

The Basics of C Programming

Adapted for PHYS2200 class from Marshall Brain's article¹

Last updated: October 22, 2015

¹<http://www.howstuffworks.com/c.htm>

Contents

I	Basics of C	1
1	C programming	3
	What is C?	3
	The simplest C program, I	4
	Spacing and indentation	4
	Compilation and run	4
	The simplest C program, II	5
	Variables and variable types	6
	Operators	8
2	Input and output	9
	printf	9
	scanf	11
	Programming exercise	13
3	Branching and looping	15
	if statement	15
	Boolean expressions	16
	Boolean: = vs ==	17
	while loop	18
	do-while loop	19
	for loop	19
	Looping: an example	20
	Programming exercise	24
4	Functions	25
	Function prototypes	27
5	Arrays	29

Programming example: sorting arrays	31
6 Structures	37
7 Pointers	41
Pointers: why?	42
Pointer Basics	43
Understanding Memory Addresses	45
Pointing to the Same Address	47
Using Pointers for Function Parameters	47
Pointers and Arrays	49
Pointers to Structures	53
 II Programming Tools	 55
8 Makefiles	57
9 Command Line Arguments	59
10 Version control	61
11 Debugging	63
12 Detecting memory problems with valgrind	65
 III More C Programming	 67
13 Dynamic Data Structures	69
Stack vs Heap	70
Malloc and Free	72
14 Strings	81
Programming exercise	83
15 Operator Precedence, II	85
16 Advanced Pointers	87
Pointer Types	87
Pointers to Arrays	88

Arrays of Pointers	88
Structures Containing Pointers	89
Pointers to Pointers	90
Pointers to Structures Containing Pointers	90
Linking	91
A Linked Stack Example	92
Programming exercise	94
17 Libraries	97
Making a library	100
Compiling and running	103
18 Text files	105
Text files: opening	107
Text files: reading	108
Main function return values	108
19 Binary Files	111
Index	117

Listings

1.1	The simplest C program	4
1.2	Integer division and typecasting	7
2.1	Using printf	9
2.2	Using scanf	11
2.3	Addition of three numbers	13
3.1	if statement	15
3.2	if else statement	16
3.3	Boolean vs integer in if statement	18
3.4	Fahrenheit-to-Celsius conversion using integers	20
3.5	Fahrenheit-to-Celsius conversion using floating point numbers, I	21
3.6	Fahrenheit-to-Celsius conversion using floating point numbers, II, with a bug	22
3.7	Fahrenheit-to-Celsius conversion using floating point numbers, II, corrected	22
5.1	Using arrays	30
5.2	Sorting arrays, I	31
5.3	Sorting arrays, I	32
6.1	Initializing structures and passing structures to functions	38
7.1	Using pointers, I	43
7.2	Using pointers for function parameters, I	47
7.3	Using pointers for function parameters, II	48
7.4	Copy the content of an array	49
7.5	Mixing arrays and pointers	50
13.1	A program that causes a stack overflow	70
13.2	Simple example of using malloc and free	72
	example-dynamicmemory2.c	74
	example-dynamicmemory3.c	75
13.3	Dereferencing of NULL pointer	76
13.4	Using malloc and free	77
13.5	Using malloc and free with structures	78
14.1	Using strncpy	82

14.2 Using <code>strncmp</code>	82
14.3 Using pointers with <code>string</code>	83

Part I

Basics of C

Chapter 1

Basics of C programming

The C programming language is a widely used language for creating computer programs. It gives maximum control and efficiency to the developer. If you are interested in becoming a C programmer, there are a couple of benefits you gain from learning C:

- You will be able to read and write code for a large number of platforms – everything from microcontrollers to the most advanced scientific systems can be written in C.
- Because of the performance and portability of C, almost all popular cross-platform programming and scripting languages are implemented in C and borrowed syntax and functions heavily from C. They share the similar operators, expressions, repetition statements, control structures, arrays, input and output, and functions. Furthermore, almost all languages can interface with C to take advantage of a large volume of existing C libraries.

What is C?

C is a *programming language*. That means that you can use C to create lists of instructions for a computer to follow. C is one of thousands of programming languages currently in use. C has been around for several decades and has won widespread acceptance. C is an easy to learn. It is a bit more cryptic in its style than some other languages, but you get beyond that fairly quickly.

C is what is called a *compiled language*. This means that once you write your C program, you must run it through a *C compiler* to turn your program into an *executable* that the computer can run (execute). The C program is the human-readable form, while the executable that comes out of the compiler is the machine-readable and executable form. What this means is that to write and run a C program, you must have access to a C compiler.

We will start at the beginning with an extremely simple C program and build up from there. I will assume that you are using the Linux command line and clang as your environment for these examples; if you are not, all of the code will still work fine – you will simply need to use whatever compiler you have available.

The simplest C program, I

Let's start with the simplest possible C program and use it both to understand the basics of C and the C compilation process. Type the following program into a standard text editor. Then save the program to a file named **hello.c**. If you leave off **.c**, you will probably get some sort of error when you compile it, so make sure you remember the **.c**. Here's the first program:

Listing 1.1: The simplest C program

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     printf ("Hello, World!\n");
6
7     return 0;
8 }
```

Spacing and indentation

When you enter this program, position **#include** so that the pound sign is in column 1. Otherwise, the spacing and indentation can be any way you like it. On many Linux systems, you will find a program called **indent**, which will format code for you. The spacing and indentation shown above is a good example to follow.

Compilation and run

When executed, this program instructs the computer to print out the line “Hello, World!” – then the program quits.

To compile this code on a Linux machine, type

```
1 clang hello.c -o hello
```

This line invokes the C compiler called clang, asks it to compile the code stored in the file `hello.c` and asks it to place the executable file it creates under the name **hello**. To run the program, type

```
1 ./hello
```

You should see the output “Hello, World!” when you run the program.

If you mistype the program, it either will not compile or it will not run. If the program does not compile or does not run correctly, edit it again and see where you went wrong in your typing. Fix the error and try again.

The simplest C program, II

Let’s walk through this program and start to see what the different lines are doing:

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     printf ("Hello, World!\n");
6
7     return 0;
8 }
```

- This C program starts with **#include <stdio.h>**. This line *includes* the “**standard input/output library**” into your program. The standard I/O library lets you read input from the keyboard (called “standard in”), write output to the screen (called “standard out”), process text files stored on the disk, and so on. It is an extremely useful library. C has a large number of standard libraries like `stdio`, including `string`, `time` and `math` libraries. A *library* is simply a package of code that someone else has written to make your life easier (we’ll discuss libraries a bit later).
- The line **int main (void)** declares the main function. Every C program must have a function named **main** somewhere in the code. We will learn more about functions shortly. At run time, program execution starts at the first line of the main function.

- In C, the { and } symbols mark the beginning and end of a block of code. In this case, the block of code making up the main function contains two lines.
- The **printf** statement in C allows you to send output to standard out (for us, the screen). The portion in quotes is called the *format string* and describes how the data is to be formatted when printed. The format string can contain string literals such as “Hello, World!”, symbols for carriage returns (\n), and operators as placeholders for variables (see below).
- The **return 0;** line causes the function to return an error code of 0 (no error) to the shell that started execution. More on this capability a bit later.

Variables and variable types

As a programmer, you will frequently want your program to “remember” a value. For example, if your program requests a value from the user, or if it calculates a value, you will want to remember it somewhere so you can use it later. The way your program remembers things is by using *variables*. For example:

```
1  int b;
```

This line says, “I want to create a space called **b** that is able to hold one integer value.” A variable has a *name* (in this case, **b**) and a *type* (in this case, **int**, an integer). You can store a value in **b** by saying something like:

```
1  b = 5;
```

You can use the value in **b** by saying something like:

```
1  printf ("%d", b);
```

In C, there are several basic types for variables:

- **int** - integer (whole number) values
- **double** - floating point values
- **char** - single character values (such as “m” or “Z”)

There are a number of derivative types:

- **long**
- **short**

- **unsigned int** - positive integers

We will see examples of these and other types as we go along.

C allows you to perform type conversions between different types of variables, so called *typecasting*. Typecasting automatically occurs during the assignment operation, e.g. when an integer value is assigned to a floating point variable.

You do explicit typecasting in C by placing the type name in parentheses and putting it in front of the value you want to change.

The following code

Listing 1.2: Integer division and typecasting

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     double a, b;
6
7     a = 10 / 3;
8     b = (double) 10 / 3;
9     printf ("a = %f      b = %f\n", a, b);
10
11     return 0;
12 }
```

prints

```
1 a = 3.000000      b = 3.333333
```

since variable **a** is assigned the result of integer division 10/3, which is integer 3, whereas variable **b** is assigned the result of floating point division, 3.33333, because 10 is converted to a floating point value before the division.

You can declare **user-defined types** in C with the *typedef* statement. For example, if you do not like the word “double” for real numbers, you can say:

```
1 typedef double real;
```

and then later in your code use:

```
1 real r1, r2, r3;
```

You can place typedef statements anywhere in a C program as long as they come prior to their first use in the code.

Operators

Many operators in C are similar to operators in most languages:

```
1 + - addition
2 - - subtraction
3 / - division
4 * - multiplication
5 % - mod
```

The `/` operator performs integer division if both operands are integers, and performs floating point division otherwise. For example:

Operator precedence in C is similar to that in most other languages. Division and multiplication occur first, then addition and subtraction. The result of the calculation `5+3*4` is 17, not 32, because the `*` operator has higher precedence than `+` in C. You can use parentheses to change the normal precedence ordering: `(5+3)*4` is 32. The `5+3` is evaluated first because it is in parentheses. We'll get into precedence later – it becomes somewhat complicated in C once pointers are introduced.

Incremental operators are important feature of C.

	Long Way	Short Way
1		
2	<code>i = i + 1;</code>	<code>i++;</code>
3	<code>i = i - 1;</code>	<code>i--;</code>
4	<code>i = i + 3;</code>	<code>i += 3;</code>
5	<code>i = i * j;</code>	<code>i *= j;</code>

Chapter 2

Input and output

printf

The `printf` statement allows you to send output to *standard out*. For us, standard out is generally the screen (although you can redirect standard out into a text file or to the input of another command).

Here is another program that will help you learn more about `printf`:

Listing 2.1: Using `printf`

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     int a, b, c;
6
7     a = 5;
8     b = 7;
9     c = a + b;
10    printf ("%d + %d = %d\n", a, b, c);
11
12    return 0;
13 }
```

Type this program into a file and save it as **add.c**. Compile it with the line **clang add.c -o add** and then run it by typing **./add**. You will see the line “5 + 7 = 12” as output.

Here is an explanation of the different lines in this program:

- The line **int a, b, c;** declares three integer variables named **a**, **b** and **c**. Integer variables hold whole numbers.
- The next line initializes the variable named **a** to the value 5.
- The next line sets **b** to 7.
- The next line adds **a** and **b** and “assigns” the result to **c**.

The computer adds the value in **a** (5) to the value in **b** (7) to form the result 12, and then places that new value (12) into the variable **c**. The variable **c** is assigned the value 12. For this reason, the = in this line is called “the assignment operator.”

- The **printf** statement then prints the line “5 + 7 = 12.” The **%d** placeholders in the printf statement act as placeholders for values. There are three **%d** placeholders, and at the end of the printf line there are the three variable names: **a**, **b** and **c**. **c** matches up the first **%d** with **a** and substitutes 5 there. It matches the second **%d** with **b** and substitutes 7. It matches the third **%d** with **c** and substitutes 12. Then it prints the completed line to the screen: 5 + 7 = 12. The **+**, the **=** and the spacing are a part of the format line and get embedded automatically between the **%d** operators as specified by the programmer.

Let’s look at some variations to understand printf completely. Here is the simplest printf statement:

```
1 printf ("Hello");
```

This call to printf has a format string that tells printf to send the word “Hello” to standard out. Contrast it with this:

```
1 printf ("Hello\n");
```

The difference between the two is that the second version sends the word “Hello” followed by a carriage return to standard out.

The following line shows how to **output the value of a variable using printf**.

```
1 printf ("%d", b);
```

The **%d** is a placeholder that will be replaced by the value of the variable **b** when the printf statement is executed. Often, you will want to embed the value within some other words. One way to accomplish that is like this:

```
1 printf ("The temperature is ");  
2 printf ("%d", b);
```

```
3 printf (" degrees\n");
```

An easier way is to say this:

```
1 printf ("The temperature is %d degrees\n", b);
```

You can also use multiple `%d` placeholders in one `printf` statement:

```
1 printf ("%d + %d = %d\n", a, b, c);
```

In the `printf` statement, it is extremely important that the number of *operators* in the format string corresponds exactly with the number and type of the variables following it. For example, if the format string contains three `%d` operators, then it must be followed by exactly three parameters and they must have the same types in the same order as those specified by the operators.

