Character, String, and Numeric Functions

C++ provides many built-in functions in addition to the `cout`, `getch()`, and `strcpy()` functions you have seen so far. These built-in functions increase your productivity and save you programming time. You don't have to write as much code because the built-in functions perform many useful tasks for you.

This chapter introduces you to

- Character conversion functions
- Character and string testing functions
- String manipulation functions
- String I/O functions
- Mathematical, trigonometric, and logarithmic functions
- Random-number processing

Character Functions

This section explores many of the character functions available in AT&T C++. Generally, you pass character arguments to the functions, and the functions return values that you can store or print. By using these functions, you off-load much of your work to C++ and allow it to perform the more tedious manipulations of character and string data.

Character Testing Functions

The character functions return True or False results based on the characters you pass to them.

Several functions test for certain characteristics of your character data. You can determine whether your character data is alphabetic, digital, uppercase, lowercase, and much more. You must pass a character variable or literal argument to the function (by placing the argument in the function parentheses) when you call it. These functions return a True or False result, so you can test their return values inside an `if` statement or a `while` loop.

NOTE: All character functions presented in this section are prototyped in the ctype.h header file. Be sure to include ctype.h at the beginning of any programs that use them.

Alphabetic and Digital Testing

The following functions test for alphabetic conditions:

- `isalpha(c)`: Returns True (nonzero) if `c` is an uppercase or lowercase letter. Returns False (zero) if `c` is not a letter.
- `islower(c)`: Returns True (nonzero) if `c` is a lowercase letter. Returns False (zero) if `c` is not a lowercase letter.
- `isupper(c)`: Returns True (nonzero) if `c` is an uppercase letter. Returns False (zero) if `c` is not an uppercase letter.

Remember that any nonzero value is True in C++, and zero is always False. If you use the return values of these functions in a relational test, the True return value is not always 1 (it can be any nonzero value), but it is always considered True for the test.

The following functions test for digits:

- `isdigit(c)`: Returns True (nonzero) if `c` is a digit 0 through 9. Returns False (zero) if `c` is not a digit.
- `isxdigit(c)`: Returns True (nonzero) if `c` is any of the hexadecimal digits 0 through 9 or A, B, C, D, E, F, a, b, c, d, e, or f. Returns False (zero) if `c` is anything else. (See Appendix A, "Memory Addressing, Binary, and Hexadecimal Review," for more information on the hexadecimal numbering system.)

NOTE: Although some character functions test for digits, the arguments are still character data and cannot be used in mathematical calculations, unless you calculate using the ASCII values of characters.

The following function tests for numeric or alphabetical arguments:

`isalnum(c)`: Returns True (nonzero) if `c` is a digit 0 through 9 or an alphabetic character (either uppercase or lowercase). Returns False (zero) if `c` is not a digit or a letter.

CAUTION: You can pass to these functions only a character value or an integer value holding the ASCII value of a character. You cannot pass an entire character array to character functions. If you want to test the elements of a character array, you must pass the array one element at a time.

Example

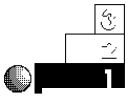

The following program asks users for their initials. If a user types anything but alphabetic characters, the program displays an error and asks again.

Identify the program and include the input/output header files. The program asks the user for his or her first initial, so declare the character variable `initial` *to hold the user's answer.*

1. *Ask the user for her or his first initial, and retrieve the user's answer.*
2. *If the answer was not an alphabetic character, tell the user this and repeat step one.*

Print a thank-you message on-screen.

```
// Filename: C22INI.CPP
// Asks for first initial and tests
// to ensure that response is correct.
#include <iostream.h>
#include <ctype.h>
void main()
{
   char initial;
   cout << "What is your first initial? ";
   cin >> initial;

   while (!isalpha(initial))
     {
       cout << "\nThat was not a valid initial! \n";
       cout << "\nWhat is your first initial? ";
       cin >> initial;
     }

   cout << "\nThanks!";
   return;
}
```

This use of the `not` operator (`!`) is quite clear. The program continues to loop as long as the entered character is not alphabetic.

Special Character-Testing Functions

A few character functions become useful when you have to read from a disk file, a modem, or another operating system device that you route input from. These functions are not used as much as the character functions you saw in the previous section, but they are useful for testing specific characters for readability.

The character-testing functions do not change characters.

