


The Hidden Language of Computer Hardware and Software



Charles Petzold

Code: The Hidden Language of Computer Hardware and Software

Charles Petzold

Copyright © 2009

Microsoft Press books are available through booksellers and distributors worldwide. For further information about international editions, contact your local Microsoft Corporation office or contact Microsoft Press International directly at fax (425) 936-7329. Visit our Web site at mspress.microsoft.com. Send comments to mspinput@microsoft.com.

Macintosh is a registered trademark of Apple Computer, Inc. Microsoft, MS-DOS, and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Other product and company names mentioned herein may be the trademarks of their respective owners.

Images of Charles Babbage, George Boole, Louis Braille, Herman Hollerith, Samuel Morse, and John von Neumann appear courtesy of Corbis Images and were modified for this book by Joel Panchot. The January 1975 cover of *Popular Electronics* is reprinted by permission of Ziff-Davis and the Ziff family. All other illustrations in the book were produced by Joel Panchot.

Unless otherwise noted, the example companies, organizations, products, people, and events depicted herein are fictitious. No association with any real company, organization, product, person, or event is intended or should be inferred.

Preface to the Paperback Edition

Code rattled around in my head for about a decade before I started writing it. As I was contemplating *Code* and then writing it, and even after the book was published, people would ask me, "What's the book about?"

I was always reluctant to answer this question. I'd mumble something about "a unique journey through the evolution of the digital technologies that define the modern age" and hope that would be sufficient.

But finally I had to admit it: "*Code* is a book about how computers work."

As I feared, the reactions weren't favorable. "Oh, I have a book like that," some people would say, to which my immediate response was, "No, no, no, you don't have a book like this one." I still think that's true. *Code* is not like other how-computers-work books. It doesn't have big color illustrations of disk drives with arrows showing how the data sweeps into the computer. *Code* has no drawings of trains carrying a cargo of zeros and ones. Metaphors and similes are wonderful literary devices but they do nothing but obscure the beauty of technology.

The other comment I heard was, "People don't want to know how computers work." And this I'm sure is true. I personally happen to enjoy learning how things work. But I also like to choose which things I learn about and which I do not. I'd be hard pressed to explain how my refrigerator works, for example.

Yet I often hear people ask questions that reveal a need to know something about the inner workings of personal computers. One such common question is, "What's the difference between storage and memory?"

That's certainly a critical question. The marketing of personal computers is based on such concepts. Even novice users are expected to know how many *megas* of the one thing and *gigas* of the other thing will be necessary for their particular applications. Novice users are also expected to master the concept of the computer "file" and to visualize how files are loaded from storage into memory and saved from memory back to storage.

The storage-and-memory question is usually answered with an analogy: "Memory is like the surface of your desk and storage is like the filing cabinet." That's not a bad answer as far as it goes. But I find it quite

unsatisfactory. It makes it sound as if computer architecture were patterned after an office. The truth is that the distinction between memory and storage is an artificial one and exists solely because we don't have a single storage medium that is both fast and vast as well as nonvolatile. What we know today as "von Neumann architecture"—the dominant computer architecture for over 50 years—is a direct result of this technical deficiency.

Here's another question that someone once asked me: "Why can't you run Macintosh programs under Windows?" My mouth opened to begin an answer when I realized that it involved many more technical issues than I'm sure my questioner was prepared to deal with in one sitting.

I want *Code* to be a book that makes you understand these things, not in some abstract way, but with a depth that just might even rival that of electrical engineers and programmers. I also hope that you might recognize the computer to be one of the crowning achievements of twentieth century technology and appreciate it as a beautiful thing in itself without metaphors and similes getting in the way.

Computers are constructed in a hierarchy, from transistors down at the bottom to the information displayed on our computer screens at the top. Moving up each level in the hierarchy—which is how *Code* is structured—is probably not as hard as most people might think. There is certainly a lot going on inside the modern computer, but it is a lot of very common and simple operations.

Although computers today are more complex than the computers of 25 years or 50 years ago, they are still fundamentally the same. That's what's so great about studying the history of technology: The further back in time you go, the simpler the technologies become. Thus it's possible to reach a point where it all makes relatively easy sense.

In *Code*, I went as far back as I could go. Astonishingly, I found that I could go back into the nineteenth century and use early telegraph equipment to show how computers are built. In theory at least, everything in the first 17 chapters of *Code* can be built entirely using simple electrical devices that have been around for over a century.

This use of antique technology gives *Code* a fairly nostalgic feel, I think. *Code* is a book that could never be titled *The Faster New Faster Thing or Business @ the Speed of a Digital Nervous System*. The "bit" isn't defined

until page 68; "byte" isn't defined until page 180. I don't mention transistors until page 142, and that's only in passing.

