



Gorgo C for C++ Programmers

April 11, 2026

Contents

Author Intro	6
0. How to use this booklet	7
Tips	7
Try It	7
Chapter Layout	7
Exercises	7
1. Introduction	8
Getting Started	8
Differences Summary	8
Printing with <code>printf</code>	9
Reading Input with <code>scanf</code>	10
Try It	11
Key Points	12
Exercises	12
2. Variables	13
Basic Types	13
Variables as Named Memory	13
Pointer Declarations	14
Arrays	14
<code>const</code>	16
Structures	16
Designated Initializers	18
<code>typedef</code>	18
Try It: Variables Starter	19
Key Points	20
Exercises	21
3. Strings	22
Declaring C Strings	22
String Literals vs. Arrays	22
String Functions	22
The Dangers of <code>strcat</code>	26
Character Classification: <code><ctype.h></code>	26
A Preview: <code>sprintf</code> and <code>scanf</code>	27
Try It: String Starter	28
Key Points	28
Exercises	29
4. Expressions	30
Assignment	30
Arithmetic Operators	30
Comparison and Logical Operators	31
Bitwise Operators	32
Compound Assignment Operators	33
Increment and Decrement	34
The Ternary Operator	34
Operator Precedence	34
Try It: Expressions Starter	35
Key Points	36

Exercises	36
5. Control Flow	38
if / else	38
while and do-while	38
break and continue	39
for Loops	40
switch Statements	41
goto	42
Try It: Control Flow Starter	43
Key Points	44
Exercises	44
6. Pointers	46
What Is a Pointer?	46
Declaring Pointers	46
The Address-Of Operator: &	46
Dereferencing: *	46
Pointers to Pointers	47
Visualizing Pointers in Memory	47
NULL Pointers	47
Pointers and Arrays	48
Pointers and Structures	48
Pass by Value (and Pointers as a Workaround)	49
Try It: Pointer Starter	49
Key Points	50
Exercises	50
7. Functions	52
Function Declarations and Definitions	52
Pass by Value	53
const Parameters	53
Passing Structures	53
Recursive Functions	54
Function Pointers	55
Try It: Functions Starter	57
Key Points	58
Exercises	58
8. Allocating Memory	60
Global Variables	60
Local Variables	60
Static Local Variables	61
Dynamic Allocation: malloc and free	61
Working with Raw Memory: memcpy and memset	62
Where Variables Live: A Summary	63
Try It: Memory Lifetimes	63
Key Points	64
Exercises	64
9. Numbers and Casting	65
Everything is a Number	65
Strings are Not Special	65
Converting Strings to Numbers	66
Integer Types	66

Floating Point Types	68
Casting	69
Try It: Numbers Starter	70
Key Points	71
Exercises	71
10. Standard I/O	73
scanf for Input	73
Scan Sets	73
The %m Modifier (POSIX)	74
stdin, stdout, and stderr	74
fprintf and fscanf	75
Opening and Closing Files	75
sprintf and sscanf	76
Binary I/O: fread and fwrite	76
Reading Lines: fgets	77
Buffering and fflush	77
Try It: Standard I/O Starter	78
Key Points	79
Exercises	79
11. Low-Level I/O	80
File Descriptors	80
read and write	80
open and close	80
Seeking: lseek	81
pread and pwrite	82
Try It: Low-Level I/O Starter	82
Key Points	83
Exercises	83
12. Odds and Ends	84
exit vs return	84
extern "C" — Calling C from C++	84
Pointer Ownership	85
Error Handling Without Exceptions	86
enum	87
union	87
qsort — Function Pointers in Action	88
Try It: Odds and Ends Starter	89
Key Points	90
Exercises	91
Conclusion	93
Appendix A: Macros	94
Object-Like Macros	94
Function-Like Macros	95
Stringification and Token Pasting	96
Variadic Macros	97
Multi-Level Expansion	97
X-Macros	98
Try It: Macro Starter	99
Key Points	100
Exercises	100

Author Intro

C was the first programming language i truly loved! i had taught myself BASIC and Z80 assembly, and later i learned Pascal and VAX assembly for college. i liked VAX assembly. it was at my internship that i was handed The C Programming Language K&R that i realized how much i could love a language. i also learned FORTRAN 77 at that internship. i do NOT like FORTRAN.

C has a simplicity about it that makes it easy to jump into. it has a rawness that lets you really control what is happening in your program and consequently shoot yourself in the foot. the really great thing about C is that it gives you visibility into what really happens when your code executes that is only surpassed by assembly but without the verbosity and awkward constructs of working with assembly.

this booklet is to be experimented with. if you want a book to read, these are not the droids you are looking for. if you want something to introduce you to key concepts of C with some pointers for experiments you can try yourself, read on and write some code! programming is learned by writing code; there is just no other way that comes close to it. not everyone starts out as a great programmer, but everyone can become one by writing code either from scratch or by modifying code of others.

i hope this booklet will be something that can help you learn to love the challenge/satisfaction/power/rewards that come with being a great programmer.

ben

ps - you may guess after reading the booklet that i really loved the 80s. you can listen to the music references in the text at the [sounds of c4c++ youtube playlist](#)

0. How to use this booklet

This is a short booklet to help a C++ programmer get up to speed with C. As a side benefit, you will probably find that you understand C++ better at the end of it.

Each chapter is meant to help you understand a topic, but you will still want to reference API descriptions for more specifics of the parameters and operating conditions. Hopefully, you'll have the context you need to understand API documentation.

Tips

Tips call out details that you need to pay special attention to. **Traps** warn you of common mistakes made. **Wut** calls out a detail that is counter-intuitive, so make sure you pay attention.

Try It

Remember:

The only way to learn a new programming language is by writing programs in it.

You need to write some code. Make sure you try writing some programs from scratch. At the end of most sections is a starter program that you can type in and experiment with. Don't use it as an excuse to avoid writing some of your own starter programs. It's the only way to master a language.

Chapter Layout

Each chapter follows a consistent structure:

- A **brief overview** at the top that sets the context and lists what the chapter covers.
- The **main content** with explanations, code examples, and callouts (Tips, Traps, and Wuts).
- A **Try It** starter program you can compile and modify.
- **Key Points** summarizing the most important takeaways.
- **Exercises** with a mix of questions to test your understanding.

Exercises

Don't skip the exercises at the end of the chapters. You can get the answer key, but don't look at the answer key before you work out the answer yourself. If you look at the answer key first, the concepts will not sink in.

1. Introduction

In the beginning, knowing C++ automatically meant you knew C. The original C++ compiler, `cfront`, literally translated C++ code into C before compiling it. But modern C++ has diverged dramatically from C. You have `std::string`, `std::vector`, smart pointers, RAII, templates, `iostream`, and exceptions — none of which exist in C. If someone hands you a C codebase today, your modern C++ instincts will not get you very far.

So why learn C? Because C is everywhere. Operating systems, embedded firmware, database engines, language runtimes — the foundational software that the world runs on is written in C. Even if you spend most of your career in C++, you will encounter C code, C libraries, and C APIs. Understanding C makes you a better programmer, period.

My go-to C textbook is *The C Programming Language (Second Edition)* by Brian Kernighan and Dennis Ritchie — often called “K&R” [1]. It is one of the most influential programming books ever written, and it opens like this:

Getting Started

The only way to learn a new programming language is by writing programs in it. The first program to write is the same for all languages: Print the words hello, world. This is the big hurdle; to leap over it you have to be able to create the program text somewhere, compile it successfully, load it, run it, and find out where your output went. With these mechanical details mastered, everything else is comparatively easy.

In C, the program to print “hello, world” is

```
#include <stdio.h>
main()
{
    printf("hello, world\n");
}
```

Just how to run this program depends on the system you are using. As a specific example, on the UNIX operating system you must create the program in a file whose name ends in “.c”, such as `hello.c`, then compile it with the command `cc hello.c`. If you haven’t botched anything, such as omitting a character or misspelling something, the compilation will proceed silently, and make an executable file called `a.out`. If you run `a.out` by typing the command `a.out`, it will print `hello, world`. On other systems, the rules will be different; check with a local expert.

Differences Summary

Notice the differences from C++. There is no `#include <iostream>`, no `std::println`, no `std::cout`. In C, you use `printf` from `<stdio.h>` for output. The file ends in `.c`, not `.cpp`. You compile with `cc` (the C compiler) rather than `c++`.

Here is a quick summary of the biggest differences you will encounter:

C++	C
<code>std::string</code>	char arrays with <code>'\0'</code>
<code>std::vector</code>	raw arrays or <code>malloc</code>
<code>std::cout</code> / <code>std::println</code>	<code>printf</code>
<code>std::cin</code>	<code>scanf</code>
<code>new</code> / <code>delete</code>	<code>malloc</code> / <code>free</code>
Smart pointers	Raw pointers (only option)
Classes and objects	Structs and functions

C++	C
References (&)	Pointers (*)
bool (built-in)	_Bool or #include <stdbool.h>
// comments	/* comments */ (C89); // allowed since C99

Tip: C source files use the .c extension and are compiled with cc. If you accidentally compile a .c file with c++, it will be treated as C++ and may accept syntax that real C compilers reject. Always use cc when writing C.

Printing with printf

```
int printf(const char *format, ...);
```

In C++, you use `std::cout` or `std::println` for output. In C, you use `printf` from `<stdio.h>`. Unlike `std::println`, `printf` does not automatically add a newline at the end of the output. If you want each call to end on its own line, you must include `\n` in the format string yourself.

The first argument to `printf` is a **format string** containing literal text and **format specifiers** that start with `%`. Each specifier is replaced by the corresponding argument that follows:

```
int year = 1984;
printf("Year: %d\n", year); // Year: 1984
```

Here are the format specifiers you will use most:

Specifier	Type	Example	Output
%d	int (decimal)	printf("%d", 42)	42
%x	int (hex, lowercase)	printf("%x", 255)	ff
%X	int (hex, uppercase)	printf("%X", 255)	FF
%f	double	printf("%f", 3.14)	3.140000
%e	double (scientific)	printf("%e", 3.14)	3.140000e+00
%c	char	printf("%c", 'A')	A
%s	char * (string)	printf("%s", "hola")	hola
%p	pointer	printf("%p", (void *)ptr)	0x7ffd...
%zu	size_t	printf("%zu", sizeof(int))	4



Trap: The first argument to `printf` should always be a string literal. Never pass a variable as the first argument. It may work, but it is a potential security vulnerability (format string attack). If you only want to print a string variable, don't do `printf(str)`, do `printf("%s", str)`.

You can control the width and precision of output by placing numbers between the `%` and the specifier letter. A number before the specifier sets the minimum field width, and a `.` followed by a number sets the precision (decimal places for floats, max characters for strings):

```
double score = 98.6;
printf("Score: %f\n", score); // 98.600000 (default: 6 decimal places)
printf("Score: %.2f\n", score); // 98.60 (2 decimal places)
printf("Score: %e\n", score); // 9.860000e+01 (scientific notation)
```

Zero-filled output is useful for track numbers, timestamps, and hex addresses. Place a `0` before the width to pad with zeros instead of spaces:

```

for (int i = 1; i <= 5; i++) {
    printf("Track %02d\n", i);
}
// Track 01
// Track 02
// Track 03
// Track 04
// Track 05

int color = 0xFF8800;
printf("Color: 0x%06X\n", color); // Color: 0xFF8800

int score = 95;
printf("Score: %d%%\n", score); // Score: 95%

```

Since % introduces a format specifier, you must write %% to print a literal percent sign.



Trap: printf does not check that your format specifiers match the types of your arguments. If you write printf("%d", 3.14), the compiler may warn you, but it will not stop you. The result is garbage. Always match specifiers to types: %d for int, %f for double, %s for char *, and so on.

Reading Input with scanf

```
int scanf(const char *format, ...);
```

In C++, you read input with `std::cin >>`. In C, you use `scanf` from `<stdio.h>`. Like `printf`, `scanf` uses a format string with specifiers — but instead of printing values, it reads them from standard input and stores them in variables.

The critical difference from `printf` is that `scanf` needs to fill in a variable rather than read a variable. C doesn't have references like C++ has. Instead, we give `scanf` the **address** (location in memory) of each variable, not the value. You obtain the address of a variable using the `&` operator. (It's kind of confusing that `&` is also used for references, right?)

```
int age;
scanf("%d", &age); // read an integer into age
```



Trap: Without the `&`, `scanf` would receive the current (uninitialized) value of `age` instead of its address. The compiler will warn you that it is wrong, but will still compile the code and you can try to run it even though it will not work.

Strings in C are an array of characters. A variable representing an array takes on the address of the first element of the array when used in expressions or passed to functions. (This is called **array decay**.) For this reason, you don't need to use the `&` operator when reading strings — the array variable already represents an address.

```
char name[50]; // make sure you have enough space
scanf("%s", name); // no & needed - name is already an address
```

`%s` reads only until the first whitespace character. If the user types `Rick James`, `scanf("%s", name)` stores only `"Rick"`. Chapter 10 covers how to read full lines and handle input more robustly.

You can also read multiple items in a single `scanf`. `scanf` returns the number of items read. If there is no more input to be read, also known as **end of file**, `scanf` will return EOF.

```
#include <stdio.h>

int main(void) {
    int track;
    char name[50];

    if (scanf("%d %s", &track, name) != 2) {
        printf("I needed a number and a name\n");
    } else {
        printf("Got %s for track %d\n", name, track);
    }
}
```

If `scanf` returns 2, we know that both variables have been read.



Trap: `scanf` does no bounds checking. If the user types more characters than your buffer can hold, it writes past the end of the array. You can limit how many characters `%s` reads by specifying a width: `scanf("%49s", name)` reads at most 49 characters (leaving room for the `'\0'`).

`scanf` will be covered more extensively in Chapter 10 (Standard I/O). For now, `%d` for integers and `%s` for strings are enough to write interactive programs.

Try It

```
#include <stdio.h>

int main(void) {
    int num;

    printf("Enter a number: ");
    scanf("%d", &num);

    // Basic format specifiers
    printf("Integer: %d\n", num);
    printf("Hex:      %x (lowercase) %X (uppercase)\n", num, num);
    printf("Char:      %c\n", num);
    printf("Float:     %f\n", 3.14);
    printf("Sci:      %e\n", 3.14);
    printf("String:    %s\n", "Hola");

    // Width and precision
    printf("\nFormatted output:\n");
    for (int i = 1; i <= 3; i++)
        printf("  Track %02d\n", i);

    printf("  Pi: %.2f\n", 3.14159);
    printf("  100% complete\n");
}
```

```
    return 0;
}
```

Key Points

- C uses `printf` from `<stdio.h>` for output — there is no `std::cout` or `std::println`.
- Format specifiers (`%d`, `%s`, `%f`, etc.) must match the types of the arguments passed to `printf`.
- C uses `scanf` for input — you must pass the **address** of each variable with `&` (except for arrays, which already decay to pointers).
- C source files end in `.c` and are compiled with `cc`, not `c++`.
- C has no classes, no templates, no exceptions, no smart pointers, and no `std::string`.
- `printf` does not add a newline automatically — you must include `\n` yourself.

Exercises

1. **Think about it:** C uses format specifiers in `printf` while C++ uses `operator<<` or `std::format`. What advantage does the format string approach give you when writing output to a log file? What is a disadvantage?

2. **What does this print?**

```
printf("%05d %x\n", 42, 255);
```

3. **Calculation:** How many bytes does the string literal "hello" occupy in memory?

4. **Where is the bug?**

```
double pi = 3.14159;
printf("Pi is %d\n", pi);
```

5. **Where is the bug?**

```
int count;
scanf("%d", count);
```

6. **Write a program** that reads the user's name (single word) and a year with `scanf`, then prints a greeting like Hello, Sting! The year is 1983. using `printf`.
7. **Write a program** that prints a 5x5 multiplication table using `printf` with width formatting so the columns are aligned.

2. Variables

In C++, you have `auto` to let the compiler figure out types, `std::string` to handle text, and classes to bundle data with behavior. In C, none of that exists. Every type is explicit, strings are raw character arrays, and if you want to group data together, you use a `struct` with no member functions. This chapter covers how C handles variables — from basic types and arrays to pointers, `const`, and structures.

Basic Types

C provides a small set of built-in types. There are no classes, no `std::string`, and no `bool` keyword (without a header). Here are the types you will use most:

Type	Typical Size	Description
<code>char</code>	1 byte	A single character (or small integer)
<code>short</code>	2 bytes	Short integer
<code>int</code>	4 bytes	Standard integer
<code>long</code>	4 or 8 bytes	Long integer (platform-dependent)
<code>long long</code>	8 bytes	At least 64-bit integer
<code>float</code>	4 bytes	Single-precision floating point
<code>double</code>	8 bytes	Double-precision floating point
<code>_Bool</code>	1 byte	Boolean (C99); use <code>bool</code> via <code><stdbool.h></code>

Each integer type has an unsigned variant that stores only non-negative values, giving you twice the positive range. For example, a signed `char` holds -128 to 127, while an unsigned `char` holds 0 to 255:

```
unsigned char brightness = 255;
unsigned int count = 4000000000U;
```



Tip: C99 added `<stdbool.h>`, which defines `bool`, `true`, and `false`. Without it, you must use `_Bool` for the type and integer values 0 and 1. Most modern C code includes `<stdbool.h>` and uses `bool` just like C++.

```
#include <stdbool.h>
```

```
bool done = false;
if (!done) {
    /* keep going */
}
```

Variables as Named Memory

A variable declaration does two things: it allocates a region of memory large enough to hold the declared type, and it gives that region a name. The amount of memory allocated depends on the type:

```
#include <stdio.h>
```

```
int main(void) {
    char letter = 'J';
    int year = 1981;
    double rating = 9.5;
```

```

printf("char:  %zu bytes\n", sizeof(letter)); // 1
printf("int:   %zu bytes\n", sizeof(year));   // 4
printf("double: %zu bytes\n", sizeof(rating)); // 8

return 0;
}

```

The `sizeof` operator returns the size in bytes of a type or variable. It evaluates at compile time, so there is no runtime cost. You can use it with a type name or with a variable:

```

printf("int is %zu bytes\n", sizeof(int));
printf("year is %zu bytes\n", sizeof(year));

```



Tip: `sizeof` returns a value of type `size_t`. Always use `%zu` to print it. Using `%d` is technically undefined behavior, even though it often appears to work.

Pointer Declarations

A pointer variable holds the address of another variable. You declare a pointer by placing `*` after the base type:

```

int score = 100;
int *p = &score; /* p holds the address of score */

```

The type before the `*` tells the compiler what kind of data lives at that address. An `int *` points to an `int`, a `char *` points to a `char`, and so on.

We will not go deeper into pointers here. The Pointers chapter covers dereferencing, pointer arithmetic, pointers to pointers, and how pointers interact with arrays and structures in detail.

Arrays

An array is a fixed-size sequence of elements of the same type, declared with `[]`:

```

int scores[5]; /* 5 uninitialized ints */
int primes[5] = {2, 3, 5, 7, 11}; /* initialized */
/* compiler counts: 5 chars (including '\0') */
char greeting[] = "Hola";

```

When you provide an initializer, the compiler can determine the size for you, so you can leave the brackets empty.

The “Value” of an Array Name

In most expressions, the name of an array evaluates to the address of its first element. This is called **decay** — the array “decays” into a pointer:

```

int primes[5] = {2, 3, 5, 7, 11};
int *p = primes; /* p points to primes[0]; no & needed */

```

This is why you can pass an array to a function that expects a pointer. The Pointers chapter will explore this relationship thoroughly.



Wut: The `sizeof` operator is one of the few contexts where an array does *not* decay to a pointer. `sizeof(primes)` gives the total size of the array, not the size of a pointer. If `int` is 4 bytes and an address is 8 bytes (common sizes on modern hardware), `sizeof(primes)` is 20 and `sizeof(p)` is 8.

Initialization

You can partially initialize an array — remaining elements are set to zero:

```
int totals[10] = {1, 2, 3}; /* totals[3] through totals[9] are 0 */
int zeros[100] = {0};      /* all 100 elements are 0 */
```

Multidimensional Arrays

C supports multidimensional arrays. A two-dimensional array is really an array of arrays, stored in **row-major** order — all elements of row 0 come first, then all elements of row 1, and so on:

```
int grid[3][4] = {
    {1, 2, 3, 4},
    {5, 6, 7, 8},
    {9, 10, 11, 12}
};
```

In memory, this is stored as: 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12 — twelve consecutive integers. The expression `grid[r][c]` accesses row `r`, column `c`.