You can print all of the normal C types with `printf` by using different placeholders:

- **int** (integer values) uses **%d**
- **double** (floating point values) uses **%f**
- **char** (single character values) uses **%c**
- **character strings** (arrays of characters, discussed later) use **%s**

You can learn more about the nuances of `printf` on a Linux machine by typing **man 3 printf**.

scanf

The previous program is good, but it would be better if it read in the values 5 and 7 from the user instead of using constants. Try this program instead:

Listing 2.2: Using `scanf`

```
1 #include <stdio.h>
2
3 int main(void)
4 {
5     int a, b, c;
6
7     printf("Enter the first value:");
8     scanf("%d", &a);
```

```
9     printf("Enter the second value:");
10     scanf("%d", &b);
11     c = a + b;
12     printf("%d + %d = %d\n", a, b, c);
13
14     return 0;
15 }
```

Make the changes, then compile and run the program to make sure it works. Note that `scanf` uses the same sort of format string as `printf` (type **man scanf** for more info). Also note the `&` in front of `a` and `b`. This is the *address operator* in C: It returns the address of the variable (this will not make sense until we discuss pointers). You must use the `&` operator in `scanf` on any variable of type `char`, `int`, or `float`, as well as structure types (which we will get to shortly). If you leave out the `&` operator, you will get an error when you run the program. Try it so that you can see what that sort of run-time error looks like.

The `scanf` function allows you to accept input from standard in, which for us is generally the keyboard. The `scanf` function can do a lot of different things, but it is generally unreliable unless used in the simplest ways. It is unreliable because it does not handle human errors very well. But for simple programs it is good enough and easy-to-use.

The simplest application of **scanf** looks like this:

```
1     scanf ("%d", &b);
```

The program will read in an integer value that the user enters on the keyboard (`%d` is for integers, as is `printf`, so `b` must be declared as an `int`) and place that value into `b`.

The **scanf** function uses the same placeholders as **printf**:

- **int** uses `%d`
- **double** uses `%f`
- **char** uses `%c`
- **character strings** (discussed later) use `%s`

You must put `&` in front of the variable used in `scanf`. The reason why will become clear once you learn about **pointers**. It is easy to forget the `&` sign, and when you forget it your program will almost always crash when you run it.

The `printf` and `scanf` functions will take a bit of practice to be completely understood, but once mastered they are extremely useful.

Programming exercise

1. Modify the following program so that it accepts three values instead of two and adds all three together:

Listing 2.3: Addition of three numbers

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     int a, b, c;
6
7     printf ("Enter the first value:");
8     scanf ("%d", &a);
9     printf ("Enter the second value:");
10    scanf ("%d", &b);
11    c = a + b;
12    printf ("%d + %d = %d\n", a, b, c);
13
14    return 0;
15 }
```

2. Try deleting or adding random characters or words in one of the previous programs and watch how the compiler reacts to these errors.

For example, delete the `b` variable in the first line of the above program and see what the compiler does when you forget to declare a variable. Delete a semicolon and see what happens. Leave out one of the braces. Remove one of the parentheses next to the main function. Make each error by itself and then run the program through the compiler to see what happens. By simulating errors like these, you can learn about different compiler errors, and that will make your typos easier to find when you make them for real.

Chapter 3

Branching and looping

if statement

Here is a simple C program demonstrating an if statement:

Listing 3.1: if statement

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     int b;
6
7     printf ("Enter a value:");
8     scanf ("%d", &b);
9     if (b < 0)
10    {
11        printf ("The value is negative\n");
12    }
13
14    return 0;
15 }
```

This program accepts a number from the user. It then tests the number using an if statement to see if it is less than 0. If it is, the program prints a message. Otherwise, the program is silent. The (**b < 0**) portion of the program is the *Boolean expression* or *logical expression*. C evaluates this expression

to decide whether or not to print the message. If the Boolean expression evaluates to **true**, then C executes the single line immediately following the if statement (or a block of lines within braces immediately following the if statement). If the Boolean expression is **false**, then C skips the line or block of lines immediately following the if statement.

Here's slightly more complex example:

Listing 3.2: if else statement

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     int b;
6
7     printf ("Enter a value:");
8     scanf ("%d", &b);
9
10    if (b < 0)
11    {
12        printf ("The value is negative\n");
13    }
14    else if (b == 0)
15    {
16        printf ("The value is zero\n");
17    }
18    else
19    {
20        printf ("The value is positive\n");
21    }
22
23    return 0;
24 }
```

In this example, the **else if** and **else** sections evaluate for zero and positive values as well.

Boolean expressions

Here is a more complicated Boolean expression:

```
1  if ((x == y) && (j > k))
2  {
3      z = 1;
4  }
5  else
6  {
7      q = 10;
8  }
```

This statement says, “If the value in variable x equals the value in variable y, and if the value in variable j is greater than the value in variable k, then set the variable z to 1, otherwise set the variable q to 10.” You will use if statements like this throughout your C programs to make decisions. In general, most of the decisions you make will be simple ones like the first example; but on occasion, things get more complicated.

Notice that C uses **==** to **test for equality**, while it uses **=** to **assign a value** to a variable. The **&&** in C represents a Boolean AND operation.

Here are all of the Boolean operators in C:

1	equality	==
2	less than	<
3	Greater than	>
4	<=	<=
5	>=	>=
6	inequality	!=
7	and	&&
8	or	
9	not	!

= vs == in boolean expressions

The **==** sign is a problem in C because every now and then you may forget and type just **=** in a Boolean expression. This is an easy mistake to make, but to the compiler there is a very important difference. C will accept either **=** and **==** in a Boolean expression – the behavior of the program changes remarkably between the two, however.

Boolean expressions evaluate to integers in C, and integers can be used inside of Boolean expressions. The integer value 0 in C is False, while any other integer value is True. The following is legal in C:

Listing 3.3: Boolean vs integer in if statement

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     int a;
6
7     printf ("Enter an integer number:");
8     scanf ("%d", &a);
9
10    if (a)
11    {
12        printf ("The value is true\n");
13    }
14
15    return 0;
16 }
```

If **a** is anything other than 0, the printf statement gets executed.

In C, a statement like **if (a = b)** means, "Assign **b** to **a**, and then test **a** for its Boolean value." So if **a** becomes 0, the if statement is False; otherwise, it is True. The value of **a** changes in the process. This is not the intended behavior if you meant to type **==** (although this feature is useful when used correctly), so be careful with your **=** and **==** usage.

while loop

You'll find that **while** statements are just as easy to use as if statements. For example:

```
1 while (a < b)
2 {
3     printf ("%d\n", a);
4     a = a + 1;
5 }
```

This causes the two lines within the braces to be executed repeatedly until **a** is greater than or equal to **b**. The while statement in general works as illustrated to the right.

do-while loop

C also provides a **do-while** structure:

```
1 do
2 {
3     printf ("%d\n", a);
4     a = a + 1;
5 }
6 while (a < b);
```

for loop

The **for loop** in C is simply a shorthand way of expressing a while statement. For example, suppose you have the following code in C:

```
1 x = 1;
2 while (x < 10)
3 {
4     blah blah blah
5     x++; /* x++ is the same as saying x = x+1 */
6 }
```

You can convert this into a for loop as follows:

```
1 for (x = 1; x < 10; x++)
2 {
3     blah blah blah
4 }
```

Note that the while loop contains an initialization step (**x=1**), a test step (**x<10**), and an increment step (**x++**). The for loop lets you put all three parts onto one line, but you can put anything into those three parts. For example, suppose you have the following loop:

```
1 a = 1;
2 b = 6;
3 while (a < b)
4 {
5     a++;
6     printf ("%d\n", a);
7 }
```

You can place this into a for statement as well:

```
1 for (a = 1, b = 6; a < b; a++)
2 {
3     printf ("%d\n", a);
4 }
```

It is slightly confusing, but it is possible. The *comma operator* lets you separate several different statements in the initialization and increment sections of the for loop (but not in the test section). Many C programmers like to pack a lot of information into a single line of C code; but a lot of people think it makes the code harder to understand, so they break it up.

Looping: an example

Let's say that you would like to create a program that prints a Fahrenheit-to-Celsius conversion table. This is easily accomplished with a for loop or a while loop:

Listing 3.4: Fahrenheit-to-Celsius conversion using integers

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     int a = 0;
6
7     while (a <= 100)
8     {
9         printf ("%4d degrees F = %4d degrees C\n",
10             a, (a - 32) * 5 / 9);
11         a = a + 10;
```

```
12     }
13
14     return 0;
15 }
```

If you run this program, it will produce a table of values starting at 0 degrees F and ending at 100 degrees F. The output will look like this:

```
1    0 degrees F = -17 degrees C
2    10 degrees F = -12 degrees C
3    20 degrees F =  -6 degrees C
4    30 degrees F =  -1 degrees C
5    40 degrees F =   4 degrees C
6    50 degrees F =  10 degrees C
7    60 degrees F =  15 degrees C
8    70 degrees F =  21 degrees C
9    80 degrees F =  26 degrees C
10   90 degrees F =  32 degrees C
11  100 degrees F =  37 degrees C
```

The table's values are in increments of 10 degrees. You can see that you can easily change the starting, ending or increment values of the table that the program produces.

If you wanted your values to be more accurate, you could use **floating point** values instead:

Listing 3.5: Farenheit-to-Celsius conversion using floating point numbers, I

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     double a = 0.;
6
7     while (a <= 100.)
8     {
9         printf ("%6.2f degrees F = %6.2f degrees C\n",
10             a, (a - 32.) * 5. / 9.);
11         a += 10.;
12     }
13
14     return 0;
```

15 }

You can see that the declaration for **a** has been changed to a float, and the **%f** symbol replaces the **%d** symbol in the printf statement. In addition, the **%f** symbol has some formatting applied to it: The value will be printed with six digits preceding the decimal point and two digits following the decimal point.

Now let's say that we wanted to modify the program so that the temperature 98.6F is inserted in the table at the proper position. That is, we want the table to increment every 10 degrees, but we also want the table to include an extra line for 98.6 degrees F because that is the normal body temperature for a human being. The following program accomplishes the goal:

Listing 3.6: Farenheit-to-Celsius conversion using floating point numbers, II, with a bug

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     double a = 0.;
6     double tbody = 98.6;
7
8     while (a <= 100.)
9     {
10         if (a > tbody)
11         {
12             printf ("%6.2f degrees F = %6.2f degrees C\n",
13                     tbody, (tbody - 32.) * 5. / 9.);
14         }
15         printf ("%6.2f degrees F = %6.2f degrees C\n",
16                 a, (a - 32.) * 5. / 9.);
17         a += 10.;
18     }
19
20     return 0;
21 }
```

This program works if the ending value is 100, but if you change the ending value to 200 you will find that the program has a **bug**. It prints the line for 98.6 degrees too many times. We can fix that problem in several different ways. Here is one way:

Listing 3.7: Farenheit-to-Celsius conversion using floating point numbers, II, corrected

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     double a, b, tbody;
6     a = 0.;
7     b = a;
8     tbody = 98.6;
9
10    while (a <= 150.)
11    {
12        if ((a > tbody) && (b < tbody))
13        {
14            printf ("%6.2f degrees F = %6.2f degrees C\n",
15                    tbody, (tbody - 32.) * 5. / 9.);
16        }
17        printf ("%6.2f degrees F = %6.2f degrees C\n",
18                a, (a - 32.) * 5. / 9.);
19        b = a;
20        a += 10;
21    }
22
23    return 0;
24 }
```

The output will look like this:

```
1 0.00 degrees F = -17.78 degrees C
2 10.00 degrees F = -12.22 degrees C
3 20.00 degrees F = -6.67 degrees C
4 30.00 degrees F = -1.11 degrees C
5 40.00 degrees F = 4.44 degrees C
6 50.00 degrees F = 10.00 degrees C
7 60.00 degrees F = 15.56 degrees C
8 70.00 degrees F = 21.11 degrees C
9 80.00 degrees F = 26.67 degrees C
10 90.00 degrees F = 32.22 degrees C
11 98.60 degrees F = 37.00 degrees C
12 100.00 degrees F = 37.78 degrees C
13 110.00 degrees F = 43.33 degrees C
```

```
14 120.00 degrees F = 48.89 degrees C
15 130.00 degrees F = 54.44 degrees C
16 140.00 degrees F = 60.00 degrees C
17 150.00 degrees F = 65.56 degrees C
```

Programming exercise

1. Try changing the Fahrenheit-to-Celsius program so that it adds a heading line to the table that is produced.
2. Try to find a different solution to the bug fixed by the previous example.
3. Create a table that converts pounds to kilograms or miles to kilometers.

Chapter 4

Functions

Functions let you chop up a long program into named sections so that the sections can be reused throughout the program. Functions accept **parameters** and **return** a result. C functions can accept an unlimited number of parameters. In general, C does not care in what order you put your functions in the program, so long as the function name is known to the compiler before it is called.

We have already talked a little about functions. **printf** and **main** are examples of functions.

