

3

SPELL OF REBIRTH: TRANSMUTING GREETINGS INTO SHELLCODE



In this chapter, we'll extend our ASM and C knowledge by writing a simple "Hello, world!" program. Then we'll modify that code to create our first hacking-related program: a simple shellcode for gaining command line access to a computer system. Yes, we're already that advanced. Along the way, you'll learn to work with pointers, use the GNU Debugger to inspect a program's allocated memory, and explore different addressing modes for memory access. First, though, we'll cover a few more general concepts about how computers handle information internally, and we'll take a closer look at how memory works.

Representing Information

We're going to start working with values larger than 1 byte in this chapter, and for you to fully comprehend what's happening, I'll need to introduce some basic concepts about how information is represented and stored in memory. This may not sound as exciting as writing a shellcode, but these fundamentals are just as important. Bear with me, and in no time we'll do our very first hacking thing.

Bits, Bytes, Words, and Beyond

A *bit* is the smallest unit of information a computer can manage. A bit can hold one of two values: 0 or 1. These values are sometimes referred to as *false* and *true*, especially in the context of logic circuitry.

If you put 8 bits together, you get a *byte*. Therefore, a byte is a sequence of eight 0s and/or 1s. One byte can represent any number up to 8 binary digits long. If you're more comfortable counting in decimal, a byte can store 256 different values, typically from 0 to 255. (I'll talk more about numeric representation shortly.)

Each half of a byte is known as a *nibble*. A byte has two nibbles, the high one and the low one, each of which is 4 bits long. Therefore, each nibble can represent 16 different values ($2^4 = 16$). When working in hexadecimal (that is, base 16), we can represent 16 values just with one digit (0 to 9 for 0 through 9, and a to f for 10 through 15). This means a byte can be represented with two hexadecimal digits, one for each nibble.

Going up from bytes, we start doubling the number of bits:

- A *word* is composed of 2 bytes, or 16 bits.
- A *doubleword* or *dword* is composed of 2 words, which is 4 bytes, or 32 bits.
- A *quadword* or *qword* is composed of 2 doublewords, which is 4 words, 8 bytes, or 64 bits.

We won't go further than this, but feel free to figure out the names of the next sizes in the list on your own.

Note that these are the typical interpretations of these terms, which mostly derive from the Intel processor definitions. In reality, the definition of a *word* can vary by architecture, based on the processor's native word size, as we'll see in the next section.

Processor Native Word Sizes

A processor's *native word size* is the number of bits the processor is most comfortable working with at one time. In a sense, you encounter native word sizes regularly: Whenever we talk about 32-bit versus 64-bit processors or programs, we're referencing word sizes. Those numbers are important for more than just choosing the right download link for your Linux distribution, however. They also have some low-level implications:

- Do you remember the registers in a processor? Sure you do. The processor's native word size defines the size of those registers. A 32-bit processor has 32-bit registers, and a 64-bit processor has 64-bit registers. (Actually, that isn't completely accurate, but for now you can just consider this size to be the size of the processor's registers.)
- The native word size is usually the width of the CPU's data bus.
- The native word size is usually the width of the address bus as well.

We'll talk about what those buses are later in the chapter.

Overall, and without going deep into the electronics within the processor, what you need to know is that each processor is optimized to work with its native word size. For example, a 32-bit processor will perform arithmetic operations and access memory faster when it's working with 32-bit values than when it's working with 16-bit values. I know that sounds a bit counter-intuitive and may require an act of faith to believe, but it would be very tedious to go through all the complex details to understand why this is the case. Just to give you a little taste, here's a fragment from the "Data Types" section of the Intel 80386 Programmer's Reference Manual from 1986:

When used in a configuration with a 32-bit bus, actual transfers of data between processor and memory take place in units of doublewords beginning at addresses evenly divisible by four; however, the processor converts requests for misaligned words or doublewords into the appropriate sequences of requests acceptable to the memory interface. Such misaligned data transfers reduce performance by requiring extra memory cycles. For maximum performance, data structures (including stacks) should be designed in such a way that, whenever possible, word operands are aligned at even addresses and doubleword operands are aligned at addresses evenly divisible by four. Due to instruction prefetching and queuing within the CPU, there is no requirement for instructions to be aligned on word or doubleword boundaries. (However, a slight increase in speed results if the target addresses of control transfers are evenly divisible by four.)

Summing up, try to program your processor using its native word size. It may look like you're wasting some space, but that's the right way to do it.

Decimal, Binary, Hexadecimal, and More

A *number* is an abstract entity that represents a quantity. The quantity is universal, but it can be represented in different ways. The most obvious example is human language: A 1 means one thing, but English speakers say *one*, Dutch speakers say *een*, Spanish speakers say *uno*, and so on. Each language has a different verbal representation for the same abstract quantity.

The same is true when we write numbers symbolically: The same abstract quantity can be represented in many different ways. For instance, you can write the number twenty-four with Western Arabic numerals as 24 or with Roman numerals as XXIV. Setting aside the symbols themselves, the

more general way of representing numbers is using *radix-based positional numeral systems*, where the position of each digit is associated with a value.

We normally use radix 10, which means that we use 10 possible symbols (0 to 9) for each digit or position in a number. If we need to represent more than 10 values, we add a new digit, which will be in the second position and will be multiplied by 10. Right? Figure 3-1 shows an example.

```

1234
^^^^
||+----- First position times 1    -> 4 x 1   =   4
|+----- Second position times 10   -> 3 x 10  =  30
|+----- Third position times 100   -> 2 x 100 = 200
+----- Fourth position times 1000  -> 1 x 1000 = 1000
                                         -----
                                         1234

```

Figure 3-1: Parsing a base-10 number

We all know this from school, but the terminology might be unfamiliar. The basis of the system, 10 in this case, is called the *radix* (or *base*), and the value of each digit is multiplied by a power of the radix: The first position is multiplied by $10^0 = 1$, the second position by $10^1 = 10$, the third position by $10^2 = 100$, and so on.

Here's the interesting thing: We can use any number as the radix, and furthermore, we can represent any integer number using any radix. Taking into account that computers use digital circuits that can only hold two values, any radix that's a multiple of 2 is interesting for us. Actually, there are three numerical representation systems that are used frequently when working with computers because they make some things much easier:

Binary A binary representation uses base 2, and therefore each digit can take the value 0 or 1. Each position is multiplied by a power of 2 (1, 2, 4, 8, 16, ...). Binary numbers are usually represented with a leading *0b*. For example, *0b101* in binary is 5 in decimal: $(1 \times 2^2) + (0 \times 2^1) + (1 \times 2^0) = 5$.

Octal This representation uses powers of 8 and symbols from 0 to 7 for each digit. Each octal digit is represented using 3 bits. The octal system was very common in the past, but it's rarely used nowadays. Unix filesystem permissions are one of the few cases where you'll still see it. Octal numbers are usually prefixed with a *0*. For example, *011* in octal is 9 in decimal: $(1 \times 8^1) + (1 \times 8^0) = 9$.

Hexadecimal This representation uses radix 16, so each digit can have 16 values. We only have 10 symbols in our normal numeral systems (0 to 9), so for the 6 missing symbols we use letters: *a* means 10, *b* means 11, and so on up to *f* for 15. Each hexadecimal digit requires 4 bits (a

nibble), which is why this system is commonly used in computing. A byte can be represented with two hex digits, a word with four, a dword with eight, and so on. Hex values are usually prefixed with 0x. For example, 0x21 is 33 in hexadecimal: $(2 \times 16^1) + (1 \times 16^0) = 32 + 1 = 33$.

Which representation should you use in your code? It doesn't really matter, since all of them represent the same values. Which one you choose depends on what you want to achieve and whether a particular system will make your life easier, but it won't have any impact on the final machine code a program will generate.

Our focus here has been on whole numbers. Computers also have ways to represent real numbers with fractional components (fixed point and floating point), but we won't worry about such numbers in this book because you rarely need this type of data for hacking or systems programming. No kernel system call expects a real number as a parameter, for example.

We've also sort of been assuming that all numbers are positive, but how does a computer represent negative numbers? It's not hard, but it isn't obvious either.

Negative Numbers

To get a taste for how negative numbers work, let's think about a single byte (eight 0s or 1s). We've seen that 8 bits can represent 256 values, from 0 to 255. That's perfect for natural numbers, but things change if we need negative numbers, too. First, we need to store the sign of the number, which will take up one of the bits. That leaves us with 7 bits, which allows us to represent 128 values. Figure 3-2 shows one way to represent a few of those numbers.

