



Gorgo Starting C++

April 11, 2026

Contents

7. Numbers	2
Bases	2
Literals in Other Bases	3
Printing in Other Bases	4
Strings and Numbers	5
Two's Complement	8
Integer Sizes and Ranges	9
Binary Addition and Subtraction	12
Bit Operators	13
Shift Operators	14
Key Points	16
Exercises	16

7. Numbers

There are only 10 kinds of people in the world: those who understand binary and those who don't.

If that joke doesn't make sense yet, it will by the end of this chapter.

At its heart, a CPU is a super-fast glorified calculator. Everything it does — drawing pixels on your screen, playing your favorite track, sending a message — ultimately boils down to operations on numbers. When we think about the number five, we might write 5, or V, or ||||, but they all represent the same thing. You already saw in Chapter 2 that the character 'A' is actually stored as the ASCII code 65. There are many ways to represent a number, and in this chapter you will learn why that matters and how C++ lets you work with them.

Bases

Let's extend that joke. *There are only 10 kinds of people in the world: those who understand binary, those who don't, and those who weren't expecting this joke to be in base 3. And those who weren't expecting base 4. And those who...*

You get the idea. The punchline changes depending on which **base** (or **radix**) you are using.

Decimal (Base 10)

You have been using decimal your whole life. It uses ten digits: 0 through 9. Each position represents a power of 10:

```
  4   7   2
  |   |   |
  |   |   |  +-- 2 * 10^0 =  2
  |   +----- 7 * 10^1 = 70
  +----- 4 * 10^2 = 400
                               ---
                               472
```

A number is divisible by 10 when its last digit is 0. You can tell if a number is divisible by 2 by checking whether its last digit is even. Divisibility by 8 is trickier — you have to check the last three digits. These rules come directly from how place values work.

Binary (Base 2)

It is often said that computers think in 1s and 0s. That is not entirely accurate, but at the lowest level, data is stored in **bits** — each one either 0 or 1. Binary is base 2, so each position represents a power of 2:

```
  1   0   1   0   1   0
  |   |   |   |   |   |
  |   |   |   |   |   +-- 0 * 2^0 =  0
  |   |   |   |   +----- 1 * 2^1 =  2
  |   |   |   +----- 0 * 2^2 =  0
  |   |   +----- 1 * 2^3 =  8
  |   +----- 0 * 2^4 =  0
  +----- 1 * 2^5 = 32
                               --
                               42
```

Counting in binary looks like this:

```
Decimal:  0   1   2   3   4   5   6   7   8
Binary:   0   1  10  11 100 101 110 111 1000
```

In binary, you can tell if a number is divisible by 2 by checking the last bit — if it is 0, the number is even. Divisible by 8? Check the last three bits. But divisibility by 10 is no longer obvious at a glance.



Tip: Notice the symmetry. In decimal, divisibility by the base (10) is trivial to check. In binary, divisibility by the base (2) is trivial. Each system makes certain things easy and others hard.

Hexadecimal (Base 16)

Binary numbers get long quickly. The number 255 is 11111111 in binary — eight digits for what decimal handles in three. **Hexadecimal** (hex) uses sixteen digits: 0–9 and A–F (where A = 10, B = 11, ... F = 15). Each hex digit represents exactly four bits:

```
Binary: 1010 1100
Hex:    A    C  -> 0xAC = 172
```

This makes hex a compact way to write binary values. Two hex digits represent one byte (8 bits), and eight hex digits represent a 32-bit integer. You will see hex used frequently for colors, memory addresses, and bit masks.

Octal (Base 8)

Octal uses eight digits: 0–7. Each octal digit represents exactly three bits:

```
Binary: 101 010
Octal:  5  2  -> 052 = 42
```

Octal is less common than hex in modern code, but you will encounter it when working with Unix file permissions (like 0755).