Another example is the **rand** function below. It accepts no parameters and returns an integer result:

```
1 unsigned int rand(void)
2 /*
3  * rand() generates a pseudo random integer between 0 and 32767
4  * inclusive.
5  */
6 {
7     rand_seed = rand_seed * 1103515245 + 12345;
8     return (unsigned int)(rand_seed / 65536) % 32768;
9 }
```

The **unsigned int rand()** line declares the function **rand** to the rest of the program and specifies that **rand** will accept no parameters and return a positive integer result. This function has no local variables, but if it needed locals, they would go right below the opening **{** (C allows you to declare variables after any **{** – they exist until the program reaches the matching **}** and then they disappear. A function’s local variables therefore vanish as soon as the matching **}** is reached in the function. While they exist, local variables live on the system stack.) Note that there is no **;** after the **()** in the first line. Also note that even though there are no parameters, you must use the **(void)**. It tells the compiler that you are declaring a function rather than simply declaring an **int**.

The **return** statement is important to any function that returns a result. It specifies the value that the function will return and causes the function to exit immediately. This means that you can place multiple return statements in the function to give it multiple exit points. If you do not place a return statement in a function, the function returns when it reaches `}` and returns a random value (many compilers will warn you if you fail to return a specific value). In C, a function can return values of any type: `int`, `double`, `char`, etc.

There are several ways to call the **rand** function. For example: `x = rand();`. The variable `x` is assigned the value returned by `rand` in this statement. Note that you *must* use `()` in the function call, even though no parameter is passed. Otherwise, `x` is given the memory address of the `rand` function, which is generally not what you intended.

You might also call `rand` this way:

```
1 if (rand() > 100)
2 {
3     ....
4 }
```

Or this way:

```
1 rand();
```

In the latter case, the function is called but the value returned by `rand` is discarded. You may never want to do this with `rand`, but many functions return some kind of error code through the function name, and if you are not concerned with the error code (for example, because you know that an error is impossible) you can discard it in this way.

Functions can use a void return type if you intend to return nothing. For example:

```
1 void print_header()
2 {
3     printf ("Program Number 1\n");
4     printf ("Version 1.0, released 12/12/2012\n");
5 }
```

This function returns no value. You can call it with the following statement:

```
1 print_header();
```

C functions can accept parameters of any type. For example:

```
1 int fact(int i)
2 {
3     int j, k;
4
5     j = 1;
6     for (k = 2; k <= i; k++)
7     {
8         j = j*k;
9     }
10    return j;
11 }
```

returns the factorial of **i**, which is passed in as an integer parameter. Separate multiple parameters with commas:

```
1 int add (int i, int j)
2 {
3     return i + j;
4 }
```

Function prototypes

It is considered good practice to use *function prototypes* for all functions in your program. A prototype declares the function name, its parameters, and its return type to the rest of the program prior to the function's actual declaration. To understand why function prototypes are useful, enter the following code and run it:

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     printf ("%d\n", add(3));
6     return 0;
7 }
8
9 int add(int i, int j)
10 {
```

```
11     return i+j;
12 }
```

This code compiles on some compilers without giving you a warning, even though **add** expects two parameters but receives only one. It works because many C compilers do not check for parameter matching either in type or count. You can waste an enormous amount of time debugging code in which you are simply passing one too many or too few parameters by mistake. The above code compiles properly, but it produces the wrong answer.

To solve this problem, C lets you place function prototypes at the beginning of (actually, anywhere in) a program. If you do so, C checks the types and counts of all parameter lists. Try compiling the following:

```
1 #include <stdio.h>
2
3 int add (int, int); /* function prototype for add */
4
5 int main (void)
6 {
7     printf ("%d\n", add(3));
8     return 0;
9 }
10
11 int add(int i, int j)
12 {
13     return i + j;
14 }
```

The prototype causes the compiler to flag an error on the **printf** statement.

Place one prototype for each function at the beginning of your program. They can save you a great deal of debugging time, and they also solve the problem you get when you compile with functions that you use before they are declared.

Chapter 5

Arrays

In this section, we will create a small C program that generates 10 random numbers and sorts them. To do that, we will use a new variable arrangement called an *array*.

An array lets you declare and work with a collection of values of the same type. For example, you might want to create a collection of five integers. One way to do it would be to declare five integers directly:

```
1 int a, b, c, d, e;
```

This is okay, but what if you needed a thousand integers? An easier way is to declare an array of five integers:

```
1 int a[5];
```

The five separate integers inside this array are accessed by an **index**. All arrays start at index zero and go to n-1 in C. Thus, **int a[5];** contains five elements. For example:

```
1 int a[5];  
2  
3 a[0] = 12;  
4 a[1] = 9;  
5 a[2] = 14;  
6 a[3] = 5;  
7 a[4] = 1;
```

One of the nice things about array indexing is that you can use a loop to manipulate the index. For example, the following code initializes all of the values in the array to 0:

```
1 int a[5];
2 int i;
3 for (i = 0; i < 5; i++)
4 {
5     a[i] = 0;
6 }
```

You declare arrays by inserting an array size after a normal declaration, as shown below:

```
1 int a[10];           /* array of integers */
2 char s[100];        /* array of characters */
3 double f[20];       /* array of floating point numbers */
```

The following code initializes the values in the array sequentially and then prints them out:

Listing 5.1: Using arrays

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     int a[5];
6     int i;
7
8     for (i = 0; i < 5; i++)
9     {
10         a[i] = i;
11     }
12     for (i = 0; i < 5; i++)
13     {
14         printf ("a[%d] = %d\n", i, a[i]);
15     }
16
17     return 0;
18 }
```

Programming example: sorting arrays

Arrays are used all the time in C. To understand a common usage, consider the following code:

Listing 5.2: Sorting arrays, I

```
1 #include <stdio.h>
2
3 #define MAX 10
4
5 static int a[MAX];
6 static unsigned int rand_seed=10;
7
8 unsigned int rand(void);
9
10 int main (void)
11 {
12     int i;
13
14     /* fill array */
15     for (i = 0; i < MAX; i++)
16     {
17         a[i] = (int) rand();
18         printf ("%d\n", a[i]);
19     }
20
21     /* more stuff will go here in a minute */
22
23     return 0;
24 }
25
26 /*
27 * rand() generates a pseudo random integer between 0 and 32767
28 inclusive.
29 */
30 unsigned int rand (void)
31 {
32     rand_seed = rand_seed * 1103515245 + 12345;
33     return (unsigned int) (rand_seed / 65536) % 32768;
34 }
```

This code contains several new concepts. The **#define** line declares a constant named **MAX** and sets it to 10. Constant names are traditionally written in all caps to make them obvious in the code. The line **int a[MAX];** shows you how to declare an array of integers in C. Note that because of the position of the array's declaration, it is global to the entire program.

The line **unsigned int rand_seed = 10** also declares a global variable, this time named **rand_seed**, that is initialized to 10 each time the program begins. This value is the starting seed for the random number code that follows. The **rand** function will produce the same values each time you run the program with the same seed.

The line **unsigned int rand(void)** is a function declaration. The **rand** function accepts no arguments and returns an unsigned integer value. We will learn more about functions later. The four lines that follow implement the **rand** function. We will ignore them for now.

The main function is normal. Four local integers are declared, and the array is filled with 10 random values using a for loop. Note that the array **a** contains 10 individual integers. You point to a specific integer in the array using square brackets. So **a[0]** refers to the first integer in the array, **a[1]** refers to the second, and so on. The line starting with **/*** and ending with ***/** is called a **comment**. The compiler completely ignores the line. You can place notes to yourself or other programmers in comments.

When compiled and run, the code produces the following output:

```
1 4543
2 28214
3 11245
4 8870
5 16887
6 2234
7 20162
8 27592
9 20255
10 30710
```

Now add more code, so that the program looks as following:

Listing 5.3: Sorting arrays, I

```
1 #include <stdio.h>
2
3 #define MAX 10
4
5 static int a[MAX];
```

```
6 static unsigned int rand_seed=10;
7
8 unsigned int rand(void);
9 void bubblesort(void);
10
11 int main (void)
12 {
13     int i;
14
15     /* fill array */
16     for (i = 0; i < MAX; i++)
17     {
18         a[i] = (int) rand();
19         printf ("%d\n", a[i]);
20     }
21
22     /*
23      * sort
24      */
25     bubblesort();
26
27     /*
28      * print sorted array
29      */
30     printf ("-----\n");
31     for (i = 0; i < MAX; i++)
32     {
33         printf ("%d\n", a[i]);
34     }
35
36     return 0;
37 }
38
39 /*
40  * generate a pseudo random integer between 0 and 32767 inclusive
41  *
42  */
43 unsigned int rand (void)
44 {
```



```
44     rand_seed = rand_seed * 1103515245 + 12345;
45     return (unsigned int) (rand_seed / 65536) % 32768;
46 }
47 /*
48  * bubble sort the array
49  */
50 void bubblesort (void)
51 {
52     int i, j, t;
53
54     for (i = 0; i < MAX-1; i++)
55     {
56         for (j = 0; j < MAX-i-1; j++)
57         {
58             if (a[j] > a[j+1])
59             {
60                 t = a[j];
61                 a[j] = a[j+1];
62                 a[j+1] = t;
63             }
64         }
65     }
66 }
```

This code **sorts** the random values and prints them in sorted order. Each time you run it, you will get the same values. If you would like to change the values that are sorted, change the value of `rand_seed` each time you run the program.

The output on the new version is the original unsorted array, followed by the sorted one.

```
1 4543
2 28214
3 11245
4 8870
5 16887
6 2234
7 20162
8 27592
9 20255
10 30710
```

```
11 -----  
12 2234  
13 4543  
14 8870  
15 11245  
16 16887  
17 20162  
18 20255  
19 27592  
20 28214  
21 30710
```

A good way to understand what this code is doing is to execute it “by hand.” That is, assume **MAX** is 4 to make it a little more manageable, take out a sheet of paper and pretend you are the computer. Draw the array on your paper and put four random, unsorted values into the array. Execute each line of the sorting section of the code and draw out exactly what happens. You will find that, each time through the inner loop, the larger values in the array are pushed toward the bottom of the array and the smaller values bubble up toward the top.

Chapter 6

Structures

Structures in C allow you to group variable into a package. Here's an example:

```
1 struct rec
2 {
3     int a, b, c;
4     double d, e, f;
5 };
6
7 struct rec r;
```

As shown here, whenever you want to declare structures of the type **rec**, you have to say **struct rec**. This line is very easy to forget, and you get many compiler errors because you absent-mindedly leave out the **struct**. You can compress the code into the form:

```
1 struct rec
2 {
3     int a;
4     double b, c;
5 } r;
```

where the type declaration for **rec** and the variable **r** are declared in the same statement. Or you can create a typedef statement for the structure name. For example, if you do not like saying **struct rec r** every time you want to declare a record, you can say:

```
1 typedef struct rec rec_type;
```

and then declare records of type **rec_type** by saying:

```
1 rec_type r;
```

You access fields of structure using a period, for example, **r.a = 5;**

You can initialize a structure as following:

```
1 struct rec r1 = (struct rec) { 1, 2., 3. };
```

XS The C language allows you to pass an entire structure to a function. In addition, a function can return a structure back to its caller.

Listing 6.1: Initializing structures and passing structures to functions

```
1 #include <stdio.h>
2 #include <math.h>
3
4 struct point2d
5 {
6     double x;
7     double y;
8 };
9
10 struct triangle
11 {
12     struct point2d a;
13     struct point2d b;
14     struct point2d c;
15 };
16
17 double area(struct triangle t);
18
19
20 /*
21  * Calculate the area of a triangle
22  */
23 double area (struct triangle t)
24 {
25     return fabs(0.5*( (t.c.x - t.a.x)*(t.b.y - t.a.y) - (t.c.y - t.
26         a.y)*(t.b.x - t.a.x))) ;
27 }
```

```
28 int main(void)
29 {
30     struct triangle t1, t2;
31     double s;
32
33     /*
34      * initialization
35      */
36     t1.a = (struct point2d) {0., 0.};
37     t1.b = (struct point2d) {3., 0.};
38     t1.c = (struct point2d) {0., 4.};
39
40     s = area(t1);
41     printf("Area: %14.6f\n", s);
42
43     /*
44      * initialization
45      */
46     t2 = (struct triangle) { {0., 0.}, {5., 0.}, {0., 12.} };
47     s = area(t2);
48
49     printf("Area: %14.6f\n", s);
50
51     return 0;
52 }
```


Chapter 7

Pointers

Pointers are used everywhere in C, so if you want to use the C language fully you have to have a very good understanding of pointers. The goal of this section and the next several that follow is to help you build a complete understanding of pointers and how C uses them. For most people it takes a little time and some practice to become fully comfortable with pointers, but once you master them you are a full-fledged C programmer.

C uses pointers in three different ways:

- C uses pointers to create **dynamic data structures** – data structures built up from blocks of memory allocated at run-time.
- C uses pointers to handle **variable parameters** passed to functions.
- Pointers in C provide a way to **access information stored in arrays**. Pointer techniques are especially valuable when you work with strings. There is an intimate link between arrays and pointers in C.