8 => 0 000 1000	-8 => 1 000 1000
7 => 0 000 0111	-7 => 1 000 0111
...	...
1 => 0 000 0001	-1 => 1 000 0001
0 => 0 000 0000	0 => 1 000 0000

Figure 3-2: A bad way to represent negative numbers

There are a few problems with this approach. The first is that we have two representations for the number zero. That isn't convenient as it can make computations ambiguous, and we're also losing the opportunity to represent an extra number. The second problem is that arithmetic becomes complicated: Multiplication is kind of easy, but addition is kind of hellish.

Fortunately for us, long ago, some smart people came up with a better representation for negative numbers. It also uses the most significant bit to indicate the sign, but the value of the number is encoded in a more practical way. Figure 3-3 shows how it works.

8 => 0 000 1000	-8 => 1 111 1000
7 => 0 000 0111	-7 => 1 111 1001
...	...
2 => 0 000 0010	-2 => 1 111 1110
1 => 0 000 0001	-1 => 1 111 1111
0 => 0 000 0000	

Figure 3-3: A good way to represent negative numbers

After the sign bit, the rest of the number is constructed by counting upward as usual for positive numbers and counting downward (backward) for negative numbers. Now there's a single representation for zero, and it's actually zero (a 0 in each bit). This has a consequence: Zero is considered a positive number because its most significant bit is 0 (that's our sign bit). This is why a signed char can take values from -128 to +127; zero is the first of the positive values.

With this system, you can change a number's sign (that is, find the positive equivalent of a negative number or vice versa) with two simple steps:

1. Invert all the bits in the number (this is called the *one's complement*).
2. Add 1 to the result.

Together, this sequence of operations is called taking the *two's complement*. Let's use the number 5 as an example and calculate its two's complement to determine the bit representation of -5:

1. The number 5 in binary is 0000101, so inverting the bits gives us 1111010.
2. Adding 1 gives us 1111011, which is indeed -5.

The biggest advantage of this system is that basic arithmetic operations just work. If you add the binary representations of 5 and -5 that were just shown, it results in 0. Subtraction and multiplication also work out of the box.

SIGN EXTENSION

There's a small caveat to consider when using negative numbers, known as *sign extension*. Suppose we store the value -1 in a given memory position, and we read that value into a 32-bit register. We'll get the following:

Memory	Register
1111 1111 (-1)	0000 0000 0000 0000 0000 0000 1111 1111

When considering the register as a whole 32-bit number, the value we have is 256, not -1. Representing the value -1 with 32 bits requires all the bits to be 1s (to see why, take the two's complement, as just described). To deal with these

cases, processors typically offer a way to perform a sign extension operation. This operation takes the sign bit and copies it over to the upper bits of the value.

As we saw in Chapter 2, most processors allow us to work with different sizes of register. Usually, you can select how many bits you want to extend the sign to, depending on whether you're using words or doublewords in your program. For example, on x86_64, EAX is the lower 32 bits of the 64-bit RAX register. If we only need 32 bits for our operations, we can just extend the sign for the first 32 bits; otherwise, we can extend it to the whole 64-bit value:

ORIGINAL	0000 0000 0000 0000 0000 0000 1111 1111
16-bit sign extension	0000 0000 0000 0000 1111 1111 1111 1111
32-bit sign extension	1111 1111 1111 1111 1111 1111 1111 1111
64-bit sign extension	1111 1111 1111 1111 1111 1111 1111 1111
	1111 1111 1111 1111 1111 1111 1111 1111

Remember, each 4 binary bits is a nibble, and two nibbles form a byte.

The sign extension operation works the same way for positive numbers, where the sign bit is 0, as for negative numbers.

We won't discuss negative number representation further, but the interested reader can see the "Two's complement" Wikipedia page (https://en.wikipedia.org/wiki/Two's_complement) for more information.

Text Strings

We've been talking about how to represent integer numbers, but there's one more type of data we'll want to work with: text. A *text string* is just a sequence of characters—that is, any message we write on the screen or any text we type on a keyboard to provide as input to our programs.

We can only store numbers in computer memory, so how can we represent the letters in a text string? The answer is to use a code, which in its simplest form is just a mapping between symbols. We can assign each letter and symbol we want to be able to represent a number, then use the code to convert those numbers into characters to display on the screen.

There are a plethora of options for assigning numbers to characters. We could come up with our own code, but several good systems already exist, with American Standard Code for Information Interchange (ASCII) being the best known. Back in the day, Extended Binary Coded Decimal Interchange Code (EBCDIC) was also popular because early IBM computers used it, and more recently Unicode has become the de facto standard. UTF-8, one of the encodings for Unicode, uses variable-length encoding and is backward compatible with ASCII. In other words, 1-byte-long UTF-8 codes are the same as the ASCII codes.

We'll use ASCII string encoding throughout this book. ASCII was originally defined as a 7-bit code, to be compatible with the teletype machines used back in the 1960s. This is also why the first 32 codes (from 0x00 to 0x1f) are control values intended to operate those machines. For example, the

code 0x0a, also known as *LF* or *line feed*, was intended to move the paper in the teletype machine to the next line. Most of those codes aren't used anymore, but a few of them are: Today you'll find *line feed*, *carriage return*, and *tabulation* represented as \n, \r, and \t, respectively.

In ASCII, the value 0x20 represents a space. Numbers are encoded using the values 0x30 to 0x39, which makes it convenient to print decimal digits in ASM programs: Just add 0x30 to the value. Uppercase letters are encoded using the values 0x41 to 0x5a, and lowercase letters are stored using the values 0x61 to 0x7a. This means we can easily convert between uppercase and lowercase just by switching bit 5 of a character on or off. When sorting strings alphabetically, the sort is actually based on these underlying numbers, so a word beginning with a capital letter will be sorted before a word beginning with a lowercase letter.

To use a text string in a program, we need one extra piece of information: We need to tell the program when those numbers in memory are no longer letters. Put another way, we need to define the length of the string. There are two main options for this:

- Store a number indicating the size of a string together with the sequence of numbers representing the string itself.
- Use a number that would never appear in a text string as a guard value or delimiter at the end of the string.

Languages like Java use option 1. (Historically, Java also used the UTF-16 encoding, an encoding for Unicode that has a minimum unit size of 2 bytes.) The C language uses option 2, with 0x00 indicating the end of a string. Figure 3-4 compares how these two languages would store the string "ABC" in memory.

Java		C	
Size	0x00	0x41	Char 1
Size	0x00	0x42	Char 2
Size	0x00	0x43	Char 3
Size	0x04	0x00	End of string
Char 1	0x00		
Char 1	0x41		
Char 2	0x00		
Char 2	0x42		
Char 3	0x00		
Char 3	0x43		

Figure 3-4: String representation in Java versus C

Java starts with a 32-bit value to store the size of the string, meaning even an empty string will take up at least 4 bytes in memory. Then it uses 16 bits to encode each character. Like UTF-8, UTF-16 matches the ASCII code for each character, but extends its size to 2 bytes. By contrast, C uses 1 byte per character, plus an extra byte for the guard value at the end. Our

three-character string therefore takes up just 4 bytes, compared to the 10 bytes required by Java.

NOTE

Java 9 introduced a new string representation called compact strings. Now, unless the string requires a special character that needs to be encoded using UTF-16, Java uses 1 byte per character. It also adds an extra byte to indicate which type of encoding is used (1 or 2 bytes per character). In this newer representation, the Java string in our example would take up 8 bytes—still bigger than the C version, but not by as much.

Why would Java code strings this way when it eats up so much more space? The answer is security. The Java approach is to check the sizes of all strings and make sure that we never write outside of the allotted bounds. The C implementation is lighter and more compact, but it leads to all kinds of security problems when not used properly. We'll see this in Chapter 4 when we discuss buffer overflow exploits.

A Closer Look at Memory

Let's use what you have just learned about how computers represent information—in particular, the processor's native word size—to take a closer look at how memory works. I mentioned earlier that a 32-bit processor will operate faster on a 32-bit value than a 16-bit value, and accessing it in memory will be faster too. An implication of this is that, even though the smallest addressable value in memory is 8 bits (remember the size of our memory drawers from Chapter 1?), a 32-bit processor can read 32 bits (4 bytes) from memory at once.