So, while *Code* goes fairly deep into the workings of the computer (few other books show how computer processors actually work, for example), the pace is fairly relaxed. Despite the depth, I tried to make the trip as comfortable as possible.

But without little drawings of trains carrying a cargo of zeros and ones.

Charles Petzold

August 16, 2000

Chapter 1. Best Friends

code (kōd) ...

- 3.a. A system of signals used to represent letters or numbers in transmitting messages.
- b. A system of symbols, letters, or words given certain arbitrary meanings, used for transmitting messages requiring secrecy or brevity.
4. A system of symbols and rules used to represent instructions to a computer...

—*The American Heritage Dictionary of the English Language*

You're 10 years old. Your best friend lives across the street. In fact, the windows of your bedrooms face each other. Every night, after your parents have declared bedtime at the usual indecently early hour, you still need to exchange thoughts, observations, secrets, gossip, jokes, and dreams. No one can blame you. After all, the impulse to communicate is one of the most human of traits.

While the lights are still on in your bedrooms, you and your best friend can wave to each other from the windows and, using broad gestures and rudimentary body language, convey a thought or two. But sophisticated transactions seem difficult. And once the parents have decreed "Lights out!" the situation seems hopeless.

How to communicate? The telephone perhaps? Do you have a telephone in your room at the age of 10? Even so, wherever the phone is you'll be overheard. If your family personal computer is hooked into a phone line, it might offer soundless help, but again, it's not in your room.

What you and your best friend *do* own, however, are flashlights. Everyone knows that flashlights were invented to let kids read books under the bed covers; flashlights also seem perfect for the job of communicating after dark. They're certainly quiet enough, and the light is highly directional and probably won't seep out under the bedroom door to alert your suspicious folks.

Can flashlights be made to speak? It's certainly worth a try. You learned how to write letters and words on paper in first grade, so transferring that knowledge to the flashlight seems reasonable. All you have to do is stand at your window and draw the letters with light. For an O, you turn on the flashlight, sweep a circle in the air, and turn off the switch. For an I, you make a vertical stroke. But, as you discover quickly, this method simply

doesn't work. As you watch your friend's flashlight making swoops and lines in the air, you find that it's too hard to assemble the multiple strokes together in your head. These swirls and slashes of light are not *precise* enough.

You once saw a movie in which a couple of sailors signaled to each other across the sea with blinking lights. In another movie, a spy wiggled a mirror to reflect the sunlight into a room where another spy lay captive. Maybe that's the solution. So you first devise a simple technique. Each letter of the alphabet corresponds to a series of flashlight blinks. An A is 1 blink, a B is 2 blinks, a C is 3 blinks, and so on to 26 blinks for Z. The word BAD is 2 blinks, 1 blink, and 4 blinks with little pauses between the letters so you won't mistake the 7 blinks for a G. You'll pause a bit longer between words.

This seems promising. The good news is that you no longer have to wave the flashlight in the air; all you have to do is point and click. The bad news is that one of the first messages you try to send ("How are you?") turns out to require a grand total of 131 blinks of light! Moreover, you forgot about punctuation, so you don't know how many blinks correspond to a question mark.

But you're close. Surely, you think, somebody must have faced this problem before, and you're absolutely right. With daylight and a trip to the library for research, you discover a marvelous invention known as Morse code. It's *exactly* what you've been looking for, even though you must now relearn how to "write" all the letters of the alphabet.

Here's the difference: In the system you invented, every letter of the alphabet is a certain number of blinks, from 1 blink for A to 26 blinks for Z. In Morse code, you have two kinds of blinks—short blinks and long blinks. This makes Morse code more complicated, of course, but in actual use it turns out to be much more efficient. The sentence "How are you?" now requires only 32 blinks (some short, some long) rather than 131, and that's *including* a code for the question mark.

When discussing how Morse code works, people don't talk about "short blinks" and "long blinks." Instead, they refer to "dots" and "dashes" because that's a convenient way of showing the codes on the printed page. In Morse code, every letter of the alphabet corresponds to a short series of dots and dashes, as you can see in the following table.

A	·—	J	·— — —	S	···
B	—···	K	—·—	T	—
C	—·—·	L	·—··	U	··—
D	—··	M	— —	V	···—
E	·	N	—·	W	·— —
F	··—·	O	— — —	X	—·— —
G	— — ·	P	·— — ·	Y	— · — —
H	····	Q	— — — —	Z	— — ··
I	··	R	·— ·		

Although Morse code has absolutely nothing to do with computers, becoming familiar with the nature of codes is an essential preliminary to achieving a deep understanding of the hidden languages and inner structures of computer hardware and software.