In some cases, C programmers also use pointers because they make the code slightly more efficient. What you will find is that, once you are completely comfortable with pointers, you tend to use them all the time.

We will start this discussion with a basic introduction to pointers and the concepts surrounding pointers, and then move on to the three techniques described above.

Pointers: why?

Imagine that you would like to create a **text editor** – a program that lets you edit normal text files, like “vim” on UNIX or “Notepad” on Windows. A text editor is a fairly common thing for someone to create because, if you think about it, a text editor is probably a programmer’s most commonly used piece of software. The text editor is a programmer’s intimate link to the computer – it is where you enter all of your thoughts and then manipulate them. Obviously, with anything you use that often and work with that closely, you want it to be just right. Therefore many programmers create their own editors and customize them to suit their individual working styles and preferences.

So one day you sit down to begin working on your editor. After thinking about the features you want, you begin to think about the “data structure” for your editor. That is, you begin thinking about how you will store the document you are editing in memory so that you can manipulate it in your program. What you need is a way to store the information you are entering in a form that can be manipulated quickly and easily. You believe that one way to do that is to organize the data on the basis of lines of characters. Given what we have discussed so far, the only thing you have at your disposal at this point is an array. You think, “Well, a typical line is 80 characters long, and a typical file is no more than 1,000 lines long.” You therefore declare a **two-dimensional array**, like this:

```
1 char doc[1000][80];
```

This declaration requests an array of 1,000 80-character lines. This array has a total size of 80,000 characters.

As you think about your editor and its data structure some more, however, you might realize three things:

- Some documents are long lists. Every line is short, but there are thousands of lines.
- Some special-purpose text files have very long lines. For example, a certain data file might have lines containing 542 characters, with each character representing the amino acid pairs in segments of DNA.
- In most modern editors, you can open multiple files at one time.

Let’s say you set a maximum of 10 open files at once, a maximum line length of 1,000 characters and a maximum file size of 50,000 lines. Your declaration now looks like this:

```
1 char doc[50000][1000][10];
```

That doesn’t seem like an unreasonable thing, until you pull out your calculator, multiply 50,000 by 1,000 by 10 and realize the array contains 500 million characters! Even if the computer would accept a request for such a large array, you can see that it is an extravagant waste of space. It seems

strange to declare a 500 million character array when, in the vast majority of cases, you will run this editor to look at 100 line files that consume at most 4,000 or 5,000 characters. The problem with an array is the fact that *you have to declare it to have its maximum size in every dimension from the beginning*. Those maximum sizes often multiply together to form very large numbers. Also, if you happen to need to be able to edit an odd file with a 2,000 character line in it, you are out of luck. There is really no way for you to predict and handle the maximum line length of a text file, because, technically, that number is infinite.

Pointers are designed to solve this problem. With pointers, you can create *dynamic data structures*. Instead of declaring your worst-case memory consumption up-front in an array, you instead *allocate* memory from *the heap* while the program is running. That way you can use the exact amount of memory a document needs, with no waste. In addition, when you close a document you can return the memory so that other parts of the program can use it. With pointers, memory can be recycled while the program is running.

Pointer Basics

To understand pointers, it helps to compare them to normal variables.

A “normal variable” is a location in memory that can hold a value. For example, when you declare a variable `i` as an integer, a chunk of memory are set aside for it. In your program, you refer to that location in memory by the name `i`. At the machine level that location has a memory address. The memory at that address are known to you, the programmer, as `i`, and it can hold one integer value.

A pointer is different. A pointer is a variable that **points** to another variable. This means that a pointer holds the memory address of another variable. Put another way, the pointer does not hold a value in the traditional sense; instead, it holds the address of another variable. A pointer “points to” that other variable by holding a copy of its address.

Because a pointer holds an address rather than a value, it has two parts. The pointer itself holds the address. That address points to a value. There is the pointer and the value pointed to. This fact can be a little confusing until you get comfortable with it, but once you get comfortable it becomes extremely powerful.

The following example code shows a typical pointer:

Listing 7.1: Using pointers, I

```
1 #include <stdio.h>
2
3 int main (void)
```

```
4 {
5     int i, j;
6     int *p;                /* a pointer to an integer */
7
8     p = &i;
9     *p = 5;
10    j = i;
11
12    printf ("%d %d %d\n", i, j, *p);
13
14    return 0;
15 }
```

The first declaration in this program declares two normal integer variables named **i** and **j**. The line **int *p** declares a pointer named **p**. This line asks the compiler to declare a variable **p** that is a **pointer** to an integer. The ***** indicates that a pointer is being declared rather than a normal variable. You can create a pointer to anything: a double, a structure, a char, and so on. Just use a ***** to indicate that you want a pointer rather than a normal variable.

The line **p = &i;** will definitely be new to you. In C, **&** is called the **address operator**. The expression **&i** means, "The memory address of the variable **i**." Thus, the expression **p = &i;** means, "Assign to **p** the address of **i**." Once you execute this statement, **p** "points to" **i**. Before you do so, **p** contains a random, unknown address, and its use will likely cause a segmentation fault or similar program crash.

In the program above the three variables **i**, **j** and **p** have been declared, but none of the three has been initialized. Once **p** points to **i**, the memory location **i** has two names. It is still known as **i**, but now it is known as ***p** as well. This is how C talks about the two parts of a pointer variable: **p** is the location holding the address, while ***p** is the location pointed to by that address. Therefore ***p=5** means that the location pointed to by **p** should be set to 5. Because the location ***p** is also **i**, **i** also takes on the value 5. Consequently, **j = i;** sets **j** to 5, and the **printf** statement produces as following:

```
1 5 5 5
```

The main feature of a pointer is its two-part nature. The pointer itself holds an address. The pointer also points to a value of a specific type - the value at the address the point holds. The pointer itself, in this case, is **p**. The value pointed to is ***p**.

Understanding Memory Addresses

The previous discussion becomes a little clearer if you understand how memory addresses work in a computer's hardware. All computers have memory, also known as RAM (Random Access Memory). RAM holds the programs that your computer is currently running along with the data they are currently manipulating (their variables and data structures). Memory can be thought of simply as an array of bytes. In this array, every memory location has its own address – the address of the first byte is 0, followed by 1, 2, 3, and so on. Memory addresses act just like the indexes of a normal array. The computer can access any address in memory at any time and any order (hence the name “random access memory”). It can also group bytes together as it needs to to form larger variables, arrays, and structures. For example, a **double** variable consumes 8 contiguous bytes in memory. You might make the following global declaration in a program:

```
1 double f;
```

This statement says, "Declare a location named **f** that can hold one floating point value." When the program runs, the computer reserves space for the variable **f** somewhere in memory. That location has a fixed address in the memory space. The variable **f** consumes eight bytes of RAM in memory.

While you think of the variable **f**, the computer thinks of a specific address in memory (for example, 248,440). Therefore, when you create a statement like this:

```
1 f = 3.14;
```

The compiler might translate that into, “Load the value 3.14 into memory location 248,440.” The computer is always thinking of memory in terms of addresses and values at those addresses.

There are, by the way, several interesting side effects to the way your computer treats memory. For example, say that you include the following code in one of your programs:

```
1 int i, s[4], t[4];
2
3 for (i = 0; i <= 4; i++)
4 {
5     s[i] = i;
6     t[i] = i;
7 }
8
9 printf ("s:t\n");
10
11 for (i = 0; i <= 4; i++)
```

```
12 {  
13     printf ("%d:%d\n", s[i], t[i]);  
14 }
```

The output that you see from the program¹ will look like this:

```
1 s:t  
2 4:0  
3 1:1  
4 2:2  
5 3:3  
6 4:4
```

Why are `s[0]` incorrect? If you look carefully at the code, you can see that the for loops are writing one element past the end of each array. In memory, the arrays are placed adjacent to one another. Therefore, when you try to write to `t[4]`, which does not exist, the system writes into `s[0]` instead because `s[0]` is where `t[4]` ought to be. As far as the computer is concerned, `t[4]` is simply an address, and it can write into it. As you can see however, even though the computer executes the program, it is not correct or valid. The program corrupts the array `s` in the process of running. If you execute the following statement, more severe consequences may result:

```
1 s[100000000] = 5;
```

The location `s[100000000]` is likely outside of your program's memory space. In other words, you are writing into memory that your program does not own. On a system with protected memory spaces, this sort of statement will cause the system to terminate execution of the program. On other systems, however, the system is not aware of what you are doing. You end up damaging the code or variables in another application. The effect of the violation can range from nothing at all to a complete system crash. In memory, `i`, `s`, `t` are all placed next to one another at specific addresses. Therefore, if you write past the boundaries of a variable, the computer will do what you say but it will end up corrupting another memory location.

Because C does not perform any sort of range checking when you access an element of an array, it is essential that you, as a programmer, pay careful attention to **array ranges** yourself and keep within the array's appropriate boundaries. Unintentionally reading or writing outside of array boundaries always leads to faulty program behavior.

¹Compiled as following: `clang -O0 example-overwrite.c -o example-overwrite`. This particular example is sensitive to the compilation options and the compiler used.

Pointing to the Same Address

Here is an important aspect of C: **Any number of pointers can point to the same address.** For example, you could declare `p`, `q`, and `r` as integer pointers and set all of them to point to `i`, as shown here:

```
1 int i;  
2 int *p, *q, *r;  
3  
4 p = &i;  
5 q = &i;  
6 r = p;
```

Note that in this code, `r` points to the same thing that `p` points to, which is `i`. You can assign pointers to one another, and the address is copied from the right-hand side to the left-hand side during the assignment. The variable `i` now has four names: `i`, `*p`, `*q` and `*r`. There is no limit on the number of pointers that can hold (and therefore point to) the same address.

Using Pointers for Function Parameters

Variable parameters in functions are one of the most common uses of pointers in C. To understand how variable parameters work, let's see how we might go about implementing a *swap* function in C. To implement a swap function, what you would like to do is pass in two variables and have the function swap their values. Here's one attempt at an implementation – enter and execute the following code and see what happens:

Listing 7.2: Using pointers for function parameters, I

```
1 #include <stdio.h>  
2  
3 void swap(int, int);  
4  
5 void swap(int i, int j)  
6 {  
7     int t;  
8  
9     t = i;  
10    i = j;  
11    j = t;
```

```
12 }
13
14 int main (void)
15 {
16     int a, b;
17
18     a = 5;
19     b = 10;
20     printf ("%d %d\n", a, b);
21
22     swap(a, b);
23
24     printf ("%d %d\n", a, b);
25
26     return 0;
27 }
```

When you execute this program, you will find that no swapping takes place:

```
1 5 10
2 5 10
```

The values of **a** and **b** are passed to **swap**, and the **swap** function does swap them, but when the function returns nothing happens.

To make this function work correctly you can use pointers, as shown below:

Listing 7.3: Using pointers for function parameters, II

```
1 #include <stdio.h>
2
3 void swap(int *, int *);
4
5 void swap(int *i, int *j)
6 {
7     int t;
8
9     t = *i;
10    *i = *j;
11    *j = t;
12 }
```

```
13
14 int main (void)
15 {
16     int a, b;
17
18     a = 5;
19     b = 10;
20     printf ("%d %d\n", a, b);
21
22     swap(&a, &b);
23
24     printf ("%d %d\n", a, b);
25
26     return 0;
27 }
```

Now swapping works:

```
1 5 10
2 10 5
```

Pointers and Arrays

Arrays and pointers are intimately linked in C. To use arrays effectively, you have to know how to use pointers with them.

Let's start with a simple example of copying arrays:

Listing 7.4: Copy the content of an array

```
1 #include <stdio.h>
2
3 #define MAX 10
4
5 int main (void)
6 {
7     int a[MAX];
8     int b[MAX];
9     int i;
```



```
10
11     for (i = 0; i < MAX; i++)
12     {
13         a[i] = i;
14     }
15
16     /*
17      * copy arrays, a -> b
18      */
19     for (i = 0; i < MAX; b[i] = a[i], i++);
20
21     printf ("a:b\n");
22     for (i = 0; i < MAX; i++)
23     {
24         printf ("%d:%d\n", a[i], b[i]);
25     }
26
27     return 0;
28 }
```

To do the copy we need to use something like this:

```
1 for (i = 0; i < MAX; b[i] = a[i], i++);
```

(Simply writing **a=b**; would produce a compilation error.)

Variables **a** and **b** are not arrays themselves. Instead they are permanent pointers to arrays. **a** and **b** permanently point to the first elements of their respective arrays – they hold the addresses of **a[0]** and **b[0]** respectively. Since they are **permanent** pointers you cannot change their addresses. The statement **a=b**; therefore does not work.