Try It: Counting in Binary

This program counts from 0 to 15 and prints each number in decimal, binary, hex, and octal so you can see the patterns side by side:

```
#include <print>

int main() {
    std::println("{:>4} {:>8} {:>4} {:>4}", "Dec", "Binary", "Hex", "Oct");
    std::println("{:->4} {:->8} {:->4} {:->4}", "", "", "", "");
    for (int i = 0; i <= 15; ++i) {
        std::println("{:4d} {:08b} {:4x} {:4o}", i, i, i, i);
    }
}
```

Literals in Other Bases

C++ lets you write integer literals in binary, hexadecimal, and octal using prefixes. You first used integer literals in Chapter 2 when you initialized variables like `int x = 42;` — that 42 is a decimal literal.

```
int dec = 42;           // decimal (no prefix)
int bin = 0b101010;    // binary (0b prefix)
int hex = 0x2A;        // hex (0x prefix)
int oct = 052;         // octal (0 prefix)
```

All four variables hold exactly the same value: 42. The prefix only affects how you write the number in your source code, not how it is stored.



Tip: Be careful with leading zeros. `052` is **not** decimal 52 — it is octal 52, which equals decimal 42. This is a common source of confusion. If you want decimal 52, write 52 without a leading zero.

Digit Separators

Large numbers can be hard to read. C++14 introduced the single-quote `'` as a digit separator, which you can place anywhere between digits for readability:

```
int billion    = 1'000'000'000;    // easier to read than 1000000000
int bits       = 0b1010'1100;      // group binary by nibbles (4 bits)
int color      = 0xFF'80'00;       // group hex by byte
```

The separator has no effect on the value — the compiler ignores it completely.

Try It: Same Value, Different Spellings

This program shows that no matter how you write a literal, the compiler stores the same number:

```
#include <print>

int main() {
    int dec = 1984;
    int bin = 0b11111000000;
    int hex = 0x7C0;
    int oct = 03700;

    std::println("The year George Orwell warned us about:");
    std::println("  Decimal:  {}", dec);
    std::println("  Binary:   {}", bin);
    std::println("  Hex:      {}", hex);
    std::println("  Octal:   {}", oct);
    std::println("  All equal? {} ", dec == bin && bin == hex && hex == oct);
}
```

Printing in Other Bases

C++23 gives you `std::format` and `std::println` with format specifiers that make it easy to print numbers in different bases:

```
int val = 42;
std::println("Decimal:   {}", val);    // 42
std::println("Binary:    {:b}", val);  // 101010
std::println("Hex:       {:x}", val);  // 2a
std::println("Hex (upper): {:X}", val); // 2A
std::println("Octal:     {:o}", val);  // 52
```

You can also add a `#` flag to include the base prefix in the output:

```
std::println("Binary:   {:#b}", val);  // 0b101010
std::println("Hex:      {:#x}", val);  // 0x2a
std::println("Octal:    {:#o}", val);  // 052
```



Tip: The format specifiers `{:b}`, `{:x}`, and `{:o}` display the value in a different base, but they do not change the value itself. The variable still holds the same number — you are just looking at it from a different angle.

Try It: Number Viewer

This program asks for a number and displays it in every base:

```
#include <iostream>
#include <print>

int main() {
    std::print("Enter a number: ");
    int val{};
    std::cin >> val;

    std::println("Decimal:  {}", val);
    std::println("Binary:   {:#b}", val);
    std::println("Hex:     {:#x}", val);
    std::println("Octal:   {:#o}", val);
}
```

Strings and Numbers

Programs frequently need to convert between strings and numbers — reading user input, parsing files, or displaying results. C++ provides several functions for this, each with different strengths. When you see "52" and 52 they may appear to be the same number, but to C++, the former is a `std::string` and not a number at all while the latter is an integer number. If we want to use "52" as a number, we need to convert the string to an integer.

Strings to Integers

The `std::stoi` function (string-to-integer) converts a `std::string` to an `int`. Its signature is:

```
int stoi(const std::string& str, std::size_t* pos = nullptr, int base = 10);
```

Used in its simplest form:

```
int a = std::stoi("42");           // 42
int b = std::stoi(" -7");          // -7 (leading whitespace is skipped)
int c = std::stoi("1984abc");      // 1984 (stops at first non-digit)
```

For larger values, use `std::stol` for long or `std::stoll` for long long:

```
// too large for int on most systems
long big = std::stol("3000000000");
long long huge = std::stoll("9000000000000"); // needs long long
```



Tip: `std::stoi` and its siblings throw `std::invalid_argument` if no conversion can be performed and `std::out_of_range` if the value won't fit. You will learn how to catch these exceptions in Chapter 11. Always be prepared to handle them when converting user input.