Iterating through a 2D array:

```
#include <stdio.h>

int main(void) {
    int grid[3][4] = {
        {1, 2, 3, 4},
        {5, 6, 7, 8},
        {9, 10, 11, 12}
    };

    for (int r = 0; r < 3; r++) {
        for (int c = 0; c < 4; c++) {
            printf("%3d", grid[r][c]);
        }
        printf("\n");
    }
    return 0;
}
```



Trap: When passing a multidimensional array to a function, you must specify all dimensions except the first. The compiler needs the column count to calculate offsets:

```
void print_grid(int grid[][4], int rows); /* OK - column size specified */
You cannot write int grid[][] — the compiler would not know how wide each row is.
```

const

The `const` qualifier marks a variable as read-only. Any attempt to modify it is a compile-time error:

```
const int MAX_TRACKS = 12;
const double PI = 3.14159265358979;
```

const with Pointers

Things get interesting when `const` meets pointers. There are three combinations, and the **read right-to-left** rule helps you decode them:

Pointer to const data — you cannot modify the data through this pointer, but you can change where the pointer points:

```
const int *p = &x;
*p = 10;      /* ERROR: cannot modify data through p */
p = &y;       /* OK: p itself can change */
```

Read right-to-left: `p` is a pointer (`*`) to `int` that is `const`. The data is `const`, not the pointer.

const pointer to data — the pointer itself cannot change, but you can modify the data it points to:

```
int *const p = &x;
*p = 10;     /* OK: data can be modified */
p = &y;       /* ERROR: p itself is const */
```

Read right-to-left: `p` is a `const` pointer (`*`) to `int`. The pointer is `const`, not the data.

const pointer to const data — neither the pointer nor the data can change:

```
const int *const p = &x;
*p = 10;      /* ERROR */
p = &y;        /* ERROR */
```



Tip: The most common form is `const int *p` — a pointer through which you promise not to modify the data. You will see this constantly in function parameters, like `const char *msg`, where the function reads the data but does not change it.

Structures

A `struct` groups related variables together under one name. If you come from C++, think of a `struct` as a class with only public data members — no member functions, no constructors, no destructors, and no access specifiers:

```
struct song {
    char title[40];
    int year;
};
```

Declaring and Initializing

To declare a variable of a `struct` type, you write `struct` followed by the tag name:

```
struct song track;
track.year = 1981;
```

You can also initialize at declaration time:

```
struct song track = {"I Love Rock 'n' Roll", 1981};
```

Accessing Members

Use the `.` operator to access members:

```
#include <stdio.h>

int main(void) {
    struct song {
        char title[40];
        int year;
    };

    struct song track = {"I Love Rock 'n' Roll", 1981};

    printf("Title: %s\n", track.title); // I Love Rock 'n' Roll
    printf("Year: %d\n", track.year); // 1981

    return 0;
}
```

Assignment Copies

Assigning one struct to another copies the entire contents — every byte:

```
struct song original = {"I Love Rock 'n' Roll", 1981};
struct song copy = original;

printf("%s (%d)\n", copy.title, copy.year); // I Love Rock 'n' Roll (1981)
```

This is a **shallow copy**. If the struct contained a pointer, both copies would point to the same memory. For structs that contain only arrays and plain values (like `struct song` above), the copy is complete and independent.



Wut: Unlike C++, there is no copy constructor or operator= to customize what happens during assignment. C copies the raw bytes, period. If your struct contains a pointer to dynamically allocated memory, the copy will share that memory, leading to double-free bugs if you are not careful.

No Member Functions

In C++, you might write:

```
class Song {
public:
    void print() { std::cout << title << " (" << year << ")\n"; }
    // ...
};
```

In C, structs cannot have member functions. Instead, you write standalone functions that take a pointer to the struct:

```
void song_print(const struct song *s) {
    printf("%s (%d)\n", s->title, s->year);
}
```

The `->` operator accesses a member through a pointer — it is shorthand for `(*s).title`. We will cover this in detail in the Pointers chapter.



Tip: A common C pattern is to prefix functions with the struct name they operate on: `song_print`, `song_init`, `song_compare`. This gives you something like namespaced methods. Function pointers (covered in a later chapter) can even be stored in structs to simulate virtual functions.

Designated Initializers

C99 introduced **designated initializers**, which let you initialize struct members by name instead of relying on declaration order:

```
struct song track = {.year = 1985, .title = "I Can't Drive 55"};
```

This is clearer than positional initialization because the reader does not need to remember the order of members. It also lets you skip members you want left at zero:

```
struct album {
    char title[40];
    int year;
    int tracks;
};
```

```
struct album a = {.title = "Hysteria", .year = 1987};
// a.tracks is initialized to 0
```

Any member you do not mention is initialized to zero (or null for pointers).



Tip: Prefer designated initializers over positional initialization. They make the code self-documenting and protect you from bugs when someone reorders the struct members later.

Designated initializers also work with arrays. You can initialize specific elements by index:

```
int flags[8] = {[0] = 1, [4] = 1, [7] = 1};
// flags is {1, 0, 0, 0, 1, 0, 0, 1}
```



Wut: C++ did not support designated initializers until C++20, and even then it requires the initializers to appear in declaration order. C has no such restriction — you can list members in any order.

typedef

The `typedef` keyword creates an alias for an existing type. It does not create a new type — it just gives you a shorter or more descriptive name:

```
typedef unsigned long ulong;
typedef unsigned char byte;

ulong population = 4000000000UL;
byte channel = 83;
```

To remember the syntax, notice that without the typedef above, byte would have been a variable named byte of type unsigned char. With the typedef, instead of declaring byte to be a variable, we are declaring it to be the name of a new type alias.

One of the most common uses of typedef is to simplify struct declarations. Without typedef, you must write struct every time you use the type:

```
struct point {
    double x;
    double y;
};

struct point origin;  /* must say "struct point" every time */
```

With typedef, you can drop the struct keyword:

```
typedef struct {
    double x;
    double y;
} Point;

Point origin = {0.0, 0.0};  /* much cleaner */
```



Tip: In C++, you can use a struct name directly as a type. In C, you cannot — you must either write struct name every time or use typedef to create an alias. Most C codebases use typedef for any struct that appears frequently.

Try It: Variables Starter

```
#include <stdio.h>
#include <stdbool.h>

struct song {
    char title[40];
    int year;
};

int main(void) {
    /* Basic types and sizeof */
    char initial = 'J';
    int year = 1981;
    double rating = 9.5;
    bool classic = true;

    printf("=== Sizes ===\n");
    printf("char:   %zu byte\n", sizeof(initial));
    printf("int:    %zu bytes\n", sizeof(year));
    printf("double: %zu bytes\n", sizeof(rating));
```

```

printf("bool:   %zu byte\n", sizeof(classic));

/* Arrays */
int scores[] = {95, 87, 92, 78, 100};
int n = sizeof(scores) / sizeof(scores[0]);
printf("\n=== Array ===\n");
for (int i = 0; i < n; i++) {
    printf("scores[%d] = %d\n", i, scores[i]);
}
printf("Total array size: %zu bytes (%d elements)\n",
       sizeof(scores), n);

/* const */
const int MAX = 100;
printf("\nMAX = %d\n", MAX);

/* Struct */
struct song track = {"I Love Rock 'n' Roll", 1981};
printf("\n=== Struct ===\n");
printf("Title: %s\n", track.title);
printf("Year:   %d\n", track.year);
printf("Size:   %zu bytes\n", sizeof(track));

/* Struct copy */
struct song backup = track;
printf("Copy:   %s (%d)\n", backup.title, backup.year);

/* 2D array */
int grid[2][3] = {{1, 2, 3}, {4, 5, 6}};
printf("\n=== 2D Array ===\n");
for (int r = 0; r < 2; r++) {
    for (int c = 0; c < 3; c++) {
        printf("%3d", grid[r][c]);
    }
    printf("\n");
}

return 0;
}

```

Key Points

- C has no auto, no std::string, and no classes. Every type is spelled out explicitly.
- A variable declaration allocates memory of the type's size and gives it a name. Use sizeof to see how many bytes each type occupies.
- typedef creates type aliases — especially useful for structs so you do not have to write struct everywhere.
- Arrays are fixed-size, and the array name decays to a pointer to the first element in most contexts.
- Multidimensional arrays are stored in row-major order. When passing them to functions, all dimensions except the first must be specified.
- const marks a variable as read-only. With pointers, read right-to-left to determine what is const — the data, the pointer, or both.
- Structs group data together but have no member functions, constructors, or access specifiers. Assignment copies the raw bytes.

Exercises

1. **Think about it:** In C++, `auto x = 42;` lets the compiler deduce the type. C has no `auto` type deduction. What advantage does requiring explicit types give to someone reading unfamiliar C code?

2. **What does this print?**

```
int a[] = {10, 20, 30, 40};
printf("%zu %zu\n", sizeof(a), sizeof(a[0]));
```

3. **Calculation:** Given the following declarations on a system where `int` is 4 bytes, what is `sizeof(grid)`?

```
int grid[3][5];
```

4. **Where is the bug?**

```
const int MAX = 100;
int *p = &MAX;
*p = 200;
printf("MAX = %d\n", MAX);
```

5. **What does this print?**

```
struct point { int x; int y; };
struct point a = {3, 7};
struct point b = a;
b.x = 99;
printf("%d %d\n", a.x, b.x);
```

6. **Think about it:** In C, assigning one struct to another copies the raw bytes. What problem could this cause if the struct contains a `char *` member that points to dynamically allocated memory (via `malloc`)? How is this different from what C++ does by default?
7. **Write a program** that declares a struct `student` with fields `name` (a char array), `id` (an `int`), and `gpa` (a `double`). Create an array of 3 students, initialize them with values, and print each student's information using a loop. Use `%s`, `%d`, and `%.2f` in your `printf`.

3. Strings

In C++, you use `std::string` and barely think about what is happening under the hood. In C, there is no string type at all. A “string” in C is just an array of `char` that ends with a null character `'\0'`. Every string function in C depends on finding that null terminator to know where the string ends.

Declaring C Strings

There are several ways to create a string in C:

```
// compiler sizes it: 12 bytes (11 + '\0')
char greeting[] = "Hola, mundo";
char band[20] = "Depeche Mode";           // 20-byte buffer, only 13 used
char empty[10] = "";                      // 10 bytes, first byte is '\0'
```

When you write `"Hola, mundo"`, the compiler automatically adds a `'\0'` at the end. The array `greeting` is 12 bytes: 11 visible characters plus the null terminator.



Trap: Always remember the null terminator when sizing your buffers. The string `"hello"` needs 6 bytes, not 5. Off-by-one errors with null terminators are one of the most common bugs in C.

String Literals vs. Arrays

There is an important difference between these two declarations:

```
char arr[] = "I Can't Drive 55";    /* modifiable copy on the stack */
const char *ptr = "I Can't Drive 55"; /* pointer to read-only memory */
```

With `char arr[]`, the compiler creates a local array and copies the string into it. You can modify the contents freely.

With a string literal like `"I Can't Drive 55"`, the compiler stores the text in read-only memory. A pointer to a literal should be declared `const char *` because modifying the literal is undefined behavior:

```
arr[0] = 'i';    /* fine - arr is a modifiable array */
ptr[0] = 'i';    /* COMPILE ERROR - ptr is const char * */
```

If you forget the `const`, the compiler will not prevent the mistake:

```
char *bad = "I Can't Drive 55"; /* compiles, but dangerous */
bad[0] = 'i';                  /* UNDEFINED BEHAVIOR - likely a crash */
```



Trap: Modifying a string literal is one of the most common mistakes C++ programmers make when moving to C. The compiler may not warn you (unless you use `-Wwrite-strings`), and the program might even appear to work — until it crashes on a different platform or optimization level. Always use `const char *` when pointing to a string literal.

String Functions

C provides a library of string manipulation functions in `<string.h>`. These are the ones you will use most:

`strlen` — Get the Length

```
size_t strlen(const char *s);
```

Returns the number of characters before the null terminator:

```
#include <stdio.h>
#include <string.h>

int main(void) {
    char song[] = "Take On Me";
    printf("%s' has %zu characters\n", song, strlen(song)); // 10
    return 0;
}
```

Note that `strlen` does not count the `'\0'`. The array holding "Take On Me" is 11 bytes, but `strlen` returns 10.

`strcpy` / `strncpy` — Copy a String

```
char *strcpy(char *dest, const char *src);
char *strncpy(char *dest, const char *src, size_t n);
```

`strcpy` copies the source string (including the null terminator) into the destination buffer. You must make sure the destination is large enough:

```
char dest[20];
strcpy(dest, "Tainted Love"); // copies 13 bytes (12 chars + '\0')
```

`strncpy` is the safer variant — it copies at most `n` characters:

```
char dest[10];
strncpy(dest, "Enjoy the Silence", 9);
dest[9] = '\0'; // strncpy does NOT guarantee null termination!
```



Wut: `strncpy` does not null-terminate the destination if the source is longer than `n`. Always set the last byte yourself: `dest[sizeof(dest) - 1] = '\0';`

`strcmp` — Compare Strings

```
int strcmp(const char *s1, const char *s2);
```

In C++, you can compare strings with `==`. In C, you cannot — using `==` on two char arrays compares the *addresses*, not the contents. Use `strcmp` instead:

```
char a[] = "Rush";
char b[] = "Rush";

if (a == b) {
    // This compares ADDRESSES, not content - almost always false
}

if (strcmp(a, b) == 0) {
    // This compares the actual characters - correct!
    printf("Same band!\n");
}
```

`strcmp` returns 0 if the strings are equal, a negative value if the first is lexicographically less, and a positive value if it is greater.



Trap: Yes, `strcmp` returns 0 for equal strings. It is a common source of confusion. Think of it as returning the “difference” between the strings — zero means no difference.

strncmp — Compare a Prefix

```
int strncmp(const char *s1, const char *s2, size_t n);
```

`strncmp` works like `strcmp` but compares at most `n` characters. It is useful for checking whether a string starts with a particular prefix:

```
char cmd[] = "walk like an Egyptian";
if (strncmp(cmd, "walk", 4) == 0) {
    printf("Walking!\n");
}
```

Unlike `strncpy`, there are no gotchas — `strncmp` simply stops comparing after `n` characters or a null terminator, whichever comes first.

strchr / strrchr — Find a Character

```
char *strchr(const char *s, int c);
char *strrchr(const char *s, int c);
```

`strchr` finds the first occurrence of a character. `strrchr` finds the last:

```
char lyric[] = "Don't You Forget About Me";
char *first_o = strchr(lyric, 'o'); // points to 'o' in "Don't"
char *last_o = strrchr(lyric, 'o'); // points to 'o' in "About"

if (first_o != NULL) {
    printf("First 'o' at position %td\n", first_o - lyric);
}
```

Both functions return a pointer to the found character, or `NULL` if the character is not in the string.

strstr — Find a Substring

```
char *strstr(const char *haystack, const char *needle);
```

Finds the first occurrence of a substring within a string:

```
char title[] = "Blade Runner 1982";
char *found = strstr(title, "Runner");
if (found != NULL) {
    printf("Found: %s\n", found); // "Runner 1982"
}
```

Like `strchr`, it returns a pointer to the start of the match, or `NULL` if not found.

strcat / strncat — Concatenate Strings

```
char *strcat(char *dest, const char *src);
char *strncat(char *dest, const char *src, size_t n);
```

`strcat` appends one string to the end of another:

```
char message[50] = "I'll be ";
strcat(message, "back");
printf("%s\n", message); // "I'll be back"
```

This works fine as long as the destination buffer is large enough. But `strcat` does not check bounds — if you run out of space, it writes past the end of the array.

Trap: `strcat` is one of the most dangerous functions in C. It has no way to know how large the destination buffer is, so it blindly appends bytes. If the combined strings exceed the buffer size, you get a **buffer overflow** — one of the most common security vulnerabilities in the history of software. Use `strncat` instead.

`strncat` takes a maximum number of characters to append as the last argument. Make sure to leave space for the null terminator, so if there are three bytes left in a buffer, the maximum number of characters to append would be 2.

```
char buf[20] = "Hello";
strncat(buf, ", World!", sizeof(buf) - strlen(buf) - 1);
```

`strdup` — Duplicate a String

```
char *strdup(const char *s);
```

`strdup` allocates new memory on the heap, copies the string into it, and returns a pointer to the copy. You are responsible for freeing the memory when you are done:

```
char *copy = strdup("Video Killed the Radio Star");
printf("%s\n", copy);
free(copy); // you must free memory allocated by strdup
```



Tip: `strdup` calls `malloc` internally. Every call to `strdup` must eventually be paired with a call to `free`. If you forget, you have a memory leak. Note that `strdup` is a POSIX function, not part of the C standard until C23. It is available on virtually every system you will use, but compiling with `-std=c99 -pedantic` or `-std=c11 -pedantic` will produce a warning.

`strtok` — Tokenize a String

```
char *strtok(char *str, const char *delim);
```

`strtok` splits a string into tokens separated by any character in a delimiter set. In C++, you might use `std::istringstream` with `>>` or `std::string::find` — in C, `strtok` is the standard approach:

```
#include <stdio.h>
#include <string.h>

int main(void) {
    char line[] = "Girls Just Want to Have Fun";

    char *tok = strtok(line, " ");
    while (tok != NULL) {
        printf("%s\n", tok);
        tok = strtok(NULL, " ");
    }
    return 0;
}

// Output:
// 'Girls'
// 'Just'
// 'Want'
// 'to'
```

```
// 'Have'  
// 'Fun'
```

The first call passes the string to tokenize. Each subsequent call passes NULL to continue tokenizing the same string. `strtok` returns NULL when there are no more tokens.

There are two important gotchas. First, `strtok` **modifies the original string** by replacing delimiter characters with `'\0'`. If you need the original string preserved, make a copy with `strdup` before tokenizing. Second, `strtok` stores its state in a hidden static variable, which means it is **not thread-safe** and you cannot tokenize two strings at the same time.

Tip: Use `strtok_r` (POSIX) or `strtok_s` (C11 Annex K / Windows) instead of `strtok`. These reentrant versions take an extra `char **saveptr` parameter to store the tokenizer state, making them safe to use in multi-threaded programs and allowing nested tokenization:

```
char *strtok_r(char *str, const char *delim, char **saveptr);  
char *saveptr;  
char *tok = strtok_r(line, " ", &saveptr);  
while (tok != NULL) {  
    printf("%s\n", tok);  
    tok = strtok_r(NULL, " ", &saveptr);  
}
```

The Dangers of `strcat`

Let's look at a concrete example of why `strcat` is dangerous:

```
#include <stdio.h>  
#include <string.h>  
  
int main(void) {  
    char buf[12] = "Buenas ";    // 7 chars + '\0', 4 bytes remaining  
  
    // "noches" is 6 chars - doesn't fit in 4 remaining bytes!  
    strcat(buf, "noches");      // BUFFER OVERFLOW - undefined behavior  
  
    printf("%s\n", buf);       // might print garbage, might crash  
    return 0;  
}
```

The buffer is 12 bytes, "Buenas " uses 8 (including the `'\0'`), and "noches" needs 7 more bytes (including its `'\0'`). That is 14 bytes total in a 12-byte buffer. The extra bytes overwrite whatever happens to be next to the buffer in memory, which can corrupt other variables, crash the program, or — worst of all — create a security exploit.

Character Classification: `<ctype.h>`

The `<ctype.h>` header provides functions for classifying and converting individual characters. These are great when processing strings character by character:

Function	Returns non-zero if
<code>isalpha(c)</code>	<code>c</code> is a letter (A–Z, a–z)
<code>isdigit(c)</code>	<code>c</code> is a digit (0–9)
<code>isalnum(c)</code>	<code>c</code> is a letter or digit
<code>isspace(c)</code>	<code>c</code> is whitespace (space, tab, newline, etc.)
<code>isupper(c)</code>	<code>c</code> is an uppercase letter

Function	Returns non-zero if
islower(c)	c is a lowercase letter

toupper and tolower convert a character's case:

```
#include <ctype.h>
#include <stdio.h>
#include <string.h>

int main(void) {
    char title[] = "hysteria";
    for (size_t i = 0; i < strlen(title); i++) {
        title[i] = toupper((unsigned char)title[i]);
    }
    printf("%s\n", title); /* HYSTERIA */
    return 0;
}
```

Here is a handy function that checks whether a string is entirely digits:

```
int all_digits(const char *s) {
    for (; *s; s++) {
        if (!isdigit((unsigned char)*s))
            return 0;
    }
    return 1;
}
```



Wut: The `<ctype.h>` functions take an `int`, not a `char`. On platforms where `char` is signed, a character with a value above 127 would be passed as a negative `int`, which is undefined behavior. Always cast to `unsigned char` first: `toupper((unsigned char)c)`.

A Preview: `sprintf` and `scanf`

```
int sprintf(char *str, const char *format, ...);
int sscanf(const char *str, const char *format, ...);
```

C has two powerful functions for building and parsing strings that we will cover in detail in the Standard I/O chapter: `sprintf` writes formatted output into a string buffer (like `printf` but to a string), and `scanf` reads formatted input from a string (like `scanf` but from a string). They are the C programmer's Swiss Army knife for string manipulation:

```
char result[50];
int year = 1985;
sprintf(result, "The year is %d. Que bueno!", year);
// result is now "The year is 1985. Que bueno!"
```

Just as `strncpy` is the safer sibling of `strcpy`, `snprintf` is the safer sibling of `sprintf`:

```
char buf[20];
snprintf(buf, sizeof(buf), "The year is %d", 2112);
/* buf is "The year is 2112" - safely truncated if it were longer */
```

snprintf guarantees it will not write more than the specified number of bytes, preventing buffer overflows. We will cover sprintf, snprintf, and sscanf in detail in the Standard I/O chapter.