Because **a** and **b** are pointers, you can do several interesting things with pointers and arrays. For example, the following code works:

Listing 7.5: Mixing arrays and pointers

```
1 #include <stdio.h>
2
3 #define MAX 10
4
5 int main (void)
6 {
```

```
7   int a[MAX];
8   int i;
9   int *p;
10
11  p = a;
12  for (i = 0; i < MAX; i++)
13  {
14      a[i] = MAX - i;
15  }
16
17  printf ("%d\n", *p);
18
19  return 0;
20 }
```

The statement **p = a;** works because **a** is a pointer. **a** points to the address of the 0th element of the actual array. This element is an integer, so **a** is a pointer to a single integer. Therefore, declaring **p** as a pointer to an integer and setting it equal to **a** works. Another way to say exactly the same thing would be to replace **p=a;** with **p=&a[0];**. Since **a** contains the address of **a[0]**, **a** and **&a[0]** mean the same thing.

Now that **p** is pointing at the 0th element of **a**, you can do some rather strange things with it. The **a** variable is a permanent pointer and can not be changed, but **p** is not subject to such restrictions. C actually encourages you to move it around using *pointer arithmetic*. For example, if you say **p++;**, the compiler knows that **p** points to an integer, so this statement increments **p** the appropriate number of bytes to move it to the next element of the array. If **p** were pointing to an array of 100-byte-long structures, **p++;** would move **p** over by 100 bytes. C takes care of the details of element size.

You can copy the array **a** into **b** using pointers as well. The following code can replace (**for i=0; i<MAX; a[i]=b[i], i++;**) :

```
1  p = a;
2  q = b;
3  for (i = 0; i < MAX; i++)
4  {
5      *q = *p;
6      q++;
7      p++;
8  }
```

You can abbreviate this code as follows:

```
1 p = a;
2 q = b;
3 for (i = 0; i < MAX; i++)
4 {
5     *q++ = *p++;
6 }
```

And you can further abbreviate it to:

```
1 for (p = a, q = b, i = 0; i < MAX; *q++ = *p++, i++);
```

What if you go beyond the end of the array **a** or **b** with the pointers **p** or **q**? C does not care – it blithely goes along incrementing **p** and **q**, copying away over other variables with abandon. You need to be careful when indexing into arrays in C, because C assumes that you know what you are doing.

You can pass an array such as **a** or **b** to a function in two different ways. Imagine a function **dump** that accepts an array of integers as a parameter and prints the contents of the array to stdout. There are two ways to code **dump**:

```
1 void dump (int a[], int nia)
2 {
3     for (int i = 0; i < nia; i++)
4     {
5         printf ("%d\n", a[i]);
6     }
7 }
```

or:

```
1 void dump (int *p, int nia)
2 {
3     for (int i = 0; i < nia; i++)
4     {
5         printf ("%d\n", *p++);
6     }
7 }
```

The **nia** (number_in_array) variable is required so that the size of the array is known. Note that only a pointer to the array, rather than the contents of the array, is passed to the function. Also note that C functions can accept variable-size arrays as parameters.

Pointers to Structures

It is possible to create a pointer to almost any type in C, including user-defined types. It is extremely common to create pointers to structures.

```
1 struct point
2 {
3     double x;
4     double y;
5     char color;
6 } pt;
7
8 pt.x = 1;
9 pt.y = 1;
10 pt.color = 'b';
11
12 struct point *p;
13 p = &pt;
14
15 printf ("%f  %f  %c\n", (*p).x, (*p).y, (*p).color);
```

The pointer **p** in the code above is a pointer to a structure. You deal with ***p** just like a normal structure variable, but you have to be careful with the precedence of operators in C. If you were to leave off the parenthesis around ***p** the code would not compile because the “.” operator has a higher precedence than the “*” operator. Because it gets tedious to type so many parentheses when working with pointers to structures, C includes a shorthand notation that does exactly the same thing:

```
1 p->x = 5;
```

The **p->** notation is exactly equivalent to **(*p).**, but takes two fewer characters.

Part II

Programming Tools

Chapter 8

Makefiles

It can be cumbersome to type all of the **gcc/clang** lines over and over again, especially if you are making a lot of changes to the code and it has several libraries. The **make** facility solves this problem. You can use the following makefile to replace the compilation sequence above:

```
1 main: main.o util.o
2     clang -o main main.o util.o
3
4 main.o: main.c util.h
5     clang -c -g main.c
6
7 util.o: util.c util.h
8     clang -c -g util.c
```

Enter this into a file named **makefile**, and type **make** to build the executable. Note that you *must* precede all **gcc** lines with a tab. (Eight spaces will not suffice – it must be a tab. All other lines must be flush left.)

This makefile contains two types of lines. The lines appearing flush left are **dependency lines**. The lines preceded by a tab are executable lines, which can contain any valid UNIX command. A dependency line says that some file is dependent on some other set of files. For example, **main.o: main.c util.h** says that the file **main.o** is dependent on the files **main.c** and **util.h**. If either of these two files changes, the following executable line(s) should be executed to recreate **main.o**.

Note that the final executable produced by the whole makefile is **main**, on line 1 in the makefile. The final result of the makefile should always go on line 1, which in this makefile says that the file **main** is dependent on **main.o** and **util.o**. If either of these changes, execute the line **gcc -o main main.o util.o** to recreate **main**.

It is possible to put multiple lines to be executed below a dependency line – they must all start with a tab. A large program may have several libraries and a main program. The makefile automatically recompiles everything that needs to be recompiled because of a change.

Chapter 9

Command Line Arguments

C provides a fairly simple mechanism for retrieving command line parameters entered by the user. It passes an **argv** parameter to the main function in the program. **argv** structures appear in a fair number of the more advanced library calls, so understanding them is useful to any C programmer.

Enter the following code and compile it:

```
1 #include <stdio.h>
2
3 int main (int argc, char *argv[])
4 {
5     int x;
6
7     printf ("%d\n",argc);
8     for (x = 0; x < argc; x++)
9         printf ("%s\n", argv[x]);
10    return 0;
11 }
```

In this code, the main program accepts two parameters, **argv** and **argc**. The **argv** parameter is an array of pointers to string that contains the parameters entered when the program was invoked at the UNIX command line. The **argc** integer contains a count of the number of parameters. This particular piece of code types out the command line parameters. To try this, compile the code to an executable file named **aaa** and type **aaa xxx yyy zzz**. The code will print the command line parameters **xxx**, **yyy** and **zzz**, one per line.

The **char *argv[]** line is an array of pointers to string. In other words, each element of the array is a pointer, and each pointer points to a string (technically, to the first character of the string). Thus,

argv[0] points to a string that contains the first parameter on the command line (the program's name), **argv[1]** points to the next parameter, and so on. The **argc** variable tells you how many of the pointers in the array are valid. You will find that the preceding code does nothing more than print each of the valid strings pointed to by **argv**.

Because **argv** exists, you can let your program react to command line parameters entered by the user fairly easily. For example, you might have your program detect the word **help** as the first parameter following the program name, and dump a help file to **stdout**. File names can also be passed in and used in your **fopen** statements.

Chapter 10

Version control

Chapter 11

Debugging

Chapter 12

Detecting memory problems with valgrind

Part III

More C Programming

Chapter 13

Dynamic Data Structures

The operating system and several applications, along with their global variables and stack spaces, all consume portions of memory. When a program completes execution, it releases its memory for reuse by other programs.

Memory holds the executable code for the different applications currently running on the machine, along with the executable code for the operating system itself. Each application has certain global variables associated with it. These variables also consume memory. Finally, each application uses an area of memory called *the stack*, which holds all local variables and parameters used by any function. The stack also remembers the order in which functions are called so that function returns occur correctly. Each time a function is called, its local variables and parameters are “pushed onto” the stack. When the function returns, these locals and parameters are “popped.” Because of this, the size of a program’s stack fluctuates constantly as the program is running, but it has some maximum size.

As a program finishes execution, the operating system unloads it, its globals and its stack space from memory. A new program can make use of that space at a later time. In this way, the memory in a computer system is constantly “recycled” and reused by programs as they execute and complete.

Dynamic data structures are data structures that grow and shrink as you need them to by allocating and deallocating memory from a place called *the heap*. They are extremely important in C because they allow the programmer to exactly control memory consumption.

Dynamic data structures allocate blocks of memory from the heap as required, and link those blocks together into some kind of data structure using pointers. When the data structure no longer needs a block of memory, it will return the block to the heap for reuse. This recycling makes very efficient use of memory.

To understand dynamic data structures completely, we need to start with the stack and the heap.

Stack vs Heap

So far we have seen how to declare basic type variables such as `int`, `double`, etc, and complex types such as arrays and structs. The way we have been declaring them so far puts these variables on the stack in C.

The Stack

What is the stack? It's a special region of your computer's memory that stores temporary variables created by each function (including the `main()` function). The stack is a "FILO" (first in, last out) data structure, that is managed and optimized by the CPU quite closely. Every time a function declares a new variable, it is "pushed" onto the stack. Then every time a function exits, all of the variables pushed onto the stack by that function, are released (that is to say, they are deleted). Once a stack variable is released, that region of memory becomes available for other stack variables.

The advantage of using the stack to store variables, is that memory is managed for you. You don't have to allocate memory by hand, or free it once you don't need it any more. What's more, because the CPU organizes stack memory so efficiently, reading from and writing to stack variables is very fast.

A key to understanding the stack is the notion that when a function exits, all of its variables are popped off of the stack (and hence lost forever). Thus stack variables are local in nature. This is related to a concept we saw earlier known as variable scope, or local vs global variables. A common bug in C programming is attempting to access a variable that was created on the stack inside some function, from a place in your program outside of that function (i.e. after that function has exited).

Another feature of the stack to keep in mind, is that there is a limit (varies with OS) on the size of variables that can be store on the stack. If the program tries to put too much information on the stack, stack overflow will result. Stack overflow happens when all the memory in the stack has been allocated – in that case, further allocations begin overflowing into other sections of memory.

Here is an example program that causes a stack overflow:

Listing 13.1: A program that causes a stack overflow

```
1 #define HUGE 1000000000L
2
3 int main(void)
```

```
4 {  
5     int stack[HUGE];  
6     stack[0] = 1;  
7  
8     return 0;  
9 }
```

This program tries to allocate a huge array on the stack. Because the stack is not large enough to handle this array, the array allocation overflows into portions of memory the program is not allowed to use. Consequently, the program crashes.

To summarize the stack:

- the stack grows and shrinks as functions push and pop local variables
- there is no need to manage the memory yourself, variables are allocated and freed automatically
- the stack has size limits
- stack variables only exist while the function that created them, is running

The Heap

The heap is a region of your computer's memory that is not managed automatically for you, and is not as tightly managed by the CPU. It is a more free-floating region of memory (and is larger). To allocate memory on the heap, you must use built-in C functions. Once you have allocated memory on the heap, you are responsible to deallocate that memory once you don't need it any more. If you fail to do this, your program will have what is known as a memory leak. That is, memory on the heap will still be set aside (and won't be available to other processes). As we will see in the debugging section, there is a tool called `valgrind` that can help you detect memory leaks.

Unlike the stack, the heap does not have size restrictions on variable size (apart from the obvious physical limitations of your computer). Heap memory is slightly slower to be read from and written to, because one has to use pointers to access memory on the heap.

Unlike the stack, variables created on the heap are accessible by any function, anywhere in your program. Heap variables are essentially global in scope.

Summary: Stack vs Heap

- Stack
 - very fast access
 - don't have to explicitly de-allocate variables
 - space is managed efficiently by CPU, memory will not become fragmented
 - local variables only
 - limit on stack size (OS-dependent)
 - variables cannot be resized
- Heap
 - variables can be accessed globally
 - no limit on memory size
 - (relatively) slower access
 - no guaranteed efficient use of space, memory may become fragmented over time as blocks of memory are allocated, then freed
 - you must manage memory (you're in charge of allocating and freeing variables)
 - variables can be resized

Dynamic Data Structures. Malloc and Free

Let's say that you would like to allocate a certain amount of memory during the execution of your application. You can call the malloc function at any time, and it will request a block of memory from the heap. The operating system will reserve a block of memory for your program, and you can use it in any way you like. When you are done with the block, you return it to the operating system for recycling by calling the free function. Then other applications can reserve it later for their own use.

For example, the following code demonstrates the simplest possible use of the heap:

Listing 13.2: Simple example of using malloc and free

```
1 #include <stdio.h>
2 #include <stdlib.h>
```

```
3
4 int main (void)
5 {
6     int *p;
7
8     p = (int *) malloc(sizeof(int));
9     if (p == NULL)
10    {
11        printf ("ERROR: Out of memory\n");
12        return 1;
13    }
14
15    *p = 5;
16    printf ("%d\n", *p);
17
18    free(p);
19
20    return 0;
21 }
```

Line 8 in this program calls the malloc function. This function does three things:

1. The malloc statement first looks at the amount of memory available on the heap and asks, “Is there enough memory available to allocate a block of memory of the size requested?” The amount of memory needed for the block is known from the parameter passed into malloc – in this case, **sizeof(int)** is 4 bytes. If there is not enough memory available, the malloc function returns the address zero to indicate the error (another name for zero is **NULL** and you will see it used throughout C code). Otherwise malloc proceeds.
2. If memory is available on the heap, the system “allocates” or “reserves” a block from the heap of the size specified. The system reserves the block of memory so that it isn’t accidentally used by more than one malloc statement.
3. The system then places into the pointer variable (p, in this case) the address of the reserved block. The pointer variable itself contains an address. The allocated block is able to hold a value of the type specified, and the pointer points to it.