The pos Parameter

The second parameter, `pos`, receives the index of the first character that was **not** part of the number. This is useful when a string contains a number followed by other data:

```
std::size_t pos;
int val = std::stoi("42px", &pos);
// val == 42, pos == 2 (index of 'p')
```

You can use `pos` to parse multiple numbers from a single string or to check whether the entire string was consumed:

```
std::string input = "100 200 300";
std::size_t pos = 0;
int first = std::stoi(input, &pos);           // first == 100, pos == 3
int second = std::stoi(input.substr(pos), &pos); // second == 200
```

If you don't need `pos`, pass `nullptr` (or just omit it):

```
int x = std::stoi("42", nullptr); // same as std::stoi("42")
```

Converting Bases with `std::stoi`

The third parameter of `std::stoi` specifies the base to use when parsing. You can use any base from 2 to 36:

```
int a = std::stoi("101010", nullptr, 2); // binary -> 42
int b = std::stoi("2A", nullptr, 16);    // hex -> 42
int c = std::stoi("52", nullptr, 8);     // octal -> 42
```

The same base parameter works with `std::stol` and `std::stoll`:

```
long d = std::stol("7C1", nullptr, 16); // hex -> 1985
```



Tip: Pass base 0 and `std::stoi` will auto-detect the base from the prefix — `0x` for hex, `0b` for binary, and a leading `0` for octal:

```
std::stoi("0x2A", nullptr, 0); // hex -> 42
std::stoi("0b101010", nullptr, 0); // binary -> 42
std::stoi("052", nullptr, 0); // octal -> 42
```

This is convenient, but beware: with base 0, the string `"010"` is parsed as **octal 8**, not decimal 10. If your data might have leading zeros that should be treated as decimal, specify base 10 explicitly.

Strings to Floating Point

The functions `std::stof`, `std::stod`, and `std::stold` convert strings to float, double, and long double respectively:

```
float f = std::stof("3.14"); // 3.14f
double d = std::stod("2.71828"); // 2.71828
double e = std::stod("1.5e3"); // 1500.0 (scientific notation)
```

These follow the same pattern as the integer functions — they skip leading whitespace, stop at the first character that doesn't fit the number format, and throw the same exceptions for bad input.

Numbers to Strings

The `std::to_string` function converts numeric types back to strings:

```
std::string s1 = std::to_string(42);      // "42"
std::string s2 = std::to_string(-7);     // "-7"
std::string s3 = std::to_string(3.14);   // "3.140000"
```

For more control over formatting, use `std::format` (covered in Chapter 10):

```
std::string s = std::format("{:.2f}", 3.14); // "3.14"
std::string h = std::format("{:#x}", 255);   // "0xff"
```

Manual Place-Value Math

You can also convert a number from one base to another by hand using place-value arithmetic — no library functions needed. To convert from another base to decimal, multiply each digit by its place value and add:

Hex "2A" to decimal:

$$\begin{array}{r} 2 * 16^1 = 32 \\ A * 16^0 = 10 \\ -- \\ 42 \end{array}$$

Binary "101010" to decimal:

$$1*32 + 0*16 + 1*8 + 0*4 + 1*2 + 0*1 = 42$$

You can write a loop to do this programmatically using Horner's method — process each digit left to right, multiplying the running total by the base before adding the next digit:

```
std::string hex = "2A";
int result = 0;
for (char c : hex) {
    int digit = (c >= 'A') ? (c - 'A' + 10) : (c - '0');
    result = result * 16 + digit;
}
// result is 42
```

This is essentially what `std::stoi` does internally. Understanding the math helps you debug base-conversion issues and work with bases that library functions don't directly support.

