really means. Because we are adding nonnegative integers, we could ask

the user to enter one extra integer that is not part of the legitimate data but is instead a signal that there *are* no more data. Such a value is called a **sentinel value**. For this problem, any negative number would be a good sentinel value. Because the numbers to be added are all nonnegative, the appearance of a negative number signals the end of the legitimate data. We don't want to process the sentinel value (because it is not a legitimate data item); we only want to use it to terminate the looping process. This might suggest the following code:

```
sum = 0; //initialize sum
while (number >= 0) //but there is a problem here,
//see following discussion
{
    number = Convert.ToInt32(Console.ReadLine());
    sum = sum + number;
}
Console.WriteLine("The total is " + sum);
```

Here's the problem. How can we test whether *number* is greater than or equal to 0 if we haven't read the value of *number* yet? We need to do a preliminary input for the first value of *number* outside of the loop and then test that value in the loop condition. If it is nonnegative, we want to add it to *sum* and then read the next value and test it. Whenever the value of *number* is negative (including the first value), we want to do nothing with it—that is, we want to avoid executing the loop body. The following statements do this; we've also added instructions to the user.

```
sum = 0; //initialize sum
Console.Write("Please enter numbers to add; ");
Console.WriteLine("terminate with a negative "
    + "number.");
number = Convert.ToInt32(Console.ReadLine());
while (number >= 0)
{
    sum = sum + number;
    number = Convert.ToInt32(Console.ReadLine());
}
Console.WriteLine("The total is " + sum);
```

The value of *number* gets changed within the loop body by reading in a new value. The new value is tested, and if it is nonnegative, the loop body executes again, adding the data value to *sum* and reading in a new value for *number*. The loop terminates when a negative value is read in. Remember the requirement that something within the loop body must be able to affect the truth value of the condition. In this case, it is reading in a new value for *number* that has the potential to change the value of the condition from true to false. Without this requirement, the condition, once true, would remain true forever, and the loop body would be endlessly executed. This results in what is called an **infinite loop**. A program that contains an infinite loop will execute forever (or until the programmer gets tired of waiting and interrupts the program, or until the program exceeds some preset time limit).

Here is a sample of the program output.

```
Please enter numbers to add; terminate with a negative
number.
```

```
5
6
10
-1
The total is 21
```

The problem we've solved here, adding nonnegative integers until a negative sentinel value occurs, is the same one solved using assembly language in Chapter 6. The preceding C# code is almost identical to the pseudocode version of the algorithm shown in Figure 6.7. Thanks to the power of the language, the C# code embodies the algorithm directly, at a high level of thinking, whereas in assembly language this same algorithm had to be translated into the lengthy and awkward code of Figure 6.8.

To process data for a number of different trips in the *TravelPlanner* program, we could use a while loop. During each pass through the loop, the program computes the time for a given speed and distance. The body of the loop is therefore exactly like our previous code. All we are adding here is the framework that provides looping. To terminate the loop, we could use a sentinel value, as we did for the program above. A negative value for *speed*, for example, is not a valid value and could serve as a sentinel value. Instead of that, let's allow the user to control loop termination by having the program ask the user whether he or she wishes to continue. We'll need a variable to hold the user's response to this question. Of course, the user could answer "N" at the first query, the loop body would never be executed at all, and the program would terminate. Figure 13 shows the complete program.

FIGURE 13

*The TravelPlanner Program
with Looping*

```
//Computes and outputs travel time
//for a given speed and distance
//Written by J. Q. Programmer, 6/20/16

using System;

namespace InvitationCSharp
{
    class TravelPlanner
    {
        static void Main(string[] args)
        {
            int speed;           //rate of travel
            double distance;     //miles to travel
            double time;         //time needed for this travel
            int hours;           //time for travel in hours
```

**FIGURE 13**

The TravelPlanner Program with Looping (continued)

```

int minutes;           //leftover time in minutes
char choice;          //choice of output as
                      //decimal hours
                      //or hours and minutes
char more;            //user's choice to do
                      //another trip

Console.Write("Do you want to plan a trip? " +
              "(Y or N): ");
more = Convert.ToChar(Console.ReadLine());

while (more == 'Y') //more trips to plan
{
    Console.Write("Enter your speed in mph: ");
    speed = Convert.ToInt32(Console.ReadLine());
    Console.Write("Enter your distance in miles: ");
    distance = Convert.ToDouble(Console.ReadLine());
    Console.WriteLine("Enter your choice of format" +
                      " for time, ");
    Console.Write("decimal hours (D) " +
                  "or hours and minutes (M): ");
    choice = Convert.ToChar(Console.ReadLine());

    if (choice == 'D')
    {
        time = distance / speed;
        Console.Write("At " + speed + " mph, ");
        Console.WriteLine("it will take ");
        Console.Write(time + " hours to travel ");
        Console.WriteLine(distance + " miles.");
    }
    else
    {
        time = distance / speed;
        hours = (int)time;
        minutes = (int)((time - hours) * 60);
        Console.Write("At " + speed + " mph, ");
        Console.WriteLine("it will take ");
        Console.WriteLine(hours + " hours and " +
                            minutes + " minutes to travel " +
                            distance + " miles.");
    }
    Console.WriteLine();
    Console.Write("Do you want to plan a trip? " +
                  "(Y or N): ");
    more = Convert.ToChar(Console.ReadLine());
} //end of while loop
}
}
}

```


PRACTICE PROBLEMS

Assume all variables have previously been declared.

1. What is the output from the following section of code?

```
number1 = 15;
number2 = 7;
if (number1 >= number2)
    Console.WriteLine(2*number1);
else
    Console.WriteLine(2*number2);
```

2. What is the output from the following section of code?

```
scores = 1;
while (scores < 20)
{
    scores = scores + 2;
    Console.WriteLine(scores);
}
```

3. What is the output from the following section of code?

```
quotaThisMonth = 7;
quotaLastMonth = quotaThisMonth + 1;
if ((quotaThisMonth > quotaLastMonth) ||
    (quotaLastMonth >= 8))
{
    Console.WriteLine("Yes");
    quotaLastMonth = quotaLastMonth + 1;
}
else
{
    Console.WriteLine("No");
    quotaThisMonth = quotaThisMonth + 1;
}
```

4. How many times is the *WriteLine* statement executed in the following section of code?

```
left = 10;
right = 20;
while (left <= right)
{
    Console.WriteLine(left);
    left = left + 2;
}
```

5. Write a C# statement that outputs "Equal" if the integer values of *night* and *day* are the same, but otherwise does nothing.

4

Another Example

Let's briefly review the types of C# programming statements we've learned. We can do input and output—reading values from the user into memory, writing values out of memory for the user to see, being sure to use meaningful variable identifiers to reference memory locations. We can assign values to variables within the program. And we can direct the flow of control by using conditional statements or looping. Although many other statement types are available in C#, you can do almost everything using only the modest collection of statements we have described. The power lies in how these statements are combined and nested within groups to produce ever more complex courses of action.

For example, suppose we write a program to assist SportsWorld, a company that installs circular swimming pools. In order to estimate their costs for swimming pool covers or for fencing to surround the pool, SportsWorld needs to know the area or circumference of a pool, given its radius. A pseudocode version of the program is shown in Figure 14.

We should be able to translate this pseudocode fairly directly into the body of the Main function. Other things we need to add to complete the program are:

- A prologue comment to explain what the program does (optional but always recommended for program documentation)
- A using directive for namespace *System* (necessary because our program uses *ReadLine*, *Write*, and *WriteLine*)
- A declaration for the constant value *pi* (3.1416)
- Variable declarations

Figure 15 gives the complete program. Figure 16 shows what actually appears on the screen when this program is executed with some sample data.


FIGURE 14

A Pseudocode Version of the SportsWorld Program

```

Get value for user's choice about continuing
While user wants to continue, do the following steps
    Get value for pool radius
    Get value for choice of task
    If task choice is circumference
        Compute pool circumference
        Print output
    Else (task choice is area)
        Compute pool area
        Print output
    Get value for user's choice about continuing
Stop
  
```

**FIGURE 15***The SportsWorld Program*

```

//This program helps SportsWorld estimate costs
//for pool covers and pool fencing by computing
//the area or circumference of a circle
//with a given radius.
//Any number of circles can be processed.
//Written by M. Phelps, 10/05/16

using System;

namespace SportsWorld
{
    class Program
    {
        static void Main(string[] args)
        {
            const double pi = 3.1416;    //value of pi
            double radius;                //radius of a pool - given
            double circumference;        //circumference of a pool -
                                        //computed
            double area;                 //area of a pool -
                                        //computed
            char taskToDo;               //holds user choice to
                                        //compute circumference
                                        //or area
            char more;                  //controls loop for
                                        //processing more pools

            Console.WriteLine("Do you want to process a pool? (Y or N): ");
            more = Convert.ToChar(Console.ReadLine());

            while (more == 'Y') //more circles to process
            {
                Console.WriteLine();
                Console.WriteLine("Enter the value of the radius of a " +
                    "pool: ");
                radius = Convert.ToDouble(Console.ReadLine());

                //See what user wants to compute
                Console.WriteLine("Enter your choice of task.");
                Console.WriteLine("C to compute circumference, " +
                    "A to compute area: ");
                taskToDo = Convert.ToChar(Console.ReadLine());

                if (taskToDo == 'C') //compute circumference
                {
                    circumference = 2 * pi * radius;
                    Console.WriteLine("The circumference for a pool " +
                        "of radius " + radius.ToString("##.##") +
                        " is " + circumference.ToString("##.##"));
                }
                else //compute area
                {
                    area = pi * radius * radius;
                    Console.WriteLine("The area for a pool " +
                        "of radius " + radius.ToString("##.##") +

```

**FIGURE 15**

*The SportsWorld Program
(continued)*

```

        " is " + area.ToString("##.##"));
    }
    Console.WriteLine();
    Console.Write("Do you want to process more pools? " +
        "(Y or N): ");
    more = Convert.ToChar(Console.ReadLine());
} //end of while loop

        //finish up
    Console.WriteLine("Program will now terminate.");
}
}
}

```

**FIGURE 16**

*A Sample Session Using the
Program of Figure 15*

```

Do you want to process a pool? (Y or N): Y
Enter the value of the radius of a pool: 2.7
Enter your choice of task.
C to compute circumference, A to compute area: C
The circumference for a pool of radius 2.70 is 16.96

Do you want to process more pools? (Y or N): Y

Enter the value of the radius of a pool: 2.7
Enter your choice of task.
C to compute circumference, A to compute area: A
The area that for a pool of radius 2.70 is 22.90

Do you want to process more pools? (Y or N): Y

Enter the value of the radius of a pool: 14.53
Enter your choice of task.
C to compute circumference, A to compute area: C
The circumference for a pool of radius 14.53 is 91.29

Do you want to process more pools? (Y or N): N

Program will now terminate.

```

PRACTICE PROBLEMS

1. Write a complete C# program to read in an integer number and write out the square of that number.
2. Write a complete C# program that asks for the price of an item and the quantity purchased, and writes out the total cost.
3. Write a complete C# program that asks for a number. If the number is less than 5, it is written out, but if it is greater than or equal to 5, twice that number is written out.
4. Write a complete C# program that asks the user for a positive integer n and then writes out all the numbers from 1 up to and including n .

Zooming Through the Universe

C# is the programming language behind Worldwide Telescope, an application built by Microsoft Research. Worldwide Telescope uses a catalogue of astronomical images and data from many ground-based and space-based observatories around the world to let users virtually view the night sky. Users can choose which telescope to look through, including the Hubble Space Telescope. They can zoom through the

solar system, the galaxy, and beyond, and view the relative positions of planets and other heavenly bodies in the past, present, and future. Then from far-flung views of the universe, it's possible to zoom in on planet Earth right down to the individual building level. Users can also choose different light wavelengths through which to view the universe and explore hidden structures in the galaxies. "Guided tours" with experts are available. Worldwide Telescope is available for free download at www.worldwidetelescope.org.

5 Managing Complexity

The programs we have written have been relatively simple. More complex problems require more complex programs to solve them. Although it is fairly easy to understand what is happening in the 40 or so lines of the SportsWorld program, imagine trying to understand a program that is 50,000 lines long. Imagine trying to write such a program! It is not possible to understand—all at once—everything that goes on in a 50,000-line program.