The program next checks the pointer p to make sure that the allocation request succeeded with the line **if (p == NULL)** (which could have also been written as **if (!p)**). If the allocation fails (if p is zero), the program terminates. If the allocation is successful, the program then initializes the block to the value 5, prints out the value, and calls the free function to return the memory to the heap

before the program terminates.

There is really no difference between this code and previous code that sets `p` equal to the address of an existing integer `i`. The only distinction is that, in the case of the variable `i`, the memory existed as part of the program's pre-allocated memory space and had the two names: `i` and `*p`. In the case of memory allocated from the heap, the block has the single name `*p` and is allocated during the program's execution. Two common questions:

It is really important to check that the pointer is non-NULL after each allocation. Since the heap varies in size constantly depending on which programs are running, how much memory they have allocated, etc., there is never any guarantee that a call to `malloc` will succeed. You should check the pointer after any call to `malloc` to make sure the pointer is valid.

When a program terminates, the operating system “cleans up after it,” releasing its executable code space, stack, global memory space and any heap allocations for recycling. Therefore, there are no long-term consequences to leaving few allocations pending at program termination. However, “memory leaks” during the execution of a program are harmful, as discussed below.

The following two programs show two different valid uses of pointers, and try to distinguish between the use of a pointer and of the pointer's value:

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 int main (void)
5 {
6     int *p, *q;
7
8     p = (int *) malloc(sizeof(int));
9     q = p;
10    *p = 10;
11    printf ("%d\n", *q);
12    *q = 20;
13    printf ("%d\n", *q);
14
15    free (p);
16
17    return 0;
18 }
```

The final output of this code would be 10 from the first print statement and 20 from second one.

```
1 10
2 20
```

The following code is slightly different:

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 int main (void)
5 {
6     int *p, *q;
7
8     p = (int *) malloc(sizeof(int));
9     q = (int *) malloc(sizeof(int));
10    *p = 10;
11    *q = 20;
12    *p = *q;
13    printf ("%d\n", *p);
14
15    free (p);
16    free (q);
17
18    return 0;
19 }
```

The final output from this code would be 20.

```
1 20
```

Notice that the compiler will allow ***p = *q**, because ***p** and ***q** are both integers. This statement says, “Move the integer value pointed to by **q** into the integer value pointed to by **p**.” The statement moves the values. The compiler will also allow **p = q**, because **p** and **q** are both pointers, and both point to the same type (if **s** is a pointer to a character, **p = s** is not allowed because they point to different types). The statement **p = q** says, “Point **p** to the same block **q** points to.” In other words, the address pointed to by **q** is moved into **p**, so they both point to the same block. This statement moves the addresses.

From all of these examples, you can see that there are four different ways to initialize a pointer. When a pointer is declared, as in **int *p**, it starts out in the program in an uninitialized state. It may point anywhere, and therefore to dereference it is an error. Initialization of a pointer variable involves pointing it to a known location in memory.

1. One way, as seen already, is to use the malloc statement. This statement allocates a block of memory from the heap and then points the pointer at the block. This initializes the pointer, because it now points to a known location. The pointer is initialized because it has been filled with a valid address – the address of the new block.
2. The second way, as seen just a moment ago, is to use a statement such as **p = q** so that p points to the same place as q. If q is pointing at a valid block, then p is initialized. The pointer p is loaded with the valid address that q contains. However, if q is uninitialized or invalid, p will pick up the same useless address.
3. The third way is to point the pointer to a known address, such as a global variable's address. For example, if i is an integer and p is a pointer to an integer, then the statement **p=&i** initializes p by pointing it to i.
4. The fourth way to initialize the pointer is to use the value NULL. NULL is a special values used with pointers:

```
1  p = NULL;
```

What this does physically is to place a zero into p. The pointer p's address is zero.

Any pointer can be set to point to zero. When p points to zero, however, it does not point to a block. The pointer simply contains the address zero, and this value is useful as a tag. You can use it in statements such as:

```
1  if (p == NULL)
2  {
3      ...
4  }
```

or:

```
1  while (p != NULL)
2  {
3      ...
4  }
```

The system also recognizes the NULL value, and will generate runtime error message, e.g. Segmentation fault, if you happen to dereference a NULL pointer, e.g. in the following code:

Listing 13.3: Dereferencing of NULL pointer

```
1  #include <stdlib.h>
```

```
2
3 int main(void)
4 {
5     int *p;
6
7     p = NULL;
8     *p = 5;
9
10    return 0;
11 }
```

The pointer `p` does not point to a block, it points to zero, so a value cannot be assigned to `*p`.

The `malloc` command is used to allocate a block of memory. It is also possible to deallocate a block of memory when it is no longer needed. When a block is deallocated, it can be reused by a subsequent `malloc` command, which allows the system to recycle memory. The command used to deallocate memory is called **free**, and it accepts a pointer as its parameter. The `free` command does two things:

1. The block of memory pointed to by the pointer is unreserved and given back to the free memory on the heap. It can then be reused by later new statements.
2. The pointer is left in an uninitialized state, and must be reinitialized before it can be used again.

The `free` statement simply returns a pointer to its original uninitialized state and makes the block available again on the heap.

The following example shows how to use the heap. It allocates an integer block, fills it, writes it, and disposes of it:

Listing 13.4: Using `malloc` and `free`

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 int main (void)
5 {
6     int *p;
7
8     p = (int *)malloc (sizeof(int));
9     *p=10;
10    printf ("%d\n", *p);
```

```
11
12     free(p);
13
14     return 0;
15 }
```

This code is really useful only for demonstrating the process of allocating, deallocating, and using a block in C. The **malloc** line allocates a block of memory of the size specified – in this case, **sizeof(int)** bytes. The **sizeof** command in C returns the size, in bytes, of any type. Using **sizeof**, makes the code portable and more readable.

The **malloc** function returns a pointer to the allocated block. This pointer is generic. Using the pointer without typecasting generally produces a type warning from the compiler. The **(int *)** typecast converts the generic pointer returned by malloc into a “pointer to an integer,” which is what **p** expects.

The second example illustrates the same functions as the previous example, but it uses a structure instead of an integer.

Listing 13.5: Using malloc and free with structures

```
1  #include <stdio.h>
2  #include <stdlib.h>
3
4  struct rec
5  {
6      int i;
7      double f;
8      char c;
9  };
10
11 int main (void)
12 {
13     struct rec *p;
14     p = (struct rec *) malloc (sizeof(struct rec));
15     (*p).i = 10;
16     (*p).f = 3.14;
17     (*p).c = 'a';
18     printf ("%d %f %c\n", (*p).i, (*p).f, (*p).c);
19     free(p);
20 }
```

```
21     return 0;  
22 }
```

Note the following line:

```
1     (*p).i = 10;
```

Many wonder why the following doesn't work:

```
1     *p.i = 10;
```

The answer has to do with the precedence of operators in C. The result of the calculation $5+3*4$ is 17, not 32, because the $*$ operator has higher precedence than $+$ in most computer languages. In C, the $.$ operator has higher precedence than $*$, so parentheses force the proper precedence.

Most people tire of typing $(*p).i$ all the time, so C provides a shorthand notation. The following two statements are exactly equivalent, but the second is easier to type:

```
1     (*p).i = 10;  
2     p->i = 10;
```

You will see the second more often than the first when reading other people's code.

Chapter 14

Strings

Strings in C are intertwined with pointers to a large extent. You must become familiar with the pointer concepts covered in the previous articles to use C strings effectively. Once you get used to them, however, you can often perform string manipulations very efficiently.

A string in C is simply an array of characters. The following line declares an array that can hold a string of up to 99 characters.

```
1 char str[100];
```

It holds characters as you would expect: **str[0]** is the first character of the string, **str[1]** is the second character, and so on. But why is a 100-element array unable to hold up to 100 characters? Because C uses *null-terminated strings*, which means that the end of any string is marked by the ASCII value 0 (the null character), which is also represented in C as **'\0'**.

Because C provides no explicit support for strings in the language itself, all of the string-handling functions are implemented in libraries. The string I/O operations (gets, puts, and so on) are implemented in `<stdio.h>`, and a set of fairly simple string manipulation functions are implemented in `<string.h>`.

Suppose you want to copy the contents of one string to another. In C you cannot simply assign one array to another. You have to copy it element by element. The string library (`<string.h>`) contains a function called **strncpy** for this task. Here is a common piece of code to find in a normal C program:

```
1 char s[100];  
2 strncpy(s, "hello", 100);
```

The following code shows how to use **strncpy** in C:

Listing 14.1: Using strncpy

```
1 #include <string.h>
2
3 int main (void)
4 {
5     char s1[100], s2[100];
6     strncpy(s1, "hello", 100); /* copy "hello" into s1 */
7     strncpy(s2, s1, 100);      /* copy s1 into s2 */
8
9     return 0;
10 }
```

strncpy is used whenever a string is initialized in C. You use the **strncmp** function in the string library to compare two strings. It returns an integer that indicates the result of the comparison. Zero means the two strings are equal, a negative value means that **s1** is less than **s2**, and a positive value means **s1** is greater than **s2**.

Listing 14.2: Using strncmp

```
1 #include <stdio.h>
2 #include <string.h>
3
4 int main (void)
5 {
6     char s1[] = "abcde", s2[] = "pqrst";
7
8     if (strncmp(s1, s2, strlen(s1)) == 0) {
9         printf ("equal\n");
10    } else if (strncmp(s1, s2, strlen(s1)) < 0) {
11        printf ("s1 is less than s2\n");
12    } else {
13        printf ("s1 is greater than s2\n");
14    }
15    return 0;
16 }
```

The output of the program:

```
1 s1 is less than s2
```

Other common functions in the string library include **strlen**, which returns the length of a string, and **strncat** which concatenates two strings.

Most of the string functions return a string pointer as a result, and **strncpy** returns the value of **s1** as its result.

Using pointers with strings can sometimes result in definite improvements in speed. For example, suppose you want to remove the leading blanks from a string. You might be inclined to shift characters over on top of the blanks to remove them. In C, you can avoid the movement altogether:

Listing 14.3: Using pointers with string

```
1 #include <stdio.h>
2 #include <string.h>
3
4 int main (void)
5 {
6     char s[] = "    asdf", *p;
7
8     p = s;
9     while (*p == ' ')
10     {
11         p++;
12     }
13     printf ("%s\n", p);
14
15     return 0;
16 }
```

This is much faster than the movement technique, especially for long strings.

Programming exercise

- Create a program that reads in a string containing a first name followed by a blank followed by a last name. Write functions to remove any leading or trailing blanks. Write another function that returns the last name.
- Write a function that converts a string to uppercase.
- Write a function that gets the first word from a string and returns the remainder of the string.

Chapter 15

Operator Precedence, II

C contains many operators, and because of the way in which operator precedence works, the interactions between multiple operators can become confusing.

```
1 x = 5 + 3 * 6;
```

x receives the value 23, not 48, because in C multiplication and division have higher precedence than addition and subtraction.