Try It: Strings and Numbers

This program demonstrates conversions in both directions:

```
#include <print>
#include <string>

int main() {
    // string to integer
    int year = std::stoi("1985");
    std::println("Year: {}", year);

    // string to double
    double bpm = std::stod("150.2");
    std::println("BPM: {}", bpm);

    // number to string
    std::string msg = "Side " + std::to_string(1) + " of the cassette";
    std::println!("{}", msg);

    // using the pos parameter
```

```

std::size_t pos;
int tempo = std::stoi("120bpm", &pos);
std::println("Tempo: {} (unit starts at index {})", tempo, pos);

// base conversions
std::println("\n1985 in different bases:");
std::println("  Hex \"7C1\"          = {}",
  std::stoi("7C1", nullptr, 16));
std::println("  Binary \"11111000001\" = {}",
  std::stoi("11111000001", nullptr, 2));
std::println("  Octal \"3701\"         = {}",
  std::stoi("3701", nullptr, 8));

// base 0 auto-detection
std::println("\nBase 0 auto-detection:");
std::println("  \"0x7C1\" = {}", std::stoi("0x7C1", nullptr, 0));
std::println("  \"010\"   = {} (octal, not 10!)",
  std::stoi("010", nullptr, 0));
std::println("  \"10\"    = {}", std::stoi("10", nullptr, 0));
}

```

Two's Complement

So far we have only talked about positive numbers. But how does the computer represent negative numbers? You can't just put a minus sign in front of a binary number — there are only 1s and 0s.

One's Complement (and Why We Don't Use It)

An early idea was **one's complement**: flip every bit to get the negative. In an 8-bit system:

```

42 = 0010 1010
-42 = 1101 0101 (every bit flipped)

```

This mostly works, but it has a fatal flaw: there are two representations of zero.

```

+0 = 0000 0000
-0 = 1111 1111

```

Having two zeros complicates hardware and arithmetic. Is $-0 == +0$? It should be, but the bit patterns differ. Engineers needed a better solution.

Two's Complement

Two's complement fixes this by adding one extra step: flip all the bits *and* add 1.

```

42 = 0010 1010
    1101 0101 (flip bits)
+ 0000 0001 (add 1)
-----
-42 = 1101 0110

```

Now there is only one zero:

```

0 = 0000 0000
   1111 1111 (flip bits)
+ 0000 0001 (add 1)
-----
0000 0000 (overflow discarded --- still 0!)

```

In two's complement, the highest bit (the **sign bit**) tells you whether the number is negative. For an 8-bit signed integer:

- 0xxx xxxx — positive (0 to 127)
- 1xxx xxxx — negative (-128 to -1)

This gives 8-bit signed integers a range of **-128 to 127**. Notice that there is one more negative value than positive — that's because zero takes one of the "positive" slots.



Tip: Nearly every modern computer uses two's complement for signed integers. When you declare `int x = -42;`, the bit pattern stored in memory is the two's complement representation.

Why Two's Complement Is Brilliant

The beauty of two's complement is that addition and subtraction **just work** with the same hardware used for unsigned numbers. The CPU does not need separate circuitry for signed math — it uses the same adder for both.

Try It: Seeing Two's Complement

This program shows the bit patterns of positive and negative numbers so you can see two's complement in action:

```
#include <cstdlib>
#include <print>

int main() {
    int8_t values[] = {42, -42, 0, -1, 127, -128};
    for (int8_t i : values) {
        // cast to unsigned to see the raw bit pattern
        uint8_t bits = static_cast<uint8_t>(i);
        std::println("{:4d} = {:08b}", i, bits);
    }
}
```

Output:

```
42 = 00101010
-42 = 11010110
0 = 00000000
-1 = 11111111
127 = 01111111
-128 = 10000000
```

Integer Sizes and Ranges

You saw that an 8-bit signed integer can hold values from -128 to 127. But `int` is not 8 bits — so how big is it, and what range can it hold?