The remaining character-testing functions follow:

- `iscntrl(c)`: Returns True (nonzero) if `c` is a *control character* (any character from the ASCII table numbered 0 through 31). Returns False (zero) if `c` is not a control character.
- `isgraph(c)`: Returns True (nonzero) if `c` is any printable character (a noncontrol character) except a space. Returns False (zero) if `c` is a space or anything other than a printable character.
- `isprint(c)`: Returns True (nonzero) if `c` is a printable character (a noncontrol character) from ASCII 32 to ASCII 127, including a space. Returns False (zero) if `c` is not a printable character.
- `ispunct(c)`: Returns True (nonzero) if `c` is any punctuation character (any printable character other than a space, a letter, or a digit). Returns False (zero) if `c` is not a punctuation character.
- `isspace(c)`: Returns True (nonzero) if `c` is a space, newline (`\n`), carriage return (`\r`), tab (`\t`), or vertical tab (`\v`) character. Returns False (zero) if `c` is anything else.

Character Conversion Functions

Both `tolower()` and `toupper()` return lowercase or uppercase arguments.

The two remaining character functions are handy. Rather than test characters, these functions change characters to their lower- or uppercase equivalents.

- `tolower(c)`: Converts `c` to lowercase. Nothing changes if you pass `tolower()` a lowercase letter or a nonalphabetic character.
- `toupper(c)`: Converts `c` to uppercase. Nothing changes if you pass `toupper()` an uppercase letter or a nonalphabetic character.

These functions return their changed character values. These functions are useful for user input. Suppose you are asking users a yes or no question, such as the following:

```
Do you want to print the checks (Y/N)?
```

Before `toupper()` and `tolower()` were developed, you had to check for both a *Y* and a *y* to print the checks. Instead of testing for both conditions, you can convert the character to uppercase, and test for a *Y*.

Example

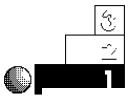

Here is a program that prints an appropriate message if the user is a girl or a boy. The program tests for *G* and *B* after converting the user's input to uppercase. No check for lowercase has to be done.

Identify the program and include the input/output header files. The program asks the user a question requiring an alphabetic answer, so declare the character variable `ans` *to hold the user's response.*

Ask whether the user is a girl or a boy, and store the user's answer in `ans`. *The user must press Enter, so incorporate and then discard the Enter keypress. Change the value of* `ans` *to uppercase. If the answer is G, print a message. If the answer is B, print a different message. If the answer is something else, print another message.*

```
// Filename: C22GB.CPP
// Determines whether the user typed a G or a B.
#include <iostream.h>
#include <conio.h>
#include <ctype.h>
void main()
{
```

```
   char ans;                    // Holds user's response.
   cout << "Are you a girl or a boy (G/B)? ";
   ans=getch();                              // Get answer.
   getch();                            // Discard newline.

cout <<ans<<"\n";
   ans = toupper(ans);     // Convert answer to uppercase.
   switch (ans)
   {   case ('G'): { cout << "You look pretty today!\n";
                     break; }
       case ('B'): { cout << "You look handsome today!\n";
                     break; }
       default :   { cout << "Your answer makes no sense!\n";
                     break; }
   }
   return;
}
```

Here is the output from the program:

```
Are you a girl or a boy (G/B)? B
You look handsome today!
```

String Functions

Some of the most powerful built-in C++ functions are the string functions. They perform much of the tedious work for which you have been writing code so far, such as inputting strings from the keyboard and comparing strings.

As with the character functions, there is no need to "reinvent the wheel" by writing code when built-in functions do the same task. Use these functions as much as possible.

Now that you have a good grasp of the foundations of C++, you can master the string functions. They enable you to concentrate on your program's primary purpose, rather than spend time coding your own string functions.

Useful String Functions

You can use a handful of useful string functions for string testing and conversion. You have already seen (in earlier chapters) the `strcpy()` string function, which copies a string of characters to a character array.

NOTE: All string functions in this section are prototyped in the string.h header file. Be sure to include string.h at the beginning of any program that uses the string functions.

The string functions work on string literals or on character arrays that contain strings.