That simple sentence has some important concepts behind it. Let's go deeper.

A Simplified Hardware Memory Model

The memory system in a PC can be quite complicated, and it's beyond the scope of this book to go into all the details. However, discussing a simplified model of the memory system will be beneficial for us and will make it easier for curious readers to dive deeper into this topic.

In its simplest form, a memory chip provides the following pins:

- A set of *address pins*, designated from A0 to AN, where N is one less than the system's native word size.
- A set of *data pins*, designated from D0 to DN, where again N is one less than the native word size.
- A set of *control pins* that allow the processor to command the memory. These include the CS (chip select), OE (output enable), and WR/ \overline{RD} (write/read) pins. CS activates the memory chip, OE controls whether data is placed on the output pins or they remain in a high-impedance state (see the "Buses" box on page 82 for more on this), and WR/ \overline{RD} indicates whether the operation is a write or a read.

NOTE

In digital electronics, a line over a signal name indicates negation, so WR/\overline{RD} means we want to write when the signal is a 1 and read when the signal is negated (that is, when it's a 0).

The CPU has similar pins that are connected to the memory. When the CPU wants to read a value from the memory, it places the address to access on its address pins, which are connected to the memory chip's address pins, usually by a bus (discussed in the following box). The appropriate control signal is then activated on the memory chip, in this case enabling the OE and CS pins and setting WR/\overline{RD} to 0 to select a read operation. This triggers the memory to place the value from the indicated memory address onto its data pins so the CPU can read it.

Remember, this is a simplified model. In reality, a lot more goes on when accessing memory.

BUSES

A simple *bus* is just a bunch of wires connecting all the chips in a computer system together. To prevent conflicts, the chips must have some internal circuitry to process input signals (activated by the CS pin) and to allow the output pins to enter a third state (in addition to the normal low and high states) where they present a very high resistance, or *impedance* (activated by the OE pin), so no current flows through them. In this high-impedance state, the pins are effectively disconnected from the bus, so they don't interfere with signals from other chips. Control logic ensures that the CS and OE signals are properly issued to all chips in the bus, so that only one (apart from the CPU) can use it at any given time.

What makes one bus different from another are the specific control signals and protocols they use to interchange information. The S-100 bus was one of the first created for microcomputers. The Versa Module Eurocard (VME) was another popular bus used on more powerful machines. The Industrial Standard Architecture (ISA) was the selected expansion bus for the first PCs, but nowadays most computers use Peripheral Component Interconnect Express (PCIe) for that purpose.

Figure 3-5 is a basic diagram of the interface between the CPU and the memory chip, without the control signals.




Figure 3-5: The connections between the CPU and memory

The address and data pins on the processor and memory are connected by tracks in the motherboard. Activating and deactivating those pins is what makes the computer work. Let's use some concrete numbers to better understand this. Imagine that the RIP register (the instruction pointer on x86_64) is set to 4, meaning the instruction at address 4 is the next one to execute. The CPU has to read that instruction from memory, so it puts the value 4, which in binary is 0b100, on the address bus. Figure 3-6 illustrates the process.




Figure 3-6: Preparing for memory access by setting the address bus and activating the read signal

Only pin A2 is set, which corresponds to the address 4 in binary. When the \overline{RD} control signal is activated, the memory chip will access the value stored at address 4. *Activating* a pin means setting a given voltage on it. In this case, we need a logic value of 0, which usually means 0 volts. The voltage for a 1 depends on the technology used; it might be 5 V, 3.3 V, or even less.

Next, the memory chip puts the value from address 4 on its data pins so the CPU can read it. This is illustrated in Figure 3-7.

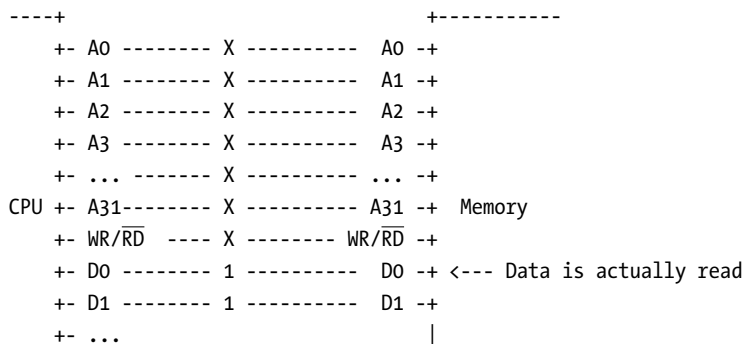


Figure 3-7: Retrieving the value from memory through the data bus

As each memory location is 8 bits, we only need 8 physical pins to send the value from address 4 to the CPU. But since there are 32 data pins, only

reading the one address would be a huge waste. As such, the memory typically doesn't put just the 1 byte that was requested on its data pins: It puts as much on as it can. In this example, that means it will output 4 bytes, starting from address 4, using the 32 data lines on the bus.

At this stage, the values on the address pins are no longer needed. Some processors take advantage of this and use the same pins to set the memory address and to write or read the values at that address. That's fine, as long as the processor is given some time to read the address and access the memory position before the pins are reset.

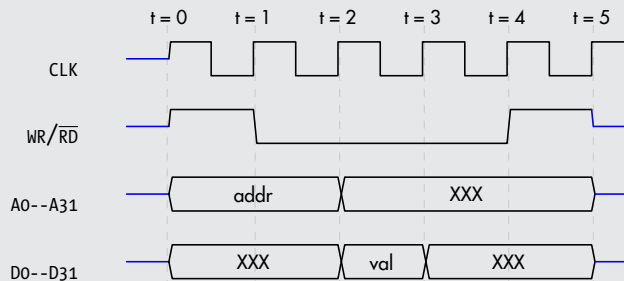
Again, note that this is a simplified example. In reality, there are a few more lines involved than we've discussed here. Also, sometimes the memory only has a data bus that's 8 bits long, and multiple memory chips are used to access 16-, 32-, or 64-bit words. In those cases, there's a real physical constraint with regard to memory-aligned accesses. These details can make a big difference between two systems.

For a more complete picture of how memory access works, you'll have to consult the datasheet of the particular processor you're using. The memory system also has its own datasheet, which you'll need to check when building your hardware. The datasheets should include chronograms diagramming the exact timing of different memory transactions; see the "Chronograms" box for more on what these look like.

CHRONOGRAMS

The bus interactions between chips are usually represented with *chronograms* in the chips' datasheets. These are graphs that show all the processor lines and how they vary with time. At the top you'll typically see the system clock signal (CLK), which orchestrates the activation and deactivation of each of the processor signals.

The chronograms don't just provide information about the sequence of signals to activate; they also indicate access time or other delays that may happen. For example, the following simplified chronogram represents the memory read access described in the previous section:



At the top, you can see the system clock. Different things happen when this signal changes, but not every edge triggers an event; some signals may be active for several clock edges. The other lines appear below the clock. The data and address buses are usually represented as a combination of lines, since

showing each line individually wouldn't be helpful and would make the diagram huge.

This particular chronogram tells us that:

- The processor puts the address value on the bus at t_0 .
- At t_1 , it sets the WR/\overline{RD} signal low. This tells the memory to read the address internally and start to retrieve the data. The amount of time for which the lines need to be set depends on the memory chip.
- After a while (that's the memory access time), the memory chip outputs the value to the data bus and the CPU can finally read the value.

Real chronograms show many more signals, to allow the hardware designer to properly interconnect the different chips in the computer.

Native Word Size and Memory Alignment

I've already hinted that there's a relationship between a processor's native word size and the functionality of the rest of the computer. For example, a 32-bit processor will interface to a 32-bit memory system. (It's not that straightforward, but that's roughly what happens; when I say *system*, that may actually mean several chips.)

Imagine you want to read just 1 byte. You put your desired address on the address lines and ask the memory chip to spit out the content of that address. The memory chip will retrieve the data at that address and the next three addresses and place it on the data lines. The CPU will then read those 4 bytes from the bus but take just the lower 8 bits from that value. In a sense, that requires more work than just reading the whole 32-bit value into a 32-bit register: The CPU has to consciously discard 3 of the 4 bytes.