Try It: String Starter

This program exercises several string functions so you can see them in action:

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void) {
    // strlen
    char song[] = "Sweet Dreams";
    printf("Song: '%s' (%zu chars)\n", song, strlen(song));

    // strcpy
    char copy[20];
    strcpy(copy, song);
    printf("Copy: '%s'\n", copy);

    // strcmp
    printf("Compare '%s' to '%s': %d\n", song, copy, strcmp(song, copy));
    printf("Compare 'A' to 'B': %d\n", strcmp("A", "B"));

    // strcat
    char greeting[30] = "Buenos ";
    strcat(greeting, "dias");
    printf("Greeting: '%s'\n", greeting);

    // strchr and strstr
    char *ch = strchr(song, 'D');
    if (ch) printf("Found 'D' at position %td\n", ch - song);

    char *sub = strstr(song, "Dreams");
    if (sub) printf("Found substring: '%s'\n", sub);

    // strdup
    char *dup = strdup("Never Gonna Give You Up");
    printf("Duplicate: '%s'\n", dup);
    free(dup);

    return 0;
}
```

Key Points

- C strings are char arrays terminated by '\0'. There is no std::string.
- Always account for the null terminator when sizing buffers.
- String literals like "hello" live in read-only memory. Use char s[] = "... " for a mutable copy and const char *p = "... " for a read-only pointer.
- Use strcmp to compare strings — the == operator compares addresses, not content. strncmp compares at most n characters — useful for prefix matching.
- strcpy and strcat do not check buffer bounds. Prefer strncpy and strncat.

- `strdup` allocates memory with `malloc` — you must free the result.
- `strtok` splits strings into tokens but modifies the original string and is not thread-safe. Use `strtok_r` or `strtok_s` instead.
- `<ctype.h>` provides character classification (`isalpha`, `isdigit`, etc.) and conversion (`toupper`, `tolower`). Cast `char` to `unsigned char` before passing.
- `snprintf` is the safe alternative to `sprintf` — it limits output to a maximum buffer size.
- `strlen` returns the number of characters *before* the `'\0'`, not the buffer size.

Exercises

1. **Think about it:** Why does `strcmp` return 0 for equal strings rather than 1? How does this relate to the function's actual purpose?

2. **What does this print?**

```
char s[] = "Ghostbusters";
printf("%zu %zu\n", strlen(s), sizeof(s));
```

3. **Calculation:** What is `sizeof(buf)` for `char buf[20] = "HoLa";`?

4. **Where is the bug?**

```
char buf[10] = "Livin'";
strcat(buf, " on a Prayer");
printf("%s\n", buf);
```

5. **Where is the bug?**

```
char *a = "Hello";
char *b = "Hello";
if (a == b) {
    printf("Equal\n");
} else {
    printf("Not equal\n");
}
```

(Hint: what does `==` actually compare here? Is the output guaranteed?)

6. **Write a program** that reads a string from the user, reverses it in place using pointer arithmetic, and prints the result.

7. **Where is the bug?**

```
char *greeting = "We Got the Beat";
greeting[0] = 'w';
printf("%s\n", greeting);
```

8. **What does this print?**

```
printf("%d\n", strncmp("Final Countdown", "Final Fantasy", 5));
```

9. **What does this print?**

```
char c = '7';
printf("%d %d %d\n", isdigit((unsigned char)c) != 0, isalpha((unsigned char)c) != 0, isalnum((unsigned char)c) != 0);
```

10. **Write a program** that takes a string and counts how many uppercase letters, lowercase letters, digits, and other characters it contains. Use `<ctype.h>` functions.

4. Expressions

C and C++ share most of their operators, and if you have been writing C++ you will find the syntax immediately familiar. But there are important differences. C has no operator overloading — + always means arithmetic addition, never something a class author decided it should mean. The << and >> operators are strictly bitwise shifts, not I/O operations. And C uses `int` for boolean results — there is no built-in `bool` type (though C99 added `_Bool` and `<stdbool.h>`).

This chapter walks through the operators you will use every day in C, highlights a few traps, and ends with the precedence table you will want to bookmark.

Assignment

The = operator assigns a value to a variable. In C, assignment is an **expression** — it produces a value, which is the value being assigned. This lets you chain assignments:

```
int a, b, c;
a = b = c = 0; // all three are now 0
```

The chain works right to left: `c` gets 0, then `b` gets the value of that assignment (also 0), then `a` gets the same.

Because assignment is an expression, you can (and sometimes will) use it inside other expressions. A common pattern is assigning and testing a return value in one step:

```
int ch;
while ((ch = getchar()) != EOF) {
    putchar(ch);
}
```



Trap: Because = is assignment and == is comparison, a common mistake is writing `if (x = 5)` when you mean `if (x == 5)`. The first assigns 5 to `x` and then evaluates as true (since 5 is nonzero). Modern compilers warn about this, but it is still one of the most famous bugs in C:

```
int x = 0;
if (x = 5) {
    printf("This always runs!\n"); // x is now 5, which is true
}
```

Some programmers write the constant on the left — `if (5 == x)` — so that `if (5 = x)` would be a compiler error. This is called a **Yoda condition**.

Arithmetic Operators

The arithmetic operators work on numeric types just like in C++:

Operator	Operation	Example	Result
+	addition	3 + 4	7
-	subtraction	10 - 3	7
*	multiplication	6 * 7	42
/	division	17 / 5	3
%	remainder	17 % 5	2

Integer division truncates toward zero. This means `17 / 5` gives 3, not 3.4. If you want a floating-point result, at least one operand must be a floating-point type:

```
printf("%d\n", 17 / 5);      // 3
printf("%f\n", 17.0 / 5);  // 3.400000
```

The % operator gives the **remainder** after integer division. In C99 and later, the result of % has the same sign as the dividend (the left operand):

```
printf("%d\n", 17 % 5);    // 2
printf("%d\n", -17 % 5);  // -2
printf("%d\n", 17 % -5);  // 2
printf("%d\n", -17 % -5); // -2
```



Wut: The % operator is often called “modulo,” but it is technically the **remainder** operator. For positive numbers, remainder and modulo are the same. For negative numbers, they differ. In mathematics, modulo always returns a non-negative result. In C, % preserves the sign of the dividend. If you need a true modulo that always returns a non-negative value, you need to adjust the result yourself:

```
/* assumes m > 0 */
int mod(int a, int m) {
    int r = a % m;
    return r < 0 ? r + m : r;
}
```

Comparison and Logical Operators

Comparison operators produce 1 for true and 0 for false. The result type is int, not bool:

Operator	Meaning
==	equal to
!=	not equal to
<	less than
>	greater than
<=	less than or equal to
>=	greater than or equal to

Logical operators combine boolean expressions:

Operator	Meaning
&&	logical AND
	logical OR
!	logical NOT

Both && and || use **short-circuit evaluation**, just like C++. With &&, if the left side is false, the right side is never evaluated. With ||, if the left side is true, the right side is skipped:

```
int *p = NULL;
if (p != NULL && *p > 0) {
    // safe - *p is only evaluated if p is not NULL
}
```

No Built-in bool

In C++, `bool` is a built-in type. In C89, there is no boolean type at all — you use `int` where `0` is false and anything nonzero is true. C99 added `_Bool` as a keyword and `<stdbool.h>` as a convenience header that defines `bool`, `true`, and `false`:

```
#include <stdbool.h>

bool is_valid = true;
if (is_valid) {
    printf("All aboard the Crazy Train!\n");
}
```

Without `<stdbool.h>`, you will see code like this:

```
int done = 0;    // 0 means false
while (!done) {
    // ... do work ...
    done = 1;    // nonzero means true
}
```



Tip: In C, any nonzero value is true. The number 42, the character 'A', and the pointer `0x7fff` are all true. Only `0` (and `NULL` for pointers) is false. This is why you can write `if (ptr)` instead of `if (ptr != NULL)` — they mean the same thing.

Bitwise Operators

Bitwise operators work on the individual bits of integer values. In C++, `<<` and `>>` are commonly used for stream I/O. In C, they are exclusively bit shift operators.

Operator	Operation	Example	Result
<code>&</code>	bitwise AND	<code>0xF0 & 0x3C</code>	<code>0x30</code>
<code>\ </code>	bitwise OR	<code>0xF0 \ 0x0F</code>	<code>0xFF</code>
<code>^</code>	bitwise XOR	<code>0xFF ^ 0x0F</code>	<code>0xF0</code>
<code>~</code>	bitwise NOT	<code>~0x00</code>	<code>0xFF...FF</code>
<code><<</code>	left shift	<code>1 << 3</code>	<code>8</code>
<code>>></code>	right shift	<code>16 >> 2</code>	<code>4</code>

Flag Manipulation

One of the most common uses of bitwise operators in C is manipulating **flags** — individual bits within an integer that each represent an on/off setting:

```
#include <stdio.h>

#define FLAG_READ    (1 << 0)    // bit 0: 0x01
#define FLAG_WRITE   (1 << 1)    // bit 1: 0x02
#define FLAG_EXEC    (1 << 2)    // bit 2: 0x04

int main(void) {
    unsigned int perms = 0;

    // Set bits with |
```

```

perms |= FLAG_READ;
perms |= FLAG_WRITE;
printf("perms = 0x%02X\n", perms); // 0x03

// Check a bit with &
if (perms & FLAG_READ) {
    printf("Read permission is set\n");
}
if (!(perms & FLAG_EXEC)) {
    printf("Execute permission is NOT set\n");
}

// Toggle a bit with ^
perms ^= FLAG_WRITE;
printf("After toggling write: 0x%02X\n", perms); // 0x01

// Clear a bit with & and ~
perms &= ~FLAG_READ;
printf("After clearing read: 0x%02X\n", perms); // 0x00

return 0;
}

```

The pattern is straightforward:

- **Set** a bit: `flags |= BIT;`
- **Clear** a bit: `flags &= ~BIT;`
- **Toggle** a bit: `flags ^= BIT;`
- **Check** a bit: `if (flags & BIT)`



Tip: Shifting 1 to create bit masks — $(1 \ll n)$ — is a common idiom in C for hardware registers, permission flags, and option bitmasks. It is clearer than writing raw hex values because you can see exactly which bit position you are targeting.

Compound Assignment Operators

Compound assignment operators combine an arithmetic or bitwise operation with assignment. They work exactly as in C++:

Operator	Equivalent to
<code>a += b</code>	<code>a = a + b</code>
<code>a -= b</code>	<code>a = a - b</code>
<code>a *= b</code>	<code>a = a * b</code>
<code>a /= b</code>	<code>a = a / b</code>
<code>a %= b</code>	<code>a = a % b</code>
<code>a &= b</code>	<code>a = a & b</code>
<code>a \ = b</code>	<code>a = a \ b</code>
<code>a ^= b</code>	<code>a = a ^ b</code>
<code>a <<= b</code>	<code>a = a << b</code>
<code>a >>= b</code>	<code>a = a >> b</code>

These are not just shortcuts — they express intent more clearly. When you write `count += 1`, the reader knows you are incrementing `count`. When you write `count = count + 1`, the reader has to verify that the same variable appears on both sides.

Increment and Decrement

The `++` and `--` operators increment or decrement a variable by one. They come in prefix and postfix forms:

```
int x = 5;
int a = ++x; // prefix: x becomes 6, then a gets 6
int b = x++; // postfix: b gets 6 (current value), then x becomes 7
```

In a standalone statement, `x++` and `++x` do the same thing — increment `x`. The difference only matters when the result is used in a larger expression.



Trap: Do not modify a variable more than once in the same expression. The result is **undefined behavior**:

```
int i = 3;
int result = i++ + ++i; // UNDEFINED BEHAVIOR - do not do this
```

The compiler is free to evaluate the sub-expressions in any order, and different compilers (or even the same compiler with different optimization levels) may produce different results. If you need multiple modifications, use separate statements.

The Ternary Operator

The ternary operator `? :` is a compact alternative to `if/else` for simple value selection:

```
int volume = 11;
const char *verdict = (volume > 10) ? "Loco" : "Tranquilo";
printf("Volume %d: %s\n", volume, verdict); // Volume 11: Loco
```

The syntax is `condition ? value_if_true : value_if_false`. The ternary operator is an expression, so it produces a value that can be used in assignments, function arguments, or anywhere a value is expected:

```
printf("Track %d is %s\n", track,
      (track % 2 == 0) ? "even" : "odd");
```



Tip: The ternary operator is great for simple one-line decisions. If your condition or either branch is complex, use a regular `if/else` instead. Readability matters more than cleverness.

Operator Precedence

When multiple operators appear in an expression, C evaluates them according to a precedence table. Here are the most important levels, from highest (evaluated first) to lowest:

Precedence	Operators	Description
1	<code>() [] -> .</code>	grouping, subscript, member access
2	<code>! ~ ++ -- + - * & (type) sizeof</code>	unary operators

Precedence	Operators	Description
3	* / %	multiplication, division, remainder
4	+ -	addition, subtraction
5	<< >>	bitwise shifts
6	< <= > >=	relational
7	== !=	equality
8	&	bitwise AND
9	^	bitwise XOR
10	\	bitwise OR
11	&&	logical AND
12	\ \	logical OR
13	? :	ternary
14	= += -= etc.	assignment
15	,	comma

Common Precedence Traps

The most dangerous precedence surprise is that **bitwise operators bind more loosely than comparison operators**:

```
// WRONG - this checks (x) & (0x04 == 0x04), which is (x) & (1)
if (x & 0x04 == 0x04) { ... }
```

```
// RIGHT - parentheses fix the precedence
if ((x & 0x04) == 0x04) { ... }
```

Similarly, `||` has lower precedence than `&&`, which matches mathematical convention (AND before OR) but can surprise you:

```
// This evaluates as: a || (b && c) - not (a || b) && c
if (a || b && c) { ... }
```



Tip: When in doubt, use parentheses. No one will fault you for writing `(a & b) == c` instead of relying on precedence rules. The few extra characters make your intent unmistakable and save the next reader (who might be you) from having to look up the precedence table.

Try It: Expressions Starter

```
#include <stdio.h>
```

```
int main(void) {
    // Assignment chaining
    int a, b, c;
    a = b = c = 1980;
    printf("a=%d b=%d c=%d\n", a, b, c);

    // Integer division and remainder
    printf("17 / 5 = %d\n", 17 / 5);
    printf("17 %% 5 = %d\n", 17 % 5);
    printf("-17 %% 5 = %d\n", -17 % 5);
}
```

```

// Boolean values are just ints
printf("(10 > 5) = %d\n", 10 > 5);
printf("(10 < 5) = %d\n", 10 < 5);

// Bitwise flag manipulation
unsigned int flags = 0;
flags |= (1 << 0); // set bit 0
flags |= (1 << 2); // set bit 2
printf("flags = 0x%02X\n", flags); // 0x05
printf("bit 1 set? %d\n", (flags & (1 << 1)) != 0); // 0
printf("bit 2 set? %d\n", (flags & (1 << 2)) != 0); // 1

// Ternary operator
int vol = 11;
printf("Volume: %s\n", (vol > 10) ? "Muy alto" : "Normal");

// Compound assignment
int total = 100;
total += 50;
total -= 20;
total *= 2;
printf("total = %d\n", total); // 260

return 0;
}

```

Key Points

- Assignment is an expression in C — it produces a value, enabling chaining (`a = b = c = 0`) and assignment within conditions.
- Integer division truncates toward zero. The `%` operator gives the remainder, which preserves the sign of the dividend.
- C uses `int` for boolean results: `0` is false, nonzero is true. Include `<stdbool.h>` for `bool`, `true`, and `false`.
- Bitwise `<<` and `>>` are shifts only — they are not overloaded for I/O as in C++.
- Use `|` to set bits, `&` to check bits, `^` to toggle bits, and `& ~` to clear bits.
- Bitwise operators have lower precedence than comparison operators. Always use parentheses when mixing them.
- Never modify a variable more than once in the same expression — the result is undefined behavior.

Exercises

1. **Think about it:** In C++, you can overload operators to give `+`, `<<`, `==`, and others custom meanings for your classes. C does not allow operator overloading. What advantage does this give you when reading unfamiliar C code? Can you think of a situation where operator overloading would have been genuinely useful in C?

2. **What does this print?**

```

int x = 10;
int y = x++ + ++x;
printf("%d %d\n", x, y);

```

(Be careful — is the answer even defined?)

3. **Calculation:** What is the result of each of these expressions?

```
25 / 4
25 % 4
-25 % 4
(1 << 4) | (1 << 1)
0xFF & 0x0F
```

4. **Where is the bug?**

```
int status = 0x07;
if (status & 0x04 == 0x04) {
    printf("Bit 2 is set\n");
}
```

5. **What does this print?**

```
int a = 5, b = 10;
a ^= b;
b ^= a;
a ^= b;
printf("a=%d b=%d\n", a, b);
```

6. **Where is the bug?**

```
int count = 0;
if (count = 0) {
    printf("El contador es cero\n");
} else {
    printf("El contador no es cero\n");
}
```

7. **Write a program** that takes an unsigned `int` and prints its value in binary (most significant bit first). Use bitwise operators to test each bit. Test it with the values 0, 1, 255, and 1024.

5. Control Flow

Control flow in C will feel very familiar if you are coming from C++. The `if`, `while`, `for`, and `switch` statements work mostly the same way. The differences are small but worth knowing: C has no range-based `for` loops, no structured bindings, and no `std::optional`. In C89, all variable declarations had to appear at the top of a block before any statements — C99 relaxed this and let you declare variables anywhere, which is how you are used to writing code in C++.

The biggest conceptual difference is that C has no native `bool` type. Conditions are just integers: zero is false, and any nonzero value is true.

`if / else`

The `if` statement in C is identical to C++:

```
int score = 85;

if (score >= 90) {
    printf("Excelente!\n");
} else if (score >= 70) {
    printf("Passing\n");
} else {
    printf("Try again\n");
}
```

The condition in an `if` is an integer expression. Zero means false, nonzero means true. There is no built-in `bool` type in C89. C99 added `_Bool` and the convenience header `<stdbool.h>`, which defines `bool`, `true`, and `false`:

```
#include <stdbool.h>

bool is_hurricane = true;
if (is_hurricane) {
    printf("Here I am\n");
}
```

Without `<stdbool.h>`, C programmers traditionally use `int` for boolean values and `0/1` for false/true, or define their own macros.



Trap: Because conditions are just integers, assignments inside `if` are legal and a common source of bugs:

```
int x = 0;
if (x = 5) {           // BUG: assigns 5 to x, then tests 5 (nonzero = true)
    printf("oops\n"); // always prints
}
```

The compiler may warn you about this, but it will not stop you. If you mean to compare, use `==`. Compiling with `-Wall` helps catch these.

`while` and `do-while`

A `while` loop tests the condition *before* executing the body. If the condition is false from the start, the body never executes:

```

int countdown = 5;
while (countdown > 0) {
    printf("%d... ", countdown);
    countdown--;
}
printf("Vamos!\n");
// 5... 4... 3... 2... 1... Vamos!

```

A do-while loop tests *after* the body, guaranteeing at least one iteration. This is useful when you need to perform an action before you can check whether to continue:

```

#include <stdio.h>

int main(void) {
    int choice;
    do {
        printf("1) Rock  2) Paper  3) Scissors  0) Quit\n");
        printf("Enter choice: ");
        scanf("%d", &choice);
        printf("You picked %d\n", choice);
    } while (choice != 0);

    printf("Adios!\n");
    return 0;
}

```

The semicolon after while (choice != 0) is required — forgetting it is a syntax error.



Tip: Use do-while when the loop body must execute at least once. Menu loops and input validation loops are classic use cases. If you find yourself duplicating code before a while loop just to set up the first test, a do-while is probably cleaner.

break and continue

break exits the nearest enclosing loop (or switch) immediately. continue skips the rest of the current iteration and jumps to the next one.