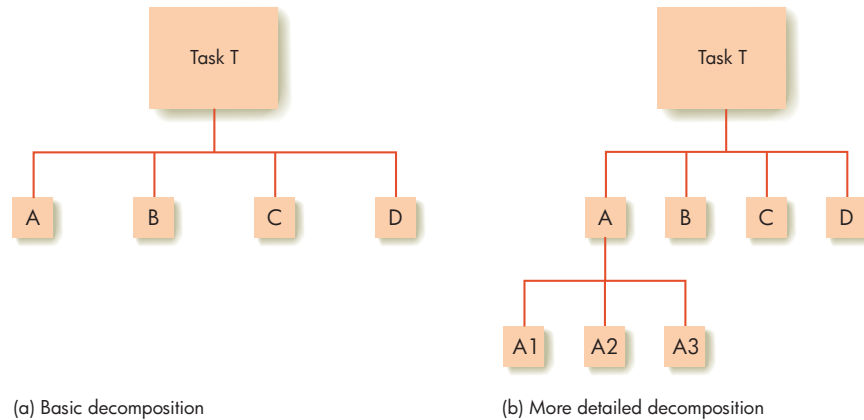
▶ 5.1 Divide and Conquer

Writing large programs is an exercise in managing complexity. The solution is a problem-solving approach called **divide and conquer**. Suppose a program is to be written to do a certain task; let's call it task T. Suppose further that we can divide this task into smaller tasks, say A, B, C, and D, such that, if we can do those four tasks in the right order, we can do task T. Then our high-level understanding of the problem need only be concerned with *what* A, B, C, and D do and how they must work together to accomplish T. We do not, at this stage, need to understand *how* A, B, C, and D can be done. Figure 17(a), an example of a **structure chart** or **structure diagram**, illustrates this situation. Task T is composed in some way of subtasks A, B, C, and D. Later we can turn our attention to, say, subtask A and see if it too can be decomposed into smaller subtasks, as in Figure 17(b). In this way, we continue to break the task down into smaller and smaller pieces, finally arriving at subtasks that are simple enough that it is easy to write the code to carry them out. By *dividing* the problem into small pieces, we can *conquer* the complexity that is overwhelming if we look at the problem as a whole.

Divide and conquer is a problem-solving approach and not just a computer programming technique. Outlining a term paper into major and minor topics is a divide-and-conquer approach to writing the paper. Doing a Form 1040 Individual Tax Return for the Internal Revenue Service can involve the subtasks of completing Schedules A, B, C, D, and so on and then reassembling the results. Designing a house can be broken down into subtasks of designing floor plans, wiring, plumbing, and the like. Large companies organize their management responsibilities using a divide-and-conquer approach; what we have called structure charts become, in the business world, organization charts.



FIGURE 17
Structure Charts



How is the divide-and-conquer problem-solving approach reflected in the resulting computer program? If we think about the problem in terms of subtasks, then the program should show that same structure; that is, part of the code should do subtask A, part should do subtask B, and so on. We divide the code into *modules* or *subprograms*, each of which does some part of the overall task. Then we empower these modules to work together to solve the original problem.

▶ 5.2 Using Functions

In C#, modules of code are called **functions**. Each function in a program should do one and only one subtask. These “subtask functions” are the optional functions listed before the mandatory Main function in the C# program outline of Figure 2. When subtask functions are used, the Main function consists primarily of invoking these subtask functions in the correct order.

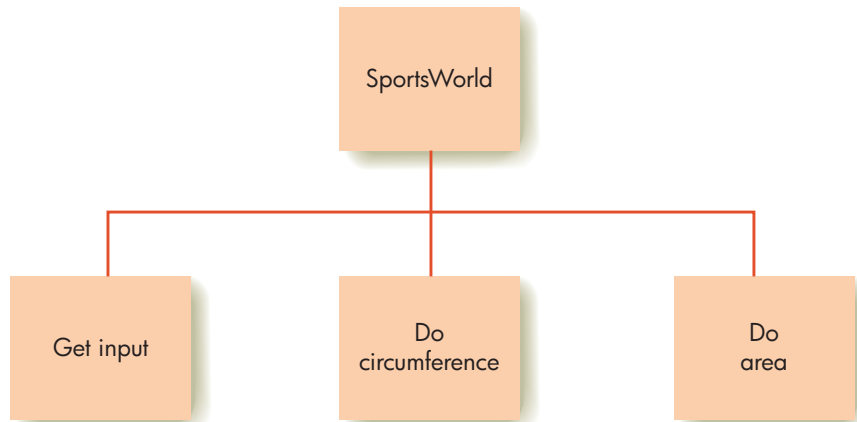
Let’s review the Main function body of the SportsWorld program (Figure 15) with an eye to further subdividing the task. There is a loop that does some operations as long as the user wants. What gets done? Input is obtained from the user about the radius of the circle and the choice of task to be done (compute circumference or compute area). Then the circumference or the area gets computed and written out. We’ve identified three subtasks, as shown in the structure chart of Figure 18.

We can visualize the Main function body of the program at a pseudocode level as shown in Figure 19. This divide-and-conquer approach to solving the problem can (and should) be planned first in pseudocode, without regard to the details of the programming language to be used. If the three subtasks (input, circumference, area) can all be done, then arranging them within the structure of Figure 19 solves the problem. We can write a function for each of the subtasks. Although we now know what form the Main function body will take, we have pushed the details of how to do each of the subtasks off into the other functions. Execution of the program begins with the Main function. Every time the flow of control reaches the equivalent of a “do subtask” instruction, it transfers execution to the appropriate function code. When execution of the function code is complete, flow of control returns to the Main function and picks up where it left off.



FIGURE 18

Structure Chart for the SportsWorld Task



Before we look at the details of how to write a function, we need to examine the mechanism that allows the functions to work with each other and with the Main function. This mechanism consists of passing information about various quantities in the program back and forth between the other functions and the Main function. Because each function is doing only one subtask of the entire task, it does not need to know the values of all variables in the program. It only needs to know the values of the variables with which its particular subtask is concerned. Allowing a function access only to pertinent variables prevents that function from inadvertently changing a value it has no business changing.

When the Main function wants a subtask function to be executed, it gives the name of the function (which is an ordinary C# identifier) and also a list of the identifiers for variables pertinent to that function. This is called an **argument list**. In our SportsWorld program, let's name the three functions *getInput*, *doCircumference*, and *doArea* (names that are descriptive of the subtasks these functions carry out). The *getInput* function collects the values for the variables *radius* and *taskToDo*. The Main function invokes the *getInput* function with the statement

```
getInput(ref double radius, ref char taskToDo);
```

which takes the place of the "Do the input subtask" line in Figure 19. (The meaning of "ref" will be explained later.) When this statement is reached,



FIGURE 19

A High-Level Modular View of the SportsWorld Program

```

Get value for user's choice about continuing
While the user wants to continue
  Do the input subtask
  If (Task = 'C') then
    do the circumference subtask
  else
    do the area subtask
  Get value for user's choice about continuing
  
```

control passes to the *getInput* function. After execution of this function, control returns to the Main function, and the variables *radius* and *taskToDo* have the values obtained for them within *getInput*.

The *doCircumference* function computes and writes out the value of the circumference, and in order to do that it needs to know the radius. Therefore, the variable *radius* is a legitimate argument for this function. The Main function contains the statement

```
doCircumference(radius);
```

in place of the “do the circumference subtask” in Figure 19. When this statement is reached, the variable *radius* conveys the value of the radius to the *doCircumference* function, and the function computes and writes out the circumference. The variable *circumference*, then, is also a variable of interest to the *doCircumference* function, but it is of interest to this function alone, in the sense that *doCircumference* does the computation and writes out the result. No other use is made of the circumference in the entire program, so no other function, including the Main function, has anything to do with *circumference*. Instead of being declared in the body of the Main function, *circumference* will be declared (and can be used) only within the *doCircumference* function; it will be **local** to that function. Any function can have its own **local constants** and **local variables**, declared within and known only to that function.

The *doCircumference* function also needs to know the value of the constant *pi*. We could declare *pi* as a constant local to *doCircumference*, but *doArea* needs the same constant, so we will declare *pi* right after the program *namespace* directive, not within any function. This will make *pi* a **global constant** whose value is known everywhere. The value of a constant cannot be changed, so there is no reason to prevent any function from having access to its value.

The *doArea* function computes and writes out the area and needs to know the radius, so the line “do the area subtask” in Figure 19 is replaced by

```
doArea(radius);
```

Within *doArea*, *area* is a local variable.

Now we can write the Main function of the modularized version of the SportsWorld program, shown in Figure 20. The Main function body is a direct translation of Figure 19. If, in starting from scratch to write this program, we had taken a divide-and-conquer approach, broken the original problem down into three subtasks, and come up with the outline of Figure 19, it would have been easy to get from there to Figure 20. The only additional task would have been determining the variables needed.

At a glance, the Main function in Figure 20 does not look a great deal different from our former Main function. However, it is conceptually quite different; the subtasks of getting the input values, computing and writing out the circumference, and computing and writing out the area have been relegated to functions. The details (such as the formulas for computing circumference and area) are now hidden and have been replaced by function invocations. If these subtasks had required many lines of code, our new Main function would indeed be shorter—and easier to understand—than before.

**FIGURE 20**

The Main Function in a Modularized Version of the SportsWorld Program

```
static void Main(string[] args)
{
    double radius = 0;           //radius of a pool - given
    char taskToDo = ' ';        //holds user choice to
                                //compute circumference
                                //or area
    char more;                  //controls loop for
                                //processing more pools

    Console.Write("Do you want to process a pool?(Y or N): ");
    more = Convert.ToChar(Console.ReadLine());

    while (more == 'Y')         //more circles to process
    {
        getInput(ref radius, ref taskToDo);

        if (taskToDo == 'C')    //compute circumference
        {
            doCircumference(radius);
        }
        else                    //compute area
        {
            doArea(radius);
        }
        Console.WriteLine();
        Console.Write("Do you want to process more pools? " +
            "(Y or N): ");
        more = Convert.ToChar(Console.ReadLine());
    } //end of while loop
    //finish up
    Console.WriteLine("Program will now terminate.");
}
```

▶ 5.3 Writing Functions

Now we know how the Main function can invoke another function. (In fact, using the same process, any function can invoke another function. A function can even invoke itself.) It is time to see how to write the code for these other, non-Main functions. The general outline for a C# function is shown in Figure 21.

The function header consists of three parts:

- A return indicator
- The function identifier
- A parameter list

**FIGURE 21**

The Outline for a C++ Function

```
function header
{
    local declarations [optional]
    function body
}
```

The **return indicator** classifies a function as a “void” or a “nonvoid” function. We’ll explain this distinction later, but the three functions for the circle program are all void functions, so the return indicator is the keyword **void**. (All of our Main functions have been void functions as well.) The **function identifier** can be any legitimate C# identifier. The parameters in the **parameter list** correspond to the arguments in the statement that invokes this function; that is, the first parameter in the list matches the first argument given in the statement that invokes the function, the second parameter matches the second argument, and so on. It is through this correspondence between parameters and arguments that information (data) flows from the Main function to other functions, and vice versa. The data type of each parameter must be given as part of the parameter list, and it must match the data type of the corresponding argument. For example, because the *getInput* function is invoked with the two arguments *radius* and *taskToDo*, the parameter list for the *getInput* function header has two parameters, the first of type *double* and the second of type *char*. Parameters may have, but do not have to have, the same identifiers as the corresponding arguments; arguments and parameters correspond by virtue of their respective positions in the argument list and the parameter list, regardless of the identifiers used. For the *getInput* function, we choose the parameter identifiers *radius* and *taskToDo*, matching the argument identifiers. No semicolon is used at the end of a function header.

One additional aspect of the parameter list in the function header concerns the use the function will make of each parameter. Consider the statement that invokes the function; an argument in the invoking statement carries a data value to the corresponding parameter in the function header. If the value is one that the function must know to do its job but should not change, then the argument is **passed by value**. The function receives a copy of the data value but never knows the memory location where the original value is stored. If the function changes the value of its copy, this change has no effect when control returns to the Main function. If, however, the value passed to the function is one that the function should change, and the Main function should know the new value, then the argument is **passed by reference**. The function receives access to the memory location where the value is stored, and any changes it makes to the value are seen by the Main function after control returns there. Included in this category are arguments whose values are unknown when the function is invoked (which really means that they are meaningless values of whatever happens to be in the memory location associated with that identifier), but the function changes those unknown values into meaningful values.

By default, arguments in C# are passed by value, which protects them from change by the function. Explicit action must be taken by the programmer to pass an argument by reference; specifically, the keyword *ref* must appear in front of the corresponding parameter in the function parameter list, and in front of the corresponding argument in the argument list.