String functions that test or manipulate strings follow:

- `strcat(s1, s2)`: Concatenates (merges) the `s2` string to the end of the `s1` character array. The `s1` array must have enough reserved elements to hold both strings.
- `strcmp(s1, s2)`: Compares the `s1` string with the `s2` string on an alphabetical, element-by-element basis. If `s1` alphabetizes before `s2`, `strcmp()` returns a negative value. If `s1` and `s2` are the same strings, `strcmp()` returns `0`. If `s1` alphabetizes after `s2`, `strcmp()` returns a positive value.
- `strlen(s1)`: Returns the length of `s1`. Remember, the length of a string is the number of characters, not including the null zero. The number of characters defined for the character array has nothing to do with the length of the string.

TIP: Before using `strcat()` to concatenate strings, use `strlen()` to ensure that the target string (the string being concatenated to) is large enough to hold both strings.

String I/O Functions

In the previous few chapters, you have used a character input function, `cin.get()`, to build input strings. Now you can begin to use the string input and output functions. Although the goal of the

string-building functions has been to teach you the specifics of the language, these string I/O functions are much easier to use than writing a character input function.

The string input and output functions are listed as follows:

- `gets(s)`: Stores input from `stdin` (usually directed to the keyboard) to the string named `s`.
- `puts(s)`: Outputs the `s` string to `stdout` (usually directed to the screen by the operating system).
- `fgets(s, len, dev)`: Stores input from the standard device specified by `dev` (such as `stdin` or `stdaux`) in the `s` string. If more than `len` characters are input, `fgets()` discards them.
- `fputs(s, dev)`: Outputs the `s` string to the standard device specified by `dev`.

Both `gets()` and `puts()` input and output strings.

These four functions make the input and output of strings easy. They work in pairs. That is, strings input with `gets()` are usually output with `puts()`. Strings input with `fgets()` are usually output with `fputs()`.

TIP: `gets()` replaces the string-building input function you saw in earlier chapters.

Terminate `gets()` or `fgets()` input by pressing Enter. Each of these functions handles string-terminating characters in a slightly different manner, as follows:

`gets()`	A newline input becomes a null zero (`\0`).
`puts()`	A null at the end of the string becomes a newline character (`\n`).
`fgets()`	A newline input stays, and a null zero is added after it.
`fputs()`	The null zero is dropped, and a newline character is not added.

Therefore, when you enter strings with `gets()`, C++ places a string-terminating character in the string at the point where you press Enter. This creates the input string. (Without the null zero, the

input would not be a string.) When you output a string, the null zero at the end of the string becomes a newline character. This is preferred because a newline is at the end of a line of output and the cursor begins automatically on the next line.

Because `fgets()` and `fputs()` can input and output strings from devices such as disk files and telephone modems, it can be critical that the incoming newline characters are retained for the data's integrity. When outputting strings to these devices, you do not want C++ inserting extra newline characters.

CAUTION: Neither `gets()` nor `fgets()` ensures that its input strings are large enough to hold the incoming data. It is up to you to make sure enough space is reserved in the character array to hold the complete input.

One final function is worth noting, although it is not a string function. It is the `fflush()` function, which flushes (empties) whatever standard device is listed in its parentheses. To flush the keyboard of all its input, you would code as follows:

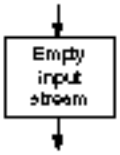


```
fflush(stdin);
```

When you are doing string input and output, sometimes an extra newline character appears in the keyboard buffer. A previous answer to `gets()` or `getc()` might have an extra newline you forgot to discard. When a program seems to ignore `gets()`, you might have to insert `fflush(stdin)` before `gets()`.

Flushing the standard input device causes no harm, and using it can clear the input stream so your next `gets()` works properly. You can also flush standard output devices with `fflush()` to clear the output stream of any characters you sent to it.

NOTE: The header file for `fflush()` is in stdio.h.

Example

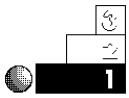

The following program shows you how easy it is to use `gets()` and `puts()`. The program requests the name of a book from the user using a single `gets()` function call, then prints the book title with `puts()`.