```

#include <stdio.h>

int main(void) {
    /* break example: stop at the first multiple of 7 */
    for (int i = 1; i <= 20; i++) {
        if (i % 7 == 0) {
            printf("Found it: %d\n", i);
            break;
        }
    }

    /* continue example: skip odd numbers */
    for (int i = 0; i < 10; i++) {
        if (i % 2 != 0)
            continue;
        printf("%d ", i);
    }
}

```

```

}
printf("\n");
// 0 2 4 6 8

return 0;
}

```

break and continue only affect the innermost loop. If you have nested loops and need to break out of an outer loop, you either use a flag variable or, in some cases, goto (discussed later in this chapter).

for Loops

The for loop has the same structure as in C++:

```

for (init; condition; update) {
    /* body */
}

```

Here is a classic example iterating over an array:

```

int scores[] = {90, 84, 77, 95, 88};
int n = sizeof(scores) / sizeof(scores[0]);

for (int i = 0; i < n; i++) {
    printf("Score %d: %d\n", i + 1, scores[i]);
}

```

C99 allows you to declare the loop variable inside the for statement, just like modern C++. In C89, you had to declare it before the loop:

```

/* C89 style */
int i;
for (i = 0; i < n; i++) {
    printf("%d\n", scores[i]);
}

/* C99 style - preferred */
for (int i = 0; i < n; i++) {
    printf("%d\n", scores[i]);
}

```



Tip: C has no range-based for loop. There is no `for (auto x : vec)`. You always iterate with an index or a pointer. The `sizeof(arr) / sizeof(arr[0])` idiom gives you the element count of a stack-allocated array, but it does not work on pointers — a pointer does not carry size information. This is why C functions that take arrays almost always take a separate size parameter.

You can iterate over an array with a pointer instead of an index. This is idiomatic C and worth getting comfortable with:

```

int nums[] = {10, 20, 30, 40, 50};
int *end = nums + sizeof(nums) / sizeof(nums[0]);

for (int *p = nums; p < end; p++) {
    printf("%d ", *p);
}

```

```

}
printf("\n");
// 10 20 30 40 50

```

Any part of the for header can be omitted. Omitting all three creates an infinite loop:

```

for (;;) {
    /* runs forever - use break to exit */
}

```

switch Statements

A switch statement selects among multiple cases based on an integer expression. If you have used switch in C++, the syntax is identical:

```

#include <stdio.h>

int main(void) {
    int wind = 5;    /* Beaufort scale */

    switch (wind) {
    case 0:
        printf("Calm\n");
        break;
    case 5:
        printf("Fresh breeze\n");
        break;
    case 12:
        printf("Huracan!\n");
        break;
    default:
        printf("Wind level: %d\n", wind);
        break;
    }

    return 0;
}

```

Each case must be an integer constant expression. You cannot use strings, floats, or variables as case labels — only values the compiler can evaluate at compile time.

Fall-through is the most important thing to understand about switch in C. If you forget a break, execution falls through to the next case. This is sometimes intentional:

```

char grade = 'B';

switch (grade) {
case 'A':
case 'B':
case 'C':
    printf("Passing\n");
    break;
case 'D':
case 'F':
    printf("Not passing\n");
    break;
default:

```

```

    printf("Invalid grade\n");
    break;
}

```

Here, cases 'A', 'B', and 'C' all fall through to the same printf. This is deliberate and a common pattern. But accidental fall-through is a frequent bug:

```

switch (x) {
case 1:
    printf("one\n");
    /* oops, forgot break - falls into case 2 */
case 2:
    printf("two\n");
    break;
}

```

If x is 1, this prints both "one" and "two."



Trap: Every case should end with break unless you intentionally want fall-through. When you do use fall-through on purpose, add a comment like `/* fall through */` so the next person reading the code knows it is deliberate. Some compilers recognize this comment and suppress warnings.

goto

In C++, you were probably taught to never use goto. In C, goto has a legitimate and widely used role: the **cleanup pattern**. When a function acquires multiple resources (files, memory, locks), and something goes wrong partway through, goto provides a clean way to release everything in the correct order.

Here is a brief example:

```

#include <stdio.h>
#include <stdlib.h>

int process(const char *path) {
    int status = -1;

    FILE *f = fopen(path, "r");
    if (!f) return -1;

    char *buf = malloc(1024);
    if (!buf) goto close_file;

    /* ... do work with f and buf ... */
    status = 0;

    free(buf);
close_file:
    fclose(f);
    return status;
}

```

If malloc fails, control jumps to close_file, which closes the file that was already opened. Without goto, you would need deeply nested if statements or duplicated cleanup code. The goto cleanup pattern is used

extensively in production C code, including the Linux kernel.



Tip: The goto cleanup pattern is covered in more detail in the Odds and Ends chapter. For now, just know that goto in C is not the taboo it is in C++. When used strictly for forward jumps to cleanup labels at the end of a function, it makes resource management clearer, not messier.

goto has two restrictions: you can only jump within the same function, and you cannot jump over a variable declaration that has an initializer (in C99+).

Try It: Control Flow Starter

```
#include <stdio.h>
```

```
int main(void) {
    /* if / else */
    int wind = 74; /* mph */
    if (wind >= 74) {
        printf("Rock you like a hurricane!\n");
    } else if (wind >= 39) {
        printf("Tropical storm\n");
    } else {
        printf("Calm seas\n");
    }
}

/* while */
int n = 1;
while (n <= 5) {
    printf("%d ", n);
    n++;
}
printf("\n");

/* do-while: repeat until valid input */
int guess;
do {
    printf("Guess (1-10): ");
    scanf("%d", &guess);
} while (guess < 1 || guess > 10);
printf("You guessed %d\n", guess);

/* for with break and continue */
printf("Even numbers up to 20: ");
for (int i = 1; i <= 100; i++) {
    if (i > 20)
        break;
    if (i % 2 != 0)
        continue;
    printf("%d ", i);
}
printf("\n");

/* switch */
```

```

int track = 3;
printf("Side B, Track %d: ", track);
switch (track) {
case 1: printf("Big City Nights\n"); break;
case 2: printf("Wind of Change\n"); break;
case 3: printf("Rock You Like a Hurricane\n"); break;
default: printf("Unknown track\n"); break;
}

return 0;
}

```

Key Points

- C's control flow statements (if, while, do-while, for, switch) are syntactically identical to C++.
- C uses integers for boolean conditions: 0 is false, nonzero is true. Include `<stdbool.h>` for `bool`, `true`, and `false` (C99+).
- `break` exits the nearest loop or switch. `continue` skips to the next iteration.
- C has no range-based for loop. Use index or pointer iteration.
- switch cases must be integer constants. Watch for accidental fall-through.
- `goto` is legitimate in C for the cleanup pattern — forward jumps to release resources in reverse order.

Exercises

1. **Think about it:** C uses 0 for false and nonzero for true, while C++ has a built-in `bool` type. What practical problems can arise from using integers as booleans? Can you think of a case where a nonzero value that is not 1 might cause a subtle bug?

2. **What does this print?**

```

for (int i = 0; i < 5; i++) {
    if (i == 3)
        continue;
    printf("%d ", i);
}
printf("\n");

```

3. **What does this print?**

```

int x = 2;
switch (x) {
case 1:
    printf("uno ");
case 2:
    printf("dos ");
case 3:
    printf("tres ");
    break;
default:
    printf("other ");
}
printf("\n");

```

4. **Where is the bug?**

```

int total = 0;
int i;

```

```
for (i = 0; i < 10; i++);
{
    total += i;
}
printf("Total: %d\n", total);
```

5. **Calculation:** How many times does the body of this loop execute?

```
int count = 0;
int i = 10;
do {
    count++;
    i--;
} while (i > 10);
```

6. **Where is the bug?**

```
int level = 5;
if (level = 10) {
    printf("Max level!\n");
}
```

7. **Write a program** that reads integers from the user (using scanf) until the user enters 0. Print the sum and average of all numbers entered (not counting the 0). Use a do-while or while loop.

6. Pointers

If you have been writing modern C++, you may have rarely (or never) used raw pointers. Smart pointers like `std::unique_ptr` and `std::shared_ptr` manage memory for you. References let you pass objects without copying them. The standard library hides pointer details behind iterators and containers.

In C, none of that exists. Pointers are everywhere, and you must be comfortable with them. Every dynamic data structure, every function that needs to modify its arguments, every interaction with the operating system — all involve pointers.

What Is a Pointer?

A pointer is a variable that holds a memory address. That's it. Instead of holding a value like 42, a pointer holds the *location* where 42 is stored.

Declaring Pointers

A pointer type is declared by placing a `*` after the base type. The type before the `*` tells you what kind of data lives at the address the pointer holds:

```
int *p;           // p is a pointer to an int
char *s;         // s is a pointer to a char
double *d;       // d is a pointer to a double
```



Trap: The `*` belongs to the variable, not the type. This declaration creates one pointer and one regular int:

```
int *p, q;       // p is a pointer to int; q is just an int
To declare two pointers, you need two stars: int *p, *q;
```

The Address-Of Operator: &

The `&` operator returns the address of a variable. You have seen `&` in C++ for references — in C, it is strictly the address-of operator:

```
int score = 100;
int *p = &score; // p now holds the address of score

printf("score = %d\n", score); // 100
// something like 0x7ffd5e8a3b2c
printf("address of score = %p\n", (void *)p);
```

Dereferencing: *

The `*` operator on a pointer gives you the value at the address the pointer holds. This is called **dereferencing**:

```
int score = 100;
int *p = &score;

printf("Value at p: %d\n", *p); // 100

*p = 200; // modify score through the pointer
printf("score is now: %d\n", score); // 200
```

Notice the dual use of `*`: in a declaration, it means “this is a pointer.” In an expression, it means “follow the pointer to the value.”

Pointers to Pointers

Since a pointer is just a variable, it has an address too. You can create a pointer to a pointer:

```
int val = 42;
int *p = &val;    // p points to val
int **pp = &p;    // pp points to p

printf("val = %d\n", val);    // 42
printf("*p = %d\n", *p);    // 42
printf("**pp = %d\n", **pp); // 42
```

You dereference `pp` twice: once to get `p`, and again to get `val`. Pointers to pointers show up frequently in C — for example, `main` can be declared as `int main(int argc, char **argv)`, where `argv` is a pointer to an array of string pointers.

Visualizing Pointers in Memory

Consider this small program:

```
int x = 1985;
int y = 80;
int *p = &x;
int **pp = &p;
```

Every variable lives at some address in memory. Here is what the layout looks like (using made-up but realistic addresses):

Variable	Address	Value	
x	0x1000	1985	
y	0x1004	80	
p	0x1008	0x1000	-----> x
pp	0x1010	0x1008	-----> p -----> x

The variable `x` lives at address `0x1000` and holds the value `1985`. The pointer `p` lives at address `0x1008` and holds the value `0x1000` — the address of `x`. The pointer-to-pointer `pp` lives at `0x1010` and holds `0x1008` — the address of `p`. Following the chain: `*pp` gives you `p` (which is `0x1008`), and `**pp` gives you `x` (which is `1985`).

Notice that `p` and `pp` are just variables that hold numbers. Those numbers happen to be memory addresses. There is nothing magical about a pointer — it is just a variable whose value is an address.



Tip: You can take the address of any variable with `&`, including the address of a pointer variable. The expression `&p` gives you the address where `p` itself is stored, not the address `p` points to.

NULL Pointers

A pointer that does not point to anything should be set to `NULL`:

```
int *p = NULL;    // p points to nothing

if (p != NULL) {
```

```

    printf("Value: %d\n", *p);
} else {
    printf("Pointer is NULL\n");
}

```

Dereferencing a NULL pointer is undefined behavior and usually crashes your program with a segmentation fault. Always check before dereferencing a pointer you did not initialize yourself.



Tip: In C, NULL is typically defined as `((void *)0)`. You may also see `0` used directly. C++11 introduced `nullptr` as a type-safe null pointer — C does not have `nullptr`, so use NULL.

Pointers and Arrays

A pointer might point to a single value in memory, or it might point to the first element of an array of values. There is nothing in the type system that tells you which — an `int *` looks the same either way:

```

int score = 100;
int *p = &score;           // points to one int

int nums[] = {10, 20, 30};
int *q = nums;             // points to the first of three ints

```

Both `p` and `q` are `int *`. The compiler does not know whether there are more `int` values after the one being pointed to. It is up to you, the programmer, to keep track of how many elements a pointer refers to and to stay within bounds.

You already saw in the Variables chapter that an array name decays to a pointer to its first element. Now let's see what that lets you do.

Pointer arithmetic works in units of the pointed-to type. If `p` is an `int *` and `int` is 4 bytes, then `p + 1` advances the address by 4 bytes to the next `int`. You never have to think about byte sizes — the compiler handles it:

```

int nums[] = {10, 20, 30, 40, 50};
int *p = nums;

printf("%d\n", *p);        // 10 (same as nums[0])
printf("%d\n", *(p + 1)); // 20 (same as nums[1])
printf("%d\n", p[2]);     // 30 - yes, you can use [] on pointers!

```



Tip: Array indexing `nums[i]` is actually syntactic sugar for `*(nums + i)`. This is why `2[nums]` technically works — it is `*(2 + nums)`, which is the same thing. Don't write code like that, but knowing this helps you understand how arrays and pointers relate.

Pointers and Structures

You already know how to declare structs and access members with `.` from the Variables chapter. Pointers to structures are extremely common in C — almost any non-trivial program passes struct pointers around rather than copying entire structs.

```

struct song {
    char title[40];

```

```

    int year;
};

struct song track = {"Karma Chameleon", 1983};
struct song *p = &track;

```

To access a field through a pointer, you must dereference the pointer first. But the `.` operator has higher precedence than `*`, so you need parentheses:

```

printf("Title: %s\n", (*p).title); // parentheses required
printf("Year: %d\n", (*p).year);

```

Writing `(*p).field` everywhere is tedious. C provides the `->` operator as a convenient shorthand:

```

printf("Title: %s\n", p->title); // same as (*p).title
printf("Year: %d\n", p->year); // same as (*p).year

```



Tip: The `->` operator is simply `(*p).field` written more clearly. You will see `->` far more often than `(*p).` in real C code. If you have a pointer to a struct, reach for `->`.

Pass by Value (and Pointers as a Workaround)

In C++, you can pass arguments by reference using `&`:

```

void increment(int &x) { x++; } // C++ - modifies the original

```

C does not have references. **All function parameters in C are pass by value** — the function receives a copy of the argument, not the original. If you want a function to modify a variable in the caller, you pass a *pointer* to it:

```

void increment(int *x) {
    (*x)++; // dereference the pointer, then increment
}

int main(void) {
    int score = 99;
    increment(&score); // pass the ADDRESS of score
    printf("%d\n", score); // 100
    return 0;
}

```

The function `increment` receives a *copy* of the pointer (the address), but since both the original and the copy point to the same memory, dereferencing either one reaches the same variable. This is how C simulates pass by reference.



Tip: Every time you see a function parameter with `*` in C, ask yourself: “Is this pointer here so the function can modify the caller’s variable, or because it needs to access a block of memory (like an array)?” Often it is both.

Try It: Pointer Starter

This program demonstrates the core pointer operations:

```

#include <stdio.h>

void swap(int *a, int *b) {
    int temp = *a;
    *a = *b;
    *b = temp;
}

int main(void) {
    // Basic pointer usage
    int val = 1985;
    int *p = &val;
    printf("val = %d, *p = %d\n", val, *p);
    printf("Address of val: %p\n", (void *)&val);

    // Modify through pointer
    *p = 1989;
    printf("After *p = 1989: val = %d\n", val);

    // Pointer to pointer
    int **pp = &p;
    printf("**pp = %d\n", **pp);

    // Pass by value with pointers
    int x = 10, y = 20;
    printf("Before swap: x=%d, y=%d\n", x, y);
    swap(&x, &y);
    printf("After swap: x=%d, y=%d\n", x, y);

    // Arrays and pointers
    const char *words[] = {"Totally", "Radical", "Tubular"};
    for (int i = 0; i < 3; i++) {
        printf("words[%d] = %s\n", i, words[i]);
    }

    return 0;
}

```

Key Points

- A pointer holds a memory address. Use & to get an address and * to dereference it.
- All pointers are the same size on a given system, regardless of the type they point to.
- Arrays decay to pointers in most expressions. `a[i]` is equivalent to `*(a + i)`.
- Pointer arithmetic moves in units of the pointed-to type, not bytes.
- Use `->` to access struct fields through a pointer. It is shorthand for `(*p).field`.
- All function parameters in C are pass by value. Pass a pointer to modify the caller's variable.

Exercises

1. **Think about it:** In C++, you can pass by reference to modify a caller's variable. Why do you think C was designed with only pass by value? What does this simplify in the language?
2. **What does this print?**

```
int a[] = {10, 20, 30, 40, 50};
```

```
int *p = a + 2;
printf("%d %d %d\n", *p, *(p - 1), p[1]);
```

3. **Calculation:** On a 64-bit system, what is `sizeof(int *)`, `sizeof(char *)`, and `sizeof(double *)`?

4. **Where is the bug?**

```
int *get_value(void) {
    int result = 42;
    return &result;
}
```

5. **What does this print?**

```
int x = 10;
int *p = &x;
int **pp = &p;
**pp = 20;
printf("%d\n", x);
```

6. **Where is the bug?**

```
struct song {
    char title[40];
    int year;
};

struct song *p = NULL;
printf("%s\n", p->title);
```

7. **Write a program** that declares an array of 5 integers, uses a pointer to iterate through the array, and prints each element along with its memory address.

7. Functions

C functions look a lot like C++ functions — same return types, same curly braces, same return statement. But several features you rely on in C++ are simply not available in C. There is no function overloading, no default arguments, no references, and no auto return type deduction. Every parameter is pass by value. If you want a function to modify a caller’s variable, you pass a pointer.

Despite these restrictions, C functions are straightforward. Once you understand the handful of differences from C++, you will find them easy to work with.

Function Declarations and Definitions

In C++, you can often define a function before it is called and skip the separate declaration. In C, you should always declare a function (provide its **prototype**) before you call it. A prototype tells the compiler the function’s return type, name, and parameter types — without the body:

```
int add(int a, int b);           // declaration (prototype)
```

The **definition** provides the body:

```
int add(int a, int b) {         // definition
    return a + b;
}
```

Prototypes typically go in header files (.h) so that other .c files can call the function. The definition lives in exactly one .c file.



Tip: In C++, `int foo()` means “takes no parameters.” In C, `int foo()` means “takes an *unspecified* number of parameters” — the compiler will not check your arguments at all. To declare a function that truly takes no parameters in C, write `int foo(void)`. Always use `void` in empty parameter lists.

Here is the difference in action:

```
int get_score(void);           // takes no parameters - compiler enforces this
int get_score();               // unspecified parameters - compiler won't check
```

You will see `int main(void)` throughout this book for exactly this reason.

No Overloading

In C++, you can have multiple functions with the same name but different parameter types:

```
int max(int a, int b);
double max(double a, double b); // C++ overload - fine
```

C does not support this. Every function must have a unique name. If you need a `max` for both `int` and `double`, you name them differently:

```
int max_int(int a, int b);
double max_double(double a, double b);
```

No Default Arguments

C++ lets you provide default values for parameters:

```
void greet(const char *name, int times = 1); // C++ - valid
```

C does not support default arguments. Every argument must be provided at every call site:

```
void greet(const char *name, int times); // C - no defaults allowed

greet("Iron Man", 3); // must pass both arguments
```

Pass by Value

As you saw in the Pointers chapter, all function parameters in C are pass by value — the function receives a copy of each argument. To modify a caller’s variable, pass a pointer. There are no & reference parameters in C; every “output” parameter is a pointer.

const Parameters

When a function takes a pointer parameter, the caller cannot tell from the call site whether the function will modify the data. The const keyword solves this by making a promise: “I will not modify what this pointer points to.”