To cut a long story short, a 32-bit processor will access *32-bit-aligned memory*, meaning chunks of memory starting from addresses that are multiples of 4: 0, 4, 8, 12, 16, and so on. Deviating from this alignment creates more work. For example, if we need to access a doubleword that starts at address 2, the processor will have to read addresses 0 through 3, then addresses 4 through 7, and then reorder the bytes from those two reads to get the single 32-bit value from addresses 2 through 5. This is why memory alignment is so important. Some processors will raise an exception when accessing unaligned memory. Others will just slow down. Still others may do both, depending on what you're trying to do. Fortunately, unless you really try to do otherwise on purpose, the compiler will generally ensure that you access memory the right way.

What I just described is a very simplified example to illustrate how a processor works more efficiently when everything aligns with its native internal word size. The reality is more complex, but unless you need to design your own computer (the motherboard, at least), it doesn't really matter, especially with today's large cache memories. The main takeaway from all this is that

it's always better to use the processor's native word size and to access data aligned to that size.

Little vs. Big Endian

We know the importance of using the processor's native word size and the right memory alignment now. So far, so good, but there's one more thing to take into account: How are the bytes of dwords, qwords, and other multibyte pieces of data mapped between the memory and a processor's registers?

Let's assume that the native word size of our processor is 4 bytes (that is, it's a 32-bit processor), and that we have the memory layout shown in Figure 3-8.

```
| ... | Drawer ...  
+-----+  
| 0x44 | Drawer p + 3  
+-----+  
| 0x33 | Drawer p + 2  
+-----+  
| 0x22 | Drawer p + 1  
+-----+  
| 0x11 | Drawer p  
+-----+
```

Figure 3-8: A hypothetical memory layout

Now, let's assume that we want to read the content of addresses p through $p + 3$ into one of our 32-bit registers. Which of the following values do you think we'll get?

- a) 0x44332211
- b) 0x11223344

The answer is: It depends on the processor. If our processor is *little endian*, we'll get the first value. If it's *big endian*, we'll get the second value. Most of the processors we're using in this book are little endian, although ARM can be configured to work as big endian. MIPS processors are usually big endian, but you can find little-endian hardware out there too. For the purposes of this book, we'll consider MIPS to be big endian and ARM to be little endian.

In general, you don't have to worry about whether your processor is big or little endian. You can just write and read values to and from memory, and the processor will do the right thing. Endianness only becomes important when you have to send data between computers, as the sender and receiver may have different endiannesses. This happens frequently in network programming when using open protocols that have to work with any kind of

machine. We'll see examples later in the book. For now, just keep in mind that different processors may read data in different ways.

But enough introduction. Let's see how all these concepts can be put to practical use by writing a traditional "Hello, world!" program.

"Hello, World!" in ASM

I'm pretty sure you're familiar with the classic "Hello, world!" programs that display the message "Hello, world!" in the console. In fact, we wrote a simple one that used libc functions in Chapter 2. Let's create another one now, entirely in ASM. Then we'll see how the program translates to C.

The way to display a message like "Hello, world!" on Linux is to write it to the *standard output* (usually the console) using the `SYS_write` system call, which expects three parameters:

File descriptor The file or device to write to. Broadly speaking, a file descriptor represents a file, and in the Unix world, everything is a file. In this case, we'll need file descriptor 1, which is always associated with `stdout`. We'll discuss file descriptors in more detail in Chapter 6. For now, you just need to know that if you pass 1 as the first parameter to `SYS_write`, you'll be writing to the console.

Buffer The memory address containing the data we want to write to the file descriptor provided in the first parameter.

Length The number of bytes we want to write, starting at the address provided in the second parameter.

Remember that you can check section 2 of the Unix manual to get all the details about any system call, using the command `man 2 syscall`. Just remove the `SYS_` from the syscall name.

On the `x86_64` architecture, the `SYS_write` system call has the identifier 1. Knowing this, and applying everything we learned in Chapter 2 about invoking system calls, Listing 3-1 shows what our little program looks like.

```
global _start
_start: mov rax, 1 ; SYS_write = 1
        mov rdi, 1 ; fd = 1
        mov rsi, msg ; buf = msg
        mov rdx, 14 ; count = 14 (the number of bytes to write)
        syscall ; (SYS_write = rax(1), fd = rdi, buf = rsi, count = rdx)

        ;; Exit program
        mov rax, 0x3c ; SYS_exit = 0x3c
        mov rdi, 0 ; status = 0
        syscall ; (SYS_exit = rax , status = rdi)

msg: db 'Hello, world!', 0x0a
```

Listing 3-1: A "Hello, world!" program in ASM for Intel `x86_64`

I hope you can identify the two system calls in there. The first one writes the message, and the second one exits the program with status 0.

Let's compile the program, following the same steps we used at the start of Chapter 2, then run it:

```
$ nasm -f elf64 -o hello.x86.o x86-hello.asm
$ x86_64-linux-gnu-ld -o hello.x86 hello.x86.o
$ qemu-x86_64-static ./hello.x86
Hello, world!
```

Once you've verified that it works, there are some key new elements in our program that we need to discuss.

Labels and Assembler Commands

The first key element in our program is the `msg` label. A *label* is a name used to reference a part of a program (actually, a memory address). We've already seen these at work in the past—for example, the `_start` label we used in Chapter 2. In this case, `msg` references our “Hello, world!” message. In general, we don't know where in memory our program will be loaded, and using symbolic names like `msg` lets us write programs without caring about those details. Wherever the data representing “Hello, world!” ends up in memory, `msg` will refer to that location. In this way, labels can make programs easier to read and modify.

The second key thing to highlight is the `db` instruction paired with the `msg` label. This is an *assembler directive*, an instruction that's understood only by the assembler and doesn't directly translate into an opcode in the program. We've also seen assembler directives before, such as `global`.

The `db` assembler directive, likely short for *define byte* or *data byte*, allows us to set some memory area with a sequence of bytes. In this case, our `db` instruction has two parts, separated by a comma. The first part is the string `'Hello, world!'`. The assembler will output one byte per character in the string, starting at position `msg`. The second part is an extra byte, expressed in hexadecimal: `0x0a`. Sure, you could just use the decimal value (10) instead, and everything would stay the same. This is the ASCII code for a line feed, meaning it will move the cursor to the next line in the console. Try removing it, and see what happens to your output. With `db`, you could also write your whole string as a list of the ASCII codes for each character, separated by commas, but that wouldn't be very practical.

Exercise Use the `SYS_unlink` system call (87) to write a program named *ghost* that will delete itself every time it's executed. For this exercise, assume the name of the program is always *ghost*.

Pointers

There's one more key feature of our “Hello, world!” program that we need to highlight: a pointer.

Okay, don't panic. You may have heard that pointers are the trickiest part of low-level programming, but they're actually a lot simpler than you might think. Furthermore, you can't do much in ASM without using them. Just repeat to yourself that you have to understand pointers to create cool things like shellcodes. Take a deep breath, repeat it again, and let's dive in.

A *pointer* is just a position in memory that contains the address of another position in memory. Remember, a memory address is just a number. The content of a given memory address is also just a number. It's therefore only natural that the number stored at a memory address could itself be interpreted as some other memory address.

One thing you may figure out from this recursive definition is that a pointer's size has to be equal to the number of address pins in the processor. In other words, it has to be the size of the address bus, which is usually the same as the processor's native word size. A 32-bit processor with a 32-bit address bus (for example, the Intel 386) will require 32 bits (4 bytes) to store any potential memory address a program could reference. A pointer on a 32-bit machine will therefore take up four consecutive positions in memory. (Remember, each position holds a single byte.) Likewise, for a 64-bit processor with a 64-bit address bus, a pointer will need 64 bits (8 bytes) to reference any possible memory address. A pointer on a 64-bit machine will therefore be stored in eight consecutive memory addresses. The entire range of memory addresses a processor can reference is known as the *addressing space*.

Let's look at an example. Suppose we have the memory layout shown in Figure 3-9.

```
| ... | p + 4 = 0x400004
+-----+
| 0x40 | p + 3 = 0x400003
+-----+
| 0x00 | p + 2 = 0x400002
+-----+
| 0x00 | p + 1 = 0x400001
+-----+
| 0x04 | p      = 0x400000
+-----+
```

Figure 3-9: A pointer in memory

On a 32-bit little-endian machine, the 4 bytes at memory addresses p (0x40000000) through $p + 3$ (0x40000003) can be read together as having the value 0x40000004. If we interpret this as a pointer, or a memory address storing another memory address, we can say that address 0x40000000 points at (stores) address 0x40000004.