Identify the program and include the input/output header files. The program asks the user for the name of a book. Declare the character array `book` *with 30 elements to hold the user's answer.*

Ask the user for the book's title, and store the user's response in the `book` *array. Display the string stored in* `book` *to an output device, probably your screen. Print a thank-you message.*

```
// C22GPS1.CPP
// Inputs and outputs strings.
#include <iostream.h>
#include <stdio.h>
#include <string.h>

void main()
{
   char book[30];

   cout << "What is the book title? ";
   gets(book);                    // Get an input string.
   puts(book);                    // Display the string.
   cout << "Thanks for the book!\n";
   return;
}
```

The output of the program follows:

```
What is the book title? Mary and Her Lambs
Mary and Her Lambs
Thanks for the book!
```

Converting Strings to Numbers

Sometimes you have to convert numbers stored in character strings to a numeric data type. AT&T C++ provides three functions that enable you to do this:

- `atoi(s)`: Converts `s` to an integer. The name stands for *a*lphabetic *to i*nteger.
- `atol(s)`: Converts `s` to a long integer. The name stands for *a*lphabetic *to l*ong integer.
- `atof(s)`: Converts `s` to a floating-point number. The name stands for *a*lphabetic *to f*loating-point.

NOTE: These three `ato()` functions are prototyped in the stdlib.h header file. Be sure to include stdlib.h at the beginning of any program that uses the `ato()` functions.

The string must contain a valid number. Here is a string that can be converted to an integer:

```
"1232"
```

The string must hold a string of digits short enough to fit in the target numeric data type. The following string could not be converted to an integer with the `atoi()` function:

```
"-1232495.654"
```

However, it could be converted to a floating-point number with the `atof()` function.

C++ cannot perform any mathematical calculation with such strings, even if the strings contain digits that represent numbers. Therefore, you must convert any string to its numeric equivalent before performing arithmetic with it.

NOTE: If you pass a string to an `ato()` function and the string does not contain a valid representation of a number, the `ato()` function returns `0`.

These functions become more useful to you after you learn about disk files and pointers.

Numeric Functions

This section presents many of the built-in C++ numeric functions. As with the string functions, these functions save you time by converting and calculating numbers instead of your having to write functions that do the same thing. Many of these are trigonometric and advanced mathematical functions. You might use some of these numeric functions only rarely, but they are there if you need them.

This section concludes the discussion of C++'s standard built-in functions. After mastering the concepts in this chapter, you are ready to learn more about arrays and pointers. As you develop more skills in C++, you might find yourself relying on these numeric, string, and character functions when you write more powerful programs.

Useful Mathematical Functions

Several built-in numeric functions return results based on numeric variables and literals passed to them. Even if you write only a few science and engineering programs, some of these functions are useful.

NOTE: All mathematical and trigonometric functions are prototyped in the math.h header file. Be sure to include math.h at the beginning of any program that uses the numeric functions.

These numeric functions return double-precision values.

Here are the functions listed with their descriptions:

- `ceil(x)`: The `ceil()`, or *ceiling*, function rounds numbers up to the nearest integer.
- `fabs(x)`: Returns the absolute value of `x`. The absolute value of a number is its positive equivalent.

TIP: Absolute value is used for distances (which are always positive), accuracy measurements, age differences, and other calculations that require a positive result.

- `floor(x)`: The `floor()` function rounds numbers down to the nearest integer.
- `fmod(x, y)`: The `fmod()` function returns the floating-point remainder of (`x/y`) with the same sign as `x`, and `y` cannot be zero. Because the modulus operator (`%`) works only with integers, this function is used to find the remainder of floating-point number divisions.
- `pow(x, y)`: Returns `x` raised to the `y` power, or x^y. If `x` is less than or equal to zero, `y` must be an integer. If `x` equals zero, `y` cannot be negative.
- `sqrt(x)`: Returns the square root of `x`; `x` must be greater than or equal to zero.

The *n*th Root

No function returns the *n*th root of a number, only the square root. In other words, you cannot call a function that gives you the 4th root of 65,536. (By the way, 16 is the 4th root of 65,536, because 16 times 16 times 16 times 16 = 65,536.)

You can use a mathematical trick to simulate the *n*th root, however. Because C++ enables you to raise a number to a fractional power—with the `pow()` function—you can raise a number to the *n*th root by raising it to the (1/*n*) power. For example, to find the 4th root of 65,536, you could type this:

```
root = pow(65536.0, (1.0/4.0));
```

Note that the decimal point keeps the numbers in floating-point format. If you leave them as integers, such as

```
root = pow(65536, (1/4));
```

C++ produces incorrect results. The `pow()` function and most other mathematical functions require floating-point values as arguments.