Bytes

A group of 8 bits is called a **byte**, and it is the basic unit of measurement for computer memory. When you buy a computer with 16 GB of RAM, you are buying roughly 16 billion bytes of memory. More precisely, 1 GB (gigabyte) is 2^{30} bytes, which equals 1,073,741,824 — just over a billion. So 16 GB is 17,179,869,184

bytes. Close enough to 16 billion for casual conversation, but the difference matters when you are counting precisely.

You will also encounter kilobytes (KB, $2^{10} = 1,024$ bytes), megabytes (MB, $2^{20} = 1,048,576$ bytes), and terabytes (TB, 2^{40}). Notice a pattern: each unit is a power of 2, not a power of 10, because binary makes powers of two the natural grouping.



Tip: The `sizeof` operator in C++ returns sizes in bytes, not bits. Since 1 byte = 8 bits, a 4-byte `int` has 32 bits. You will see `sizeof` used frequently when working with memory and data structures.

How Many Bits?

C++ has several integer types, each with a minimum guaranteed size. On most modern systems, the sizes are:

Type	Typical Size	Bits
<code>char</code>	1 byte	8
<code>short</code>	2 bytes	16
<code>int</code>	4 bytes	32
<code>long</code>	4 or 8 bytes	32 or 64
<code>long long</code>	8 bytes	64

You can check the size of any type with the `sizeof` operator:

```
std::println("char:      {} bytes", sizeof(char));    // 1
std::println("short:     {} bytes", sizeof(short));   // 2
std::println("int:       {} bytes", sizeof(int));     // 4
std::println("long long: {} bytes", sizeof(long long)); // 8
```

The Range Formula

With n bits, you can represent 2^n distinct values. How those values are divided depends on whether the type is signed or unsigned:

- **Unsigned** (no negatives): 0 to $2^n - 1$
- **Signed** (two's complement): $-2^{(n-1)}$ to $2^{(n-1)} - 1$

For example, with 8 bits:

- Unsigned: 0 to 255 ($2^8 - 1$)
- Signed: -128 to 127 (-2^7 to $2^7 - 1$)

Here are the ranges for the common types:

Type	Range
<code>unsigned char</code>	0 to 255
<code>signed char</code>	-128 to 127
<code>unsigned short</code>	0 to 65,535
<code>short</code>	-32,768 to 32,767
<code>unsigned int</code>	0 to 4,294,967,295 (about 4.3 billion)
<code>int</code>	-2,147,483,648 to 2,147,483,647 (about +/- 2.1 billion)
<code>unsigned long long</code>	0 to 18,446,744,073,709,551,615

Type	Range
long long	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807



Tip: If you are ever unsure about a type's range, use `std::numeric_limits` from `<limits>`:

```
std::println("int max: {}", std::numeric_limits<int>::max());
std::println("int min: {}", std::numeric_limits<int>::min());
```

Signed vs. Unsigned

Every integer type has a signed and unsigned variant.

```
unsigned int positive_only = 42;
int can_be_negative = -42;
```

An unsigned type gives up negative values in exchange for a larger positive range. A 32-bit unsigned `int` can hold values up to about 4.3 billion, while a signed `int` tops out around 2.1 billion. You will see unsigned again in the next chapter as `size_t`, the type returned by `.size()` on standard containers.



Tip: Mixing signed and unsigned types in comparisons can cause surprising bugs. The expression `-1 < 0u` is `false` because `-1` is implicitly converted to a very large unsigned value. When possible, stick to signed types for general arithmetic and use unsigned only when you have a specific reason (like bit manipulation or interfacing with APIs that require it).

What Happens When You Overflow and Underflow?

If you try to store a value that does not fit, the behavior depends on whether the type is signed or unsigned:

- **Unsigned overflow/underflow** wraps around. Adding 1 to the maximum value gives 0:

```
unsigned char x = 255;
x = x + 1; // x is now 0 (wraps around)
x = x - 2; // x is now 254 (wraps around)
```

- **Signed overflow/underflow** is undefined behavior. The compiler can do anything:

```
int y = 2'147'483'647; // INT_MAX
y = y + 1; // undefined behavior!
```



Tip: "Undefined behavior" is not just a theoretical concern. Compilers actively exploit it for optimization. A signed overflow might silently wrap, or it might cause your program to behave in completely unexpected ways. Never rely on signed overflow.