```
#include <stdio.h>

void print_name(const char *name) {
    printf("I am %s\n", name);
    // name[0] = 'X'; // ERROR - name points to const data
}

int main(void) {
    print_name("Iron Man");
    return 0;
}
```

Using const serves two purposes. First, it documents intent — anyone reading the prototype knows the function will not modify the data. Second, it catches bugs at compile time. If you accidentally try to write through a const pointer, the compiler will stop you.



Tip: Make every pointer parameter const unless the function genuinely needs to modify the pointed-to data. This is one of the cheapest and most effective ways to prevent bugs in C.

You will see const char * everywhere in C — it is the standard way to accept a string that the function will read but not modify. Functions like printf, strlen, and strcmp all take const char * parameters.

Passing Structures

Structures can be passed by value just like any other type. The function receives a complete copy of the struct:

```
#include <stdio.h>

struct hero {
    char name[40];
    int power;
};

void print_hero(struct hero h) {
    printf("%s (power: %d)\n", h.name, h.power);
}
```

```
int main(void) {
    struct hero tony = {"Iron Man", 100};
    print_hero(tony); // passes a COPY of tony
    return 0;
}
```

This works, but copying a large struct every time you call a function is wasteful. If `struct hero` had thousands of bytes of data, every call to `print_hero` would copy all of it onto the stack.

The solution is to pass a pointer to the struct instead. Since the function only needs to read the data, use `const`:

```
void print_hero(const struct hero *h) {
    printf("%s (power: %d)\n", h->name, h->power);
}

int main(void) {
    struct hero tony = {"Iron Man", 100};
    print_hero(&tony); // passes only an 8-byte pointer
    return 0;
}
```



Tip: For small structs (a few bytes), passing by value is fine and sometimes clearer. For anything larger, prefer `const struct type *param`. When the function needs to modify the struct, drop the `const`.

If the function needs to modify the struct, you pass a non-const pointer:

```
void level_up(struct hero *h) {
    h->power += 10;
}
```

This pattern — `const` pointer for read-only access, non-const pointer for modification — is the C equivalent of `const` references and non-const references in C++.

Recursive Functions

C supports recursion just like C++. A function can call itself, and each call gets its own set of local variables on the stack.

Factorial

```
#include <stdio.h>

long factorial(int n) {
    if (n <= 1)
        return 1;
    return n * factorial(n - 1);
}

int main(void) {
    printf("5! = %ld\n", factorial(5)); // 5! = 120
    printf("10! = %ld\n", factorial(10)); // 10! = 3628800
}
```

```
    return 0;
}
```

Fibonacci

```
#include <stdio.h>

int fibonacci(int n) {
    if (n <= 0) return 0;
    if (n == 1) return 1;
    return fibonacci(n - 1) + fibonacci(n - 2);
}

int main(void) {
    for (int i = 0; i < 10; i++) {
        printf("%d ", fibonacci(i));
    }
    printf("\n");
    // 0 1 1 2 3 5 8 13 21 34
    return 0;
}
```



Trap: This Fibonacci implementation has exponential time complexity because it recomputes the same values over and over. It works for small inputs, but try `fibonacci(50)` and you will be waiting a long time. In practice, you would use an iterative approach or memoization.

Stack Depth

Every function call pushes a new frame onto the call stack. Recursive functions can exhaust the stack if the recursion goes too deep. There is no built-in protection — C has no exception mechanism, so there is no stack overflow exception to catch. The program will simply crash (usually with a segmentation fault). Keep your recursion depth reasonable, or convert deep recursion to iteration.

Function Pointers

In C++, you might use `std::function`, lambdas, or template parameters to pass behavior around. C has none of those — but it does have **function pointers**. A function pointer is a variable that holds the address of a function.

Declaring Function Pointers

The syntax takes some getting used to. A pointer to a function that takes two `int` parameters and returns an `int` looks like this:

```
int (*fp)(int, int);
```

Read it from the inside out: `fp` is a pointer (`*fp`) to a function that takes (`int, int`) and returns `int`.



Wut: The parentheses around `*fp` are critical. Without them, `int *fp(int, int)` declares a *function* that returns an `int *` — a completely different thing. The parentheses force `fp` to be a pointer first.

Using Function Pointers

You can assign a function’s name to a function pointer (the function name decays to a pointer, just like array names do) and then call the function through the pointer:

```
#include <stdio.h>

int add(int a, int b) { return a + b; }
int multiply(int a, int b) { return a * b; }

int main(void) {
    int (*op)(int, int);

    op = add;
    printf("3 + 4 = %d\n", op(3, 4));    // 7

    op = multiply;
    printf("3 * 4 = %d\n", op(3, 4));    // 12

    return 0;
}
```

Simplifying with typedef

The raw function pointer syntax is noisy. A typedef gives it a clean name:

```
#include <stdio.h>

typedef int (*binop_fn)(int, int);

int add(int a, int b) { return a + b; }
int subtract(int a, int b) { return a - b; }

void apply(binop_fn fn, int x, int y) {
    printf("result = %d\n", fn(x, y));
}

int main(void) {
    apply(add, 10, 3);    // result = 13
    apply(subtract, 10, 3);    // result = 7
    return 0;
}
```

Now `binop_fn` is a type that means “pointer to a function taking two `int` parameters and returning `int`.” You can use it for parameters, local variables, struct members, and arrays of function pointers.

Callbacks

Function pointers are often used as **callbacks** — you pass a function pointer to another function, which calls it at the right moment. This is the C equivalent of passing a lambda or functor in C++. The Odds and Ends chapter shows a practical example using `qsort`, the standard library's sorting function that takes a comparison callback.

Try It: Functions Starter

```
#include <stdio.h>

// Prototype with const pointer parameter
void shout(const char *msg);

// Struct and a function that takes a const pointer to it
struct song {
    char title[50];
    int year;
};

void print_song(const struct song *s) {
    printf("%s (%d)\n", s->title, s->year);
}

// Recursive factorial
long factorial(int n) {
    if (n <= 1) return 1;
    return n * factorial(n - 1);
}

// Function pointer demo
int sumar(int a, int b) { return a + b; }
int restar(int a, int b) { return a - b; }

typedef int (*math_fn)(int, int);

void compute(math_fn fn, int x, int y) {
    printf(" result = %d\n", fn(x, y));
}

// shout definition
void shout(const char *msg) {
    printf(">>> %s <<<\n", msg);
}

int main(void) {
    // const parameter
    shout("I Am Iron Man");

    // Passing a struct by pointer
    struct song fav = {"Iron Man", 1970};
    print_song(&fav);

    // Recursion
```

```

printf("7! = %ld\n", factorial(7)); // 5040

// Function pointers
printf("sumar:\n");
compute(sumar, 8, 3); // result = 11
printf("restar:\n");
compute(restar, 8, 3); // result = 5

return 0;
}

```

Key Points

- C functions have no overloading, no default arguments, and no references. Every parameter is pass by value.
- Use void in empty parameter lists: `int foo(void)`. In C, `int foo()` means “unspecified parameters,” not “no parameters.”
- Use `const` on pointer parameters when the function does not modify the pointed-to data. It documents intent and catches bugs.
- For large structs, prefer `const struct type *param` over pass by value to avoid expensive copies.
- Function pointers let you store and pass functions as values — C’s replacement for lambdas and `std::function`.
- Use `typedef` to make function pointer types readable.

Exercises

1. **Think about it:** C does not have function overloading. How does the C standard library handle providing similar functionality for different types? Look at `abs` (for `int`) and `fabs` (for `double`) as examples. What naming convention do you see?
2. **What does this print?**

```

void mystery(int x) {
    x = x * 2;
    printf("inside: %d\n", x);
}

int main(void) {
    int val = 5;
    mystery(val);
    printf("outside: %d\n", val);
    return 0;
}

```

3. **Where is the bug?**

```

int count_chars(const char *s) {
    int count;
    while (*s != '\0') {
        count++;
        s++;
    }
    return count;
}

```

4. **Calculation:** Given the struct below, approximately how many bytes are copied each time `process_data` is called with pass by value? Assume `int` is 4 bytes.

```
struct data {
    int values[1000];
    int count;
};

void process_data(struct data d) { /* ... */ }
```

5. **What does this print?**

```
int apply(int (*fn)(int, int), int a, int b) {
    return fn(a, b);
}

int mul(int a, int b) { return a * b; }

int main(void) {
    printf("%d\n", apply(mul, 6, 7));
    return 0;
}
```

6. **Where is the bug?**

```
void swap(int a, int b) {
    int tmp = a;
    a = b;
    b = tmp;
}

int main(void) {
    int x = 10, y = 20;
    swap(x, y);
    printf("x=%d y=%d\n", x, y);
    return 0;
}
```

7. **Write a program** that defines a function `void transform(int *arr, int n, int (*fn)(int))` which applies the function `fn` to each element of `arr`, modifying the array in place. Test it with a function that doubles each element and another that negates each element.

8. Allocating Memory

Every variable in your program lives somewhere in memory, but not all memory is created equal. Understanding where variables live — and how long they last — is important for writing correct C programs. It's such a challenging task that new languages such as Java and Rust were designed to help developers manage variable lifetimes.

Global Variables

A **global variable** is declared outside of any function. It is created when the program starts and exists until the program exits:

```
#include <stdio.h>

int high_score = 0; // global - lives for the entire program

void update_score(int points) {
    if (points > high_score) {
        high_score = points;
    }
}

int main(void) {
    update_score(1000);
    update_score(500);
    printf("High score: %d\n", high_score); // 1000
    return 0;
}
```

Global variables are visible to every function in the file. They are convenient but can make programs harder to reason about, since any function can change them.



Tip: Use global variables sparingly. When every function can read and write the same variable, bugs become harder to track down. Prefer passing data through function parameters.

Local Variables

A **local variable** is declared inside a function (or block). It is created when the function is called and destroyed when the function returns:

```
void greet(void) {
    char message[] = "Hola, amigo"; // local - exists only during greet()
    printf("%s\n", message);
}
// message is gone once greet() returns
```

Local variables live on the **stack** — a region of memory that grows and shrinks automatically as functions are called and return. You do not need to free stack memory; it is reclaimed automatically.



Trap: Never return a pointer to a local variable. The memory is freed when the function returns, and the pointer becomes a **dangling pointer** — it points to memory that no longer belongs to you:

```
int *bad(void) {
    int x = 42;
    return &x;    // BUG: x is destroyed when bad() returns
}
```

Static Local Variables

A **static local variable** has the scope of a local variable but the lifetime of a global. It is declared inside a function with the `static` keyword, created once when the program starts, and retains its value between calls:

```
#include <stdio.h>

void count_calls(void) {
    static int count = 0;    // initialized once, persists between calls
    count++;
    printf("Called %d time(s)\n", count);
}

int main(void) {
    count_calls();    // Called 1 time(s)
    count_calls();    // Called 2 time(s)
    count_calls();    // Called 3 time(s)
    return 0;
}
```

Without `static`, `count` would be reset to 0 on every call. With `static`, it lives in the data segment (like a global) but is only accessible inside `count_calls`.

Dynamic Allocation: `malloc` and `free`

```
void *malloc(size_t size);
void free(void *ptr);
```

Sometimes you need memory that outlives the function that created it, or memory whose size you do not know at compile time. For this, C provides `malloc` and `free` from `<stdlib.h>`.

`malloc` allocates a block of memory on the **heap** and returns a pointer to it. The heap is a region of memory that persists until you explicitly release it:

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

int main(void) {
    int *nums = malloc(5 * sizeof(int));
    if (nums == NULL) {
        printf("Allocation failed!\n");
        return 1;
    }

    for (int i = 0; i < 5; i++) {
        nums[i] = (i + 1) * 10;
    }
}
```

```

}

for (int i = 0; i < 5; i++) {
    printf("nums[%d] = %d\n", i, nums[i]);
}

free(nums); // release the memory back to the system
return 0;
}

```

`malloc` returns a `void *` — a generic pointer that can be assigned to any pointer type without a cast in C. It returns `NULL` if the allocation fails. Always check for `NULL` after calling `malloc`. Having said that, there is a school of thought held by some very good engineers that `NULL` checks just clutter the code. To handle `NULL` gracefully, there will be a check and logic every place where it could be `NULL`, so you end up with checks and logic that are rarely used and almost never tested. If you are out of memory, you'll probably need to shut down the program, so the reasoning goes that if you try to use a `NULL` pointer the CPU will do the check for you. For safety-critical systems, the above argument does not hold, although those systems often forbid the use of dynamically allocated memory entirely.

Tip: There are no smart pointers in C. There is no RAII. There is no garbage collector. If you call `malloc`, you *must* call `free` when you are done. If you forget, you leak memory. If you call `free` twice on the same pointer, you get undefined behavior. If you use a pointer after freeing it, you get undefined behavior. Memory management in C is entirely your responsibility.

`calloc` is a variant that allocates memory and initializes it to zero:

```

void *calloc(size_t count, size_t size);

int *nums = calloc(5, sizeof(int)); // 5 ints, all initialized to 0

```

And `realloc` lets you resize a previously allocated block:

```

void *realloc(void *ptr, size_t size);

nums = realloc(nums, 10 * sizeof(int)); // grow to 10 ints

```



Trap: Never assign the result of `realloc` directly back to the same pointer. If `realloc` fails, it returns `NULL` and the original memory is not freed — so `nums = realloc(nums, ...)` loses your only pointer to the original block, causing a memory leak. Use a temporary pointer instead:

```

int *tmp = realloc(nums, 10 * sizeof(int));
if (tmp == NULL) {
    // handle error - nums is still valid
} else {
    nums = tmp;
}

```

Working with Raw Memory: `memcpy` and `memset`

Two functions from `<string.h>` operate on raw bytes rather than strings. You will see them constantly in C code:

`memset` fills a block of memory with a byte value. It is commonly used to zero out a buffer:

```

void *memset(void *s, int c, size_t n);

```

```
int nums[10];
memset(nums, 0, sizeof(nums)); // set all bytes to 0
```

memcpy copies a block of bytes from one location to another. Unlike strcpy, it does not stop at a '\0' — you tell it exactly how many bytes to copy:

```
void *memcpy(void *dest, const void *src, size_t n);

int src[] = {10, 20, 30};
int dest[3];
memcpy(dest, src, sizeof(src)); // copy all 12 bytes (3 ints × 4 bytes)
```

Trap: memcpy requires that the source and destination do not overlap. If they might overlap (e.g., shifting elements within the same array), use memmove instead, which handles overlapping regions correctly.

```
void *memmove(void *dest, const void *src, size_t n);
```

Where Variables Live: A Summary

Kind	Where	Lifetime	Example
Global	Data segment	Entire program	int count = 0; (outside functions)
Local	Stack	Until function returns	int x = 5; (inside a function)
Static local	Data segment	Entire program	static int n = 0; (inside a function)
Dynamic	Heap	Until you call free	int *p = malloc(...)

Try It: Memory Lifetimes

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

int total = 0; // global - lives for the whole program

void add_to_total(int n) {
    int local = n; // local - gone when add_to_total returns
    total += local;
}

int main(void) {
    add_to_total(10);
    add_to_total(20);
    printf("Total: %d\n", total); // 30

    // dynamic - lives until we free it
    int *data = malloc(3 * sizeof(int));
    if (data == NULL) return 1;

    data[0] = 1985;
    data[1] = 1986;
    data[2] = 1987;

    for (int i = 0; i < 3; i++) {
```

```

        printf("data[%d] = %d\n", i, data[i]);
    }

    free(data);
    return 0;
}

```

Key Points

- Global variables live for the entire program; local variables live only until the function returns.
- Static local variables have the scope of a local but the lifetime of a global.
- malloc allocates memory on the heap. You must call free when done.
- calloc allocates and zeroes memory. realloc resizes an allocation.
- memcpy copies bytes between non-overlapping regions. Use memmove for overlapping regions.
- memset fills a block of memory with a byte value.

Exercises

1. **Think about it:** Why would you choose calloc over malloc followed by memset to zero?
2. **What does this print?**

```

#include <stdio.h>

void counter(void) {
    static int n = 0;
    n++;
    printf("%d ", n);
}

int main(void) {
    counter(); counter(); counter();
    return 0;
}

```

3. **Calculation:** On a system where int is 32 bits, how many bytes does malloc(5 * sizeof(int)) allocate?
4. **Where is the bug?**

```

int *p = malloc(10 * sizeof(int));
for (int i = 0; i < 10; i++) {
    p[i] = i;
}
free(p);
printf("%d\n", p[0]);

```

5. **Where is the bug?**

```

int *a = malloc(5 * sizeof(int));
int *b = a;
free(a);
free(b);

```

6. **Write a program** that uses malloc to allocate an array of n integers (where n is provided by the user via scanf), fills the array with squares (0, 1, 4, 9, ...), prints them, and frees the memory.

9. Numbers and Casting

To the CPU, everything is just a number. It has no concept of characters, strings, pointers, or objects. The CPU only knows about the numeric values in registers and memory addresses.

The types we assign to variables in C do not change the underlying bits; they simply tell the *compiler* how we want to interpret and use those numbers. Different numeric types provide different sizes (which determine the range of values they can hold) and different semantics (like whether they are signed or unsigned, integer or floating-point).

Everything is a Number

A char is just a small integer (usually 8 bits). Assigning 'A' is exactly the same as assigning the number 65 (in ASCII). You can perform math on characters just as you would on any integer:

```
#include <stdio.h>

int main(void) {
    char grade = 'A';
    int score = 65;

    // Both variables hold the exact same numeric value
    printf("grade as char: %c, as int: %d\n", grade, grade);
    printf("score as char: %c, as int: %d\n", score, score);

    char next_grade = grade + 1; // 65 + 1 = 66
    printf("Next grade: %c\n", next_grade); // 'B'

    return 0;
}
```

A pointer is also just a number. Specifically, it is an integer that represents a memory address. When you provide a pointer type like `int *`, you are telling the compiler: "Treat this address number as the location of an `int`." C allows you to print the address just like a number, usually formatted as hexadecimal using the `%p` specifier for pointers:

```
int val = 1986; // Year "Danger Zone" charted
int *p = &val;

printf("Address: %p\n", (void *)p); // e.g., 0x7ffd9b8
```

Strings are Not Special

Because everything is just a number, C has no native understanding of "strings." A string in C is merely an array of characters (small integers) that ends with a special 0 byte (the null terminator, `'\0'`).

When you write a string literal like "Africa", the compiler simply lays out an array of numbers (65, 102, 114, 105, 99, 97, 0) in read-only memory and gives you a `char *` pointer to the first number (65). The standard libraries (`<string.h>`) are what give us the illusion of strings, by processing these arrays of numbers until they hit that 0 byte.

```
#include <stdio.h>

int main(void) {
    char word[] = "Hola";

    printf("String: %s\n", word);
}
```

```

printf("Bytes: ");
for (int i = 0; i < (int)sizeof(word); i++)
    printf("%d ", word[i]);
printf("\n");

return 0;
}
// Output:
// String: Hola
// Bytes: 72 111 108 97 0

```

The five bytes are the ASCII values for H, o, l, a, and the null terminator.

Converting Strings to Numbers

Because strings are really arrays of characters, the text "1986" is entirely different from the integer 1986. Unsurprisingly, converting between the text representation of a number and its purely numeric form is a very common task.

To translate a string representation into an actual integer, use `strtol` (string to long) from `<stdlib.h>`:

```
long strtol(const char *str, char **endptr, int base);
```

The first argument is the string to parse. The second argument, `endptr`, is an optional pointer that `strtol` sets to point to the first character after the parsed number (pass `NULL` if you don't need it). The third argument is the numeric base (10 for decimal, 16 for hex, 0 to auto-detect from the prefix).

```

#include <stdio.h>
#include <stdlib.h>

int main(void) {
    char *year_str = "1986";

    // 10 is the base (base-10 decimal)
    long year = strtol(year_str, NULL, 10);

    printf("Year: %ld\n", year);
    return 0;
}

```

For floating-point conversions, use `strtod` (string to double) in the same way — it takes a string, an optional `endptr`, but no base argument since floating-point literals are always decimal.

Later on when we discuss Standard I/O, we will cover `sprintf` and `scanf`, which provide convenient formatting routines to convert back and forth between strings and numbers.

Integer Types

Integers are the most common type of number you will use in C. The CPU can do arithmetic on integers very quickly, and integers are also used for pointer arithmetic and array indexing. Integers can be signed or unsigned and can be of different sizes. Here is a table of the common sizes, ranges, and literal suffixes of integers on most 64-bit systems:

Type	Bytes	Range	Suffix
signed char	1	−128 to 127	
unsigned char	1	0 to 255	
short	2	−32,768 to 32,767	

Type	Bytes	Range	Suffix
unsigned short	2	0 to 65,535	
int	4	-2^{31} to $2^{31} - 1$	
unsigned int	4	0 to $2^{32} - 1$	U
long	8	-2^{63} to $2^{63} - 1$	L
unsigned long	8	0 to $2^{64} - 1$	UL
long long	8	-2^{63} to $2^{63} - 1$	LL
unsigned long long	8	0 to $2^{64} - 1$	ULL

Note that the sizes and ranges specified here can vary. The C standard's rules about sizes are so vague they aren't worth quoting here :'. Notice that the range is 1 number larger for negative numbers than for positive numbers. This is because most systems use **two's complement** representation for signed integers. The top bit indicates the sign: if it is set, the number is negative. But the remaining bits are not a simple magnitude — two's complement encodes negative values so that addition and subtraction work the same way for signed and unsigned numbers. A consequence of this encoding is that there is one more negative value than positive values, and zero is represented only one way (with the sign bit clear). While this may not make your math teacher happy, it gets worse when we talk about floating-point numbers!

You can explore the actual sizes and ranges on your machine using `<limits.h>`:

```
#include <stdio.h>
#include <limits.h>

int main(void) {
    printf("char:      %zu byte,  range %d to %d\n",
           sizeof(char), CHAR_MIN, CHAR_MAX);
    printf("short:     %zu bytes, range %d to %d\n",
           sizeof(short), SHRT_MIN, SHRT_MAX);
    printf("int:        %zu bytes, range %d to %d\n",
           sizeof(int), INT_MIN, INT_MAX);
    printf("long:       %zu bytes, range %ld to %ld\n",
           sizeof(long), LONG_MIN, LONG_MAX);
    printf("long long: %zu bytes, range %lld to %lld\n",
           sizeof(long long), LLONG_MIN, LLONG_MAX);

    return 0;
}
```



Trap: Most integer types are always signed by default, but `char` is different — its signedness is implementation-defined. Sadly, on `x86_64` CPUs, `char` is signed by default, and on `ARM` CPUs, `char` is unsigned by default! Watch out!

Integer Promotion

When you use an integer type in an expression that is equal to or less than the size of `int`, it is automatically promoted to `int` if `int` is large enough to hold all the values of the type. Otherwise, it is promoted to unsigned `int`. This is called integer promotion.

For expressions that involve larger integer types, the rules generally promote to the signedness and the size of the larger type.

Trap: One place where this can cause problems is when using `strlen`. `strlen` returns a `size_t` which is an unsigned integer type. If you subtract two `size_t` values, the result will be a `size_t`. If the first value is smaller than the second value, the result will be a large positive number. This can cause problems when using the result in other expressions.