How do we decide whether to pass an argument by value or by reference? If the Main function needs to obtain a new value back from a function when execution of that function terminates, then the argument must be passed by reference (by inserting the *ref* into the parameter and argument lists). Otherwise, the argument should be passed by value, the default arrangement.

In the *getInput* function, both *radius* and *taskToDo* are values that *getInput* obtains from the user and that the Main function needs to know when *getInput* terminates, so both of these are passed by reference. The header for the *getInput* function is shown next along with the invoking statement from the Main function. Note that the parameters *radius* and *taskToDo* are in the

right order, have been given the correct data types, and are both marked for passing by reference. Also remember that, although the arguments are named *radius* and *taskToDo* because those are the variable identifiers declared in the Main function, the parameters could have different identifiers, and it is the parameter identifiers that are used within the body of the function.

```
void getInput(ref double radius, ref char taskToDo)
//function
//header
getInput(ref radius, ref taskToDo); //function
//invocation
```

In C# an argument, whether passed by value or by reference, must have been given a value before the invocation. This explains the statements

```
double radius = 0;
char taskToDo = ' ';
```

at the beginning of the Main function above. These combine declaration with initialization of the variables.

The body of the *getInput* function comes from the corresponding part of Figure 15. If we hadn't already written this code, we could have done a pseudocode plan first. The complete function appears in Figure 22, where a comment has been added to document the purpose of the function.

The *doCircumference* function needs to know the value of *radius* but does not change that value. Therefore, *radius* is passed by value. Why is the distinction between arguments passed by value and those passed by reference important? If functions are to effect any changes at all, then clearly reference parameters are necessary, but why not just make everything a reference parameter? Suppose that in this example *radius* is made a reference parameter. If an instruction within *doCircumference* were to inadvertently change the value of *radius*, then that new value would be returned to the Main function, and any subsequent calculations using this value (there are none in this example) would be in error. Making *radius* a value parameter prevents this. How could one possibly write a program statement that changes the value of a variable inadvertently? In something as short and simple as our example, this probably would not happen, but in a more complicated program, it might. Distinguishing

FIGURE 22

The *getInput* Function

```
static void getInput(ref double radius, ref char taskToDo)
//gets radius and choice of task from the user
{
    Console.WriteLine();
    Console.Write("Enter the value of the radius of a " +
        "pool: ");
    radius = Convert.ToDouble(Console.ReadLine());

    //See what user wants to compute
    Console.WriteLine("Enter your choice of task.");
    Console.Write("C to compute circumference, " +
        "A to compute area: ");
    taskToDo = Convert.ToChar(Console.ReadLine());
}
```



FIGURE 23

The *doCircumference* Function

```
static void doCircumference(double radius)
//computes and writes out the circumference of
//a circle with given radius
{
    double circumference;           //circumference of a pool -
                                   //computed
    circumference = 2 * pi * radius;
    Console.WriteLine("The circumference for a pool " +
        "of radius " + radius.ToString("##.##") +
        " is " + circumference.ToString("##.##"));
}
```

between passing by value and passing by reference is just a further step in controlling a function's access to data values, to limit the damage the function might do. The code for the *doCircumference* function appears in Figure 23.

The *doArea* function is very similar. Let's reassemble everything and give the complete modularized version of the program. In Figure 24, only the Main function needs to know the value of *more*. No other function needs access to this value, so this variable is never passed as an argument. The Main function header

```
void Main(string[] args)
```

also follows the form for any function header. In other words, the Main function truly is a C# function.

Because it seems to have been a lot of effort to arrive at this complete, modularized version of our SportsWorld program (which, after all, does the same thing as the program in Figure 15), let's review why this effort is worthwhile.

The modularized version of the program is compartmentalized in two ways. First, it is compartmentalized with respect to task. The major task is accomplished by a series of subtasks, and the work for each subtask takes place within a separate function. This leaves the Main function free of details and consisting primarily of invoking the appropriate function at the appropriate point. As an analogy, think of the president of a company calling on various assistants to carry out tasks as needed. The president does not need to know *how* a task is done, only the name of the person responsible for carrying it out. Second, the program is compartmentalized with respect to data, in the sense that the data values known to the various functions are controlled by parameter lists and by the use of value parameters instead of reference parameters where appropriate. In our analogy, the president gives each assistant the information he or she needs to do the assigned task, and expects relevant information to be returned—but not all assistants know all information.

This compartmentalization is useful in many ways. It is useful when we *plan the solution* to a problem, because it allows us to use a divide-and-conquer approach. We can think about the problem in terms of subtasks. This makes it easier for us to understand how to achieve a solution to a large and complex problem. It is also useful when we *code the solution* to a problem, because it allows us to concentrate on writing one section of the code at a time. We can write a function and then fit it into the program, so that the program gradually expands rather than having to be written all at once. Developing a large software project is a team effort, and different parts of the team can be writing different functions at the same time. It is useful when we *test the program*, because we

**FIGURE 24**

*The Complete Modularized
SportsWorld Program*

```
//This program helps SportsWorld estimate costs
//for pool covers and pool fencing by computing
//the area or circumference of a circle
//with a given radius.
//Any number of circles can be processed.
//Written by M. Phelps, 10/23/16

using System;

namespace SportsWorld
{
    class Program
    {
        const double pi = 3.1416; //value of pi

        static void getInput(ref double radius, ref char taskToDo)
        //gets radius and choice of task from the user
        {
            Console.WriteLine();
            Console.Write("Enter the value of the radius of a " +
                "pool: ");
            radius = Convert.ToDouble(Console.ReadLine());

            //See what user wants to compute
            Console.WriteLine("Enter your choice of task.");
            Console.Write("C to compute circumference, " +
                "A to compute area: ");
            taskToDo = Convert.ToChar(Console.ReadLine());
        }

        static void doCircumference(double radius)
        //computes and writes out the circumference of
        //a circle with given radius
        {
            double circumference; //circumference of a pool -
                //computed
            circumference = 2 * pi * radius;
            Console.WriteLine("The circumference for a pool " +
                "of radius " + radius.ToString("##.##") +
                " is " + circumference.ToString("##.##"));
        }

        static void doArea(double radius)
        //computes and writes out the area of
        //a circle with given radius
        {
            double area; //area of a pool -
                //computed
            area = pi * radius * radius;
            Console.WriteLine("The area for a pool " +
                "of radius " + radius.ToString("##.##") +
                " is " + area.ToString("##.##"));
        }

        static void Main(string[] args)
        {
            double radius = 0; //radius of a pool - given

```



FIGURE 24

The Complete Modularized
SportsWorld Program
(continued)

```

char taskToDo= ' ';           //holds user choice to
                              //compute circumference
                              //or area

char more;                    //controls loop for
                              //processing more pools

Console.Write("Do you want to process a pool? (Y or N): ");
more = Convert.ToChar(Console.ReadLine());

while (more == 'Y')          //more circles to process
{
    getInput(ref radius, ref taskToDo);

    if (taskToDo == 'C')     //compute circumference
    {
        doCircumference(radius);
    }
    else                      //compute area
    {
        doArea(radius);
    }
    Console.WriteLine();
    Console.Write("Do you want to process more pools? " +
        "(Y or N): ");
    more = Convert.ToChar(Console.ReadLine());
} //end of while loop

                              //finish up
Console.WriteLine("Program will now terminate.");
}
}
}

```

can test one new function at a time as the program grows, and any errors are localized to the function being added. (The Main function can be tested early by writing appropriate headers but empty bodies for the remaining functions.) Compartmentalization is useful when we *modify the program*, because changes tend to be within certain subtasks and hence within certain functions in the code. And finally it is useful for anyone (including the programmer) who wants to *read* the resulting program. The overall idea of how the program works, without the details, can be gleaned from reading the Main function; if and when the details become important, the appropriate code for the other functions can be consulted. In other words, modularizing a program is useful for its

- Planning
- Coding
- Testing
- Modifying
- Reading

A special type of C# function can be written to compute a single value rather than to carry out a subtask. For example, *doCircumference* does everything connected with the circumference, both calculating the value and writing it

out. We can write a different *doCircumference* function that only computes the value of the circumference and then returns that value to the Main function, which writes it out. A function that returns a single value to the section of the program that invoked it is a **nonvoid** function. Instead of using the word *void* as the return indicator in the function header, a non void function uses the data type of the single returned value. In addition, a nonvoid function must contain a return statement, which consists of the keyword **return** followed by an expression for the value to be returned. (This explains why we have always written the Main function as a void function; it is never invoked anywhere else in the program and does not return a value.)

The code for this new *doCircumference* function would be simply

```
static double doCircumference(double radius)
//computes the circumference of a circle with given
//radius
{
    return 2*pi*radius;
}
```

A nonvoid function is invoked wherever the returned value is to be used, rather than in a separate statement. For example, the statement,

```
Console.Write(doCircumference(radius));
```

invokes the *doCircumference* function by giving its name and argument, and this invocation actually becomes the value returned by the *doCircumference* function, which is then written out.

Figure 25 shows a third version of the SportsWorld program using nonvoid *doCircumference* and *doArea* functions. There are no variables

FIGURE 25

The SportsWorld Program Using Nonvoid Functions

```
//This program helps SportsWorld estimate costs
//for pool covers and pool fencing by computing
//the area or circumference of a circle
//with a given radius.
//Any number of circles can be processed.
//Written by M. Phelps, 10/23/16

using System;

namespace SportsWorld
{
    class Program
    {
        const double pi = 3.1416;        //value of pi

        static void getInput(ref double radius, ref char taskToDo)
        //gets radius and choice of task from the user
        {
            Console.WriteLine();
            Console.Write ("Enter the value of the radius of a " +
                "pool: ");
            radius = Convert.ToDouble(Console.ReadLine());
            //See what user wants to compute
```


**FIGURE 25**

*The SportsWorld Program Using
Nonvoid Functions
(continued)*

```

Console.WriteLine("Enter your choice of task.");
Console.Write("C to compute circumference, " +
              "A to compute area: ");
taskToDo = Convert.ToChar(Console.ReadLine());
}

static double doCircumference(double radius)
//computes the circumference of a circle with given radius
{
    return 2 * pi * radius;
}

static double doArea(double radius)
//computes the area of a circle with given radius
{
    return pi * radius * radius;
}

static void Main(string[] args)
{
    double radius = 0;           //radius of a pool - given
    char taskToDo = ' ';        //holds user choice to
                                //compute circumference
                                //or area
    char more;                  //controls loop for
                                //processing more pools

    Console.Write("Do you want to process a pool? (Y or N): ");
    more = Convert.ToChar(Console.ReadLine());

    while (more == 'Y')        //more circles to process
    {
        getInput(ref radius, ref taskToDo);

        if (taskToDo == 'C')    //compute circumference
        {
            Console.WriteLine("The circumference for a pool " +
                              "of radius " + radius.ToString("###.##") +
                              " is " + doCircumference(radius).ToString("###.##"));
        }
        else                    //compute area
        {
            Console.WriteLine("The area for a pool " +
                              "of radius " + radius.ToString("###.##") +
                              "is " + doArea(radius).ToString("###.##"));
        }
        Console.WriteLine();
        Console.Write("Do you want to process more pools? " +
                     "(Y or N): ");
        more = Convert.ToChar(Console.ReadLine());
    } //end of while loop

                                //finish up
    Console.WriteLine("Program will now terminate.");
}
}
}

```



FIGURE 26

Some C# Terminology

TERM	MEANING	TERM	MEANING
Local variable	Declared and known only within a function	Global constant	Declared outside any function and known everywhere
Argument passed by value	Function receives a copy of the value and can make no permanent changes in the value	Argument passed by reference	Function gets access to memory location where the value is stored; changes it makes to the value persist after control returns to main function
Void function	Performs a task, function invocation is a complete C# statement	Nonvoid function	Computes a value; must include a return statement; function invocation is used within another C# statement

anywhere for the circumference and the area of the circle. The *doCircumference* and *doArea* functions use the usual formulas for their computations, but instead of using local variables for circumference and area, we've compressed the code into a single return statement. These functions are now invoked within an output statement, so the values get printed out without being stored anywhere.

Figure 26 summarizes several sets of terms introduced in this section.