Try It: Exploring Sizes and Limits

This program prints the size and range of each integer type:

```
#include <limits>
#include <print>

int main() {
    std::println("{:<12} {:>5} {:>22} {:>22}",
                "Type", "Bytes", "Min", "Max");
```

```

std::println("{:<12} {:>5} {:>22} {:>22}", "char",
    sizeof(char),
    std::numeric_limits<signed char>::min(),
    std::numeric_limits<signed char>::max());

std::println("{:<12} {:>5} {:>22} {:>22}", "short",
    sizeof(short),
    std::numeric_limits<short>::min(),
    std::numeric_limits<short>::max());

std::println("{:<12} {:>5} {:>22} {:>22}", "int",
    sizeof(int),
    std::numeric_limits<int>::min(),
    std::numeric_limits<int>::max());

std::println("{:<12} {:>5} {:>22} {:>22}", "long long",
    sizeof(long long),
    std::numeric_limits<long long>::min(),
    std::numeric_limits<long long>::max());
}

```

Binary Addition and Subtraction

Binary Addition

Binary addition works just like decimal addition, but you carry at 2 instead of 10:

```

0 + 0 = 0
0 + 1 = 1
1 + 0 = 1
1 + 1 = 10 (0, carry 1)

```

Let's add 42 + 15 in 8-bit binary:

```

  0010 1010 (42)
+ 0000 1111 (15)
-----
  0011 1001 (57)

```

Subtraction with Two's Complement

To subtract, you negate the second number (using two's complement) and add. Let's compute 42 - 15:

Step 1 — find -15 using two's complement:

```

15 = 0000 1111
    1111 0000 (flip bits)
+   0000 0001 (add 1)
-----
-15 = 1111 0001

```

Step 2 — add 42 + (-15):

```

  0010 1010 (42)
+ 1111 0001 (-15)
-----
 1 0001 1011

```

^
overflow bit (discarded in 8 bits)

The result is 0001 1011 = 27. Correcto! The overflow bit is discarded because we are working with 8-bit values, and the addition produces the right answer automatically.

Let's also try a negative result: 15 - 42:

Step 1 — find -42:

```
42 = 0010 1010
    1101 0101 (flip bits)
+ 0000 0001 (add 1)
-----
-42 = 1101 0110
```

Step 2 — add 15 + (-42):

```
0000 1111 (15)
+ 1101 0110 (-42)
-----
1110 0101 (-27)
```

The sign bit is 1, so the result is negative. To verify, convert back: flip the bits (0001 1010), add 1 (0001 1011 = 27), so the answer is -27. Perfecto.



Tip: Overflow can happen when the result is too large (or too small) for the number of bits. Adding two large positive numbers can wrap around to a negative value. In C++, signed integer overflow is **undefined behavior** — the compiler is free to do anything. Be careful with arithmetic near the limits of a type.

Try It: Unsigned Wraparound

This program demonstrates unsigned addition wrapping around, the same way the CPU adds binary numbers:

```
#include <cstdlib>
#include <print>

int main() {
    uint8_t a = 200;
    uint8_t b = 100;
    uint8_t result = a + b; // 300 doesn't fit in 8 bits

    std::println("{} + {} = {} (wrapped)", a, b, result);
    // 200 + 100 = 44 (wrapped), because 300 - 256 = 44

    uint8_t x = 42;
    uint8_t y = 15;
    std::println("{} + {} = {}", x, y, static_cast<int>(x + y));
    std::println("{} - {} = {}", x, y, static_cast<int>(x - y));
}
```

Bit Operators

We saw the logical operators in the previous chapter, in this chapter we will look at the version of those operators that operate on bits of a number.

Operator	Operation	Example (binary)	Result
&	AND	0b1100 & 0b1010 0b1000	
	OR	0b1100 0b1010 0b1110	
^	XOR	0b1100 ^ 0b1010 0b0110	
~	NOT (complement)	~0b1100	flips all bits
<<	left shift	0b0001 << 2 0b0100	
>>	right shift	0b1000 >> 2 0b0010	

The 0b prefix lets you write binary literals, which makes it easier to see what is happening bit by bit.