In this book, the word *code* usually means a system for transferring information among people and machines. In other words, a code lets you communicate. Sometimes we think of codes as secret. But most codes are not. Indeed, most codes must be well understood because they're the basis of human communication.

In the beginning of *One Hundred Years of Solitude*, Gabriel Garcia Marquez recalls a time when "the world was so recent that many things lacked names, and in order to indicate them it was necessary to point." The names that we assign to things usually seem arbitrary. There seems to be no reason why cats aren't called "dogs" and dogs aren't called "cats." You could say English vocabulary is a type of code.

The sounds we make with our mouths to form words are a code intelligible to anyone who can hear our voices and understands the language that we speak. We call this code "the spoken word," or "speech." We have other code for words on paper (or on stone, on wood, or in the air, say, via skywriting). This code appears as handwritten characters or printed in newspapers, magazines, and books. We call it "the written word," or "text." In many languages, a strong correspondence exists between speech and

text. In English, for example, letters and groups of letters correspond (more or less) to spoken sounds.

For people who can't hear or speak, another code has been devised to help in face-to-face communication. This is sign language, in which the hands and arms form movements and gestures that convey individual letters of words or whole words and concepts. For those who can't see, the written word can be replaced with Braille, which uses a system of raised dots that correspond to letters, groups of letters, and whole words. When spoken words must be transcribed into text very quickly, stenography or shorthand is useful.

We use a variety of different codes for communicating among ourselves because some codes are more convenient than others. For example, the code of the spoken word can't be stored on paper, so the code of the written word is used instead. Silently exchanging information across a distance in the dark isn't possible with speech or paper. Hence, Morse code is a convenient alternative. A code is useful if it serves a purpose that no other code can.

As we shall see, various types of codes are also used in computers to store and communicate numbers, sounds, music, pictures, and movies.

Computers can't deal with human codes directly because computers can't duplicate the ways in which human beings use their eyes, ears, mouths, and fingers. Yet one of the recent trends in computer technology has been to enable our desktop personal computers to capture, store, manipulate, and render all types of information used in human communication, be it visual (text and pictures), aural (spoken words, sounds, and music), or a combination of both (animations and movies). All of these types of information require their own codes, just as speech requires one set of human organs (mouths and ears) while writing and reading require others (hands and eyes).

Even the table of Morse code shown on page 4 is itself a code of sorts. The table shows that each letter is represented by a series of dots and dashes. Yet we can't actually send dots and dashes. Instead, the dots and dashes correspond to blinks.

When sending Morse code with a flashlight, you turn the flashlight switch on and off very quickly (a fast blink) for a dot. You leave the flashlight turned on somewhat longer (a slower on-off blink) for a dash. To send an A,

for example, you turn the flashlight on and off very quickly and then on and off at a lesser speed. You pause before sending the next character. By convention, the length of a dash should be about three times that of a dot. For example, if a dot is one second long, a dash is three seconds long. (In reality, Morse code is transmitted much faster than that.) The receiver sees the short blink and the long blink and knows it's an A.

Pauses between the dots and dashes of Morse code are crucial. When you send an A, for example, the flashlight should be off between the dot and the dash for a period of time equal to about one dot. (If the dot is one second long, the gap between dots and dashes is also a second.) Letters in the same word are separated by longer pauses equal to about the length of one dash (or three seconds if that's the length of a dash). For example, here's the Morse code for "hello," illustrating the pauses between the letters:

●●●● ● ●—●● ●—●● — — —

Words are separated by an off period of about two dashes (six seconds if a dash is three seconds long). Here's the code for "hi there":

●●●● ●● — ●●●● ● ●—● ●

The lengths of time that the flashlight remains on and off aren't fixed. They're all relative to the length of a dot, which depends on how fast the flashlight switch can be triggered and also how quickly a Morse code sender can remember the code for a particular letter. A fast sender's dash may be the same length as a slow sender's dot. This little problem could make reading a Morse code message tough, but after a letter or two, the receiver can usually figure out what's a dot and what's a dash.

At first, the definition of Morse code—and by *definition* I mean the correspondence of various sequences of dots and dashes to the letters of the alphabet—appears as random as the layout of a typewriter. On closer inspection, however, this is not entirely so. The simpler and shorter codes are assigned to the more frequently used letters of the alphabet, such as E and T. Scrabble players and *Wheel of Fortune* fans might notice this right away. The less common letters, such as Q and Z (which get you 10 points in Scrabble), have longer codes.

Almost everyone knows a little Morse code. Three dots, three dashes, and three dots represent SOS, the international distress signal. SOS isn't an abbreviation for anything—it's simply an easy-to-remember Morse code