```
#include <stdio.h>
#include <string.h>

int main(void) {
    char *a = "Jump";           // strlen = 4
    char *b = "Jump!!!!";      // strlen = 8

    // size_t is unsigned, so 4 - 8 wraps around!
    size_t diff = strlen(a) - strlen(b);
    printf("strlen(a) - strlen(b) = %zu\n", diff);

    // cast to a signed type to get the correct result
    long sdiff = (long)strlen(a) - (long)strlen(b);
    printf("Signed difference:      %ld\n", sdiff);

    return 0;
}
// Output:
// strlen(a) - strlen(b) = 18446744073709551612
// Signed difference:      -4
```

Floating Point Types

Floating point numbers are used for decimal numbers that may have fractional components. They also vary in size and precision, but they are always signed. Here is a table of the common sizes, ranges, and literal suffixes of floating point numbers on most 64-bit systems:

Type	Size	Range	Literal Suffix
float	4 bytes	-3.4e+38 to 3.4e+38	F
double	8 bytes	-1.7e+308 to 1.7e+308	(none, default)
long double	16 bytes	-1.2e+4932 to 1.2e+4932	L

Floating point numbers also have a sign bit, but unlike integers, the sign bit can be set on zero. This means floating point numbers have two zeros: positive zero and negative zero. (They compare as equal to each other though.) Floating point numbers can also be used to represent infinity and NaN (Not a Number).

Floating point numbers cannot always exactly represent a value. The fractional part of a number is stored as a sum of negative powers of 2 (1/2, 1/4, 1/8, ...). Values like 0.5 (which is 1/2) and 0.25 (which is 1/4) can be represented exactly, but values like 0.1 and 0.2 cannot. Rounding is used to get as close as possible to the value you want. This can cause surprising results: for example, 0.1 + 0.2 is not exactly equal to 0.3.

Observe the following code:

```
#include <stdio.h>
int main(void) {
    float f = 1.2;
    if (f != 1.2) printf("what?!?\n");
    if (f == 1.2f) printf("ok\n");
    return 0;
}
```

You would expect only the ok message to print, but the what message prints too! Why? The literal 1.2 is a double. When it is assigned to f, the value is rounded to fit in a float, losing some precision. In the first if, f is promoted back to double for the comparison, but the precision lost during the assignment is not recovered, so f and 1.2 are not equal. In the second if, 1.2f is a float literal, which was rounded the same way, so the comparison succeeds.

Casting

A **cast** is a way of forcing the compiler to treat a value of one type as another type. The syntax is (type) value.

Casts in C are much simpler than in C++. They are also much less magical! C has a single unified cast syntax (type)value. You can only cast between scalar types (integers, floating-point, and pointers). Remember that char is an integer type. The following table summarizes the allowed casts:

From / To	Integer	Floating-Point	Pointer
Integer	Yes	Yes	Yes
Floating-Point	Yes	Yes	No
Pointer	Yes	No	Yes

Casting from floating point to integer drops the decimal part (it does not round). As long as the numbers fit into the target type, the conversion works well. But if the number is too large, you get the dreaded “undefined behavior” — anything can happen.

Casting from a floating point to a pointer doesn’t make sense, but you can do it if you cast to an integer type first. Why would you ever want to do that though? Keep your coworkers happy and don’t do it.

By using a cast, you are telling the compiler: “I know what I am doing, suppress any warnings, and just treat this as the type I specified.”

```
double pi = 3.14159;  
int roughly_pi = (int)pi; // truncates to 3
```

Because C trusts you implicitly, casting can be dangerous. Magic *does not* happen when you cast.



Trap: A classic beginner mistake is trying to convert a string to an integer by casting the pointer.

```
char *movie_year = "1985"; // The Goonies  
int bad_year = (int)movie_year; // THIS IS A BUG! not 1985!!
```

Casting a char * string to an int does not convert the text "1985" into the number 1985. It tells the compiler to take the *memory address* where the string is stored and chop or pad it to fit inside an int. On a 64-bit system, the pointer is 8 bytes and the int is 4 bytes, so compiling this code will actually result in a warning that you are casting a pointer to an integer of different size! Always use functions like strtol to parse strings into numbers.

Casting pointers to other pointers

void * is a pointer that can point to anything. It is a pointer to a generic memory location. This is why you don’t have to cast pointers returned by malloc to a specific pointer type.

Take care when casting pointers to other pointer types. You must understand the memory layout of the structures you are working with. One common pattern is using a char * for byte-level arithmetic on a base address plus an offset, and then casting the result to the desired structure pointer type.

```

#include <stdio.h>

int main(void) {
    int nums[] = {1984, 1985, 1986, 1987};

    void *vp = nums;           // any pointer converts to void *
    int *ip = (int *)vp;      // cast back to use it
    printf("First: %d\n", ip[0]);

    // byte-level access with char *
    char *bp = (char *)nums;
    printf("First byte of nums[0]: 0x%02x\n", (unsigned char)bp[0]);

    return 0;
}
// Output:
// First: 1984
// First byte of nums[0]: 0xc0

```

Try It: Numbers Starter

This program exercises the key concepts from this chapter: characters as numbers, string-to-number conversion, integer sizes, and casting.

```

#include <stdio.h>
#include <stdlib.h>
#include <limits.h>

int main(void) {
    // Characters are just numbers
    char ch = 'A';
    printf("'%'c' is %d\n", ch, ch);
    printf("'%'c' + 3 = '%'c' (%d)\n", ch, ch + 3, ch + 3);

    // Strings are arrays of numbers
    char title[] = "Rio";
    printf("\'%s\' bytes: ", title);
    for (int i = 0; i < (int)sizeof(title); i++)
        printf("%d ", title[i]);
    printf("\n");

    // String to number conversion
    char *bpm_str = "120";
    long bpm = strtol(bpm_str, NULL, 10);
    printf("\'%s\' as a number: %ld\n", bpm_str, bpm);

    // Hex string to number
    long color = strtol("FF8000", NULL, 16);
    printf("0x%lX = %ld\n", color, color);

    // Integer sizes on this machine
    printf("\'sizeof(char) = %zu\n", sizeof(char));
    printf("sizeof(short) = %zu\n", sizeof(short));
    printf("sizeof(int) = %zu\n", sizeof(int));
}

```

```

printf("sizeof(long) = %zu\n", sizeof(long));
printf("UINT_MAX = %u\n", UINT_MAX);

// Casting: float to int truncates
double tempo = 120.7;
int whole = (int)tempo;
printf("\n(int)%0.1f = %d\n", tempo, whole);

// The classic trap: casting a pointer
char *year_str = "1982";
long bad = (long)year_str; // address, not 1982!
long good = strtol(year_str, NULL, 10);
printf("(long)\n\"1982\" = %ld (an address!)\n", bad);
printf("strtol(\n\"1982\", ...) = %ld\n", good);

// void * round-trip
int val = 42;
void *vp = &val;
int *ip = (int *)vp;
printf("\nvoid * round-trip: %d\n", *ip);

int hex_val = 0xbadd00d;
printf("%d %x\n", hex_val, hex_val);
// let's look at the bytes of hex_val (int has 4 bytes)
unsigned char *raw = (unsigned char *)&hex_val;
printf("little-endian: %02x %02x %02x %02x\n",
       raw[0], raw[1], raw[2], raw[3]);

return 0;
}

```

Key Points

- Under the hood, everything (characters, pointers) is just a number.
- C does not have native “strings,” only arrays of numbers ended with a 0.
- The type of a variable tells the compiler how you intend to interpret its numeric bits.
- C casts (type) value assert your intent to the compiler. It assumes you know what you are doing.
- Casting a pointer to an integer simply gives you the raw numeric memory address, not the parsed contents of a string. Use `strtol` to parse strings.

Exercises

1. **Think about it:** C++ provides several different cast operators (`static_cast`, `reinterpret_cast`, etc.) whereas C provides only one. What are the advantages of C++’s approach over C’s single cast syntax?
2. **What does this print?**

```
char letter = 'C';
printf("%c %d\n", letter + 2, letter + 2);
```
3. **Calculation:** Assuming a 64-bit system where pointers are 8 bytes and `int` is 4 bytes, what is the output of `sizeof("Danger")` and what is the output of `sizeof((int)0)`?
4. **Where is the bug?**

```
#include <stdio.h>
```

```
int main(void) {  
    char *score_str = "100";  
    int score = (int)score_str;  
    printf("You got a %d percent!\n", score);  
    return 0;  
}
```

5. **Write a program** that declares a double variable with a fractional component and uses casting to separate the integer part from the fractional part. Print both pieces.

10. Standard I/O

C's `<stdio.h>` library is your replacement for C++ `iostream`. It provides `printf` and `scanf` for formatted output and input, file operations with `fopen` and `fclose`, and binary I/O with `fread` and `fwrite`. Everything flows through the `FILE *` type — an opaque pointer to a structure that tracks the state of an I/O stream.

scanf for Input

```
int scanf(const char *format, ...);
```

You have already seen `printf` for output. For input, C uses `scanf`, which reads formatted data from standard input:

```
#include <stdio.h>

int main(void) {
    int year;
    printf("Enter a year: ");
    scanf("%d", &year);
    printf("You entered: %d\n", year);
    return 0;
}
```

Notice the `&` before `year`. Since C is pass by value, `scanf` needs the *address* of the variable so it can store the result there. Forgetting the `&` is a classic bug — the program compiles but crashes or produces garbage at runtime.

`scanf` uses similar format specifiers to `printf`, but not identical ones. Notably, `scanf` uses `%lf` for double while `printf` uses `%f`:

```
char name[50];
double gpa;
scanf("%s %lf", name, &gpa);
```

Note that `name` does not need `&` because an array name already points to the bytes we want to read into. But `gpa` does need a `&`, because `scanf` needs to know where `gpa` is stored to fill it in.



Trap: `scanf("%s", ...)` reads a single word (stopping at whitespace). It also has no bounds checking — it will happily overflow your buffer. Use a width specifier like `%49s` to limit input to 49 characters (plus `'\0'`).

Scan Sets

`scanf` supports **scan set specifiers** with `%[...]`, which let you define exactly which characters to accept. The scan set reads characters as long as they are in the set, and stops at the first character that is not:

```
char vowels[20];
// reads only vowels, stops at first non-vowel
scanf("%19[aeiouAEIOU]", vowels);
```

A caret `^` at the start of the set **negates** it — read everything *except* the listed characters. This gives you a way to read an entire line with `scanf`, since `%[^\n]` reads everything up to (but not including) the newline:

```
char line[80];
scanf("%79[^\n]", line); // reads a full line (up to 79 chars)
printf("You said: %s\n", line);
```

Always include a width limit to prevent buffer overflow, just like with %s.

Scan sets work with sscanf too. Here is an example that parses a structured string:

```
char buf[] = "Track 03: 99 Luftballons";
int track;
char title[50];
sscanf(buf, "Track %d: %49[^\n]", &track, title);
// track is 3, title is "99 Luftballons"
```



Tip: %^l is the scanf way to read a line, but fgets is generally safer and simpler for line-oriented input. Use scan sets when you need to parse structured input where only certain characters are valid.

The %m Modifier (POSIX)

With %s and %^l, you must always provide a buffer that is large enough. The POSIX %m modifier (called the **assignment-allocation** modifier) tells scanf to malloc the buffer for you. Instead of passing a char[], you pass a char ** and scanf allocates exactly enough memory:

```
char *line = NULL;
scanf("%m[^\n]", &line); // scanf mallocs the buffer
printf("You said: %s\n", line);
free(line); // you must free it
```

Notice &line — scanf needs a pointer to your char * so it can fill it in with the address of the newly allocated buffer. This eliminates buffer overflow risk entirely, since the buffer is always the right size.

%ms works the same way for single words:

```
char *word = NULL;
scanf("%ms", &word); // reads one word, malloc'd to fit
free(word);
```



Tip: %m is a POSIX extension (available on Linux, macOS, and most Unix systems) and is not part of the C standard. It will not work with MSVC on Windows. When portability is not a concern, %m is an excellent way to avoid buffer sizing headaches.

stdin, stdout, and stderr

When your C program starts, three streams (of type FILE *) are already open:

Stream	Purpose	C++ equivalent
stdin	Standard input (keyboard)	std::cin
stdout	Standard output (screen)	std::cout
stderr	Standard error (screen)	std::cerr

printf(...) is actually shorthand for fprintf(stdout, ...). You can write to stderr for error messages: fprintf(stderr, "Error: file not found\n");

Error messages sent to `stderr` are not affected by output redirection (`./program > output.txt` only redirects `stdout`), so error messages still appear on the screen. `./program 2> err.txt` will redirect errors to `err.txt` and `stdout` will appear on the screen.

fprintf and fscanf

```
int fprintf(FILE *stream, const char *format, ...);
int fscanf(FILE *stream, const char *format, ...);
```

`fprintf` and `fscanf` are the file versions of `printf` and `scanf`. They take a `FILE *` as the first argument:

```
fprintf(stdout, "Hello\n");           // same as printf("Hello\n")
fprintf(stderr, "Something broke\n"); // write to stderr
```

More usefully, you can use them with files you have opened yourself. Here is `fscanf` reading from a file:

```
FILE *f = fopen("scores.txt", "r");
if (f != NULL) {
    char name[50];
    int score;
    while (fscanf(f, "%49s %d", name, &score) == 2) {
        printf("%s scored %d\n", name, score);
    }
    fclose(f);
}
```

`fscanf` returns the number of items successfully read, so checking the return value tells you whether the read succeeded.

Opening and Closing Files

```
FILE *fopen(const char *path, const char *mode);
int fclose(FILE *stream);
```

To read or write a file, you open it with `fopen` and close it with `fclose`:

```
#include <stdio.h>

int main(void) {
    FILE *f = fopen("log.txt", "w");
    if (f == NULL) {
        fprintf(stderr, "Cannot open file\n");
        return 1;
    }

    fprintf(f, "Under Pressure\n");
    fprintf(f, "Year: %d\n", 1981);

    fclose(f);
    return 0;
}
```

The second argument to `fopen` is the **mode string**:

Mode	Meaning
"r"	Read (file must exist)
"w"	Write (creates or truncates)

Mode	Meaning
"a"	Append (creates or appends)
"r+"	Read and write (file must exist)
"w+"	Read and write (creates or truncates)
"a+"	Read and append

To open a file in **binary mode**, add `b` to the mode string: `"rb"`, `"wb"`, `"ab"`, etc. On Unix systems, binary and text modes behave identically. On Windows, text mode translates `\r\n` to `\n` on input and vice versa on output — binary mode does not.

sprintf and sscanf

`sprintf` writes formatted output into a string buffer instead of a stream. `sscanf` reads formatted input from a string:

```
char buf[100];
sprintf(buf, "Track %02d: %s", 3, "99 Luftballons");
// buf is now "Track 03: 99 Luftballons"

int track;
char title[50];
sscanf(buf, "Track %d: %49[^\n]", &track, title);
// track is 3, title is "99 Luftballons"
```

Trap: `sprintf` has the same buffer overflow risk as `strcpy` — it does not check the size of the destination buffer. Use `snprintf` for safety:

```
int snprintf(char *str, size_t size, const char *format, ...);
snprintf(buf, sizeof(buf), "Track %02d: %s", 3, "99 Luftballons");
snprintf guarantees it will not write more than sizeof(buf) bytes, including the null terminator.
```

`asprintf` (POSIX) goes one step further — it `mallocs` a buffer that is exactly the right size, so you never have to guess:

```
int asprintf(char **strp, const char *format, ...);

char *msg;
asprintf(&msg, "Track %02d: %s", 3, "Mis Ojos Lloran Por Ti");
printf("%s\n", msg); // "Track 03: Mis Ojos Lloran Por Ti"
free(msg);           // you must free it
```

Like `%m` in `scanf`, you pass a pointer to a `char *` and `asprintf` fills it in with the address of the newly allocated string. It returns the number of characters written, or `-1` on failure.



Tip: `asprintf` is a POSIX/GNU extension, not part of the C standard. It is available on Linux, macOS, and most Unix systems but not MSVC. When it is available, it is the safest and most convenient way to build formatted strings — no buffer sizing, no truncation, no overflow.

Binary I/O: fread and fwrite

```
size_t fread(void *ptr, size_t size, size_t count, FILE *stream);
size_t fwrite(const void *ptr, size_t size, size_t count, FILE *stream);
```

For reading and writing raw binary data (not text), use `fread` and `fwrite`:

```
#include <stdio.h>

int main(void) {
    int nums[] = {10, 20, 30, 40, 50};

    // Write binary data
    FILE *f = fopen("data.bin", "wb");
    fwrite(nums, sizeof(int), 5, f);
    fclose(f);

    // Read it back
    int result[5];
    f = fopen("data.bin", "rb");
    fread(result, sizeof(int), 5, f);
    fclose(f);

    for (int i = 0; i < 5; i++) {
        printf("%d ", result[i]); // 10 20 30 40 50
    }
    printf("\n");
    return 0;
}
```

Both functions take four arguments: a pointer to the data, the size of each element, the number of elements, and the file stream. `fwrite` returns the number of elements successfully written; `fread` returns the number of elements successfully read.

Reading Lines: `fgets`

```
char *fgets(char *s, int size, FILE *stream);
```

`fgets` reads a line from a stream into a buffer. It stops when it has read `size - 1` characters, encounters a newline (which it includes in the buffer), or reaches end of file. It always null-terminates the result:

```
char line[80];
while (fgets(line, sizeof(line), f) != NULL) {
    printf("%s", line); // line already includes '\n'
}
```

`fgets` returns `NULL` at end of file or on error, making it easy to use in a loop. It is generally safer than `scanf` for reading lines because it always respects the buffer size.

Buffering and `fflush`

`stdio` does not write directly to the output device on every call. Instead, it accumulates data in an internal buffer and writes it in larger chunks for efficiency. There are three buffering modes:

- **Full buffering:** Output is written when the buffer is full (default for files).
- **Line buffering:** Output is written when a `\n` is encountered (default for `stdout` when connected to a terminal).
- **Unbuffered:** Output is written immediately (default for `stderr`).

This means that `printf("Working...")` (no newline) will not appear on screen until the next newline when `stdout` goes to a terminal, and will not appear when redirected to a file until the buffer fills or the program exits. Use `fflush` to force the buffer to be written:

```
int fflush(FILE *stream);

printf("Working...");
fflush(stdout); // force output to appear now
// ... long computation ...
printf(" done!\n");
```



Trap: When `stdout` is connected to a terminal, output is **line buffered** — a `\n` triggers a flush. When `stdout` is redirected to a file or pipe, it is **fully buffered** — output may not appear until the buffer fills up or the program exits. If you need output to appear immediately (e.g., progress indicators), call `fflush(stdout)` after printing. `stderr` is always unbuffered, which is why error messages appear immediately.

Try It: Standard I/O Starter

```
#include <stdio.h>
#include <string.h>

int main(void) {
    // sprintf: format into a string
    char buf[100];
    sprintf(buf, "Track %02d: %s", 7, "Hungry Like the Wolf");
    printf("sprintf: %s\n", buf);

    // snprintf: safe version with size limit
    char small[15];
    snprintf(small, sizeof(small), "Year: %d", 1984);
    printf("snprintf: %s\n", small);

    // sscanf: parse from a string
    int track;
    char title[50];
    sscanf(buf, "Track %d: %49[^\n]", &track, title);
    printf("sscanf: track=%d title='%s'\n", track, title);

    // fprintf to stderr
    fprintf(stderr, "This goes to stderr\n");

    // fwrite/fread round-trip
    int nums[] = {10, 20, 30};
    FILE *f = fopen("/tmp/tryit_data.bin", "wb");
    fwrite(nums, sizeof(int), 3, f);
    fclose(f);

    int result[3];
    f = fopen("/tmp/tryit_data.bin", "rb");
    fread(result, sizeof(int), 3, f);
    fclose(f);

    printf("fread: %d %d %d\n", result[0], result[1], result[2]);

    return 0;
}
```

```
}
```

Key Points

- `printf` writes to `stdout`; `fprintf` writes to any `FILE *`.
- `scanf` needs the address (`&`) of each variable — arrays are the exception since they decay to pointers.
- `fopen` returns `NULL` on failure — always check before using the file pointer.
- Add `"b"` to the mode string for binary files. This matters on Windows.
- `fread` and `fwrite` transfer raw bytes — no format conversion.
- `stdout` is line buffered at a terminal and fully buffered when redirected. Use `fflush` when you need output immediately.

Exercises

1. **Think about it:** Why does `scanf` need the `&` operator for scalar variables but not for arrays?
2. **What does this print?**

```
char buf[50];
sprintf(buf, "%s: %d", "Score", 100);
printf("%zu\n", strlen(buf));
```

3. **Calculation:** If `buf` is declared as `char buf[20]` and you call `snprintf(buf, sizeof(buf), "Year: %d", 1984)`, how many characters (excluding the null terminator) are written to `buf`?
4. **Where is the bug?**