The bit operators work similarly to logical operators, except they work with the individual bits. When using the & operator a bit in the result will be 1 only if both corresponding bits are 1. When using the | operator a bit in the result will be 1 only if either corresponding bits are 1. There is also a XOR (^) operator that is true only if one or the other, but not both are true. Like, you can have fish or chicken but not both — you can have fish XOR chicken.

AND (&) is useful for checking if a specific bit is set.

```
int flags = 0b1010;
if (flags & 0b0010) {
    std::cout << "bit 1 is set" << std::endl;
}
```

OR (|) is useful for setting bits.

```
int flags = 0b1000;
flags |= 0b0010; // flags is now 0b1010
```

XOR (^) flips bits — if a bit is 1, it becomes 0, and vice versa.

```
// flip bit 2 of a number
void toggle_bit_2(int &num) {
    num ^= 0b100;
}
```

Shift Operators

The shift operators << and >> move bits left or right by a specified number of positions. You first saw << and >> used for stream I/O in Chapter 1 — here we are talking about their original purpose as bitwise operators.

Left Shift: <<

Left shift moves every bit to the left and fills the empty positions on the right with zeros:

```
0000 0101 (5)
<< 1
0000 1010 (10)

0000 0101 (5)
<< 3
0010 1000 (40)
```

Each left shift by 1 **multiplies the value by 2**. Shifting left by n is the same as multiplying by 2ⁿ:

```
int x = 5;
int doubled = x << 1; // 10 (5 * 2)
int times8 = x << 3; // 40 (5 * 2 * 2 * 2)
```

Right Shift: >>

Right shift moves every bit to the right. For unsigned values, zeros fill in from the left. For signed values, the sign bit is copied (called **arithmetic shift**), so negative numbers stay negative after a right shift.

```
0010 1000 (40)
>> 1
0001 0100 (20)

0010 1000 (40)
>> 3
0000 0101 (5)
```

Each right shift by 1 **divides the value by 2** (discarding any remainder). Shifting right by n is the same as dividing by 2^n :

```
int y = 40;
int halved = y >> 1; // 20 (40 / 2)
int div8 = y >> 3; // 5 (40 / 8)
```

What About Odd Numbers?

When you right-shift an odd number, the lowest bit is lost — just like integer division discards the remainder:

```
int odd = 7;
int result = odd >> 1; // 3, not 3.5 (same as 7 / 2)
```

Compound Assignment

As with other operators, there are compound assignment forms:

```
int flags = 1;
flags <<= 4; // flags is now 16 (1 shifted left 4)
flags >>= 2; // flags is now 4 (16 shifted right 2)
```



Tip: Shifting by a negative amount or by more bits than the type has is **undefined behavior**. For a 32-bit `int`, valid shift amounts are 0 to 31.



Tip: Modern compilers optimize multiplication and division by powers of two into shift operations automatically. Write `x * 4` rather than `x << 2` unless you are doing actual bit manipulation — it is clearer, and the compiler will generate the same code.

Try It: Powers of Two with Shifts

This program uses shifts to compute and display powers of two:

```
#include <print>

int main() {
    std::println("Powers of two using left shift:");
    for (int i = 0; i < 16; ++i) {
        int power = 1 << i;
        std::println(" 1 << {:2} = {:5} ({:#018b})", i, power, power);
    }
}
```

```

std::println("\nDividing 1000 by powers of two using right shift:");
int val = 1000;
for (int i = 0; i <= 4; ++i) {
    std::println("  {} >> {} = {}", val, i, val >> i);
}
}