```
1 char *a[[]10{[]}];
```

Is **a** a single pointer to an array of 10 characters, or is it an array of 10 pointers to character? Unless you know the precedence conventions in C, there is no way to find out. Similarly, because of precedence statements such as ***p.i = 10;** do not work. Instead, the form **(*p).i = 10;** must be used to force correct precedence.

The following table from *C Programming Language* by Kernighan and Ritchie, shows the precedence hierarchy in C. The top line has the highest precedence.

Operators	Associativity
([- .	Left to right
! - ++ -{- + * & (type-cast) sizeof	Right to left
(in the above line, +, - and * are the unary forms)	
* / %	Left to right
+ -	Left to right
<< >>	Left to right
< <= > >=	Left to right
== !=	Left to right

10	&	Left to right
11	^	Left to right
12		Left to right
13	&&	Left to right
14		Left to right
15	? :	Left to right
16	= += -= *= /= %= &= ^= = <<= >>=	Right to left
17	,	Left to right

Using this table, you can see that **char *a[10];** is an array of 10 pointers to character. You can also see why the parentheses are required if **(*p).i** is to be handled correctly. After some practice, you will memorize most of this table, but every now and again something will not work because you have been caught by a subtle precedence problem.

Chapter 16

Advanced Pointers

You will normally use pointers in somewhat more complicated ways than those shown in some of the previous examples. For example, it is much easier to create a normal integer and work with it than it is to create and use a pointer to an integer. In this section, some of the more common and advanced ways of working with pointers will be explored.

Pointer Types

It is possible, legal, and beneficial to create pointer types in C, as shown below:

```
1 typedef int *IntPtrter;  
2 ...  
3 IntPtrter p;
```

This is the same as saying:

```
1 int *p;
```

This technique will be used in many of the examples on the following pages. The technique often makes a data declaration easier to read and understand, and also makes it easier to include pointers inside of structures or pass pointer parameters in functions.

Pointers to Arrays

It is also possible to create pointers to arrays, as shown below:

```
1  int *p;  
2      int i;  
3  
4      p = (int *)malloc(sizeof(int[10]));  
5      for (i=0; i<10; i++)  
6          p[i] = 0;  
7      free(p);
```

or:

```
1  int *p;  
2      int i;  
3  
4      p = (int *)malloc(sizeof(int[10]));  
5      for (i=0; i<10; i++)  
6          *(p+i) = 0;  
7      free(p);
```

Note that when you create a pointer to an integer array, you simply create a normal pointer to `int`. The call to `malloc` allocates an array of whatever size you desire, and the pointer points to that array's first element. You can either index through the array pointed to by `p` using normal array indexing, or you can do it using pointer arithmetic. C sees both forms as equivalent.

This particular technique is extremely useful when working with strings. It lets you allocate enough storage to exactly hold a string of a particular size.

Arrays of Pointers

Sometimes a great deal of space can be saved, or certain memory-intensive problems can be solved, by declaring an array of pointers. In the example code below, an array of 10 pointers to structures is declared, instead of declaring an array of structures. If an array of the structures had been created instead, $243 * 10 = 2,430$ bytes would have been required for the array. Using the array of pointers allows the array to take up minimal space until the actual records are allocated with `malloc`

statements. The code below simply allocates one record, places a value in it, and disposes of the record to demonstrate the process:

```
1  typedef struct
2  {
3      char s1[81];
4      char s2[81];
5      char s3[81];
6  } Rec;
7  Rec *a[10];
8
9  a[0] = (Rec *)malloc(sizeof(Rec));
10 strcpy(a[0]->s1, "hello");
11 free(a[0]);
```

Structures Containing Pointers

Structures can contain pointers, as shown below:

```
1  typedef struct
2  {
3      char name[21];
4      char city[21];
5      char phone[21];
6      char *comment;
7  } Addr;
8  Addr s;
9  char comm[100];
10
11 gets(s.name, 20);
12 gets(s.city, 20);
13 gets(s.phone, 20);
14 gets(comm, 100);
15 s.comment =
16     (char *)malloc(sizeof(char[strlen(comm)+1]));
17 strcpy(s.comment, comm);
```


This technique is useful when only some records actually contained a comment in the comment field. If there is no comment for the record, then the comment field would consist only of a pointer (4 bytes). Those records having a comment then allocate exactly enough space to hold the comment string, based on the length of the string typed by the user.

Pointers to Pointers

It is possible and often useful to create pointers to pointers. This technique is sometimes called a **handle**, and is useful in certain situations where the operating system wants to be able to move blocks of memory on the heap around at its discretion.

```
1  int **p;  
2      int *q;  
3  
4      p = (int **)malloc(sizeof(int *));  
5      *p = (int *)malloc(sizeof(int));  
6      **p = 12;  
7      q = *p;  
8      printf ("%d\n", *q);  
9      free(q);  
10     free(p);
```

Windows and the Mac OS use this structure to allow memory compaction on the heap. The program manages the pointer `p`, while the operating system manages the pointer `*p`. Because the OS manages `*p`, the block pointed to by `*p` (****p**) can be moved, and `*p` can be changed to reflect the move without affecting the program using `p`. Pointers to pointers are also frequently used in C to handle pointer parameters in functions.

Pointers to Structures Containing Pointers

It is also possible to create pointers to structures that contain pointers. The following example uses the **Addr** record from the previous section:

```
1  typedef struct  
2      {  
3          char name[21];  
4          char city[21];
```

```
5      char phone[21];
6      char *comment;
7  } Addr;
8  Addr *s;
9  char comm[100];
10
11  s = (Addr *)malloc(sizeof(Addr));
12  gets(s->name, 20);
13  gets(s->city, 20);
14  gets(s->phone, 20);
15  gets(comm, 100);
16  s->comment =
17      (char *)malloc(sizeof(char[strlen(comm)+1]));
18  strcpy(s->comment, comm);
```

The pointer `s` points to a structure that contains a pointer that points to a string.

In this example, it is very easy to create lost blocks if you aren't careful. For example, here is a different version of the AP example.

```
1  s = (Addr *)malloc(sizeof(Addr));
2  gets(comm, 100);
3  s->comment =
4      (char *)malloc(sizeof(char[strlen(comm)+1]));
5  strcpy(s->comment, comm);
6  free(s);
```

This code creates a lost block because the structure containing the pointer pointing to the string is disposed of before the string block is disposed of, as shown to the right.

Linking

Finally, it is possible to create structures that are able to point to identical structures, and this capability can be used to link together a whole string of identical records in a structure called a linked list.

```
1  typedef struct
2  {
3      char name[21];
```

```
4      char city[21];
5      char state[21];
6      Addr *next;
7  } Addr;
8  Addr *first;
```

The compiler will let you do this, and it can be used with a little experience to create structures like the one shown to the right.

A Linked Stack Example

A good example of dynamic data structures is a simple stack library, one that uses a dynamic list and includes functions to init, clear, push, and pop. The library's header file looks like this:

```
1  /* Stack Library - This library offers the
2     minimal stack operations for a
3     stack of integers (easily changeable) */
4
5  typedef int stack_data;
6
7  extern void stack_init();
8  /* Initializes this library.
9     Call first before calling anything. */
10
11 extern void stack_clear();
12 /* Clears the stack of all entries. */
13
14 extern int stack_empty();
15 /* Returns 1 if the stack is empty, 0 otherwise. */
16
17 extern void stack_push(stack_data d);
18 /* Pushes the value d onto the stack. */
19
20 extern stack_data stack_pop();
21 /* Returns the top element of the stack,
22    and removes that element.
23    Returns garbage if the stack is empty. */
```

The library's code file follows:

```
1 #include "stack.h"
2 #include <stdio.h>
3
4 /* Stack Library - This library offers the
5    minimal stack operations for a stack of integers */
6
7 struct stack_rec
8 {
9     stack_data data;
10    struct stack_rec *next;
11 };
12
13 struct stack_rec *top=NULL;
14
15 void stack_init()
16 /* Initializes this library.
17    Call before calling anything else. */
18 {
19     top=NULL;
20 }
21
22 void stack_clear()
23 /* Clears the stack of all entries. */
24 {
25     stack_data x;
26
27     while (!stack_empty())
28         x=stack_pop();
29 }
30
31 int stack_empty()
32 /* Returns 1 if the stack is empty, 0 otherwise. */
33 {
34     if (top==NULL)
35         return(1);
36     else
37         return(0);
38 }
```

```
39
40 void stack_push(stack_data d)
41 /* Pushes the value d onto the stack. */
42 {
43     struct stack_rec *temp;
44     temp = (struct stack_rec *) malloc(sizeof(struct stack_rec));
45     temp->data=d;
46     temp->next=top;
47     top=temp;
48 }
49
50 stack_data stack_pop()
51 /* Returns the top element of the stack,
52    and removes that element.
53    Returns garbage if the stack is empty. */
54 {
55     struct stack_rec *temp;
56     stack_data d=0;
57     if (top!=NULL)
58     {
59         d=top->data;
60         temp=top;
61         top=top->next;
62         free(temp);
63     }
64     return(d);
65 }
```

Programming exercise

- Add a **dup**, a **count**, and an **add** function to the stack library to duplicate the top element of the stack, return a count of the number of elements in the stack, and add the top two elements in the stack.
- Build a driver program and a makefile, and compile the stack library with the driver to make sure it works.

Note how this library practices information hiding: Someone who can see only the header file cannot tell if the stack is implemented with arrays, pointers, files, or in some other way. Note also that C uses **NULL**. **NULL** is defined in **stdio.h**, so you will almost always have to include **stdio.h** when you use pointers. **NULL** is the same as zero.

Chapter 17

Libraries

Libraries are very important in C because the C language supports only the most basic features that it needs. C does not even contain I/O functions to read from the keyboard and write to the screen. Anything that extends beyond the basic language must be written by a programmer. The resulting chunks of code are often placed in *libraries* to make them easily reusable. We have seen the standard I/O, or **stdio**, library already: Standard libraries exist for standard I/O, math functions, string handling, time manipulation, and so on. You can use libraries in your own programs to split up your programs into modules. This makes them easier to understand, test, and debug, and also makes it possible to reuse code from other programs that you write.

You can create your own libraries easily. As an example, we will take some code from a previous article in this series and make a library out of two of its functions. Here's the code we will start with:

```
1 #include <stdio.h>
2
3 #define MAX 10
4
5 int a[MAX];
6 int rand_seed = 10;
7
8 int rand(void)
9 /*
10  * from K&R - produces an integer random number
11  * between 0 and 32767.
12  */
13 {
```



```
14     rand_seed = rand_seed * 1103515245 + 12345;
15     return (unsigned int)(rand_seed / 65536) % 32768;
16 }
17
18 int main (void)
19 {
20     int i, t, x, y;
21     /* fill array */
22     for (i = 0; i < MAX; i++)
23     {
24         a[i] = rand();
25         printf ("%d\n",a[i]);
26     }
27     /* bubble sort the array */
28     for (x = 0; x < MAX-1; x++)
29     {
30         for (y = 0; y < MAX-x-1; y++)
31         {
32             if (a[y] > a[y+1])
33             {
34                 t = a[y];
35                 a[y] = a[y+1];
36                 a[y+1] = t;
37             }
38         }
39     }
40     /* print sorted array */
41     printf ("-----\n");
42     for (i = 0; i < MAX; i++)
43     {
44         printf ("%d\n",a[i]);
45     }
46     return 0;
47 }
```

This code fills an array with random numbers, sorts them using a bubble sort, and then displays the sorted list.

Take the bubble sort code, and use what you learned in the previous article to make a function from

it. Since both the array `a` and the constant `MAX` are known globally, the function you create needs no parameters, nor does it need to return a result. However, you should use local variables for `x`, `y`, and `t`.

Once you have tested the function to make sure it is working, pass in the number of elements as a parameter rather than using `MAX`:

```
1 #include <stdio.h>
2
3 #define MAX 10
4
5 int a[MAX];
6 int rand_seed = 10;
7
8 /*
9  * from K&R - produces an integer random number
10  * between 0 and 32767.
11  */
12 int rand(void)
13 {
14     rand_seed = rand_seed * 1103515245 + 12345;
15     return (unsigned int)(rand_seed / 65536) % 32768;
16 }
17
18 void bubble_sort(int m)
19 {
20     int x, y, t;
21     for (x = 0; x < m-1; x++)
22     {
23         for (y = 0; y < m-x-1; y++)
24         {
25             if (a[y] > a[y+1])
26             {
27                 t = a[y];
28                 a[y] = a[y+1];
29                 a[y+1] = t;
30             }
31         }
32     }
33 }
```

```
34
35 int main (void)
36 {
37     int i, t, x, y;
38     /* fill array */
39     for (i = 0; i < MAX; i++)
40     {
41         a[i] = rand();
42         printf ("%d\n",a[i]);
43     }
44     bubble_sort(MAX);
45     /* print sorted array */
46     printf ("-----\n");
47     for (i = 0; i < MAX; i++)
48     {
49         printf ("%d\n",a[i]);
50     }
51     return 0;
52 }
```

You can also generalize the **bubble_sort** function even more by passing in **a** as a parameter:

```
1 bubble_sort(int m, int a[])
```

This line says, “Accept the integer array **a** of any size as a parameter.” Nothing in the body of the **bubble_sort** function needs to change. To call **bubble_sort**, change the call to:

```
1 bubble_sort (MAX, a);
```

Note that **&a** has not been used in the function call even though the sort will change **a**. The reason for this will become clear once you understand pointers.

Making a library

Since the **rand** and **bubble_sort** functions in the previous program are useful, you will probably want to reuse them in other programs you write. You can put them into a utility library to make

their reuse easier.

Every library consists of two parts: a header file and the actual code file. The header file, normally denoted by a **.h** suffix, contains information about the library that programs using it need to know. In general, the header file contains constants and types, along with prototypes for functions available in the library. Enter the following header file and save it to a file named **util.h**.

```
1 /* util.h */
2 extern int rand();
3 extern void bubble_sort(int, int []);
```

These two lines are function prototypes. The word “extern” in C represents functions that will be linked in later.

Enter the following code into a file named **util.c**.