What does all this have to do with our "Hello, world!" program? Taking into account that a register can be seen as a special kind of memory, whenever we store a memory address in a register, we can consider that

register to be a pointer. With that in mind, take a look at this line from our ASM code:

```
mov rsi, msg
```

Here, we store the value of `msg` in the `RSI` register. As we've already discussed, `msg` is just a label for the memory address where our message is stored. In this context, `RSI` contains a pointer to the message: It's a memory location (a register, in this case) storing another memory location that we're interested in.

In general, whenever we need to pass some information that's too big to be stored in a register to a system call, we need to store that information somewhere else in memory and give the system call a pointer (which can fit in a register) to the relevant memory address. In our example, the string "Hello, world!" requires 13 bytes, plus an extra byte for the line feed, and our registers can only store 8 bytes on a 64-bit platform. To make the system call work for all cases and not just for data blocks of 8 bytes, `sys_write` expects the data to write to be somewhere in memory, and it expects to receive that location (a pointer) as a parameter via a register.

Strictly speaking, `RSI` is the pointer here, but it's more common to refer to the memory address stored in `RSI` as the pointer. The indirection is implicit in the *pointer* definition, and it's really the target address that we're interested in. In other words, even though technically the pointer is the address containing an address, in general, people call the address itself the pointer.

Most of the time, pointers will be true memory addresses, not registers, but the concept is the same. The names of some registers actually reflect this role: The instruction *pointer*, the stack *pointer*, and the frame *pointer* are three of the most famous cases.

"Hello, World!" in C

Now let's see how to write the same "Hello, world!" program in C. We'll learn a few more things about the C language and its associated tools in the process. Listing 3-2 shows the code.

```
#include <unistd.h>

int main ()
{
    register void *p = "Hello, world!\n";
    write (1, p, 14);
    _exit (0);
}
```

Listing 3-2: A "Hello, world!" program in C

Again, notice the two system calls in the program, `write` and `_exit`. We already know that the second parameter to `write` has to be a pointer, a

memory address containing the address in memory of the string to print. We pass the variable `p`, which we declare with the register keyword. As mentioned in Chapter 1, you don't typically use this keyword in C programs, but we include it here to force the compiler to produce machine code that uses the processor's registers. Without it, the compiler will simply allocate a regular memory block of 8 bytes (for a 64-bit machine) in which it will store the pointer—the address where our message is located, in this case.

Let's take a look at the assembly generated by `gcc`. We'll create a non-position-independent executable, as that makes it easier to spot the pointer:

```
$ x86_64-linux-gnu-gcc -no-pie -fno-pic -o hello hello.c
$ x86_64-linux-gnu-objdump -d -M intel hello
--snip--
0000000000400544 <main>:
   400544: 55                push   rbp
   400545: 48 89 e5          mov    rbp, rsp
   400548: 53                push   rbx
   400549: 48 83 ec 08       sub   rsp, 0x8
   40054d: bb 5c 06 40 00    ❶ mov    ebx, 0x40065c
   400552: ba 0d 00 00 00    mov    edx, 0xd
   400557: 48 89 de          ❷ mov    rsi, rbx
   40055a: bf 01 00 00 00    mov    edi, 0x1
   40055f: e8 dc fe ff ff   call  400440 <write@plt>
   400564: bf 00 00 00 00    mov    edi, 0x0
   400569: e8 c2 fe ff ff   call  400430 <_exit@plt>
--snip--
```

Skipping over the first four instructions in the resulting ASM (that's stack stuff, which we'll discuss in Chapter 4), try to find the pointer. Have you spotted it? Look at how we set `EBX` (the 32-bit part of `RBX`) to `0x40065c` ❶, then copy `RBX` into `RSI` ❷. And what do you think `0x40065c` is? A memory address, of course! That once again makes `RSI` a pointer: a memory location holding the address of another memory location.

NOTE

The exact memory address you see being assigned to `EBX` in your output may vary, depending on your system configuration. Take note of that address, as you'll need it in a minute.

To find out what our pointer is pointing to, let's introduce a new friend, the GNU Debugger (`gdb`).

Using the GNU Debugger

A *debugger* like `gdb` is a program that allows us to execute other programs and look inside them. Do you remember in Chapter 2, when I said that the operating system makes sure that a process can't see inside the memory assigned to another process? Well, that's still correct. A debugger uses a special system call to do its work, so technically it's the OS looking into the other process's memory and reporting what it finds to the debugger.

A debugger is able to load a binary in memory and process all its metadata (the ELF details added during compiling and linking), including the code itself and any data stored in the program. It can then examine the contents of the memory associated with the program. Later, we'll learn how to figure out the address where the program data actually gets loaded, but for now we can just use the debugger to inspect the contents of arbitrary locations.

Technically, the debugger launches the program to debug as a new process, then attaches itself to that process using the `SYS_ptrace` syscall. This allows the debugger to control the target process and access its memory, as well as the values of its registers. In other words, the `SYS_ptrace` syscall is what enables debugging. Prior to this, the debugger can only read the program's file on disk and use the metadata stored in it to help us inspect its structure.

Let's use `gdb` to see what's at the memory address our pointer is pointing to (remember that the exact memory address you need to look at may vary):

```
$ gdb ./hello
(gdb) x/s 0x40065c
0x40065c:  "Hello, world!\n"
```

In `gdb`, the `x` command allows us to dump memory contents to the console. Here, it shows us that the memory location the pointer is referencing holds our "Hello, world!" message.

The `x` command must be followed by the type of data we want to see. In this case, we pass `s` for strings, but you can, for instance, pass `i` for processor instructions, `a` for addresses, or `t` for binary values. The command is a bit more powerful than this, but that's all you need to know for now. See the `gdb` manual if you want more details.

NOTE

If you aren't working on an `x86_64` machine, you may have a problem using `gdb` directly on an `x86_64` binary. See Appendix B for information on how to debug programs for other processors using `qemu` and `gdb-multiarch`.

Declaring Pointers in C

As you may have already guessed by looking at the code in Listing 3-2, pointers in C are declared using the `*` symbol. In our code, `*p` creates a pointer called `p` that references the memory location where the "Hello, world!" message is stored:

```
register void *p = "Hello, world!\n";
```

There's a little more to pointers than that, since in C we also need to specify the data type. In this case, the data type doesn't really make a difference, so we've used `void`. In most cases, however, the pointer type is important and useful. The standard syntax is:

```
type *pointer_name;
```

This declares a pointer called *pointer_name* to a memory address containing a value of a certain data type (*type*). Table 3-1 lists the basic types C provides.

Table 3-1: C Data Types

C type	Data type	Description
char	Byte	Minimal addressable element (not necessarily 8 bits)
int	Integer	Default integer type
short	Integer	Usually half the size of the default int
long	Integer	Usually double the size of the default int
float	Floating point	Single-precision floating point
double	Floating point	Double-precision floating point
void	Nothing	Nothing or anything

C also supports compound types and allows us to define our own new types, but for the time being, the ones in Table 3-1 will be enough for our purposes.

CHAR VS. BYTE

In C, a char is defined as a byte, and a byte is defined by the C standard as the size required to store a character on a given platform. That may sound circular, but it opens up the possibility that a byte could be something other than 8 bits. In fact, 6-bit bytes were common in mainframes and minicomputers during the 1960s and early 1970s, and some machines used 5-, 7-, or 9-bit bytes.

Nowadays, the size of the byte type is 8 bits (an *octet*) on nearly all systems. For the rare exceptions, C defines the CHAR_BIT constant, which indicates how many bits define a character for the platform in question. Digital signal processors (DSPs) are one type of system where you might find different byte sizes. For example, the Texas Instruments C54x and TMS320 DSPs define CHAR_BIT as 16 bits.

Listing 3-3 shows a simple program you can use to figure out the size of each type on your system and better understand the differences between those types.

```
#include <stdio.h>
int main ()
{
❶ printf ("Size of void*   : %ld\n", sizeof(void*));
   printf ("Size of short  : %ld\n", sizeof(short));
   printf ("Size of int    : %ld\n", sizeof(int));
   printf ("Size of long   : %ld\n", sizeof(long));
   printf ("Size of float  : %ld\n", sizeof(float));
   printf ("Size of double : %ld\n", sizeof(double));
```

```
    return 0;
}
```

Listing 3-3: Showing the sizes of the main C types in bytes

This program uses the `sizeof` operator, which returns the size, in bytes, of a given data type or variable. We display the results with the `printf` function, short for *print formatted* (see the “Format Strings and Type Conversion” box for more on how this works). Notice in particular how we check the size of `void*` ❶. This is a pointer to `void`, meaning a pointer to anything. That’s what we used in our “Hello, world!” program, and it’s the C equivalent of the pointer we used in our original ASM code. If you compile and run the test program in Listing 3-3, it will show you that `void*` has a size of 4 bytes on a 32-bit platform, or 8 bytes on a 64-bit platform.