```
int x;
scanf("%d", x);
```

5. **Where is the bug?**

```
FILE *f = fopen("noexist.txt", "r");
fprintf(f, "Hello\n");
fclose(f);
```

6. **Think about it:** You run `./program > output.txt` and your program contains both `printf` and `fprintf(stderr, ...)` calls. Which messages appear in `output.txt` and which appear on the screen? Why?
7. **Write a program** that opens a text file, writes five lines to it (your choice of content), closes it, reopens it for reading, reads and prints each line using `fgets`, then closes it again.

11. Low-Level I/O

The `<stdio.h>` functions you learned in the previous chapter are built on top of a lower-level I/O interface provided by the operating system. These system calls — `read`, `write`, `open`, and `close` — work directly with **file descriptors** rather than `FILE *` pointers. You will encounter them in systems programming, and understanding them helps you see what `stdio` is actually doing under the hood.

File Descriptors

A file descriptor is a small non-negative integer that the operating system uses to identify an open file (or pipe, socket, device, etc.). When your program starts, three file descriptors are already open:

File Descriptor	POSIX Name	Purpose
-----------------	------------	---------

```
| 0 | STDIN_FILENO | Standard input | | 1 | STDOUT_FILENO | Standard output | | 2 | STDERR_FILENO | Standard error |
```

These constants are defined in `<unistd.h>`. They correspond to `stdin`, `stdout`, and `stderr` from `<stdio.h>`, but at a lower level.

read and write

```
ssize_t read(int fd, void *buf, size_t count);
ssize_t write(int fd, const void *buf, size_t count);
```

The `read` and `write` system calls transfer raw bytes between a file descriptor and a buffer:

```
#include <unistd.h>

// write(fd, buffer, count) - returns bytes written
write(1, "Blue Monday\n", 12); // write 12 bytes to stdout

// read(fd, buffer, count) - returns bytes read
char buf[100];
ssize_t n = read(0, buf, sizeof(buf)); // read from stdin
write(1, buf, n); // echo it back
```

`read` returns the number of bytes actually read (which may be less than requested), 0 at end of file, or -1 on error. `write` returns the number of bytes actually written, or -1 on error.

Unlike `printf` and `scanf`, these functions perform no formatting — they transfer raw bytes. There are no format specifiers, no newline handling, no buffering.

open and close

```
int open(const char *path, int flags, ... /* mode_t mode */);
int close(int fd);
```

To open a file at the system call level, use `open` from `<fcntl.h>`:

```
#include <fcntl.h>
#include <unistd.h>

int fd = open("data.txt", O_RDONLY);
if (fd == -1) {
    write(2, "Cannot open file\n", 17);
}
```

```

    return 1;
}

char buf[256];
ssize_t n = read(fd, buf, sizeof(buf));
write(1, buf, n);

close(fd);

```

The second argument to `open` is a set of flags combined with bitwise OR:

Flag	Purpose
<code>O_RDONLY</code>	Open for reading only
<code>O_WRONLY</code>	Open for writing only
<code>O_RDWR</code>	Open for reading and writing
<code>O_CREAT</code>	Create the file if it does not exist
<code>O_TRUNC</code>	Truncate the file to zero length
<code>O_APPEND</code>	Append to the file

To **create a new file** (or truncate an existing one), combine flags:

```
int fd = open("output.txt", O_WRONLY | O_CREAT | O_TRUNC, 0644);
```

The third argument (`0644`) is the **file permissions** — only used with `O_CREAT`. The value `0644` means the owner can read and write, and everyone else can only read.

There is also `creat`, which is equivalent to `open` with `O_WRONLY | O_CREAT | O_TRUNC`:

```
int creat(const char *path, mode_t mode);

int fd = creat("output.txt", 0644);
// same as: open("output.txt", O_WRONLY | O_CREAT | O_TRUNC, 0644)

```



Wut: Yes, it is `creat` with no `e`. Ken Thompson was once asked what he would do differently if he were redesigning Unix. His answer: "I'd spell `creat` with an `e`."

Seeking: `lseek`

```
off_t lseek(int fd, off_t offset, int whence);
```

`lseek` repositions the file offset for an open file descriptor:

```
#include <unistd.h>
```

```
lseek(fd, 0, SEEK_SET);    // go to beginning
lseek(fd, 0, SEEK_END);   // go to end
lseek(fd, 100, SEEK_SET); // go to byte 100

```

```
off_t pos = lseek(fd, 0, SEEK_CUR); // get current position (no move)
```

The three `SEEK_` constants control where the offset is relative to:

Constant	Meaning
SEEK_SET	Relative to the beginning of the file
SEEK_CUR	Relative to the current position
SEEK_END	Relative to the end of the file

`lseek` returns the new offset from the beginning of the file, or -1 on error.

Tip: In `<stdio.h>`, the equivalent functions are `fseek` and `ftell`. The low-level `lseek` combines both: calling `lseek(fd, 0, SEEK_CUR)` returns the current position without moving, just like `ftell`.

pread and pwrite

```
ssize_t pread(int fd, void *buf, size_t count, off_t offset);
ssize_t pwrite(int fd, const void *buf, size_t count, off_t offset);
```

`pread` and `ppwrite` are like `read` and `write` but take an explicit offset instead of using (or modifying) the file's current position:

```
// Read 100 bytes starting at offset 200, without changing the file position
ssize_t n = pread(fd, buf, 100, 200);
```

```
// Write 50 bytes at offset 0, without changing the file position
ppwrite(fd, data, 50, 0);
```

These are useful in multi-threaded programs where multiple threads share a file descriptor — since they do not modify the file position, there is no race condition.

Try It: Low-Level I/O Starter

```
#include <fcntl.h>
#include <unistd.h>
#include <string.h>
#include <stdio.h>

int main(void) {
    // write to stdout using file descriptor 1
    const char *msg = "Low-level I/O starter\n";
    write(STDOUT_FILENO, msg, strlen(msg));

    // open, write, close
    int fd = open("/tmp/tryit_lowio.txt",
                 O_WRONLY | O_CREAT | O_TRUNC, 0644);
    if (fd == -1) {
        write(STDERR_FILENO, "open failed\n", 12);
        return 1;
    }

    const char *lines[] = {"Come As You Are\n", "Year: 1992\n"};
    for (int i = 0; i < 2; i++)
        write(fd, lines[i], strlen(lines[i]));
    close(fd);

    // open, read, print
```

```

fd = open("/tmp/tryit_lowio.txt", O_RDONLY);
char buf[256];
ssize_t n = read(fd, buf, sizeof(buf));
write(STDOUT_FILENO, buf, n);

// lseek: go back to start and read again
lseek(fd, 0, SEEK_SET);
n = read(fd, buf, sizeof(buf));
printf("Read %zd bytes on second pass\n", n);

close(fd);
return 0;
}

```

Key Points

- File descriptors are small integers: 0 is stdin, 1 is stdout, 2 is stderr.
- read and write transfer raw bytes — no formatting, no buffering.
- open returns a file descriptor; fopen returns a FILE *. They are different levels of abstraction.
- Use O_CREAT with open to create files. Always provide a permissions argument when using O_CREAT.
- lseek repositions the read/write offset. Use SEEK_SET, SEEK_CUR, or SEEK_END.
- pread and pwrite read/write at a specific offset without changing the file position.

Exercises

1. **Think about it:** Why would you use low-level read/write instead of fprintf/fscanf? When would stdio be the better choice?
2. **What does this print?**

```

write(1, "ABC", 3);
write(1, "DEF\n", 4);

```

3. **Calculation:** If read(fd, buf, 1024) returns 512, what does that tell you? Does it mean there was an error?
4. **Where is the bug?**

```

int fd = open("newfile.txt", O_WRONLY | O_CREAT);
write(fd, "Hello\n", 6);
close(fd);

```

5. **Think about it:** Explain the difference between lseek(fd, 0, SEEK_END) and lseek(fd, -1, SEEK_END). What does each return?
6. **Write a program** that uses low-level I/O (open, read, write, close) to copy the contents of one file to another. The source and destination filenames should be taken from argv.

12. Odds and Ends

This chapter covers a few remaining topics that do not fit neatly into the previous chapters but are important for writing real C programs and for working with C code from C++.

exit vs return

```
void exit(int status);
```

You already know that `return` in `main` ends the program. The `exit` function from `<stdlib.h>` does the same thing, but it can be called from *any* function — not just `main`:

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

void check_file(const char *path) {
    FILE *f = fopen(path, "r");
    if (f == NULL) {
        fprintf(stderr, "Fatal: cannot open %s\n", path);
        exit(1); // end the program immediately
    }
    // ... work with the file ...
    fclose(f);
}
```

`exit` is useful when an error deep inside a call chain is unrecoverable and there is no reasonable way to propagate the error back through multiple layers of callers. In C++, you would throw an exception; in C, `exit` is sometimes the pragmatic choice.

`exit` also flushes all open `stdio` streams and calls any functions registered with `atexit` before terminating the program.

```
int atexit(void (*func)(void));
```



Tip: Use `exit` sparingly. It is a blunt instrument — it ends the entire program immediately, skipping any cleanup code in calling functions. If you can reasonably propagate an error code back to `main` and let `main` return, prefer that approach. Reserve `exit` for truly fatal errors.

extern "C" — Calling C from C++

If you are writing C++ code that needs to call functions from a C library, you need `extern "C"`. The reason: C++ **mangles** function names to support overloading (so `void foo(int)` and `void foo(double)` have different symbol names), but C does not. Without `extern "C"`, the C++ linker looks for the mangled name and cannot find the C function.

```
// In your C++ code:
extern "C" {
    #include "my_c_library.h" // treat these declarations as C
}
```

Or for a single function:

```
extern "C" void c_function(int x);
```

Many C headers protect themselves with this pattern so they work from both C and C++:

```

#ifdef __cplusplus
extern "C" {
#endif

void some_function(int x);
int another_function(const char *s);

#ifdef __cplusplus
}
#endif

```

The `__cplusplus` macro is only defined when compiling with a C++ compiler, so the `extern "C"` wrapper only appears in C++ compilation.



Tip: The reason C++ mangles names is to support **function overloading** — having multiple functions with the same name but different parameter types. C has no function overloading. Every function name must be unique. If you need two functions that do the same thing for different types, you give them different names — for example, `abs` for `int`, `fabs` for `double`, and `labs` for `long`. The upside is that when you see a function call in C, you know exactly which function will be invoked — there is only one version.

Pointer Ownership

In C++, smart pointers make ownership clear: a `std::unique_ptr` owns the memory, and when it goes out of scope, the memory is freed. In C, there are no smart pointers. When a function returns a pointer, you must ask: **who owns this memory?**

There are three common patterns:

1. The caller owns it (you must free). The function allocates memory and hands ownership to you:

```

char *copy = strdup("Everybody Wants to Rule the World");
// You own this memory. You must free it.
free(copy);

```

2. The library owns it (do not free). The function returns a pointer to memory it manages internally:

```

struct hostent *h = gethostbyname("example.com");
// The library owns this. Do NOT free it.

```

3. You own it (you passed it in). You allocated the memory and passed a pointer to the function. The function used it but did not take ownership:

```

char buf[100];
fgets(buf, sizeof(buf), stdin);
// You still own buf. Nothing to free (it's on the stack).

```

Trap: Always read the documentation of a C function that returns a pointer. Look for words like “the caller must free the returned pointer” or “the returned pointer points to a static buffer.” If the documentation does not say, look at the source code. Getting ownership wrong leads to either memory leaks (never freeing) or double-free bugs (freeing what you do not own).

Error Handling Without Exceptions

In C++, you can throw an exception and let a catch block handle it several call levels up. C has no exceptions. Error handling is done through return codes, and cleanup is your responsibility.

The simplest pattern is to check return values and bail out:

```
void perror(const char *s);  
  
FILE *f = fopen("La Isla Bonita.txt", "r");  
if (!f) {  
    perror("fopen");  
    return -1;  
}
```

But what happens when a function acquires multiple resources? You need to release them in the correct order when something goes wrong. The idiomatic C pattern uses goto to jump to cleanup labels:

```
int process(const char *path) {  
    int status = -1;  
  
    FILE *f = fopen(path, "r");  
    if (!f) return -1;  
  
    char *buf = malloc(1024);  
    if (!buf) goto close_file;  
  
    char *line = malloc(256);  
    if (!line) goto free_buf;  
  
    /* do work with f, buf, and line ... */  
    status = 0;  
  
    free(line);  
free_buf:  
    free(buf);  
close_file:  
    fclose(f);  
    return status;  
}
```

Each resource acquired gets a corresponding cleanup label below it. If any allocation fails, control jumps to the label that releases everything acquired so far, in reverse order. This pattern is used extensively in real C code including the Linux kernel.



Wut: C++ programmers are taught “never use goto.” In C, goto for cleanup is an accepted and widely used idiom. It is the closest thing C has to RAII — a structured way to ensure resources are always released.

The other common strategy is to return an error code (often -1 or NULL) and let the caller decide what to do. Many C library functions set the global variable `errno` (declared in `<errno.h>`) to indicate what went wrong. You can use `perror` (from `<stdio.h>`) or `strerror(errno)` (from `<string.h>`) to get a human-readable message:

```
char *strerror(int errnum);
```

```
FILE *f = fopen("No Existe.txt", "r");
if (!f) {
    perror("fopen");    // prints: fopen: No such file or directory
}
```

enum

C has `enum` just like C++, but there are a few differences. In C, `enum` constants are plain `int` values — there is no `enum` class and no scoping:

```
enum direction { NORTH, SOUTH, EAST, WEST };
```

Unlike C++, you cannot use the `enum` name as a type directly without the `enum` keyword:

```
enum direction heading = NORTH; /* need the 'enum' keyword */
```

You can avoid this with a `typedef`, just like with `struct` (see Chapter 2):

```
typedef enum { NORTH, SOUTH, EAST, WEST } Direction;
Direction heading = NORTH;
```

`Enum` values are just integers. You can assign specific values and mix `enums` with integers freely — the compiler will not complain:

```
enum status { OK = 0, ERR_FILE = -1, ERR_MEM = -2 };
int code = OK;    /* fine - enum to int */
enum status s = 42; /* also fine in C, no warning */
```



Wut: C++ `enum` class prevents implicit conversions and scopes the constants. C has neither protection. `Enum` constants are global and freely convert to `int`.

union

A `union` looks like a `struct`, but all members share the same memory. Only one member can hold a value at a time, and the size of the `union` is the size of its largest member:

```
union value {
    int i;
    float f;
    char s[20];
};
```

```
printf("sizeof(union value) = %zu\n", sizeof(union value));
/* prints 20 - the size of the largest member */
```

`Unions` are useful for saving memory when a variable can hold one of several types. A common C pattern is a **tagged union** — a `struct` with an `enum` tag that tracks which member is active:

```
enum val_type { VAL_INT, VAL_FLOAT, VAL_STR };
```

```
struct tagged_value {
    enum val_type type;
    union {
        int i;
        float f;
        char s[20];
    };
};
```

```

    } data;
};

void print_value(struct tagged_value v) {
    switch (v.type) {
        case VAL_INT:    printf("%d\n", v.data.i); break;
        case VAL_FLOAT:  printf("%f\n", v.data.f); break;
        case VAL_STR:    printf("%s\n", v.data.s); break;
    }
}

```



Tip: Tagged unions are C's equivalent of `std::variant`. Always check the tag before accessing a union member — reading the wrong member is undefined behavior.

qsort — Function Pointers in Action

You learned about function pointers in Chapter 7 (Functions). The most common place you will encounter them in practice is the standard library function `qsort` from `<stdlib.h>`. In C++, you would use `std::sort` with a lambda or comparator. In C, `qsort` takes a comparison function pointer:

```

void qsort(void *base, size_t nel, size_t width,
           int (*compar)(const void *, const void *));

#include <stdio.h>
#include <stdlib.h>
#include <string.h>

int compare_ints(const void *a, const void *b) {
    int ia = *(const int *)a;
    int ib = *(const int *)b;
    return (ia > ib) - (ia < ib);
}

int main(void) {
    int years[] = {1989, 1982, 1985, 1980, 1987};
    int n = sizeof(years) / sizeof(years[0]);

    qsort(years, n, sizeof(int), compare_ints);

    for (int i = 0; i < n; i++) {
        printf("%d ", years[i]);
    }
    printf("\n");
    // Output: 1980 1982 1985 1987 1989
    return 0;
}

```

`qsort` takes four arguments: the array, the number of elements, the size of each element, and a pointer to a comparison function. The comparison function receives `const void *` pointers — you must cast them to the correct type inside the function. It returns a negative value if the first argument is less than the second, zero if equal, and a positive value if greater.

You can sort anything with `qsort` by writing different comparison functions. Here is one that sorts strings:

```

int compare_strings(const void *a, const void *b) {
    const char *sa = *(const char **)a;
    const char *sb = *(const char **)b;
    return strcmp(sa, sb);
}

int main(void) {
    const char *songs[] = {
        "Maniac", "Footloose", "Flashdance", "Fame"
    };
    int n = sizeof(songs) / sizeof(songs[0]);

    qsort(songs, n, sizeof(char *), compare_strings);

    for (int i = 0; i < n; i++) {
        printf("%s\n", songs[i]);
    }
    // Output:
    // Fame
    // Flashdance
    // Footloose
    // Maniac
    return 0;
}

```

Notice the double cast in `compare_strings`: `qsort` passes a pointer to each array element, and each element is already a `char *`, so you receive a `char **` disguised as `const void *`.



Wut: A common mistake is to write `return a - b` in integer comparison functions. This can overflow when `a` and `b` have very different signs (e.g., `INT_MAX - (-1)` overflows). The pattern `(a > b) - (a < b)` is safe and returns `-1, 0, or 1`.

Try It: Odds and Ends Starter

```

#include <stdio.h>
#include <stdlib.h>
#include <string.h>

// A comparison function for qsort
// a normal numeric sort, but we want 1982 to always be first
// because we are big thriller fans!
int compare_ints(const void *a, const void *b) {
    int ia = *(const int *)a;
    int ib = *(const int *)b;
    // 1982 always appears first (Thriller came out in 1982)
    if (ia == 1982 && ib == 1982) return 0;
    if (ia == 1982) return -1;
    if (ib == 1982) return 1;
    // Overflow-safe alternative to (ia - ib). The three cases:
    // ia > ib → (1) - (0) = 1
    // ia == ib → (0) - (0) = 0
    // ia < ib → (0) - (1) = -1
}

```

```

    return (ia > ib) - (ia < ib);
}

void goodbye(void) {
    printf("atexit: Adios!\n");
}

int main(void) {
    atexit(goodbye);

    // Function pointer
    int (*cmp)(const void *, const void *) = compare_ints;
    int x = 10, y = 20;
    printf("compare(10, 20) = %d\n", cmp(&x, &y));

    // qsort
    int years[] = {1987, 1983, 1982, 1989, 1980, 1985};
    int n = sizeof(years) / sizeof(years[0]);
    qsort(years, n, sizeof(int), compare_ints);

    printf("Sorted: ");
    for (int i = 0; i < n; i++)
        printf("%d ", years[i]);
    printf("\n");

    // Pointer ownership: strdup allocates, you must free
    char *copy = strdup("Master of Puppets");
    printf("strdup: '%s'\n", copy);
    free(copy);

    // goto cleanup pattern
    printf("Demonstrating goto cleanup...\n");
    char *buf = malloc(100);
    if (!buf) return 1;

    char *msg = malloc(50);
    if (!msg) goto free_buf;

    snprintf(buf, 100, "Resource 1 OK");
    snprintf(msg, 50, "Resource 2 OK");
    printf(" %s, %s\n", buf, msg);

    free(msg);
free_buf:
    free(buf);

    return 0;
}

```

Key Points

- `exit` terminates the program from any function. Use it for unrecoverable errors.
- `exit` flushes `stdio` streams and calls `atexit` handlers before terminating.
- `extern "C"` tells the C++ compiler not to mangle function names, so it can link to C libraries.

- C headers often use `#ifdef __cplusplus` to wrap declarations in extern "C" automatically.
- When you receive a pointer from a function, always determine who owns the memory: you, the function, or a library.
- C has no exceptions. Use return codes for errors and goto cleanup for releasing resources in the correct order.
- `qsort` is the most common use of function pointers — it takes a comparison callback to sort any type.
- C enum constants are plain `int` values with no scoping — there is no `enum` class.
- A union stores all members in the same memory. Use a tagged union (struct with an enum tag) to track which member is active.

Exercises

1. **Think about it:** In C++ you would use exceptions for error handling. In C there are no exceptions. What strategies can you use to handle errors in deeply nested function calls? When is `exit` appropriate and when is it not?