```
1 /* util.c */
2 #include "util.h"
3
4 int rand_seed = 10;
5
6 /*
7  * from K&R - produces an integer random number
8  * between 0 and 32767.
9  */
10 int rand(void)
11 {
12     rand_seed = rand_seed * 1103515245 + 12345;
13     return (unsigned int)(rand_seed / 65536) % 32768;
14 }
15
16 void bubble_sort(int m, int a[])
17 {
18     int x, y, t;
19     for (x = 0; x < m-1; x++)
20     {
21         for (y = 0; y < m-x-1; y++)
22         {
23             if (a[y] > a[y+1])
24             {
25                 t = a[y];
```

```
26         a[y] = a[y+1];
27         a[y+1] = t;
28     }
29 }
30 }
31 }
```

Note that the file includes its own header file (*util.h*) and that it uses quotes instead of the symbols `<and>`, which are used only for system libraries. As you can see, this looks like normal C code. Note that the variable **rand_seed**, because it is not in the header file, cannot be seen or modified by a program using this library. This is called information hiding. Adding the word **static** in front of **int** enforces the hiding completely.

Enter the following main program in a file named **main.c**.

```
1 #include <stdio.h>
2 #include "util.h"
3
4 #define MAX 10
5
6 int a[MAX];
7
8 int main (void)
9 {
10     int i, t, x, y;
11     /* fill array */
12     for (i = 0; i < MAX; i++)
13     {
14         a[i] = rand();
15         printf ("%d\n", a[i]);
16     }
17
18     bubble_sort (MAX, a);
19
20     /* print sorted array */
21     printf ("-----\n");
22     for (i = 0; i < MAX; i++)
23     {
24         printf ("%d\n", a[i]);
```

```
25     }  
26     return 0;  
27 }
```

This code includes the utility library. The main benefit of using a library is that the code in the main program is much shorter.

Compiling and running

To compile the library, type the following at the command line (assuming you are using UNIX) (replace `gcc` with `cc` if your system uses `cc`):

```
1 gcc -c -g util.c
```

The `-c` causes the compiler to produce an object file for the library. The object file contains the library's machine code. It cannot be executed until it is linked to a program file that contains a main function. The machine code resides in a separate file named *util.o*.

To compile the main program, type the following:

```
1 gcc -c -g main.c
```

This line creates a file named *main.o* that contains the machine code for the main program. To create the final executable that contains the machine code for the entire program, link the two object files by typing the following:

```
1 gcc -o main main.o util.o
```

This links **main.o** and **util.o** to form an executable named **main**. To run it, type **main**.

Chapter 18

Text files

Text files in C are straightforward and easy to understand. All text file functions and types in C come from the **stdio** library.

When you need text I/O in a C program, and you need only one source for input information and one sink for output information, you can rely on **stdin** (standard in) and **stdout** (standard out). You can then use input and output redirection at the command line to move different information streams through the program. There are six different I/O commands in `<stdio.h>` that you can use with **stdin** and **stdout**:

- **printf** - prints formatted output to **stdout**
- **scanf** - reads formatted input from **stdin**
- **puts** - prints a string to **stdout**
- **gets** - reads a string from **stdin**
- **putc** - prints a character to **stdout**
- **getc, getchar** - reads a character from **stdin**

The advantage of **stdin** and **stdout** is that they are easy to use. Likewise, the ability to redirect I/O is very powerful. For example, maybe you want to create a program that reads from **stdin** and counts the number of characters:

```
1 #include <stdio.h>
2 #include <string.h>
3
4 int main (void)
5 {
```



```
6 char s[1000];
7 int count = 0;
8 while (gets(s))
9 {
10     count += strlen(s);
11 }
12 printf ("%d\n", count);
13 return 0;
14 }
```

Enter this code and run it. It waits for input from stdin, so type a few lines. When you are done, press CTRL-D to signal end-of-file (eof). The gets function reads a line until it detects eof, then returns a 0 so that the while loop ends. When you press CTRL-D, you see a count of the number of characters in stdout (the screen). (Use **man gets** or your compiler's documentation to learn more about the gets function.)

Now, suppose you want to count the characters in a file. If you compiled the program to an executable named **xxx**, you can type the following:

```
1 xxx < filename
```

Instead of accepting input from the keyboard, the contents of the file named **filename** will be used instead. You can achieve the same result using **pipes**:

```
1 cat < filename | xxx
```

You can also redirect the output to a file:

```
1 xxx < filename > out
```

This command places the character count produced by the program in a text file named **out**.

Sometimes, you need to use a text file directly. For example, you might need to open a specific file and read from or write to it. You might want to manage several streams of input or output or create a program like a text editor that can save and recall data or configuration files on command. In that case, use the text file functions in stdio:

- **fopen** - opens a text file
- **fclose** - closes a text file
- **feof** - detects end-of-file marker in a file
- **fprintf** - prints formatted output to a file

- **fscanf** - reads formatted input from a file
- **fputs** - prints a string to a file
- **fgets** - reads a string from a file
- **fputc** - prints a character to a file
- **fgetc** - reads a character from a file

Text files: opening

You use **fopen** to open a file. It opens a file for a specified mode (the three most common are r, w, and a, for read, write, and append). It then returns a file pointer that you use to access the file. For example, suppose you want to open a file and write the numbers 1 to 10 in it. You could use the following code:

```
1 #include <stdio.h>
2 #define MAX 10
3
4 int main (void)
5 {
6     FILE *f;
7     int x;
8     f = fopen("out", "w");
9     if (!f)
10         return 1;
11     for(x = 1; x <= MAX; x++)
12         fprintf (f, "%d\n", x);
13     fclose(f);
14     return 0;
15 }
```

The **fopen** statement here opens a file named **out** with the w mode. This is a destructive write mode, which means that if **out** does not exist it is created, but if it does exist it is destroyed and a new file is created in its place. The fopen command returns a pointer to the file, which is stored in the variable f. This variable is used to refer to the file. If the file cannot be opened for some reason, f will contain NULL.

The fprintf statement should look very familiar: It is just like printf but uses the file pointer as its first parameter. The fclose statement closes the file when you are done.

Text files: reading

To read a file, open it with `r` mode. In general, it is not a good idea to use **fscanf** for reading: Unless the file is perfectly formatted, `fscanf` will not handle it correctly. Instead, use **fgets** to read in each line and then parse out the pieces you need.

The following code demonstrates the process of reading a file and dumping its contents to the screen:

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     FILE *f;
6     char s[1000];
7
8     f = fopen("infile", "r");
9     if (!f)
10         return 1;
11     while (fgets(s, 1000, f) != NULL)
12         printf ("%s", s);
13     fclose(f);
14     return 0;
15 }
```

The `fgets` statement returns a `NULL` value at the end-of-file marker. It reads a line (up to 1,000 characters in this case) and then prints it to `stdout`. Notice that the `printf` statement does not include `\n` in the format string, because `fgets` adds `\n` to the end of each line it reads. Thus, you can tell if a line is not complete in the event that it overflows the maximum line length specified in the second parameter to `fgets`.

Main function return values

This program is the first program in this series that returns an error value from the main program. If the **fopen** command fails, `f` will contain a `NULL` value (a zero). We test for that error with the `if` statement. The `if` statement looks at the True/False value of the variable `f`. Remember that in C, 0 is False and anything else is true. So if there were an error opening the file, `f` would contain zero,

which is False. The **!** is the NOT operator. It inverts a Boolean value. So the if statement could have been written like this:

```
1 if (f == 0)
```

That is equivalent. However, **if (!f)** is more common.

If there is a file error, we return a 1 from the main function. In Linux, you can actually test for this value on the command line. See the shell documentation for details.

Chapter 19

Binary Files

Binary files are very similar to arrays of structures, except the structures are in a disk file rather than in an array in memory. Because the structures in a binary file are on disk, you can create very large collections of them (limited only by your available disk space). They are also permanent and always available. The only disadvantage is the slowness that comes from disk access time.

Binary files have two features that distinguish them from text files:

- You can jump instantly to any structure in the file, which provides random access as in an array.
- You can change the contents of a structure anywhere in the file at any time.

Binary files also usually have faster read and write times than text files, because a binary image of the record is stored directly from memory to disk (or vice versa). In a text file, everything has to be converted back and forth to text, and this takes time.

C supports the file-of-structures concept very cleanly. Once you open the file you can read a structure, write a structure, or seek to any structure in the file. This file concept supports the concept of a **file pointer**. When the file is opened, the pointer points to record 0 (the first record in the file). Any **read operation** reads the currently pointed-to structure and moves the pointer down one structure. Any **write operation** writes to the currently pointed-to structure and moves the pointer down one structure. **Seek** moves the pointer to the requested record.

Keep in mind that C thinks of everything in the disk file as blocks of bytes read from disk into memory or read from memory onto disk. C uses a file pointer, but it can point to any byte location in the file. You therefore have to keep track of things.

The following program illustrates these concepts:

```
1 #include <stdio.h>
2
3 /* random record description - could be anything */
4 struct rec
5 {
6     int x, y, z;
7 };
8
9 /*
10  * writes and then reads 10 arbitrary records
11  * from the file "junk".
12  */
13 int main (void)
14 {
15     int i, j;
16     FILE *f;
17     struct rec r;
18
19     /* create the file of 10 records */
20     f = fopen("junk", "w");
21     if (!f)
22     {
23         return 1;
24     }
25     for (i = 1; i <= 10; i++)
26     {
27         r.x = i;
28         fwrite(&r, sizeof(struct rec), 1, f);
29     }
30     fclose(f);
31
32     /* read the 10 records */
33     f = fopen("junk", "r");
34     if (!f)
35     {
36         return 1;
37     }
38     for (i = 1; i <= 10; i++)
```

```
39     {
40         fread(&r, sizeof(struct rec), 1, f);
41         printf ("%d\n", r.x);
42     }
43     fclose(f);
44     printf ("\n");
45
46     /* use fseek to read the 10 records
47        in reverse order */
48     f = fopen("junk", "r");
49     if (!f)
50     {
51         return 1;
52     }
53     for (i = 9; i >= 0; i--)
54     {
55         fseek(f, sizeof(struct rec)*i, SEEK_SET);
56         fread(&r, sizeof(struct rec), 1, f);
57         printf ("%d\n", r.x);
58     }
59     fclose(f);
60     printf ("\n");
61
62     /* use fseek to read every other record */
63     f = fopen("junk", "r");
64     if (!f)
65     {
66         return 1;
67     }
68     fseek(f, 0, SEEK_SET);
69     for (i = 0; i < 5; i++)
70     {
71         fread(&r, sizeof(struct rec), 1, f);
72         printf ("%d\n", r.x);
73         fseek(f, sizeof(struct rec), SEEK_CUR);
74     }
75     fclose(f);
76     printf ("\n");
77
```



```
78     /* use fseek to read 4th record,
79        change it, and write it back */
80     f = fopen("junk", "r+");
81     if (!f)
82     {
83         return 1;
84     }
85     fseek(f, sizeof(struct rec)*3, SEEK_SET);
86     fread(&r, sizeof(struct rec), 1, f);
87     r.x = 100;
88     fseek(f, sizeof(struct rec)*3, SEEK_SET);
89     fwrite(&r, sizeof(struct rec), 1, f);
90     fclose(f);
91     printf ("\n");
92
93     /* read the 10 records to insure
94        4th record was changed */
95     f = fopen("junk", "r");
96     if (!f)
97     {
98         return 1;
99     }
100    for (i = 1; i <= 10; i++)
101    {
102        fread(&r, sizeof(struct rec), 1, f);
103        printf ("%d\n", r.x);
104    }
105    fclose(f);
106    return 0;
107 }
```

In this program, a structure description **rec** has been used, but you can use any structure description you want. You can see that **fopen** and **fclose** work exactly as they did for text files.

The new functions here are **fread**, **fwrite** and **fseek**. The **fread** function takes four parameters:

- A memory address
- The number of bytes to read per block
- The number of blocks to read

- The file variable

Thus, the line **fread(&r,sizeof(struct rec),1,f);** says to read 12 bytes (the size of **rec**) from the file **f** (from the current location of the file pointer) into memory address **&r**. One block of 12 bytes is requested. It would be just as easy to read 100 blocks from disk into an array in memory by changing 1 to 100.

The **fwrite** function works the same way, but moves the block of bytes from memory to the file. The **fseek** function moves the file pointer to a byte in the file. Generally, you move the pointer in **sizeof(struct rec)** increments to keep the pointer at record boundaries. You can use three options when seeking:

- **SEEK_SET**
- **SEEK_CUR**
- **SEEK_END**

SEEK_SET moves the pointer **x** bytes down from the beginning of the file (from byte 0 in the file). **SEEK_CUR** moves the pointer **x** bytes down from the current pointer position. **SEEK_END** moves the pointer from the end of the file (so you must use negative offsets with this option).

Several different options appear in the code above. In particular, note the section where the file is opened with **r+** mode. This opens the file for reading and writing, which allows records to be changed. The code seeks to a record, reads it, and changes a field; it then seeks back because the read displaced the pointer, and writes the change back.

Index

address operator, 12
array, 29

Boolean expression, 15

comma operator, 20
compiled language, 3
compiler, 3
computer memory, 45

Fahrenheit to Celsius conversion, 20
function prototypes, 27

heap, 69

if statement, 15
Incremental operators, 8

library, 5
logical expression, 15

Operator precedence, 8

printf, 9
programming language, 3

RAM, 45

stack, 69

type, 6
typecasting, 7
typedef, 7

variables, 6