To store the 7th root of 78,125 in a variable called `root`, for example, you would type

```
root = pow(78125.0, (1.0/7.0));
```

This stores 5.0 in `root` because 5^7 equals 78,125.

Knowing how to compute the nth root is handy in scientific programs and also in financial applications, such as time-value-of-money problems.

Example

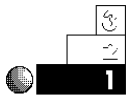

The following program uses the `fabs()` function to compute the difference between two ages.

```
// Filename: C22ABS.CPP
// Computes the difference between two ages.
#include <iostream.h>
#include <math.h>
void main()
{
   float age1, age2, diff;
   cout << "\nWhat is the first child's age? ";
   cin >> age1;
   cout << "What is the second child's age? ";
   cin >> age2;

   // Calculates the positive difference.
   diff = age1 - age2;
   diff = fabs(diff);     // Determines the absolute value.

   cout << "\nThey are " << diff << " years apart.";
   return;
}
```

The output from this program follows. Due to `fabs()`, the order of the ages doesn't matter. Without absolute value, this program would produce a negative age difference if the first age was less than the second. Because the ages are relatively small, floating-point variables are used in this example. C++ automatically converts floating-point arguments to double precision when passing them to `fabs()`.

```
What is the first child's age? 10
What is the second child's age? 12

They are 2 years apart.
```

Trigonometric Functions

The following functions are available for trigonometric applications:

- `cos(x)`: Returns the cosine of the angle `x`, expressed in radians.
- `sin(x)`: Returns the sine of the angle `x`, expressed in radians.
- `tan(x)`: Returns the tangent of the angle `x`, expressed in radians.

These are probably the least-used functions. This is not to belittle the work of scientific and mathematical programmers who need them, however. Certainly, they are grateful that C++ supplies these functions! Otherwise, programmers would have to write their own functions to perform these three basic trigonometric calculations.

Most C++ compilers supply additional trigonometric functions, including hyperbolic equivalents of these three functions.

TIP: If you have to pass an angle that is expressed in degrees to these functions, convert the angle's degrees to radians by multiplying the degrees by $\pi/180.0$ (π equals approximately 3.14159).

Logarithmic Functions

Three highly mathematical functions are sometimes used in business and mathematics. They are listed as follows:

- `exp(x)`: Returns the base of natural logarithm (*e*) raised to a power specified by `x` (e^x); *e* is the mathematical expression for the approximate value of 2.718282.
- `log(x)`: Returns the natural logarithm of the argument `x`, mathematically written as `ln(x)`. `x` must be positive.
- `log10(x)`: Returns the base-10 logarithm of argument `x`, mathematically written as `log10(x)`. `x` must be positive.

Random-Number Processing

Random events happen every day. You wake up and it is sunny or rainy. You have a good day or a bad day. You get a phone call from an old friend or you don't. Your stock portfolio might go up or down in value.

Random events are especially important in games. Part of the fun in games is your luck with rolling dice or drawing cards, combined with your playing skills.

Simulating random events is an important task for computers. Computers, however, are finite machines; given the same input, they always produce the same output. This fact can create some boring games!

The `rand()` function produces random integer numbers.

The designers of C++ knew this computer setback and found a way to overcome it. They wrote a random-number generating function called `rand()`. You can use `rand()` to compute a dice roll or draw a card, for example.

To call the `rand()` function and assign the returned random number to test, use the following syntax:

```
test = rand();
```

The `rand()` function returns an integer from 0 to 32,767. Never use an argument in the `rand()` parentheses.

Every time you call `rand()` in the same program, you receive a different number. If you run the same program over and over,

however, rand() returns the same set of random numbers. One way to receive a different set of random numbers is to call the srand() function. The format of srand() follows:

```
srand(seed);
```

where seed is an integer variable or literal. If you don't call srand(), C++ assumes a seed value of 1.

NOTE: The rand() and srand() functions are prototyped in the stdlib.h header file. Be sure to include stdlib.h at the beginning of any program that uses rand() or srand().

The seed value reseeds, or resets, the random-number generator, so the next random number is based on the new seed value. If you call srand() with a different seed value at the top of a program, rand() returns a different random number each time you run the program.