PRACTICE PROBLEMS

1. What is the output of the following C# program fragment?

```
static void doIt(ref int number)
{
    number = number + 4;
}
static void Main(string[] args)
{
    int number;

    number = 7;
    doIt(ref number);
    Console.WriteLine(number);
}
```

2. What is the output of the following C# program fragment?

```
static void doIt(int number)
{
    number = number + 4;
}
static void Main(string[] args)
```

PRACTICE PROBLEMS

```
{
    int number;

    number = 7;
    doIt(number);
    Console.WriteLine(number);
}
```

3. Write a C# function that performs an input task for the Main function, collecting two integer values *one* and *two* from the user.
4. Suppose a nonvoid function called *tax* gets a value *subTotal* from the Main function, multiplies it by a global constant tax rate called *rate*, and returns the resulting tax value. All quantities are type *double*.
 - a. Write the function header.
 - b. Write the return statement in the function body.
 - c. Write the statement in the Main function that writes out the tax.

6 Object-Oriented Programming

▶ 6.1 What Is It?

The divide-and-conquer approach to programming is a “traditional” approach. The focus is on the overall task to be done: how to break it down into subtasks, and how to write algorithms for these subtasks that are carried out by communicating modules (in the case of C#, by functions). The program can be thought of as a giant statement executor designed to carry out the major task, even though the Main function may simply call on, in turn, the various other modules that do the subtask work. **Object-oriented programming (OOP)** takes a somewhat different approach. A program is considered a simulation of some part of the world that is the domain of interest. “Objects” populate this domain. Objects in a banking system, for example, might be savings accounts, checking accounts, and loans. Objects in a company personnel system might be employees. Objects in a medical office might be patients and doctors. Each object is an example drawn from a class of similar objects. The savings account “class” in a bank has certain properties associated with it, such as name, Social Security number, account type, and account balance. Each individual savings account at the bank is an example of (an object of) the savings account class, and each has specific values for these common properties; that is, each savings account has a specific value for the name of the account holder, a specific value for the account balance, and so forth. Each object of a class therefore has its own data values.

A class also has one or more subtasks associated with it, and all objects from that class can perform those subtasks. In carrying out a subtask, each object can be thought of as providing some service. A savings account, for example, can compute compound interest due on the balance. When an object-oriented program is executed, the program generates requests for services that go to the various objects. The objects respond by performing the requested service—that is, carrying out the subtask. Thus a program that is using the savings account class might request a particular savings account object to perform the service of computing interest due on the account balance. An object always knows its own data values and may use them in performing the requested service.

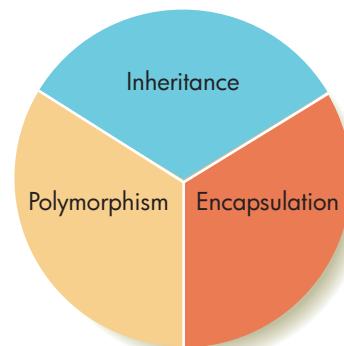
Some of this sounds familiar. We know about subtasks (functions) associated with a class. The new idea is that instead of directly asking a class to carry out a subtask, we ask an object of that class to carry out a subtask. The even bigger new idea is that such objects have data values for the class properties. Instead of storing data in variables that are available to the whole program and then passing them as arguments to subtasks, the program can simply ask an object to use its own data when it carries out a subtask.

There are three terms often associated with object-oriented programming, as illustrated in Figure 27. The first term is **encapsulation**. Each class has its own program module to perform each of its subtasks. Any user of the class (which might be some other program) can ask an object of that class to invoke the appropriate module and thereby perform the subtask service. The class user needs to know what services objects of the class can provide and how to request an object to perform any such service. The details of the module code belong to the class itself, and this code may be modified in any manner, as long as the way the user interacts with the class remains unchanged. (In the savings account example, the details of the algorithm used to compute interest due belong only to the class, and need not be known by any user of the class. If the bank wants to change how it computes interest, only the code for the interest module in the savings account class needs to be modified; any programs that use the services of the savings account class can remain unchanged.) Furthermore, the class properties represent data values that will exist as part of each object of the class. A class therefore consists of two components, its subtask modules and its properties, and both components are “encapsulated”—bundled—with the class.

A second term associated with object-oriented programming is **inheritance**. Once a class A of objects is defined, a class B of objects can be defined as a “subclass” of A. Every object of class B is also an object of class

FIGURE 27

Three Key Elements of OOP



A; this is sometimes called an “is a” relationship. Objects in the B class will “inherit” all of the properties and be able to perform all the services of objects in A, but they may also be given some special property or ability. The benefit is that class B does not have to be built from the ground up, but rather can take advantage of the fact that class A already exists. In the banking example, a senior citizens savings account would be a subclass of the savings account class. Any senior citizens savings account object is also a savings account object, but may have special properties or be able to provide special services.

The third term is **polymorphism**. *Poly* means “many.” Objects of different classes may provide services that should logically have the same name because they do roughly the same thing, but the details differ. In the banking example, both savings account objects and checking account objects should provide a “compute interest” service, but the details of how interest is computed differ in these two cases. Thus, one name, the name of the service to be performed, has several meanings, depending on the class of the object providing the service. It may even be the case that more than one service with the same name exists for the same class, although there must be some way to tell which service is meant when it is invoked by an object of that class.

Let’s change analogies from the banking world to something more fanciful, and consider a football team. Every member of the team’s backfield is an “object” of the “backfield” class. The quarterback is the only “object” of the “quarterback” class. Each backfield object can perform the service of carrying the ball if he (or she) receives the ball from the quarterback; ball carrying is a subtask of the backfield class. The quarterback who hands the ball off to a backfield object is requesting that the backfield object perform that subtask because it is “public knowledge” that the backfield class carries the ball and that this service is invoked by handing off the ball to a backfield object. The “program” to carry out this subtask is *encapsulated* within the backfield class, in the sense that it may have evolved over the week’s practice and may depend on specific knowledge of the opposing team, but at any rate, its details need not be known to other players. *Inheritance* can be illustrated by the halfback subclass within the backfield class. A halfback object can do everything a backfield object can but may also be a pass receiver. And *polymorphism* can be illustrated by the fact that the backfield may invoke a different “program” depending on where on the field the ball is handed off. Of course our analogy is imperfect, because not all human “objects” from the same class behave in precisely the same way—fullbacks sometimes receive passes and so on.

6.2 C# and OOP

How do these ideas get translated into real programs? The details, of course, vary with the programming language used, and not every language supports object-oriented programming. C# is very much an object-oriented programming language. Virtually all C# code is inside some class. From Figure 3, you

can see that the structure for the *TravelPlanner* program is dominated by the definition of the *TravelPlanner* class, which contains only the *Main* function:

```

9. class TravelPlanner
10. {
    //Main function
29. }

```

When we write a class, we specify the properties (called **member variables**) common to any object of that class. We also specify the services (called **member functions**) that any object of that class can perform. Let's rewrite the *SportsWorld* program one more time, this time using a more object-oriented approach. What are the objects of interest within the scope of this problem? *SportsWorld* deals with circular swimming pools, but they are basically just circles. So let's create a *Circle* class, and have the *SportsWorld* program create objects of (**instances** of) that class. The objects are individual circles. A *Circle* object has a radius. A *Circle* object, which knows the value of its own radius, should be able to perform the services of computing its own circumference and its own area. To do this, however, the *Circle* object also needs to know the value of pi and, because everything in C# must belong to a class, we'll also make pi a "property" of every *Circle* object. At this point, we are well on the way to answering the two major questions about our *Circle* class:

- What are the properties common to any object of this class? (In this case, the radius and the value of pi.)
- What are the services that any object of the class should be able to perform? (In this case, compute its circumference and compute its area, although as we will see shortly, we will need other services as well.)

Figure 28 shows the complete object-oriented version of *SportsWorld*, with its two classes. The class *Circle* has member variables *pi* and *radius* (well, OK, one is a constant). It has member functions *setRadius*, *getRadius*, *doCircumference*, and *doArea*. The first member function is a void function, and the remaining three return values. The class *Program* has two member functions, *getInput* and *Main*.

FIGURE 28

An Object-Oriented SportsWorld Program

```

//This program helps SportsWorld estimate costs
//for pool covers and pool fencing by computing
//the area or circumference of a circle
//with a given radius.
//Any number of circles can be processed.
//Uses class Circle
//Written by M. Phelps, 10/23/16

using System;

namespace SportsWorld
{
    class Circle

```


**FIGURE 28**

An Object-Oriented
SportsWorld Program
(continued)

```

{
    private const double pi = 3.1416;
    private double radius;

    public void setRadius(double value)
    //sets radius equal to value
    {
        radius = value;
    }

    public double getRadius()
    //returns current radius
    {
        return radius;
    }

    public double doCircumference()
    //computes and returns circumference of a circle
    {
        return 2 * pi * radius;
    }

    public double doArea()
    //computes and returns area of a circle
    {
        return pi * radius * radius;
    }
}

class Program
{
    static void getInput(ref double newRadius, ref char taskToDo)
    //gets radius and choice of task from the user
    {
        Console.WriteLine();
        Console.Write("Enter the value of the radius of a " +
            "pool: ");
        newRadius = Convert.ToDouble(Console.ReadLine());

        //See what user wants to compute
        Console.WriteLine("Enter your choice of task.");
        Console.Write("C to compute circumference, " +
            "A to compute area: ");
        taskToDo = Convert.ToChar(Console.ReadLine());
    }

    static void Main(string[] args)
    {
        double newRadius = 0;           //radius of a pool - given
        char taskToDo = ' ';           //holds user choice to
                                        //compute circumference
                                        //or area
        char more;                       //controls loop for
                                        //processing more pools
        Circle swimmingPool = new Circle();
    }
}

```



FIGURE 28

An Object-Oriented
SportsWorld Program
(continued)

```

Console.WriteLine("Do you want to process a pool? (Y or N): ");
more = Convert.ToChar(Console.ReadLine());

while (more == 'Y')           //more circles to process
{
    getInput(ref newRadius, ref taskToDo);
    swimmingPool.setRadius(newRadius);

    if (taskToDo == 'C')      //compute circumference
    {
        Console.WriteLine("The circumference for a pool " +
            "of radius " +
            swimmingPool.getRadius().ToString("###.##") +
            " is " +
            swimmingPool.doCircumference().ToString(
                "###.##"));
    }
    else                       //compute area
    {
        Console.WriteLine("The area for a pool " +
            "of radius " +
            swimmingPool.getRadius().ToString("###.##") +
            " is " +
            swimmingPool.doArea().ToString("###.##"));
    }
    Console.WriteLine();
    Console.Write("Do you want to process more pools? " +
        "(Y or N): ");
    more = Convert.ToChar(Console.ReadLine());
} //end of while loop

//finish up
Console.WriteLine("Program will now terminate.");
}
}
}

```

The question is, if these are two separate and distinct classes, how do they “talk to each other”? The answer is, if a class wants to make use of another class, it must “instantiate” (declare) an object from that class. Look at this line of code from *Main*:

```
Circle swimmingPool = new Circle();
```

This instantiates an object from the *Circle* class and calls it *swimmingPool*. Now, *Main* can ask the *swimmingPool* object to invoke methods from its class. Consider the statement

```

Console.WriteLine("The circumference for a pool " +
    "of radius " +
    swimmingPool.getRadius().ToString("###.##") +
    " is " +
    swimmingPool.doCircumference().ToString(
        "###.##"));

```

This statement asks the *swimmingPool* object first to invoke the *getRadius* function, and later to invoke the *doCircumference* function, both member functions of the *Circle* class of which *swimmingPool* is an instance.

Although we glossed over this point earlier, in the two modularized versions of the *SportsWorld* program the *doCircumference* and *doArea* functions were declared as static functions, as in

```
static void doCircumference(double radius)
```

and

```
static double doCircumference(double radius)
```

A **static function** doesn't need to be invoked by an object; it can be invoked by just giving the name of the function with an appropriate list of arguments. (The runtime system invokes the static *Main* function when the program is executed.) But now that *doCircumference* and *doArea* are member functions of the *Circle* class, and are not declared static, an object of the *Circle* class must invoke them.

The member functions of the *Circle* class are all declared using the keyword **public**. Public functions can be used anywhere, including within the *Main* function and indeed in any C# program that wants to make use of this class. Think of the *Circle* class as handing out a business card that advertises these services: Hey, you want a *Circle* object that can find its own area? Find its own circumference? Set the value of its own radius? I'm your class! (Class member functions can also be **private**, but a private member function is a sort of helping task that can be used only within the class in which it occurs.)