FORMAT STRINGS AND TYPE CONVERSION

The `printf` function used in Listing 3-3 lets us print messages using *format strings* to compose complex outputs. In this case, we use the `%ld` format string to print the long value returned by `sizeof`. This basically tells `printf`, “I have a number here that I want you to convert into a string. Please do that for me.”

The idea of type conversion may be confusing if you’re used to high-level programming languages like JavaScript, where many transformations are done automatically, but it’s a necessary part of low-level programming. Anything printed to the console, even a number, needs to be provided as a string—that is, a sequence of characters. Internally, though, numbers and characters are represented differently, so some conversion needs to take place prior to printing.

Take the number 123. As an integer, it can be stored in just 1 byte, but to print it out you need to convert it to individual characters: 1, 2, and 3. For our purposes, C represents characters using ASCII codes, each of which requires 1 byte of storage. As a string, the number 123 therefore becomes the byte sequence `0x31`, `0x32`, `0x33`, `0x00`. That’s 4 bytes: one for each digit, plus the final `0x00` to indicate the end of the string. Luckily, the `%ld` format string handles that conversion for us.

There are lots of other `%` format codes you can use with `printf`. You can find details about the available options on the `printf` man page. Note also that `printf` is defined in `stdio.h`, which is why we include that file at the start of Listing 3-3.

Writing Your First Shellcode

A *shellcode*, in its simplest form, is a piece of code that fires up a shell. It’s usually fed into a vulnerable program using an exploit, effectively enabling the attacker to acquire a shell with the same privileges as that program. Attackers typically target processes running as root, to get full access to the machine.

We'll leave exploiting a vulnerable program to Chapter 4, but believe it or not, you've already learned everything you need to write a very basic shellcode. Just like a "Hello, world!" program, all it takes is a pointer and the right system call.

Our shellcode won't be usable in the wild, but it's still instructive to learn about it. If you're curious about why it's not usable, the problem is that shellcodes usually make it into memory through a standard C function that copies strings. As we've discussed, C strings are delimited with the value 0x00 (also known as the *null byte*), but that value will also appear as part of our shellcode, which will prevent the string copy function from copying the whole program. Once you understand how the shellcode works, you can read the classic article "Smashing the Stack for Fun and Profit" by Aleph One (*Phrack Magazine*, 1996) to learn more about this issue and some tricks for avoiding it.

In the coming sections, we'll look at how to construct a shellcode for each of our hardware platforms of interest, starting with x86_64.

x86_64

In Linux, you can execute a process (for example, starting a command shell) using the `SYS_execve` system call, which on the x86_64 platform has the number 59 (0x3b). This system call takes three parameters, but for now we need only the first one, a pointer to the name of the program to run. We'll set the other parameters to 0. Listing 3-4 shows our basic shellcode.

```
section .text
global _start

_start:
    mov rax, 0x3b    ; SYS_execve
    mov rdi, cmd     ; char *cmd
    mov rsi, 0       ; No argv
    mov rdx, 0       ; No env

    syscall

cmd:    db '/bin/sh', 0
```

Listing 3-4: A shellcode for Intel x86_64

Notice how similar this is to our "Hello, world!" program. We assign a label (`cmd`) to a string, but this time our string is `'/bin/sh'`, the name of the shell we want to launch, instead of a message to print. Then we create a pointer to this string in a register (`RDI`) and invoke a system call—this time `SYS_execve` instead of `SYS_write`. We no longer need the `SYS_exit` system call, since `SYS_execve` is another one of those system calls that doesn't give control back to the calling process. See its man page (`man 2 execve`) for more on this.

HOW PROCESSES ARE EXECUTED

When we run a program from the command line, we usually specify the name of the binary to run plus a set of arguments. Such a request doesn't go straight to the kernel but rather to the command line interpreter, which uses the kernel's services to get a new process running. This involves creating the process (we will see how to do that in Chapter 6), then loading the requested binary. This second step is done with the `SYS_execve` system call, which has the following libc wrapper:

```
int execve (const char *pathname, char *const _Nullable argv[],
            char *const _Nullable envp[]);
```

When the kernel runs the `SYS_execve` system call, it loads the binary in memory according to the ELF metadata and sets up a special memory area known as the *stack* (which will become your best friend in Chapter 4) with the last two parameters, which represent the command line arguments and the environment variables. Both of these parameters are in the command line interpreter's process, and only the kernel can copy those values into the address space of the new process.

When the new process starts execution at its entry point, the command line arguments and environment variables are made available to that process via the stack. From this point on, it's up to the process what to do with that information. For C programs, the C runtime code in `crt1.o` will process this information and provide it to the `main` function as three parameters using the proper ABI, so the C program can easily access it if needed. With that said, do you remember that I told you in Chapter 2 that there's a third way of declaring `main`? Here it is:

```
int main (int argc, char *argv[], char *envp[])
```

This style of declaration gives `main` access to the environment variables as well as the command line arguments. And now you also know where that third parameter comes from: Yep, the `SYS_execve` system call.

You may be wondering why the shell we launch is called `/bin/sh`. That doesn't match the name of any of the shells you may have heard of, like `bash`, `dash`, `zsh`, or `ksh`. And that's just it: There are so many different shells out there that just about any Unix system keeps a symbolic link to the preferred shell, whatever it may be, at `/bin/sh`.

The system has to run shell scripts all the time. For example, there are shell scripts that are executed during the boot process, whenever you start or stop a service, and when you launch some applications. If those scripts explicitly referenced a particular shell, imagine what would happen if the user wanted to change their default shell. For them to keep working, the system would have to update all of those scripts, and you'd have to update all the scripts you've written that the system knows nothing about. This would also force all system users to use the same shell. This is why most Unix systems

have a binary at `/bin/sh` that runs the user's preferred shell. Shell scripts rely on the existence of that file to work, regardless of the particular shell used.

ARM32

We should be able to port our x86_64 ASM shellcode to ARM32 very easily. Listing 3-5 shows what it looks like.

```
.text
.globl _start

_start:
    mov r7, #11    @ SYS_execve

    ldr r0, =cmd   @ Command
    mov r1, #0     @ No env
    mov r2, #0     @ No argv

    svc #0

cmd:
    .asciz "/system/bin/sh"
```

Listing 3-5: A shellcode for ARM32

This code is for Android, which deviates from the standard Linux folder structure in that the default shell is no longer at `/bin/sh` but instead at `/system/bin/sh`. If you're going to test the code on another ARM platform, such as a BeagleBone Black, a BananaPi, or an Olinuxino running a standard Linux distro (usually Debian), or if you're going to use `qemu` in your development container, just change the string to the well-known location `/bin/sh`. The rest of the code should work fine.

We can compile the program like this:

```
$ arm-linux-gnueabi-as -o arm_sc.o arm_sc.s
$ arm-linux-gnueabi-ld -o arm_sc arm_sc.o
$ arm-linux-gnueabi-readelf -S arm_sc | grep ".text"
[ 1] .text          PROGBITS          00010054 000054 000020 00 AX 0 0 4
```

This produces a shellcode of 32 bytes (0x20 in hexadecimal).

Let's take a closer look at the code. First, notice how we use `.asciz` to add the shell location to memory instead of `db`. This is a GNU Assembler instruction, and it's available independent of the processor. It's similar to `nasm`'s `db` instruction, but it automatically adds the requisite zero at the end of the string so we don't have to do that explicitly (that's what the `z` in `asciz` means).

Second, notice that we have to use `ldr` (load register) to load our pointer into register `r0`. When receiving a 32-bit immediate value (like the memory address referenced by `msg`), `ldr` is an ARM pseudo-instruction. The bottom line is that you can't directly load a 32-bit value into a register on an ARM

processor. I won't go into the details, but roughly, ARM produces 32-bit machine code for all instructions (there's another mode for 16 bits, but we won't talk about that for now). This limits the size of the values that can be directly loaded into a register by one of those 32-bit machine code instructions: We would need the whole 32 bits for the value, so there wouldn't be any bits left to indicate which instruction we want to use the value with.