2. **What happens here?**

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

void cleanup(void) {
    printf("Adios!\n");
}

int main(void) {
    atexit(cleanup);
    printf("Starting...\n");
    exit(0);
}
```

3. **Where is the bug?**

```
char *get_greeting(void) {
    char buf[50];
    sprintf(buf, "Hola, mundo");
    return buf;
}
```

4. **Think about it:** You call a function `char *get_name(int id)` from a library. How would you determine whether you need to free the returned pointer?

5. **Where is the bug?** (Hint: ownership)

```
char *name = strdup("Walking on Sunshine");
char *alias = name;
free(name);
printf("%s\n", alias);
```

6. **Calculation:** Given `int nums[] = {5, 10, 15, 20};`, what is the value of `sizeof(nums) / sizeof(nums[0])`?

7. **What does this print?**

```
int compare_desc(const void *a, const void *b) {
    int ia = *(const int *)a;
    int ib = *(const int *)b;
    return (ib > ia) - (ib < ia);
}
```

```

int main(void) {
    int vals[] = {3, 1, 4, 1, 5};
    qsort(vals, 5, sizeof(int), compare_desc);
    printf("%d %d %d %d %d\n", vals[0], vals[1], vals[2], vals[3], vals[4]);
    return 0;
}

```

8. **Write a program** that uses `qsort` to sort an array of strings in reverse alphabetical order. Write a custom comparison function that calls `strcmp` with the arguments swapped.
9. **Write a program** in C++ that uses `extern "C"` to call the C function `strlen` from `<string.h>`, passes it a string, and prints the result. Compile it with `c++` to verify it works.
10. **Calculation:** What is `sizeof(union { int i; double d; char s[3]; })` on a system where `int` is 4 bytes and `double` is 8 bytes?
11. **Where is the bug?**

```

union value {
    int i;
    float f;
};

union value v;
v.i = 42;
printf("%.2f\n", v.f);

```

12. **Write a program** that defines a tagged union representing a shape (circle with a radius, or rectangle with width and height). Write a function that prints the area of the shape using the tag to determine which union member to read.

Conclusion

You have covered a lot of ground — from `printf` format specifiers to file descriptors to pointer ownership. Here are the key takeaways:

- **C and C++ are different languages.** Modern C++ has evolved far from C. Knowing one does not mean you automatically know the other.
- **`printf` and `scanf` replace `iostream`.** Format specifiers must match argument types. `scanf` needs `&` for scalar variables.
- **C types are explicit.** No `auto`, no `std::string`, no classes. You have basic types, `typedef`, arrays, and `struct`.
- **C shares most operators with C++** but there is no operator overloading, `<<` and `>>` are strictly bitwise, and boolean results are plain `int`.
- **Control flow is nearly identical to C++** except there are no range-based `for` loops and `goto` is an accepted idiom for cleanup.
- **Pointers hold memory addresses.** Use `&` to get an address, `*` to follow one, and `->` to access `struct` fields through a pointer. Arrays decay to pointers, and pointer arithmetic moves in units of the pointed-to type.
- **All function arguments are pass by value.** To modify a caller's variable, pass a pointer to it. Use `const` parameters to document read-only intent.
- **Know where your memory lives.** Global variables last the whole program, local variables live on the stack, and dynamic memory from `malloc` lives on the heap until you `free` it.
- **Strings in C are char arrays** terminated by `'\0'`. You must manage buffer sizes manually and use functions like `strlen`, `strcpy`, `strcmp`, and `strcat` instead of `std::string` methods.
- **`stdio` provides buffered I/O** through `FILE *` pointers. Low-level I/O uses file descriptors and system calls like `read`, `write`, and `open`.
- **C has no exceptions.** Use return codes for errors and `goto` cleanup to release resources in reverse order.
- **Function pointers replace lambdas.** `qsort` is the classic example — pass a comparison function to sort any type.
- **`exit` terminates from anywhere.** Use it for fatal errors. `extern "C"` bridges C and C++. Always know who owns a pointer.

Es un mundo nuevo, but you have the C++ foundation to build on. The syntax will feel familiar even when the idioms are different. Write small programs, compile them with `cc`, and get comfortable with the compiler's warnings — they are your best amigo.

Buena suerte — you have got this.

Content outline and editorial support from Ben. Words by Claude, the Opus.

Appendix A: Macros

In C++, you have `constexpr` for compile-time constants, templates for generic code, and `inline` functions to avoid call overhead. C has none of those. Instead, C leans heavily on the **preprocessor** — the `#define` macro system that rewrites your source code *before* the compiler sees it.

Macros are pure textual substitution. The preprocessor does not know about types, scope, or expressions — it just replaces text. This makes macros powerful and flexible, but also a source of subtle bugs if you are not careful.

Object-Like Macros

The simplest macros define named constants:

```
#define MAX_BUF    1024
#define PI         3.14159265
#define GREETING  "Hola, amigo"
```

Everywhere the preprocessor sees `MAX_BUF`, it replaces it with `1024`. No semicolons — a common mistake is writing `#define MAX_BUF 1024;`, which would paste `1024;` everywhere, breaking expressions like `malloc(MAX_BUF * sizeof(int))`.



Trap: Do not put a semicolon at the end of a `#define`. The semicolon becomes part of the replacement text and will cause surprising errors.

Conditional Compilation

Macros also control which code the compiler sees:

```
#define DEBUG

#ifdef DEBUG
    printf("x = %d\n", x);
#endif
```

`#ifdef` checks whether a macro is defined (regardless of its value). Its complement `#ifndef` checks that a macro is *not* defined. You can also use `#if`, `#elif`, and `#else` for more complex conditions:

```
#if VERBOSE_LEVEL >= 2
    printf("Detailed trace...\n");
#elif VERBOSE_LEVEL == 1
    printf("Basic trace...\n");
#else
    /* no tracing */
#endif
```

Include Guards

The most common use of `#ifndef` is protecting header files from being included more than once:

```
/* myheader.h */
#ifndef MYHEADER_H
#define MYHEADER_H

struct point {
```

```

    int x, y;
};

void draw_point(struct point p);

#endif /* MYHEADER_H */

```

The first time `myheader.h` is included, `MYHEADER_H` is not defined, so the contents are processed and `MYHEADER_H` gets defined. Any subsequent include finds `MYHEADER_H` already defined and skips the entire file.

Tip: Many compilers support `#pragma once` as a non-standard alternative to include guards. It is simpler to write but not portable to all compilers. When in doubt, use the `#ifndef` guard — it works everywhere.

Function-Like Macros

Macros can take parameters, making them look like functions:

```

#define SQUARE(x) ((x) * (x))
#define MAX(a, b) ((a) > (b) ? (a) : (b))
#define ABS(x) ((x) < 0 ? -(x) : (x))

```

But they are *not* functions — they are text substitution with parameter placeholders. This distinction matters.

The Parenthesization Rules

Always parenthesize every parameter use *and* the entire macro body:

```

/* Wrong: */
#define SQUARE(x) x * x

/* SQUARE(1 + 2) expands to: 1 + 2 * 1 + 2 = 5 (not 9!) */

/* Right: */
#define SQUARE(x) ((x) * (x))

/* SQUARE(1 + 2) expands to: ((1 + 2) * (1 + 2)) = 9 */

```

Without parentheses, operator precedence in the surrounding expression can silently rearrange the computation.

The Double-Evaluation Trap

Since macros substitute text, each parameter reference evaluates the argument again:

```

#define SQUARE(x) ((x) * (x))

int i = 3;
int result = SQUARE(i++);
/* Expands to: ((i++) * (i++)) - i is incremented TWICE */
/* Undefined behavior: two unsequenced modifications of i */

```

A real function evaluates its argument once. A macro evaluates it once per appearance in the replacement text. This is the most important difference between macros and functions.



Trap: Never pass expressions with side effects (like `i++`, `f()`, or assignment) to function-like macros. The expression will be evaluated multiple times, producing unexpected results or undefined behavior.

Multi-Statement Macros: `do { ... } while (0)`

If a macro needs to execute multiple statements, wrap them in `do { ... } while (0)`:

```
#define SWAP(a, b) do { \  
    int tmp = (a);    \  
    (a) = (b);        \  
    (b) = tmp;        \  
} while (0)
```

Why not just use braces? Consider:

```
if (x > y)  
    SWAP(x, y);  
else  
    printf("Already sorted\n");
```

If `SWAP` expanded to a bare `{ ... }`, the semicolon after `SWAP(x, y)` would terminate the `if` statement, and the `else` would become a syntax error. The `do { ... } while (0)` idiom creates a single statement that works correctly with semicolons and control flow.



Tip: The `do { ... } while (0)` pattern is everywhere in C codebases. It looks odd at first, but it is the standard way to make multi-statement macros behave like ordinary statements.

Stringification and Token Pasting

The preprocessor has two special operators for macro arguments.

Stringification: `#`

The `#` operator turns a macro argument into a string literal:

```
#define PRINT_VAR(x) printf("#x " = %d\n", x)  
  
int score = 42;  
PRINT_VAR(score);  
/* Expands to: printf("score" " = %d\n", score); */  
/* Adjacent string literals are concatenated: "score = %d\n" */  
/* Output: score = 42 */
```

This is useful for debug macros where you want to print both the variable name and its value.

Token Pasting: `##`

The `##` operator joins two tokens into one:

```
#define DECLARE_PAIR(type) \  
    type type##_first;    \  

```

```

    type type##_second;

DECLARE_PAIR(int)
/* Expands to:
   int int_first;
   int int_second;
*/

```

Token pasting is commonly used to generate families of related variables or functions from a single macro.

Variadic Macros

Macros can accept a variable number of arguments using `...` and `__VA_ARGS__`:

```

#define LOG(fmt, ...) fprintf(stderr, "[LOG] " fmt "\n", __VA_ARGS__)

LOG("score is %d", 42);
/* Expands to: fprintf(stderr, "[LOG] " "score is %d" "\n", 42); */

```

This is commonly used to wrap `printf`-style functions with extra decoration like timestamps or log levels.



Tip: When `__VA_ARGS__` is empty, the trailing comma before it can cause a compilation error. GNU C provides `##_VA_ARGS__` which swallows the comma when the argument list is empty:

```

#define LOG(fmt, ...) fprintf(stderr, "[LOG] " fmt "\n", ##__VA_ARGS__)
LOG("started"); /* No extra args - comma is removed */

```

This is a GCC/Clang extension. C23 standardizes this behavior with `__VA_OPT__`.

Multi-Level Expansion

Macros can expand to other macros, and the preprocessor **rescans** the result to expand again. But the `#` and `##` operators are special — they operate on the raw argument text *before* any expansion happens.

```

#define MAX_BUF 1024
#define STRINGIFY(x) #x
#define XSTRINGIFY(x) STRINGIFY(x)

printf("%s\n", STRINGIFY(MAX_BUF));
/* # operates before expansion: prints "MAX_BUF" */

printf("%s\n", XSTRINGIFY(MAX_BUF));
/* First pass: XSTRINGIFY(MAX_BUF) → STRINGIFY(1024) */
/* Rescan:    STRINGIFY(1024) → "1024" */
/* Prints "1024" */

```

`STRINGIFY(MAX_BUF)` gives `"MAX_BUF"` because `#` stringifies its argument before expansion. `XSTRINGIFY(MAX_BUF)` first expands `MAX_BUF` to `1024` (since the outer macro does not use `#` directly), then passes `1024` to `STRINGIFY`, producing `"1024"`.

This two-level indirect pattern is used whenever you need the *expanded* value of a macro as a string.



Tip: Whenever you need a macro's expanded value as a string, use the two-level indirect pattern. It comes up often when embedding version numbers or configuration values in strings.

X-Macros

X-macros are a technique for defining a list of items once and expanding it in multiple ways. The idea: define the list as a macro that calls an unspecified “action” macro on each item, then define that action differently for each use.

Here is a concrete example that generates both an enum and a string table from a single list of log levels:

```
#include <stdio.h>

/* Define the list once */
#define LOG_LEVELS(X) \
    X(LOG_DEBUG)      \
    X(LOG_INFO)       \
    X(LOG_WARN)       \
    X(LOG_ERROR)      \
    X(LOG_FATAL)

/* Generate the enum */
#define AS_ENUM(name) name,
enum log_level { LOG_LEVELS(AS_ENUM) LOG_COUNT };

/* Generate the string table */
#define AS_STRING(name) #name,
const char *log_level_names[] = { LOG_LEVELS(AS_STRING) };

int main(void) {
    for (int i = 0; i < LOG_COUNT; i++) {
        printf("%d = %s\n", i, log_level_names[i]);
    }
    return 0;
}
```

Output:

```
0 = LOG_DEBUG
1 = LOG_INFO
2 = LOG_WARN
3 = LOG_ERROR
4 = LOG_FATAL
```

Add a new log level? Add one line to LOG_LEVELS and the enum and string table stay in sync automatically. Without X-macros, you would need to update both the enum and the string array separately — and hope you never forget one.



Tip: X-macros are one of the preprocessor's most powerful patterns. You will see them in real codebases for error codes, command tables, and state machines. The key advantage: a single source of truth for a list of items.

Try It: Macro Starter

```
#include <stdio.h>

// Object-like macros
#define MAX_TRACKS 10
#define LABEL      "Sire Records"

// Function-like macro with proper parenthesization
#define SQUARE(x)  ((x) * (x))
#define MAX(a, b)  ((a) > (b) ? (a) : (b))

// Stringification: print variable name and value
#define PRINT_INT(var) printf(#var " = %d\n", var)

// Multi-statement macro using do { ... } while (0)
#define SWAP(a, b) do { \
    int tmp = (a);      \
    (a) = (b);          \
    (b) = tmp;          \
} while (0)

// Variadic macro
#define LOG(fmt, ...) fprintf(stderr, "[LOG] " fmt "\n", ##__VA_ARGS__)

int main(void) {
    // Object-like
    printf("Label: %s, Max tracks: %d\n", LABEL, MAX_TRACKS);

    // Function-like
    printf("SQUARE(5) = %d\n", SQUARE(5));
    printf("SQUARE(1+2) = %d\n", SQUARE(1 + 2));
    printf("MAX(3, 7) = %d\n", MAX(3, 7));

    // Stringification
    int year = 1984;
    PRINT_INT(year);

    // SWAP
    int a = 10, b = 20;
    printf("Before swap: a=%d, b=%d\n", a, b);
    SWAP(a, b);
    printf("After swap: a=%d, b=%d\n", a, b);

    // Conditional compilation
#ifdef DEBUG
    printf("Debug mode is on\n");
#else
    printf("Debug mode is off\n");
#endif

    // Variadic macro
    LOG("started");
    LOG("year is %d", 1985);
}
```

```

return 0;
}

```

Key Points

- Macros are **textual substitution** performed before compilation. They are not functions and do not respect scope or type rules.
- Object-like macros define constants and feature flags. Never end a `#define` with a semicolon.
- Function-like macros must have every parameter use and the entire body parenthesized to avoid precedence bugs.
- Macro arguments are evaluated each time they appear — do not pass expressions with side effects.
- Use `do { ... } while (0)` for multi-statement macros so they work correctly with `if/else` and semicolons.
- The `#` operator stringifies a macro argument; `##` pastes tokens together.
- `#` and `##` prevent argument expansion. Use the two-level indirect pattern (e.g., `XSTRINGIFY/STRINGIFY`) when you need the expanded value.
- `__VA_ARGS__` enables variadic macros for wrapping `printf`-style functions.
- X-macros define a list once and expand it multiple ways, keeping enums and string tables in sync.

Exercises

1. **Think about it:** C++ uses `constexpr` and `inline` functions to replace many uses of macros. What specific problems do macros have that these C++ features solve? Why does C still rely on macros despite these problems?

2. **What does this produce?**

```

#define DOUBLE(x) ((x) + (x))

int i = 5;
printf("%d\n", DOUBLE(i++));

```

3. **Calculation:** Given the macro `#define BUFSIZE 256`, how many bytes does `char buf[BUFSIZE + 1]` allocate? Why is the `+ 1` a common pattern?

4. **Where is the bug?**

```

#define MUL(a, b) a * b

int result = MUL(2 + 3, 4 + 5);
printf("%d\n", result);

```

5. **What does this produce?**

```

#define STRINGIFY(x) #x
#define XSTRINGIFY(x) STRINGIFY(x)
#define VERSION 3

printf("[%s] [%s]\n", STRINGIFY(VERSION), XSTRINGIFY(VERSION));

```

6. **Where is the bug?**

```

#define LOG_IF(cond, msg) \
    if (cond) \
        printf("[WARN] %s\n", msg);

if (x > 100)

```

```
LOG_IF(x > 200, "very high");  
else  
printf("normal\n");
```

7. **Write a program** that defines an X-macro list of at least four colors, then uses it to generate both an enum and a function that returns the string name for a given enum value. Print each color's enum value and name.

References

- [1] B. W. Kernighan and D. M. Ritchie, *The C programming language*, 2nd ed. Englewood Cliffs, NJ: Prentice Hall, 1988.

Index

- #define, 94
- #ifdef, 94
- #ifndef, 94
- #pragma once, 95
- _Bool, 13
- __VA_ARGS__, 97

- address-of operator (&), 46
- argc, 47
- argv, 47
- arithmetic operators, 30
- array, 14
 - decay to pointer, 14
 - designated initializers, 18
 - multidimensional, 15
 - row-major order, 15
- arrow operator (->), 49
- asprintf, 76
- assignment
 - vs equality, 30
- assignment operator, 30
- atexit, 84

- bitwise operators, 32
 - flags, 32
- bool, 32
- break, 39
- buffer overflow, 25
- buffering, 77

- C89, 38
- C99, 38
- callback, 57
- calloc, 62
- cast, 69
- cc (compiler), 9
- char, 65
- close, 80
- comparison operators, 31
- compound assignment, 33
- conditional compilation, 94
- const, 16
 - documentation, 53
 - parameter, 53
 - with pointers, 16
- continue, 39
- creat, 81
- ctype.h, 26

- dangling pointer, 85
- data segment, 61

- decrement operator, 34
- dereference operator (*), 46
- designated initializers, 18
- do-while loop, 38
- double, 68
- double free, 85

- enum, 87
- EOF, 30
- errno, 86
- error handling
 - return codes, 86
- exit, 84
- extern "C", 84

- fclose, 75
- fcntl.h, 80
- fflush, 77
- fgets, 77
- FILE, 73
- file descriptor, 80
- file permissions, 81
- float, 68
- floating-point
 - infinity, 68
- fopen, 75
 - binary mode, 76
- for loop, 40
 - C89 vs C99, 40
- format specifier, 9
 - width and precision, 9
 - zero-fill, 9
- fprintf, 75
- fread, 76
- free, 61
- fscanf, 75
- fseek, 82
- ftell, 82
- function
 - declaration, 52
 - default arguments, 52
 - definition, 52
 - overloading, 52
- function pointer, 55
- fwrite, 76

- getchar, 30
- global variable, 60
- goto, 42, 86

- header file, 52

- heap, 61
- if statement, 38
- include guard, 94
- increment operator, 34
- integer division, 30
- integer promotion, 67
- integer types, 66
- isalpha, 26
- isdigit, 26
- limits.h, 67
- local variable, 60
- logical operators, 31
- long double, 68
- lseek, 81
- macro
 - do-while(0) idiom, 96
 - function-like, 95
 - multi-level expansion, 97
 - object-like, 94
- malloc, 61
- memcpy, 62
- memmove, 63
- memory leak, 62
- memset, 62
- name mangling, 84
- NaN, 68
- NULL, 47
- null terminator, 22
- open, 80
 - flags, 81
- operator precedence, 34
 - traps, 35
- pass by value, 49
- perror, 86
- pointer, 46
 - arithmetic, 48
 - declaration, 14
 - ownership, 85
 - to pointer, 47
- pread, 82
- preprocessor, 94
- printf, 9
- prototype, 52
- ptrdiff_t, 24
- putchar, 30
- pwrite, 82
- qsort, 88
- read, 80
- realloc, 62
- recursion, 54
- remainder operator (%), 31
- scan set, 73
- scanf, 10, 73
 - %m modifier, 74
- SEEK_CUR, 81
- SEEK_END, 81
- SEEK_SET, 81
- segmentation fault, 48
- shallow copy, 17
- short-circuit evaluation, 31
- signed, 13
- size_t, 68
- sizeof, 14, 67
- snprintf, 27, 76
- sprintf, 27, 76
- sscanf, 27, 76
- ssize_t, 80
- stack, 60
- stack overflow, 55
- static, 61
- stdbool.h, 13, 32, 38
- stderr, 74
- STDERR_FILENO, 80
- stdin, 74
- STDIN_FILENO, 80
- stdio.h, 8
- stdlib.h, 61
- stdout, 74
- STDOUT_FILENO, 80
- strcat, 24
- strchr, 24
- strcmp, 23
- strcpy, 23
- strdup, 25
- strerror, 86
- string, 22
- string literal, 22
- string.h, 62
- stringification, 96
- strlen, 22
- strncat, 24
- strncmp, 24
- strncpy, 23
- strrchr, 24
- strstr, 24
- strtod, 66
- strtok, 25
- strtok_r, 26
- strtol, 66
- struct, 16
 - assignment, 17

- designated initializers, 18
- member access, 17
- passing, 53
- switch statement, 41
 - default, 41
 - fall-through, 41
- ternary operator, 34
- token pasting, 96
- tolower, 26
- toupper, 26
- two's complement, 67
- typedef, 18
 - function pointer, 56
- types, 13
- undefined behavior, 69
- union, 87
- unistd.h, 80
- unsigned, 13
- use-after-free, 85
- variable, 13
- variadic macro, 97
- void parameter list, 52
- void pointer, 69
- while loop, 38
- write, 80
- X-macro, 98
- Yoda condition, 30