The *radius* member variable of the class is declared using the keyword **private** (as is the constant *pi*). Only functions in the *Circle* class itself can use this variable. Note that *doCircumference* and *doArea* have no parameter for the value of the radius; as methods of this class, they know at all times the current value of *radius* for the object that invoked them, and it does not have to be passed to them as an argument. Because *radius* has been declared private, however, the *Main* function cannot use the value of *radius*. It cannot write out that value or directly change that value by some assignment statement. It can, however, request a *Circle* object to invoke the *getRadius* member function to return the current value of the radius in order to write it out. It can also request a *Circle* object to invoke the *setRadius* member function to change the value of its radius; *setRadius* does have a parameter to receive a new value for *radius*. Member variables are generally declared private instead of public, to protect the data in an object from reckless changes some application program might try to make. Changes in the values of member variables should be performed only under the control of class objects through functions such as *setRadius*.

The *Main* function, as before, handles all of the user interaction and now makes use of the *Circle* class. It creates a *Circle* object, an instance of the *Circle* class, by means of the following statement:

```
Circle swimmingPool = new Circle();
```

This begins like an ordinary variable declaration such as

```
int number;
```

but includes the syntax to create a “new” instance of `Circle`. Why isn’t there a need for “new” in the line of code just above? Variables of simple types, e.g., integers, are instantiated for you by the runtime system. (Recall, “new” did appear in the declaration for an array given earlier—arrays are not simple types.)

After

```
Circle swimmingPool = new Circle();
```

the object *swimmingPool* exists, and the `Main` function can ask *swimmingPool* to perform the various services of which instances of the *Circle* class are capable.

The syntax to request an object to invoke a member function is to give the name of the object, followed by a dot, followed by the name of the member function, followed by any arguments the function may need.

```
object-identifier.function-identifier(argument list)
```

The object that invokes a method is the **calling object**. Therefore the expression

```
swimmingPool.doCircumference()
```

in the `Main` function uses *swimmingPool* as the calling object to invoke the *doCircumference* method of the *Circle* class. No arguments are needed because this method has no parameters, but the empty parentheses must be present.

Looking at the code for the member functions in Figure 28, we see that the *setRadius* member function uses an assignment statement to change the value of *radius* to whatever quantity is passed to the parameter *value*. The *getRadius* function body is a single return statement. The *doCircumference* and *doArea* functions again consist of single statements that compute and return the proper value.

There is no declaration in the `Main` function for a variable called *radius*. There is a declaration for *newRadius*, and *newRadius* receives the value entered by the user for the radius of the circle. Therefore, isn’t *newRadius* serving the same purpose as *radius* did in the old program? No—this is rather subtle, so pay close attention: While *newRadius* holds the number the user wants for the circle radius, it is not itself the radius of *swimmingPool*. The radius of *swimmingPool* is the member variable *radius*, and only methods of the class can change the member variables of an object of that class. The *Circle* class provides the *setRadius* member function for this purpose. The `Main` function must ask the object *swimmingPool* to invoke *setRadius* to set the value of its radius equal to the value contained in *newRadius*. The *newRadius* argument corresponds to the *value* parameter in the *setRadius* function, which then gets assigned to the member variable *radius*.

```

swimmingPool.setRadius(newRadius);

public void setRadius(double value)
//sets radius equal to value
{
    radius = value;
}

```

The *setRadius* function is a void function because it returns no information to the invoking statement; it contains no return statement. The invocation of this method is a complete C# statement.

Finally, the output statements in *Main* that print the values of the circumference and area also have *swimmingPool* invoke the *getRadius* member function to return its current *radius* value so it can be printed as part of the output. We could have used the variable *newRadius* here instead. However, *newRadius* is what we THINK has been used in the computation, whereas *radius* is what has REALLY been used.

In C#, the functionality of *setRadius* and *getRadius* can be replaced by something called a **property**, although one might wish it were called a “property manager” instead. Consider the following code for the declaration of the property named *radiusValue*.

```

public double radiusValue
{
    get { return radius; }
    set { radius = value; }
}

```

Code to set the value for the radius in *Main* using the property above would be

```
swimmingPool.radiusValue = newRadius;
```

Notice the similarity to an ordinary assignment statement.

The “get” portion of the property is used in the output statements, e.g.:

```

Console.WriteLine("The circumference for a pool" +
    " of radius " +
    swimmingPool.radiusValue.ToString("##.##") +
    " is " +
    swimmingPool.doCircumference().ToString("##.##"));

```

6.3 One More Example

The object-oriented version of our SportsWorld program illustrates encapsulation. All data and calculations concerning circles are encapsulated in the *Circle* class. Let’s look at one final example that illustrates the other two watchwords of OOP—polymorphism and inheritance.

In Figure 29, the domain of interest is that of geometric shapes. Four different classes are given: *Circle*, *Rectangle*, *Square*, and *Square2*. Each class consists of a public part and a private or protected part. The public part provides, in the form of C# functions, the services or subtasks that an object from

the class can perform. The private or protected part describes the properties that any object of the class possesses. A Circle object, as before, has a radius property and a property for the value of pi, whereas a Rectangle object has a width property and a height property. Any Circle object can set the value of its radius and can compute its area. A Square object has a side property, as one might expect, but a Square2 object doesn't seem to have any properties or, for that matter, any way to compute its area. We will explain the difference between the *Square* class and the *Square2* class shortly.

**FIGURE 29**

A C# Program with
Polymorphism and Inheritance

```
using System;

namespace Shapes
{
    public class Circle
    {
        const double pi = 3.1416;
        private double radius;

        public void setRadius(double value)
        { //sets radius of the circle equal to value
            radius = value;
        }

        public double getRadius()
        { //return radius
            return radius;
        }

        public double doArea()
        { //computes and writes out area of circle
            return pi * radius * radius;
        }
    }

    public class Rectangle
    {
        protected double width;
        protected double height;

        public void setWidth(double value)
        { //sets width of rectangle equal to value
            width = value;
        }

        public void setHeight(double value)
        { //sets height of rectangle equal to value
            height = value;
        }

        public double getWidth()
        { //returns width
            return width;
        }
    }
}
```


**FIGURE 29**

*A C# Program with
Polymorphism and Inheritance
(continued)*

```

public double getHeight()
{ //returns height
  return height;
}

public double doArea()
{ //computes and writes out area of rectangle
  return width * height;
}
}

public class Square
{
  private double side;

  public void setSide(double value)
  { //sets the side of the square equal to value
    side = value;
  }

  public double getSide()
  { //returns side
    return side;
  }

  public double doArea()
  { //computes and writes out the area of the square
    return side * side;
  }
}

public class Square2 : Rectangle
//Square is derived class of Rectangle,
//uses the inherited height and width
//properties and the inherited doArea function
{
  public void setSide(double value)
  { //sets the side of the square equal to value
    height = value;
    width = value;
  }

  public double getSide()
  { //returns side (width)
    return width;
  }
}

class Program
{
  static void Main(string[] args)
  {
    Circle joe = new Circle();
    joe.setRadius(23.5);
  }
}

```



FIGURE 29

A C# Program with
Polymorphism and Inheritance
(continued)

```

Console.WriteLine("The area of a circle with "
+ "radius " + joe.getRadius() + " is "
+ joe.doArea());

Rectangle luis = new Rectangle();
luis.setWidth(12.4);
luis.setHeight(18.1);
Console.WriteLine("The area of a rectangle with "
+ "dimensions " + luis.getWidth() + " and "
+ luis.getHeight() + " is " + luis.doArea());

Square anastasia = new Square();
anastasia.setSide(3);
Console.WriteLine("The area of a square with "
+ "side " + anastasia.getSide() + " is "
+ anastasia.doArea());

Square2 tyler = new Square2();
tyler.setSide(4.2);
Console.WriteLine("The area of a square with "
+ "side " + tyler.getSide() + " is "
+ tyler.doArea());
}
}
}

```

The Main function uses these classes. It creates objects from the various classes. After each object is created, the Main function requests the object to set its dimensions, using the values given, and to compute its area as part of an output statement giving information about the object. For example, the statement

```
joe.setRadius(23.5);
```

instructs the circle named *joe* to invoke the *setRadius* function of *joe's* class, thereby setting *joe's* radius to 23.5. Figure 30 shows the output after the program in Figure 29 is run.

Here we see polymorphism at work, because there are lots of *doArea* functions; when the program executes, the correct function is used, on the basis of the class to which the object invoking the function belongs. After all, computing the area of a circle is quite different from computing the area of a rectangle. The algorithms themselves are straightforward; they employ assignment statements to set the dimensions and the usual formulas to compute the area of a circle, rectangle, and square. The functions can use the properties of the objects that invoke them without having the values of those properties passed as arguments.



FIGURE 30

Output from the Program of
Figure 29

```

The area of a circle with radius 23.5 is 1734.9486
The area of a rectangle with dimensions 12.4 and 18.1 is 224.44
The area of a square with side 3 is 9
The area of a square with side 4.2 is 17.64

```

Square is a stand-alone class with a *side* property and a *doArea* function. The *Square2* class, however, recognizes the fact that squares are special kinds of rectangles. The *Square2* class is a subclass of the *Rectangle* class, as is indicated by the reference to *Rectangle* in the class description for *Square2*. It inherits the *width* and *height* properties from the “parent” *Rectangle* class; the “protected” status of these properties in the *Rectangle* class indicates that they can be extended to any subclass. *Square2* also inherits the *setWidth*, *setHeight*, and *doArea* functions. In addition, *Square2* has its own function, *setSide*, because setting the value of the “side” makes sense for a square but not for an arbitrary rectangle. What the user of the *Square2* class doesn’t know is that there really isn’t a “side” property; the *setSide* function merely sets the inherited *width* and *height* properties to the same value. To compute the area, then, the *doArea* function inherited from the *Rectangle* class can be used, and there’s no need to redefine it or even to copy the existing code. Here we see inheritance at work.

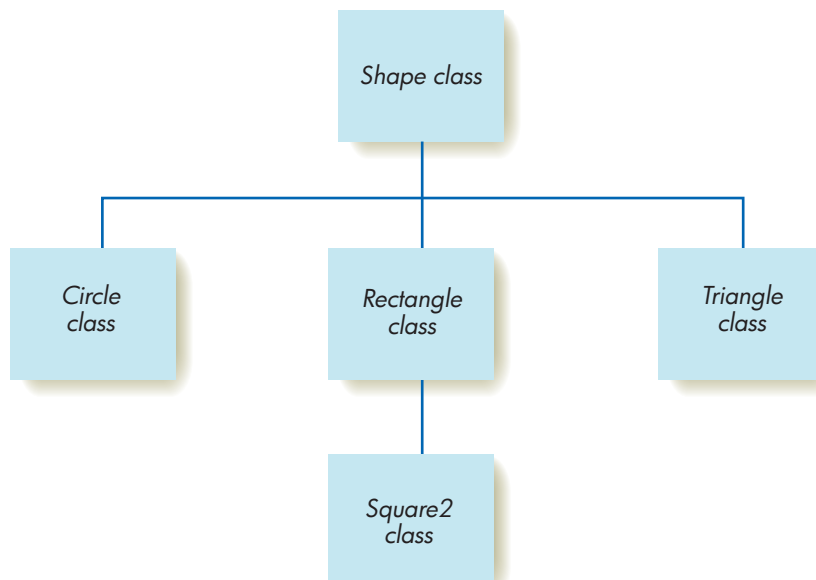
Inheritance can be carried through multiple “generations.” We might redesign the program so that there is one “superclass” that is a general *Shape* class, of which *Circle* and *Rectangle* are subclasses, *Square2* being a subclass of *Rectangle* (see Figure 31 for a possible class hierarchy).

Although the program of Figure 29 can be kept in one file, it can also be split into separate files. Each of the classes could be in its own file with a distinct namespace name. To gain access to the class, a “using” statement with the appropriate namespace name would need to be added to the main program file. This is just what happens with the “using *System*” directive to gain access to items like *Console.WriteLine()*.

▶ 6.4 What Have We Gained?

Now that we have some idea of the flavor of object-oriented programming, we should ask what we gain by this approach. There are two major reasons why OOP is a popular way to program:

FIGURE 31
A Hierarchy of Geometric Classes



- Software reuse
- A more natural “worldview”

SOFTWARE REUSE. Manufacturing productivity took a great leap forward when Henry Ford invented the assembly line. Automobiles could be assembled using identical parts so that each car did not have to be treated as a unique creation. Computer scientists are striving to make software development more of an assembly-line operation and less of a handcrafted, start-over-each-time process. Object-oriented programming is a step toward this goal: A useful class that has been implemented and tested becomes a component available for use in future software development. Anyone who wants to write an application program involving circles, for example, can use the already written, tried, and tested *Circle* class. As the “parts list” (the class library) grows, it becomes easier and easier to find a “part” that fits, and less and less time has to be devoted to writing original code. If the class doesn’t quite fit, perhaps it can be modified to fit by creating a subclass; this is still less work than starting from scratch. Software reuse implies more than just faster code generation. It also means improvements in *reliability*; these classes have already been tested, and if properly used, they will work correctly. And it means improvements in *maintainability*. Thanks to the encapsulation property of object-oriented programming, changes can be made in class implementations without affecting other code, although such change requires retesting the classes.