To get around this, the assembler expands the `ldr` pseudo-instruction into the right sequence of actual instructions to load the 32-bit value in a register. You'll find a similar workaround in other RISC processors since, like ARM, they're designed to use a fixed size for instructions. CISC processors like Intel's `x86_64` allow a variable number of bytes per instruction, so they don't have this problem.

NOTE

It's possible to directly load small values (up to 12 bits) using `ldr`. The assembler will decide, depending on the value you want to load, whether to use the immediate value or treat `ldr` as a pseudo-instruction and expand it.

If you're curious, you can use `objdump` to see the actual machine code `ldr` produces:

```
$ arm-linux-gnueabi-objdump -d arm_sc

arm_sc:      file format elf32-littlearm
```

Disassembly of section `.text`:

```
00010074 <_start>:
   10074: e3a0700b      mov     r7, #11
   10078: e59f0008      ldr     r0, [pc, #8] ; 10088 <_start+0x14>
   1007c: e3a01000      mov     r1, #0
   10080: e3a02000      mov     r2, #0
   10084: ef000000      svc     0x00000000
   10088: 0002008c      .word  0x0002008c
```

The `ldr` instruction we wrote that loads a label (a 32-bit pointer) directly into the register is changed into an `ldr` instruction that uses indirect PC-relative addressing to get the pointer. The assembler will store the pointer to `msg` near the instruction (at address `0x10088` in this example) and then will load the register `r0` by reading the 32-bit value from that address based on an offset from the current address. Because the address is nearby, the offset value will be a small number—much smaller than the 32-bit number we want to load into the register, and small enough to fit as part of the 32-bit machine code (in our example, it's 8). You'll see this pattern all over the place in code for the ARM platform.

Exercise Use `gdb` to check that the pointer in the ARM code actually points to the shell path.

A Short Digression on Addressing Modes

Processors provide lots of different methods for accessing values and memory locations. These are known as *addressing modes*. So far, we've mostly been using *immediate addressing*, where we just provide literal values in the assembly code. Here are examples for x86_64 and ARM:

```
mov rax, 0 ; x86_64
```

```
mov r1, #0 ; ARM32
```

These load the literal value 0 into register RAX or r1, respectively.

Immediate addressing is the most basic addressing mode, but we need more than that. For now, I'll introduce three other modes. You'll learn about different ones later, as we need them.

When the second operand to `mov` is a register rather than a literal value, this is called *direct addressing* or *register direct addressing*. In effect, this copies the value from one register into another register.

The next often-used addressing mode is *indirect addressing*. This mode is used when we want to access the content of a memory address. In this case, the address is provided between parentheses or square brackets (it depends on the assembly syntax used). Indirect addressing allows us to work with pointers (memory positions that contain the address of another memory position). Here are some examples:

```
mov rax, (0x600200) ; Loads the content of address 0x600200 in RAX (x86_64)
```

```
mov r0, [0x20008c] ; Loads the content of address 0x2008c in r0 (ARM)
```

Here we specify the memory address literally, but it's more common to reference a register containing the desired memory address (or pointer), like so:

```
mov rdi, 0x600200
```

```
mov rax, (rdi) ; Loads the content of address 0x600200 in RAX (x86_64)
```

```
ldr r1, =0x2008c
```

```
mov r0, [r1] ; Loads the content of address 0x2008c in r0 (ARM)
```

Finally, certain processors support *PC-relative addressing*, where the address we want to access is referenced as an offset from the address stored in the program counter (PC). This is the mode ARM's `ldr` pseudo-instruction uses to load a 32-bit value into a register with a 32-bit opcode, as we discussed in the last section. Let's go back to that example:

```
10078: e59f0008 ldr    r0, [pc, #8] ; 10088 <_start+0x14>
```

```
--snip--
```

```
10088: 0002008c .word  0x0002008c
```

Remember that brackets indicate we're reading the content of a memory address, so the first line says, "Read the value stored at the address

located 8 bytes (#8) away from where the program counter is right now.” For ARM processors, “where the program counter is right now” is always 8 bytes beyond the address of the current instruction. In other words, since all ARM32 instructions are 32 bits, or 4 bytes, pc is pointing two instructions ahead. Most processors just point to the next instruction, but ARM was originally designed as a three-stage pipeline (fetch–decode–execute), similar to the one we saw in Chapter 1, and the designers decided that pc should point to the instruction in the fetch stage. This means that, when executing an instruction, we’ve already fetched the next two.

In our example, the instruction is at address 0x10078, so pc is at 0x10078 + 8, or 0x10080. The value we want is 8 bytes ahead of pc, and 0x10080 + 8 is 0x10088. Therefore, 0x10088 is where the actual pointer we want to load into r0 is stored (0x0002008c, in this case).

In ARM jargon, the list of values used together with the ldr instruction to load full 32-bit values is known as the *literal pool*, and it’s generated automatically by the assembler somewhere close to the code that needs access to those values. You usually won’t have to worry about this, as assemblers and compilers will take care of it for you, but it’s good to know how it works.

AArch64

The code for AArch64 is basically the same as for ARM32. We just have to use the correct registers and a different syscall number, as shown in Listing 3-6.

```
.text
.globl _start
_start:
    mov w8, #221
    mov x1, #0
    mov x2, #0
    ldr x0, =cmd
    svc #0

cmd:
    .asciz "/bin/bash"
```

Listing 3-6: A shellcode for AArch64

No surprises here. You can compile the program the usual way:

```
$ aarch64-linux-gnu-as -o aarch64_sc.o aarch64_sc.s
$ aarch64-linux-gnu-ld -o aarch64_sc aarch64_sc.o
$ aarch64-linux-gnu-readelf -S aarch64_sc | grep -A1 ".text"
[ 1] .text          PROGBITS          00000000000400078 00000078
          0000000000000028 0000000000000000 AX          0          0          8
```

This produces a shellcode of 40 bytes (0x28 in hexadecimal).

MIPS32

The shellcode for MIPS32 is also pretty similar to the ARM32 version. Listing 3-7 shows what it looks like.

```
.text
.globl __start
__start:
    li $2, 4011

    la $4, cmd
    move $5, $0
    and $6, $0, $0

    syscall

cmd:
    .asciz "/bin/bash"
```

Listing 3-7: A shellcode for MIPS32

We already know that MIPS uses different register names and syscall numbers, but there are some other details to highlight. First, as with ARM, MIPS opcodes are all 32 bits long, and therefore you can't directly encode an instruction that loads a 32-bit value. Also like ARM, MIPS provides pseudo-instructions as a workaround: `li` (load immediate) and `la` (load address). They load smaller, more manageable parts of the 32-bit value into different registers, then combine those parts together.

To illustrate, say we have the following `li` pseudo-instruction:

```
li $2, 0x11223344
```

The assembler will convert this to the following two instructions:

```
lui $1, 0x1122
ori $2, $1, 0x3344
```

The `lui` (load upper word immediate) instruction loads the upper word of register `$1` with `0x1122`, so `$1` will hold `0x11220000`. If you look back at Table 2-4, you'll see that register `$1` is reserved for pseudo-instructions. Next, the `ori` (OR immediate) instruction takes the bitwise OR of register `$1` and the value `0x3344`, storing the result in register `$2`. This operation compares the two values bit by bit, yielding a 1 for bits where either value (or both) has a 1. In the end, we get the 32-bit value we want, `0x11223344`, in register `$2`. The `la` pseudo-instruction is similar to `li`, but instead of using `ori`, it uses `addiu` (add immediate unsigned).

In our shellcode, we use `la` to store our pointer in register `$4`. Then we have to set the other two syscall parameters to 0. We could just use `li`, but instead we zero registers `$5` and `$6` using MIPS's special `$0` or `$zero` register, which is always set to 0. I've shown two ways of doing this: First, we simply use `move` to copy register `$zero` into register `$5`, then we perform an AND of

register \$zero and itself, storing the result in register \$6. (You could also XOR register \$6 with itself.) Remember that the \$zero register exists when you analyze MIPS code.

Just to get you comfortable with the compilation process, here's a reminder of the command line instructions to compile the MIPS version:

```
$ mips-linux-gnu-as -o mips_sc.o mips_sc.s
$ mips-linux-gnu-ld -o mips_sc mips_sc.o
$ mips-linux-gnu-readelf -S mips_sc | grep ".text"
[ 3] .text          PROGBITS          004000d0 0000d0 000020 00 AX 0 0 16
```

The MIPS shellcode is 32 bytes long (0x20 in hexadecimal).