```

Key Points

- **A number is a number**, regardless of how you represent it. 42, 0b101010, 0x2A, and 052 are all the same value.
- **Decimal, binary, hex, and octal** are just different bases. Each one makes certain patterns easier to see.
- **C++ supports multiple bases** in literals (0b, 0x, 0), output ({:b}, {:x}, {:o}), and conversion (std::stoi with a base parameter).
- **Converting between strings and numbers** is straightforward with std::stoi, std::stod, and std::to_string. The pos parameter tells you where parsing stopped, and the base parameter lets you parse hex, binary, and octal.
- **A byte is 8 bits**, and integer types come in different sizes — from 1-byte char to 8-byte long long. The number of bits determines the range of values a type can hold.
- **Two's complement** is how computers represent negative integers. It eliminates the double-zero problem and lets addition and subtraction share the same hardware.
- **Signed overflow is undefined behavior**, but unsigned overflow wraps around predictably. Be careful with arithmetic near the limits of a type.
- **Shift operators** (<<, >>) move bits and are equivalent to multiplying and dividing by powers of two.

A number can wear many different outfits, but underneath, es el mismo numero. No te preocupes — you have got this. Nos vemos en el proximo capitulo!

Exercises

1. Convert the decimal number 200 to binary, hexadecimal, and octal by hand. Verify your answers by writing a C++ program that prints 200 in each base using std::println.

2. **What does this print?**

```

int x = 0b1100;
int y = 052;
std::println("{} ", x + y);

```

3. **Think about it:** In an 8-bit two's complement system, the most negative value is -128 but the most positive value is only 127. Why isn't the range symmetric?

4. **Where is the bug?** This loop is supposed to count down from 10 to 0, but it never terminates. Why?

```

unsigned int count = 10;
while (count >= 0) {
    std::println("{} ", count);
    --count;
}

```

5. Using two's complement with 8 bits, compute 100 - 75 by hand. Show the binary representation of 100, the two's complement of 75, and the binary addition.

6. **What values do a, b, and c hold** after these statements execute?

```

int a = 1 << 10;
int b = 100 >> 3;
int c = (1 << 4) - 1;

```

7. **Where is the bug?** A programmer wrote this code and expected it to print 700. What value does it actually print, and why?

```
int permissions = 0700;
std::println("Permissions: {}", permissions);
```

8. **What does this print?** What is wrong with this code?

```
int big = 2'000'000'000;
int doubled = big * 2;
std::println("{} * 2 = {}", big, doubled);
```

9. **Write a program** that reads a hexadecimal color code (like "FF8000") from the user, converts it to its red, green, and blue components (each 0–255), and prints each component in decimal and binary. Use `std::stoi` with the base parameter and `substr` to extract each pair of hex digits.

10. **What does this print?**

```
uint8_t a = 250;
uint8_t b = 20;
uint8_t sum = a + b;
std::println("{} + {} = {}", a, b, sum);
```

11. **Write a program** to implement a function `bool is_set(int num, int bit)` that returns true if bit number `bit` of `num` is set. For example, if `is_set(13, 1)` is false but `is_set(13, 2)` is true.

12. **What does this print?**

```
int a = 1'000'000;
int b = 0xFF'00'FF;
int c = 0b1111'0000'1111'0000;
std::println("{} {} {}", a, b, c);
```

Are the digit separators part of the value? What does the program output?

13. **What does this print?**

```
std::string input = "42 100 255";
std::size_t pos = 0;

int a = std::stoi(input, &pos);
int b = std::stoi(input.substr(pos), &pos);
int c = std::stoi(input.substr(pos));

std::println("{} {} {}", a, b, c);
```

Walk through each call and explain what `pos` ends up as after each one.

14. **Calculation:** Without running a program, compute each of these in 8-bit binary, then give the result in decimal:

```
0b1010'1100 & 0b1111'0000
0b1010'1100 | 0b0000'1111
0b1010'1100 ^ 0b1111'1111
```

What does XORing with all-ones do?

15. **Calculation:** On a typical 64-bit Linux system, what does each of these print?

```
sizeof(char)
sizeof(short)
sizeof(int)
```

```
sizeof(long)
sizeof(long long)
```

For each type, what is the largest value an *unsigned* version of that type can hold? (You may answer in terms of $2^N - 1$ instead of writing the full decimal number.)

16. **Write a program** that asks the user for an integer and prints it in decimal, hex, octal, binary, and again as a `std::string` produced by `std::to_string`. Use `std::println` (or `std::format`) for the formatted output.