A MORE NATURAL “WORLDVIEW.” The traditional view of programming is procedure-oriented, with a focus on tasks, subtasks, and algorithms. But wait—didn’t we talk about subtasks in OOP? Haven’t we said that computer science is all about algorithms? Does OOP abandon these ideas? Not at all. It is more a question of *when* these ideas come into play. Object-oriented programming recognizes that in the “real world,” tasks are done by entities (objects). Object-oriented program design begins by identifying those objects that are important in the domain of the program because their actions contribute to the mix of activities present in the banking enterprise, the medical office, or wherever. Then it is determined what data should be associated with each object and what subtasks the object contributes to this mix. Finally, an algorithm to carry out each subtask must be designed. We saw in the modularized version of the SportsWorld program in Figure 24 how the overall algorithm could be broken down into pieces that are isolated within functions. Object-oriented programming repackages those functions by encapsulating them within the appropriate class of objects.

Object-oriented programming is an approach that allows the programmer to come closer to modeling or simulating the world as we see it, rather than to mimic the sequential actions of the Von Neumann machine. It provides another buffer between the real world and the machine, another level of abstraction in which the programmer can create a virtual problem solution that is ultimately translated into electronic signals on hardware circuitry.

Finally, we should mention that a graphical user interface, with its windows, icons, buttons, and menus, is an example of object-oriented programming at work. A general button class, for example, can have properties of height, width, location on the screen, text that may appear on the button, and so forth. Each individual button object has specific values for those properties. The button class can perform certain services by responding to messages, which are generated by events (for example, the

user clicking the mouse on a button triggers a “mouseclick” event). Each particular button object individualizes the code to respond to these messages in unique ways. We will not go into details of how to develop graphical user interfaces in C#, but in the next section you will see a bit of the programming mechanics that can be used to draw the graphics items that make up a visual interface.

PRACTICE PROBLEMS

1. What is the output from the following section of code if it is added to the Main function of the C# program in Figure 29?

```
Square one = new Square();
one.setSide(10);
Console.WriteLine("The area of a square with "
    + "side " + one.getSide() + " is "
    + one.doArea());
```

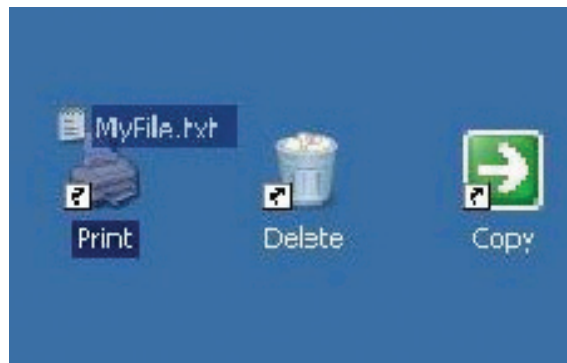
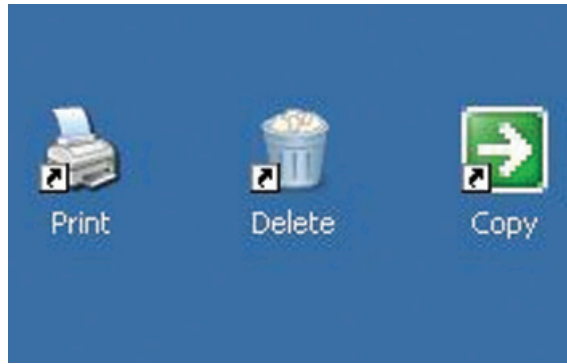
2. In the hierarchy of Figure 31, suppose that the *Triangle* class is able to perform a *doArea* function. What two properties should any triangle object have?

7 Graphical Programming

The programs that we have looked at so far all produce *text output*—output composed of the characters {A . . .Z, a . . .z, 0 . . .9} along with a few punctuation marks. For the first 30 to 35 years of software development, text was virtually the only method of displaying results in human-readable form, and in those early days it was quite common for programs to produce huge stacks of alphanumeric output. These days an alternative form of output—*graphics*—has become much more widely used. With graphics, we are no longer limited to 100 or so printable characters; instead, programmers are free to construct whatever shapes and images they desire.

The intelligent and well-planned use of graphical output can produce some phenomenal improvements in software. We discussed this issue in Chapter 6, where we described the move away from the text-oriented operating systems of the 1970s and 1980s, such as MS-DOS and VMS, to operating systems with more powerful and user-friendly graphical user interfaces (GUIs), such as Windows 7, Windows 8, and Mac OS X. Instead of requiring users to learn dozens of complex text-oriented commands for such things as copying, editing, deleting, moving, and printing files, GUIs can present users with simple and easy-to-understand visual metaphors for these operations. In the first image on the next page, the operating system presents the user with icons for printing, deleting, or copying a file. In the second image on the next page, dragging a file to the printer icon prints the file.

Not only does graphics make it easier to manage the tasks of the operating system, it can help us visualize and make sense of massive amounts of output produced by programs that model complex physical, social, and mathematical systems. (We discuss modeling and visualization in Chapter 13.) Finally, there are many applications of computers that would simply be impossible without the



ability to display output visually. Applications such as virtual reality, computer-aided design/computer-aided manufacturing (CAD/CAM), games and entertainment, medical imaging, and computer mapping would not be anywhere near as important as they are without the enormous improvements that have occurred in the areas of graphics and visualization.

So, we know that graphical programming is important. The question is: What features must be added to a programming language like C# to produce graphical output?

▶ 7.1 Graphics Hardware

Modern computer terminals use what is called a **bitmapped display**, in which the screen is made up of thousands of individual picture elements, or **pixels**, laid out in a two-dimensional grid. These are the same pixels used in visual images, as discussed in Chapter 4. In fact, the display is simply one large visual image. The number of pixels on the screen varies from system to system; typical values range from 800×600 up to 1560×1280 . Terminals with a high density of pixels are called **high-resolution** terminals. The higher the resolution—that is, the more pixels available in a given amount of space—the sharper the visual image because each individual pixel is smaller. However, if the screen size itself is small, then a high resolution image can be too tiny to read. A 30" wide-screen monitor might support a resolution of 2560×1600 , but that would not be suitable for a laptop screen. In Chapter 4 you learned that a color display

requires 24 bits per pixel, with 8 bits used to represent the value of each of the three colors red, green, and blue. The memory that stores the actual screen image is called a **frame buffer**. A high-resolution color display might need a frame buffer with (1560×1280) pixels \times 24 bits/pixel = 47,923,000 bits, or about 6 MB, of memory for a single image. (One of the problems with graphics is that it requires many times the amount of memory needed for storing text.)

The individual pixels in the display are addressed using a two-dimensional coordinate grid system, the pixel in the upper-left corner being (0, 0). The overall pixel-numbering system is summarized in Figure 32. The specific values for *maxX* and *maxY* in Figure 32 are, as mentioned earlier, system-dependent. (Note that this coordinate system is not the usual mathematical one. Here, the origin is in the upper-left corner, and *y* values are measured downward.)

The terminal hardware displays on the screen the frame buffer value of every individual pixel. For example, if the frame buffer value on a color monitor for position (24, 47) is RGB (0, 0, 0), the hardware sets the color of the pixel located at column 24, row 47 to black, as shown in Figure 33. The operation diagramed in Figure 33 must be repeated for all of the 500,000 to 2 million pixels on the screen. However, the setting of a pixel is not permanent; on the contrary, its color and intensity fade quickly. Therefore, each pixel must be “repainted” often enough so that our eyes do not detect any “flicker,” or change in intensity. This requires the screen to be completely updated, or refreshed, 30–50 times per second. By setting various sequences of pixels to different colors, the user can have the screen display any desired shape or image. This is the fundamental way in which graphical output is achieved.

▶ 7.2 Graphics Software

To control the setting and clearing of pixels, programmers use a collection of functions that are part of a special software package called a **graphics library**. Virtually all modern programming languages, including C#, come with


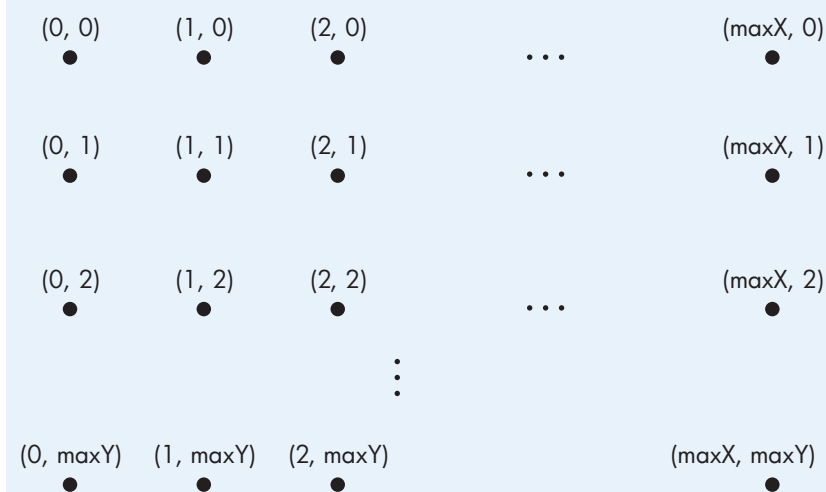



FIGURE 32

Pixel-Numbering System in a Bitmapped Display




**FIGURE 33**

Display of Information on the Terminal




an extensive and powerful graphics library for creating different shapes and images. Typically, an “industrial strength” graphics library includes hundreds of functions for everything from drawing simple geometric shapes like lines and circles, to creating and selecting colors, to more complex operations such as displaying scrolling windows, pull-down menus, and buttons. We restrict our discussion to a modest set of functions. Although the set is unrealistically small, it will still give you a good idea of what visual programming is like, and enable you to produce some interesting, nontrivial images on the screen.

To create a C# graphics program in Visual Studio 2013, you must start with a Windows Form Application Project (as opposed to a Console Application project). Then the C# IDE will automatically create a window, called Form 1. This window will hold a “drawing canvas” on which the actual graphics will appear. There is a lot of code generated by the IDE to produce Form 1, but again, it’s all done automatically by virtue of the fact that the project is a windows form application project. Form 1 shows up in the Form1.cs[Design] view in the C# IDE; it looks like this:



To see some of the code that has been generated, click on the < > symbol in the Solution Explorer window of the IDE



This brings up the Form1.cs code, which is

```
using System;
using System.Collections.Generic;
using System.ComponentModel;
using System.Data;
using System.Drawing;
using System.Linq;
using System.Text;
using System.Threading.Tasks;
using System.Windows.Forms;

namespace GraphicsProgram
{
    public partial class Form1 : Form
    {
        public Form1()
        {
            InitializeComponent();
        }
    }
}
```

where the namespace reference is the project name.

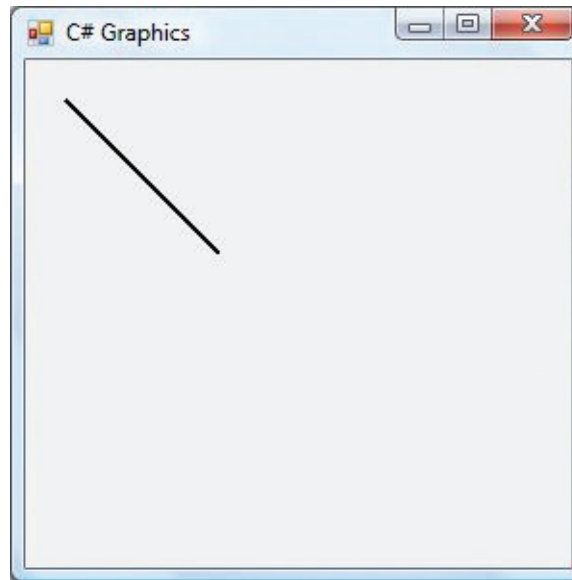
If you run the program with just the above basic code, you'll see the Form1 window. Now, where do we add the graphics code? It goes in the Form1.cs file, in the Form1 function, right below the InitializeComponent function invocation. From now on, we'll just show the public Form1() function.

```
public Form1()
{
    InitializeComponent();
    //code to actually draw
    //on the form goes here
}
```

The code to draw a line from point (20,20) to point (100,100) is shown next. Some additional comments have been added, and, as shown in the result, the text appearing on the Form1 window title bar has been changed.

```
public Form1()
{ // form on which drawing will occur
    InitializeComponent(); // initialize form
    this.Text = "C# Graphics";
    Graphics drawingCanvas = CreateGraphics();
    Show();
    Pen BlackPen = new Pen(Color.Black, 2);
    drawingCanvas.DrawLine
        (BlackPen, 20, 20, 100, 100);
}
```

The result of executing the program containing this code is:



How does this code work? The first statement invokes the `InitializeComponent()` function. This function, part of the automatically generated code, actually generates the form. The second statement changes the text that appears in the title bar of the form window.