MIPS64

The shellcode for MIPS64, shown in Listing 3-8, is almost identical to the MIPS32 version. Instead of using the register numbers, we use the register names, which greatly improves readability. The syscall number is also different.

```
.text
.globl __start
__start:
    dli $v0, 5057
    dla $a0, cmd
    dli $a1, 0
    dli $a2, 0

    syscall

cmd:
    .asciz "/bin/sh"
```

Listing 3-8: A shellcode for MIPS64

We've swapped li and la for dli and dla (the d stands for *double*) because now our registers can load up to 64 bits. These are also pseudo-instructions, and they work the same way as the others.

You can compile the program using these commands:

```
$ mips64-linux-gnuabi64-as -o mips64_sc.o mips64_sc.s
$ mips64-linux-gnuabi64-ld -o mips64_sc mips64_sc.o
$ mips64-linux-gnuabi64-readelf -S mips64_sc | grep -A1 ".text"
[ 3] .text          PROGBITS          00000001200000f0 000000f0
          0000000000000030 0000000000000000 AX      0 0 16
```

This produces a slightly bigger file than the 32-bit version: 48 bytes (0x30 in hexadecimal). Let's take a look at the resulting machine code:

```
$ mips64-linux-gnuabi64-objdump -D mips64_sc
--snip--
```

Disassembly of section .text:

```
00000001200000f0 <__start>:
    1200000f0: 240213c1      li    v0,5057
    1200000f4: 3c040000     lui   a0,0x0
    1200000f8: 3c012000     lui   at,0x2000
    1200000fc: 64840001     daddiu a0,a0,1
    120000100: 64210118     daddiu at,at,280
    120000104: 0004203c     dsll32 a0,a0,0x0
    120000108: 0081202d     daddu  a0,a0,at
    12000010c: 24050000     li    a1,0
    120000110: 24060000     li    a2,0
    120000114: 0000000c     syscall

0000000120000118 <cmd>:
    120000118: 2f62696e     sltiu v0,k1,26990
    12000011c: 2f736800     sltiu s3,k1,26624
```

In this case, `dli` is just translated to an `li` because the value is small enough to fit in the instruction. The `dla` pseudo-instruction gets expanded to a few operations, though. Let's figure out what's going on here.

The `cmd` string is stored at address `0x120010118`, as you can see in the `objdump` output. If you don't believe it, try the following:

```
$ xxd -r -p 2f62696e2f736800
/bin/sh
```

The `xxd` tool allows us to do hex dumps of data, but it can also do the opposite operation using `-r` (for *reverse*). The `-p` option shows the result immediately: Just type the hexadecimal sequence you want to convert and press ENTER, then press CTRL-C when you're done.

Here again is everything the `dla` pseudo-instruction actually does:

```
lui a0, 0x00      ; a0 = 0x0
lui at, 0x2000    ; at = 0x0000000020000000
daddiu a0, a0, 1  ; a0 = 0x0000000000000001
daddiu at, at, 280 ; at = 0x0000000020000118
dsll32 a0, a0, 0  ; a0 = 0x0000000010000000
daddu a0, a0, at  ; a0 = 0x0000000120000118
```

We'll break this down step by step:

1. As we already know, `lui` loads an upper word (16 bits). In the second line, it allows us to set bits 31–16 in the lower part of the 64-bit register `at`.
2. The `daddiu` instruction performs an unsigned addition with a given value (recall that the `i` is for *immediate*).
3. The `dsll32` instruction is a little tricky. On MIPS32, `sll` performs a logical left shift of a given number of bits (replacing the shifted bits

with zeros), but the 64-bit version performs a left shift of 32 plus the value of the third operand. In this case, since the third operand is 0, we shift left by 32 bits, meaning the lower 32 bits of the 64-bit a0 register are moved into the higher 32 bits of the register.

4. The `daddu` instruction is similar to `daddiu`, but it uses only registers and no immediate values.

After all that, we get the memory address we want, `0x120010118`, in register `a0`. Jumping through all these hoops is common for a RISC processor like MIPS. The number of available instructions is reduced (that's what RISC means), so it's normal to use several instructions to do relatively simple things. In general, you don't need to care about everything happening under the hood when you use a pseudo-instruction. However, if you're looking to save some bytes here and there, be aware that a single `dla` instruction actually takes up six 32-bit words, or 24 bytes.

As an extra step, try compiling the MIPS 32-bit code with the MIPS 64-bit compiler. You may get some warnings about not using the 64-bit versions of instructions (`daddu`, for example), but you'll get the same result in the end. MIPS was designed that way: MIPS32 code can be used directly on MIPS64 machines. The main difference between the architectures (leaving aside coprocessor and supervised mode) is that registers are 64 bits on MIPS64 and 32 bits on MIPS32. Instructions are even the same length on both (32 bits). As such, I won't include any more code for MIPS64 in this book, since it will be almost identical to the MIPS32 version. I'll just talk about MIPS32, which I'll refer to simply as MIPS from this point on.

RISC-V

The RISC-V version of our shellcode should appear quite straightforward now that we've gone through the ARM and MIPS versions. Listing 3-9 shows the code.

```
.text
.globl _start
_start:
    li a7, 221 # SYS_execve
    la a0, cmd
    li a1, 0
    li a2, 0

    ecall

cmd:
    .asciz "/bin/sh"
```

Listing 3-9: A shellcode for RISC-V

You can compile this program using the following commands:

```

$ riscv64-linux-gnu-as -o riscv64_sc.o riscv64_sc.s
$ riscv64-linux-gnu-ld -o riscv64_sc riscv64_sc.o
$ riscv64-linux-gnu-readelf -S riscv64_sc | grep -A1 ".text"
[ 1] .text          PROGBITS          000000000000100b0 000000b0
          0000000000000020 0000000000000000 AX          0          0          4

```

As usual, the code is almost identical to the MIPS version (it even has the same size of 32 bytes), and as with the ARM and MIPS versions, `la` is a pseudo-instruction. Let's take a quick look at the generated assembly:

```

$ riscv64-linux-gnu-objdump -D riscv64_sc

riscv64_sc:      file format elf64-littleriscv

```

Disassembly of section `.text`:

```

000000000000100b0 <_start>:
   100b0:  odd00893          li    a7,221
   100b4:  00000517          auipc a0,0x0
   100b8:  01450513          add  a0,a0,20 # 100c8 <cmd>
   100bc:  00000593          li    a1,0
   100c0:  00000613          li    a2,0
   100c4:  00000073          ecall

000000000000100c8 <cmd>:
   100c8:  6e69622f          .word 0x6e69622f
   100cc:  0068732f          .word 0x0068732f

```

The `li` instructions we put in the source code come through unchanged, but the `la` pseudo-instruction is substituted with a combination of `auipc` and `addi`. The first instruction, `auipc` (add upper immediate to program counter), is the way RISC-V implements PC-relative addressing. It sets the target register to the sum of the current program counter and an immediate value shifted 12 bits to the left. In other words, it adds the immediate value shifted 3 nibbles (3 hexadecimal digits) left:

```

0x100b4 auipc a0, 0x0    (pc = 0x000100b4)
                        a0 = pc + (0x0 << 12) = pc = 0x000100b4
0x100b8 add a0, a0,20    a0 = 0x000100b4 + 0x14 (20) = 0x000100c8

```

Note that when we reference data that's nearby, `auipc` will use `0x00` as a parameter, in which case it just becomes a way to load the value of `pc` into a register. The small offset (0x14, or 20, in this case) is then handled by the subsequent `add` instruction.

In general, the program counter points to the next instruction to be executed, but for RISC-V's `auipc`, it actually contains the address of the current instruction. That is, when we're about to execute `auipc`, the program counter

hasn't been incremented yet. Usually, we don't need to worry about these 4 bytes—the assembler will handle that for us—but it's good to know that there may be small differences between platforms when calculating offsets for different instructions.

Summary

In this chapter, you had your first encounter with pointers in ASM and C, and you learned how to use them together with a system call. You also had your first contact with addressing modes. Using these concepts, we wrote a simple “Hello, world!” program, which coincidentally was almost identical to a very basic shellcode. We took advantage of this to write our first shellcode too.