Why Do You Have To Do This?

There is considerable debate among C++ programmers concerning the random-number generator. Many think that the random numbers should be truly random, and that they should not have to seed the generator themselves. They think that C++ should do its own internal seeding when you ask for a random number.

However, many applications would no longer work if the random-number generator were randomized for you. Computers are used in business, engineering, and research to simulate the pattern of real-world events. Researchers have to be able to duplicate these simulations, over and over. Even though the events inside the simulations might be random from each other, the running of the simulations cannot be random if researchers are to study several different effects.

Mathematicians and statisticians also have to repeat random-number patterns for their analyses, especially when they work with risk, probability, and gaming theories.

Because so many computer users have to repeat their random-number patterns, the designers of C++ have wisely chosen to give you, the programmer, the option of keeping the same random patterns or changing them. The advantages far outweigh the disadvantage of including an extra `srand()` function call.

If you want to produce a different set of random numbers every time your program runs, you must determine how your C++ compiler reads the computer's system clock. You can use the seconds count from the clock to seed the random-number generator so it seems truly random.

Review Questions

The answers to the review questions are in Appendix B.

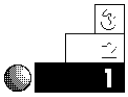

1. How do the character testing functions differ from the character conversion functions?
2. What are the two string input functions?
3. What is the difference between `floor()` and `ceil()`?

4. What does the following nested function return?

```
isalpha(islower('s'));
```

5. If the character array `str1` contains the string `Peter` and the character array `str2` contains `Parker`, what does `str2` contain after the following line of code executes?

```
strcat(str1, str2);
```

6. What is the output of the following `cout`?

```
cout << floor(8.5) << " " << ceil(8.5);
```

7. True or false: The `isxdigit()` and `isgraph()` functions could return the same value, depending on the character passed to them.

8. Assume you declare a character array with the following statement:

   ```
   char ara[5];
   ```

   Now suppose the user types `Programming` in response to the following statement:

   ```
   fgets(ara, 5, stdin);
   ```

   Would `ara` contain `Prog`, `Progr`, or `Programming`?

9. True or false: The following statements print the same results.

   ```
   cout << pow(64.0, (1.0/2.0)) ;
   cout << sqrt(64.0);
   ```

Review Exercises

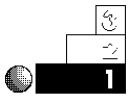

1. Write a program that asks users for their ages. If a user types anything other than two digits, display an error message.
2. Write a program that stores a password in a character array called `pass`. Ask users for the password. Use `strcmp()` to inform users whether they typed the proper password. Use the string I/O functions for all the program's input and output.
3. Write a program that rounds up and rounds down the numbers –10.5, –5.75, and 2.75.

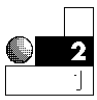

4. Ask users for their names. Print every name in *reverse* case; print the first letter of each name in lowercase and the rest of the name in uppercase.
5. Write a program that asks users for five movie titles. Print the longest title. Use only the string I/O and manipulation functions presented in this chapter.
6. Write a program that computes the square root, cube root, and fourth root of the numbers from 10 to 25, inclusive.

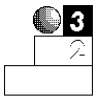

7. Ask users for the titles of their favorite songs. Discard all the special characters in each title. Print the words in the title, one per line. For example, if they enter `My True Love Is Mine, Oh, Mine!`, you should output the following:

```
My
True
Love
Is
Mine
Oh
Mine
```

8. Ask users for the first names of 10 children. Using `strcmp()` on each name, write a program to print the name that comes first in the alphabet.

Summary

You have learned the character, string, and numeric functions that C++ provides. By including the ctype.h header file, you can test and convert characters that a user types. These functions have many useful purposes, such as converting a user's response to uppercase. This makes it easier for you to test user input.

The string I/O functions give you more control over both string and numeric input. You can receive a string of digits from the keyboard and convert them to a number with the `ato()` functions. The string comparison and concatenation functions enable you to test and change the contents of more than one string.

Functions save you programming time because they take over some of your computing tasks, leaving you free to concentrate on your programs. C++'s numeric functions round and manipulate numbers, produce trigonometric and logarithmic results, and produce random numbers.

Now that you have learned most of C++'s built-in functions, you are ready to improve your ability to work with arrays. Chapter 23, "Introducing Arrays," extends your knowledge of character arrays and shows you how to produce arrays of any data type.