The next line of code in `public Form1()` produces the drawing canvas on which to draw, actually just the surface of the form.

```
Graphics drawingCanvas = CreateGraphics();
```

The statement

```
Show();
```

causes the form to be displayed on the screen.

The statement

```
Pen BlackPen = new Pen(Color.Black, 2);
```

creates the tool, in this case a pen, that will “paint” the lines on the screen. It also sets the pen color and the line thickness. Brushes can also be used as drawing tools.

The final line in the code actually draws the line on the screen.

```
drawingCanvas.DrawLine(BlackPen, 20, 20, 100, 100);
```

The parameter definitions for `DrawLine` are

```
DrawLine(Pen, x1, y1, x2, y2)
```

where (x_1, y_1) are the coordinates (in pixels) of the start of the line, and (x_2, y_2) are the coordinates of the end of the line.


The code to draw a rectangle touching the four points (25, 60), (75, 60), (25,100), and (75,100) is:

```
public Form1()
{ // form on which drawing will occur
  InitializeComponent(); // initialize form
  this.Text = "C# Graphics";
  Graphics drawingCanvas = CreateGraphics();
  Show();
  Pen BlackPen = new Pen(Color.Black, 2);
  drawingCanvas.DrawRectangle
    (BlackPen, 25, 60, 50, 40);
}
```

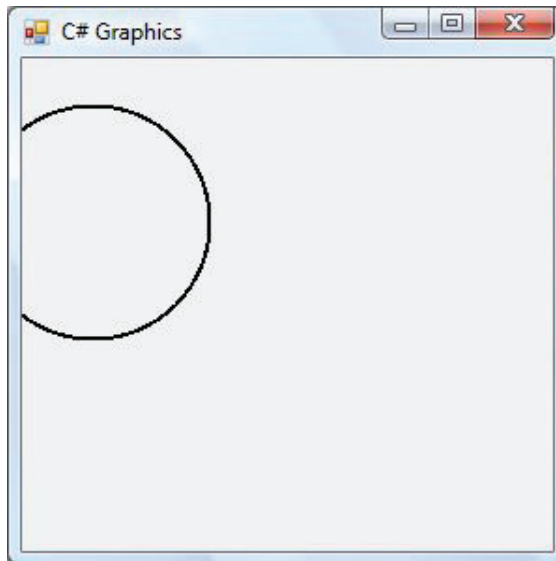
Note that the parameters for *DrawRectangle* are

```
DrawRectangle(Pen, x, y, width, height)
```

where (x, y) are the coordinates of the upper-left corner. The result is



The *DrawEllipse* method is used to draw a circle. Here is the result of drawing a circle with radius 125 pixels centered at the point (100, 150):



The circle (ellipse) is specified by giving the size of its bounding rectangle. The parameters are:

```
DrawEllipse(Pen, x, y, width, height)
```

where (x, y) are the coordinates of the upper-left corner of the bounding rectangle. In the image above, this corner will have to be “off the screen.” That’s not a problem. The code is shown below (note the negative value for the x coordinate). Also note that width and height of the bounding rectangle must be equal to get a circular shape rather than an oval.


```
public Form1()
{ // form on which drawing will occur
  InitializeComponent(); // initialize form
  this.Text = "C# Graphics";
  Graphics drawingCanvas = CreateGraphics();
  Show();
  Pen BlackPen = new Pen(Color.Black, 2);
  drawingCanvas.DrawEllipse
    (BlackPen, -25, 25, 125, 125);
}
```

How can we get text annotations on the screen? There is a method named *DrawString* that will do the job. The parameter list is

```
DrawString(string, font, brush, x, y)
```

The *string* is the string to be output, using the *font* specified, drawing with the *brush* specified, beginning at pixel position (x, y) for the upper-left corner of the bounding rectangle for the text.

Here is an example with the text “Stop” drawn inside of a circle.



The code to produce this image is

```
public Form1()
{ // form on which drawing will occur
  InitializeComponent(); // initialize form
  this.Text = "C# Graphics";
  Graphics drawingCanvas = CreateGraphics();
  Show();
  drawingCanvas.DrawEllipse
    (BlackPen, 160, 160, 40, 40);
  Brush redBrush = Brushes.Red;
  Font f = new Font("Arial", 10);
  drawingCanvas.DrawString
    ("Stop", f, redBrush, 160, 170);
}
```

In summary, we have the following graphics functions at our disposal.

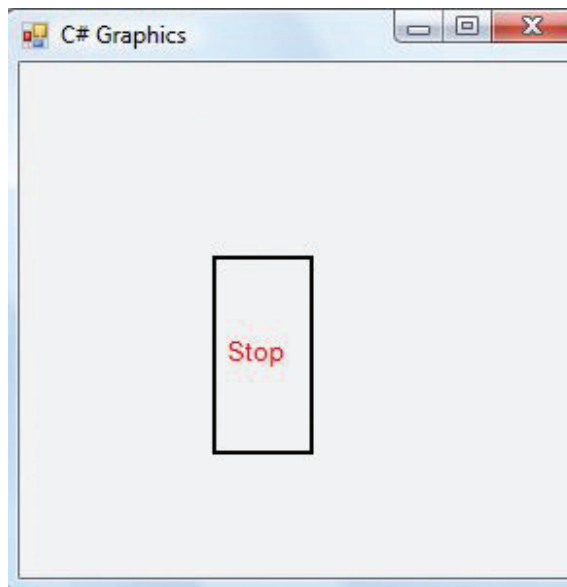
```
DrawLine(Pen, x1, y1, x2, y2)
DrawRectangle(Pen, x, y, width, height)
DrawEllipse(Pen, x, y, width, height)
DrawString(string, font, brush, x, y)
```

Now that forms and graphics are available, we seem close to producing elements of a typical GUI. Can we draw a button that acts like a button on a real GUI form—that is, can we write code to sense a mouse click on that button and respond with some action?

Here is code to draw a “Stop” button graphic.

```
public Form1()
{ // form on which drawing will occur
  InitializeComponent(); // initialize form
  this.Text = "C# Graphics";
  Graphics drawingCanvas = CreateGraphics();
  Show();
  Pen BlackPen = new Pen(Color.Black, 2);
  drawingCanvas.DrawRectangle
    (BlackPen, 100, 100, 50, 100);
  Brush redBrush = Brushes.Red;
  Font f = new Font("Arial", 10);
  drawingCanvas.DrawString
    ("Stop", f, redBrush, 105, 140);
}
```

The display is:



To sense a mouse click in a C# windows form, an “event handler” must be created. An **event handler** is code that can sense an event and respond to it. In this case we want to sense a “mousedown” event, which occurs when the user clicks the mouse button.

The C# IDE can generate the “shell” of the code to respond to a mousedown event. To make this happen, go to the Form1.cs[Design] view, right-click on the form image, and select Properties. This will open the Properties window; click on the little lightning-bolt icon near the top of this window to see the possible events on the form to which the system can respond. Double-click in the box next to the MouseClick event, which will then display “Form1_MouseClick”.

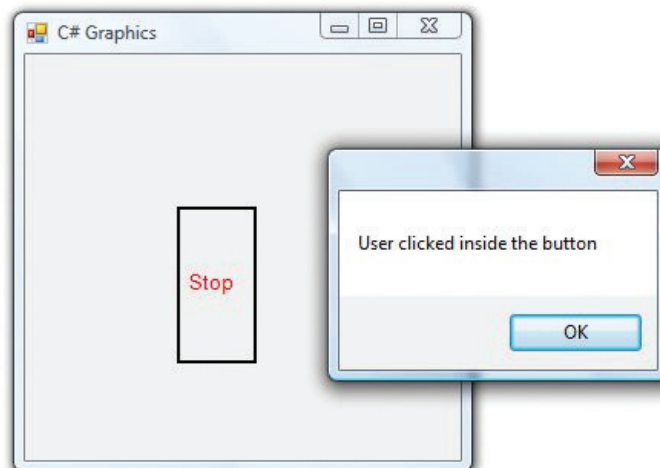
Now the Form1.cs code window shows the following additional code that has been added:

```
private void Form1_MouseClick
    (object sender, MouseEventArgs e)
{
}
}
```

Next we fill in the body of the event handler to check whether the cursor was in the rectangle when the mouseclick event occurred and to display a corresponding message.

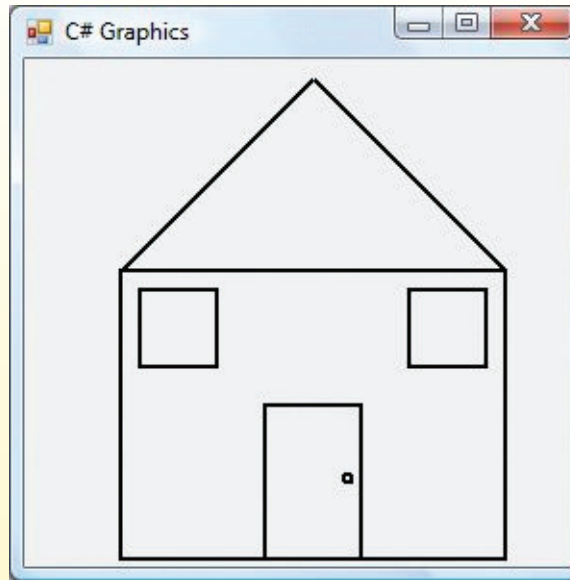
```
private void Form1_MouseClick
    (object sender, MouseEventArgs e)
{
    int x;
    int y;
    x = e.X;
    y = e.Y;
    if((x >= 100) && (x <= 150) &&
        (y >= 100) && (y <= 200))
    {
        MessageBox.Show("User clicked inside" +
            " the button");
    }
    else
    {
        MessageBox.Show("User clicked outside" +
            " the button");
    }
}
```

Integer variables *x* and *y* receive the coordinates of the cursor when the event occurs, and the above code just checks whether these (*x*, *y*) values fit within the rectangle. Here is output for the case where the user clicked the button:



PRACTICE PROBLEMS

Write the sequence of commands to draw the following “house” on the screen.



Create the house using four rectangles (for the base of the house, the door, and the two windows), two line segments (for the roof), and one circle (for the doorknob). Locate the house anywhere you want on the screen.

Of course, a real GUI interface would produce much more sophisticated responses to user mouse clicks, but this is the general idea of how **event-driven programming** works.

This brief introduction to graphical programming allows you to produce some interesting images and, even more important, gives you an appreciation for how visually oriented software is developed.

8 Conclusion

In this module we looked at one representative high-level programming language, C#. Of course, there is much about this language that has been left unsaid, but we have seen how the use of a high-level language overcomes many of the disadvantages of assembly language programming, creating a more comfortable and useful environment for the programmer. In a high-level language, the programmer need not manage the storage or movement of data values in memory. The programmer can think about the problem at a higher level, can use program instructions that are both more powerful and more

natural language–like, and can write a program that is much more portable among various hardware platforms. We also saw how modularization, through the use of functions and parameters, allows the program to be more cleanly structured and how object-oriented programming allows a more intuitive view of the problem solution and provides the possibility for reuse of helpful classes. We even had a glimpse of graphical programming.

C# is not the only high-level language. You might be interested in looking at the other online modules for languages similar to C# (Java, Python, C++, Ada). Still other languages take quite a different approach to problem solving. In Chapter 10 of *Invitation to Computer Science*, we look at some other languages and language approaches and also address the question of why there are so many different programming languages.

EXERCISES

- Write a C# declaration for one real number quantity to be called *rate*.
- Write a C# declaration for two integer quantities called *orderOne* and *orderTwo*.
- Write a C# declaration for a constant quantity called *evaporationRate*, which is to have the value 6.15.
- A C# Main function needs one constant *stockTime* with a value of 4, one integer variable *inventory*, and one real number variable *sales*. Write the necessary declarations.
- You want to write a C# program to compute the average of three quiz grades for a single student. Decide what variables your program needs, and write the necessary declarations.
- Given the declaration

```
int[] list = new int[10];
```

how do you refer to the eighth number in the array?

- An array declaration such as

```
int[,] table = new int[5,3];
```

represents a two-dimensional table of values with 5 rows and 3 columns. Rows and columns are numbered in C# starting at 0, not at 1. Given this declaration, how do you refer to the marked cell below?

- Write C# statements to prompt for and collect values for the time in hours and minutes (two integer quantities).
- An output statement may contain more than one variable identifier. Say a program computes two integer quantities *inventoryNumber* and *numberOrdered*. Write a single

output statement that prints these two quantities along with appropriate text information.

- The integer quantities *A*, *B*, *C*, and *D* currently have the values 13, 4, 621, and 18, respectively. Write the exact output generated by the following statement, using *b* to denote a blank space.

```
Console.WriteLine(A.ToString(
    " 00000") + B.ToString(" 000") +
    C.ToString(" 000") + D.ToString(
    " 000"));
```

- Write C# formatting and output statements to generate the following output, assuming that *density* is a type *double* variable with the value 63.78.

```
The current density is 63.8,
to within one decimal place.
```

- What is the output after the following sequence of statements is executed? (Assume that the integer variables *A* and *B* have been declared.)

```
A = 12;
B = 20;
B = B + 1;
A = A + B;
Console.WriteLine(2*A);
```

- Write the body of a C# Main function that gets the length and width of a rectangle from the user and computes and writes out the area. Assume that the variables have all been declared.
- In the SportsWorld program of Figure 15, the user must respond with "C" to choose the circumference task. In such a situation, it is preferable to accept either uppercase or lowercase letters. Rewrite the condition in the program to allow this.
 - In the SportsWorld program, rewrite the condition for continuation of the program to allow either an uppercase or a lowercase response.
- Write a C# Main function that gets a single character from the user and writes out a congratulatory message if the character is a vowel (a, e, i, o, or u), but otherwise writes out a "You lose, better luck next time" message.
- Insert the missing line of code so that the following adds the integers from 1 to 10, inclusive.

```
value = 0;
top = 10;
score = 1;
while (score <= top)
{
    value = value + score;
    - - - - //the missing line
}
```


EXERCISES

17. What is the output after the following Main function is executed?

```
static void Main(string[] args)
{
    int low, high;
    low = 1;
    high = 20;
    while (low < high)
    {
        Console.WriteLine(low + " " +
            high);
        low = low + 1;
        high = high - 1;
    }
}
```

18. Write a C# Main function that outputs the even integers from 2 through 30, one per line. Use a while loop.
19. In a while loop, the Boolean condition that tests for loop continuation is done at the top of the loop, before each iteration of the loop body. As a consequence, the loop body might not be executed at all. Our pseudocode language of Chapter 2 contains a do-while loop construction in which a test for loop termination occurs at the bottom of the loop rather than at the top, so that the loop body always executes at least once. C# contains a **do-while** statement that tests for loop continuation at the bottom of the loop. The form of the statement is


```
do
    S1;
while (Boolean condition);
```

where, as usual, S1 can be a compound statement. Write a C# Main function to add up a number of nonnegative integers that the user supplies and to write out the total. Use a negative value as a sentinel, and assume that the first value is nonnegative. Use a do-while statement.


20. Write a C# program that asks for a duration of time in hours and minutes, and writes out the duration only in minutes.
21. Write a C# program that asks for the user's age in years. If the user is under 35, then quote an insurance rate of \$2.23 per \$100 for life insurance; otherwise, quote a rate of \$4.32.
22. Write a C# program that reads integer values until a 0 value is encountered and then writes out the sum of the positive values read and the sum of the negative values read.
23. Write a C# program that reads in a series of positive integers and writes out the product of all the integers less than 25 and the sum of all the integers greater than or equal to 25. Use 0 as a sentinel value.
24. a. Write a C# program that reads in 10 integer quiz grades and computes the average grade. (*Hint*: Remember the peculiarity of integer division.)
b. Write a C# program that asks the user for the number of quiz grades, reads them in, and computes the average grade.
25. Write a (void) C# function that receives two integer arguments and writes out their sum and their product.
26. Write a (void) C# function that receives an integer argument representing the number of DVDs rented so far this month and a real number argument representing the sales amount for DVDs sold so far this month. The function asks the user for the number of DVDs rented today and the sales amount for DVDs sold today, and then returns the updated figures to the Main function.
27. Write a (nonvoid) C# function that receives three integer arguments and returns the maximum of the three values.
28. Write a (nonvoid) C# function that receives miles driven as a type *double* argument and gallons of gas used as a type *int* argument, and returns miles per gallon.
29. Write a C# program that uses an input function to get the miles driven (type *double*) and the gallons of gas used (type *int*), then writes out the miles per gallon, using the function from Exercise 28.
30. Write a C# program to balance a checkbook. The program needs to get the initial balance, the amounts of deposits, and the amounts of checks. Allow the user to process as many transactions as desired; use separate functions to handle deposits and checks.
31. Write a C# program to compute the cost of carpeting three rooms. Make the carpet cost a constant of \$8.95 per square yard. Use four separate functions to collect the dimensions of a room in feet, convert feet into yards, compute the area, and compute the cost per room. The Main function should use a loop to process each of the three rooms, then add the three costs, and write out the total cost. (*Hint*: The function to convert feet into yards must be used twice for each room, with two different arguments. Hence, it does not make sense to try to give the parameter the same name as the argument.)
32. a. Write a C# *doPerimeter* function for the *Rectangle* class of Figure 29.
b. Write C# code that creates a new *Rectangle* object called *yuri*, then writes out information about this object and its perimeter using the *doPerimeter* function from part (a).
33. Draw a class hierarchy diagram similar to Figure 31 for the following classes: *Student*, *UndergraduateStudent*, *GraduateStudent*, *Sophomore*, *Senior*, *PhDStudent*.

EXERCISES

34. Imagine that you are writing a program using an object-oriented programming language. Your program will be used to maintain records for a real estate office. Decide on one class in your program and a service that objects of that class might provide.
35. Determine the resolution of the screen on your computer (ask your instructor or the local computer center how to do this). Using this information, determine how many bytes of memory are required for the frame buffer to store:
- A black-and-white image (1 bit per pixel)
 - A grayscale image (8 bits per pixel)
 - A color image (24 bits per pixel)
36. Using the *DrawLine* command described in Section 7.2, draw an isosceles triangle with the following configuration:




37. Discuss what problem the display hardware might encounter while attempting to execute the following operations, and explain how this problem could be solved.
- ```
DrawLine(blackPen, 1, 1, 4, 5);
```
38. Draw a square with sides 100 pixels in length. Then inscribe a circle of radius 50 inside the square. Position the square and the inscribed circle in the middle of the screen.
39. Create the following three labeled rectangular buttons in the output window.



Have the space between the Start and Stop buttons be the same as the space between the Stop and Pause buttons.

40. Create the following image of a "teeter-totter":



41. Write a program that inputs the coordinates of three mouse clicks from the user and then draws a triangle in the output window using those three points. Here are some hints: Unlike the *mouseclick* event handler shown in the text, this event handler will do some actual drawing on the canvas. This means that the graphics canvas and the pen need to be visible to the event handler, so they need to be declared in the Form partial class that contains both the Form1 code and the event handler code. There must also be a *counter* variable to keep track of which of the three mouseclicks has just been made, as well as variables to hold the three sets of coordinates. The values for these variables have to persist through all three invocations of the event handler, so their declarations will also go in the Form1 partial class. Here is the code for the start of the program.

```
public partial class Form1 : Form
{
 Graphics drawingCanvas;
 Pen BlackPen;
 int counter = 0;
 int x1, y1, x2, y2, x3, y3;

 public Form1()
 {
 // form on which drawing
 // will occur
 // initialize form
 InitializeComponent();
 this.Text = "C# Graphics";
 drawingCanvas =
 CreateGraphics();
 Show();
 BlackPen = new Pen
 (Color.Black, 2);
 } //end Public Form 1
}
```

Here is the start of the *MouseClicked* event handler:

```
private void Form1_MouseClick
(object sender,
 MouseEventArgs e)
{
 counter = counter + 1;
 if (counter == 1)
 {
 x1 = e.X;
 y1 = e.Y;
 }
}
```

There will need to be two more *if* statements to capture the *x*, *y* values for the second and third clicks. The drawing code for the triangle will be in the body of the third *if* statement.

## ANSWERS TO PRACTICE PROBLEMS

- Section 2**
- The first three.  
 martinBradley (camel case)  
 C3P\_OH (although best not to use underscore character)  
 Amy3 (Pascal case)  
 3Right (not acceptable, begins with digit)  
 const (not acceptable, C# reserved word)
  - `int number;`
  - `const double taxRate = 5.5;`
  - `hits[7]`

- Section 3.1**
- `Console.Write("Enter a value for quantity: ");  
 quantity = Convert.ToInt32(Console.ReadLine());`
  - `Console.WriteLine(height.ToString("000000"));`
  - This isgoodbye

- Section 3.2**
- `next = newNumber;`
  - 55

- Section 3.3**
- 30
  - 3  
5  
7  
9  
11  
13  
15  
17  
19  
21
  - Yes
  - 6
  - `if (night == day)  
 Console.WriteLine("Equal");`

- Section 4**
- ```
//program to read in a number
//and write out its square
using System;

namespace PracticeProblem
{
```

```

class Problem
{
    static void Main(string[] args)
    {
        int number, square;
        Console.Write("Enter a number: ");
        number = Convert.ToInt32(Console.ReadLine());
        square = number * number;
        Console.WriteLine("The square is " + square);
    }
}

```

2.

```

//program to compute cost based on price per item
//and quantity purchased

using System;

namespace PracticeProblem
{
    class Problem
    {
        static void Main(string[] args)
        {
            int quantity;
            double price, totalCost;
            Console.Write("Enter the price: ");
            price = Convert.ToDouble(Console.ReadLine());
            Console.Write("Enter the quantity: ");
            quantity = Convert.ToInt32(Console.ReadLine());
            totalCost = price * quantity;
            Console.WriteLine("The total cost is " +
                totalCost.ToString("####.00"));
        }
    }
}

```

3.

```

//program to test a number relative to 5
//and write out the number or its double

using System;

namespace PracticeProblem
{
    class Problem
    {
        static void Main(string[] args)
        {
            int number;
            Console.Write("Enter the number: ");
            number = Convert.ToInt32(Console.ReadLine());
            if (number < 5)
            {
                Console.WriteLine("The number is " + number);
            }
        }
    }
}

```

```

else
{
    Console.WriteLine("Twice the number is " +
        number * 2);
}
}
}
}

```

4.

```

//program to collect a number, then write all
//the values from 1 to that number

using System;

namespace PracticeProblem
{
    class Problem
    {
        static void Main(string[] args)
        {
            int number;
            int i;
            Console.Write("Enter a positive integer: ");
            number = Convert.ToInt32(Console.ReadLine());
            i = 1;
            while(i <= number)
            {
                Console.WriteLine("The number is " + i);
                i = i + 1;
            }
        }
    }
}

```

Section 5.3 1. 11

2. 7

3.

```

static void getInput(ref int one, ref int two)
{
    Console.Write("Input the value for One: ");
    one = Convert.ToInt32(Console.ReadLine());
    Console.Write("Input the value for Two: ");
    two = Convert.ToInt32(Console.ReadLine());
}

```

4.

```

a. static double tax(double subTotal)
b. return subTotal * rate;
c. Console.WriteLine("Total: " +
    tax(subTotal).ToString("0.00"));

```

Section 6.4 1. The area of a square with side 10 is 100

2. height and base

Section 7.2

```
public Form1()
{ // form on which drawing will occur
  InitializeComponent(); // initialize form
  this.Text = "C# Graphics";
  Graphics drawingCanvas = CreateGraphics();
  Show();
  Pen BlackPen = new Pen(Color.Black, 2);
  drawingCanvas.DrawRectangle(BlackPen, 50, 110, 200,
    150);
  drawingCanvas.DrawRectangle(BlackPen, 60, 120, 40,
    40);
  drawingCanvas.DrawRectangle(BlackPen, 200, 120, 40,
    40);
  drawingCanvas.DrawRectangle(BlackPen, 125, 180, 50,
    80);
  drawingCanvas.DrawLine(BlackPen, 50, 110, 150, 10);
  drawingCanvas.DrawLine(BlackPen, 250, 110, 150, 10);
  drawingCanvas.DrawEllipse(BlackPen, 165, 215, 5, 5);
}
```