Computer Science E-1
Spring 2010
Scribe Notes

Lecture 1: January 25, 2010
Andrew Sellergren

Contents

Computer Science E-1
Spring 2010
Scribe Notes

Lecture 1: January 25, 2010
Andrew Sellergren

1 Introduction (0:00–7:00)

- Welcome to E-1!
- 1 handout.
- 200 lollipops.
- Although we welcome students of all different backgrounds, chances are you're one of those "less comfortable" with computers. This, actually, is a great thing because it gives you the opportunity to learn so much in this course. Our goal is to make you one of those "more comfortable" with computers and to instill in you a sense of confidence. Of course, if you are already one of those "more comfortable," you'll still find plenty to keep you interested. For example, even though you're familiar with how to use DHCP, you might not know exactly how it works. This course will hopefully fill in many of the holes in your computer science knowledge.
- Your instructors this semester will be David Malan and Dan Armendariz. I, Andrew Sellergren, am your faithful scribe.
- To begin, take a look at Apple's newest gadget.

2 Binary (7:00–42:00)

- It may surprise you that the fundamental language of computers is *binary*, which consists exclusively of zeroes and ones. Amazingly, computers are able to play music, download movies, and write business reports only by understanding these two digits.
- If we consider a decimal number like 123, we recognize that the 3 is in the 1's column, the 2 is in the 10's column, and the 1 is in the 100's column. Each of these columns thus represents a power of 10. Likewise, in binary, we have columns which hold digits. In binary, however, each of these columns represents a power of 2. Let's start counting in binary and you'll see what we mean:

zero	0
one	1
two	10
three	11
four	100
five	101

  So the rightmost column is still the 1's column, but now the column second from the right is the 2's column. The column third from the right is the 4's column. You'll notice that 1, 2, and 4 are equivalent to 2^0, 2^1, and 2^2. In decimal, the column values 1, 10, and 100 were equivalent to 10^0, 10^1, and 10^2. The same logic applies to both binary and decimal.

- In the FoxTrot comic David showed you, Jason claims to have done 100 push-ups although his mom argues that he has only done 4. As you should now understand, the joke is that 100 is the binary representation of the decimal number 4. Take a look at this second FoxTrot comic which contains a hidden joke in binary.[1] In the second panel, FoxTrot plays on the difference between bits and bytes. A *bit* is a binary digit, i.e. a 0 or a 1. A *byte*, on the other hand, consists of 8 bits. Next is a kilobyte, which consists of 1024 bytes.[2] A megabyte is approximately 1 million bytes while a gigabyte is approximately 1 billion bytes. To contextualize these units of measure, consider that Microsoft Word documents are generally a few kilobytes in size, a music MP3 is generally a few megabytes in size, and a video file is generally a few gigabytes in size. Your hard drive might be a few hundred gigabytes in size or, if you're a geek like Dan or David, it might be a few terabytes (1 trillion bytes) in size. You might think a terabyte is an unreasonable amount of data, but these days you can fill it up fast. The video file of this lecture is a half a gigabyte alone!

- To put this into perspective, consider that only a few years ago, a floppy disk, which can hold 1.44 megabytes, was the standard for data storage. When David was a child, he used to back up his hard drive (about 20 megabytes in size) with floppy disks.

- Returning to the first panel of our FoxTrot cartoon, we see a few bytes written back to back. Each byte is simply a sequence of 8 bits. To represent the number 3 as a byte, for example, we would write 00000011. But these bytes are actually representing letters in the English alphabet. How do they do this? We already know that each byte can be translated into a decimal number. This number can then be translated into an English letter using ASCII, a mapping of numbers to characters. You can see the full ASCII mapping here. In this chart, we see that the number 65 maps to the uppercase character A.

- In the old days, computers displayed simple text messages on a black screen at startup. These messages were loaded up in memory as sequences of bytes which represented alphabetic characters.

- If we open up a binary translator such as this, we can copy the bytes from the FoxTrot comic and translate them one by one. When we do these, we reveal the hidden message "YOUNERD." Ha. Ha. Ha.

- Question: CPUs are generally either 32-bit or 64-bit. What might be the advantage of a 64-bit CPU? For one, it is faster because it can do operations on data in larger chunks. Second, it can do computations with much larger numbers, which is increasingly necessary these days.

[1] Please don't hesitate at any point to raise your hand and ask a question, particularly if David or Dan uses a piece of jargon or an acronym with which you're unfamiliar!

[2] Notice that all of these are powers of 2, which are conventionally used in measurement in computer systems.

- Question: what does it mean to rewrite software for a 64-bit CPU? We'll discuss this more in a few weeks, but for now know that it means feeding program code (written by humans in an English-like programming language) into a different compiler which spits out a series of 0's and 1's that can be understood by a 64-bit CPU.
- In the video David showed you, the narrator made reference to a circular cookie inside a floppy disk on which data is stored. When the disk is inserted in the drive, levers hold back the shield and read-write heads enclose themselves around the cookie, which spins just like a hard drive platter or a CD. In the bottom corner there is a write-protect tab which, if open, signals to the computer that data should not be written to the disk. As with a hard drive platter, data is written to the surface of a floppy disk cookie via a tiny electromagnet that polarizes the bits one direction or the opposite.
- Whew. If any of that was over your head, don't worry. The take-away is just that even data storage on floppy disks can be reduced to writing 0's and 1's. Each bit is actually a tiny magnet on the disk which when aligned in one direction represents a 0 and when aligned in the polar opposite direction represents a 1.
- Floppy disks get their name from their larger predecessors which were actually floppy. These predecessors had a smaller density meaning that fewer bits could be written in the same amount of physical space.
- Interestingly, we can remove the metal shutter from a floppy disk and pry apart its plastic coating to reveal the cookie inside.

3 Course Info (42:00–56:00)

- This course will consist of lectures, take-home assignments, two exams and two movie nights (popcorn courtesy of Johnny from Kendall Square Cinema)! There is also a final project which will charge you with buying and setting up your own domain name.
- You may have noticed in the syllabus the picture of MIT's hack whereby students hooked up a fire hydrant to a water fountain. Not pictured is the sign which read "Getting an education from MIT is like drinking from a fire hose." What they meant was that the deluge of information was often overwhelming. In this course, we'll similarly be presenting you with a deluge of information, but we don't want you to feel overwhelmed. The purpose is simply to expose you to many different facets of computer science so that you can understand them well enough to explore them on your own.

- You'll find the course website (http://www.computerscience1.net) to be an invaluable resource. It gives you access to the assignments, the lecture videos (and notes!), as well as the syllabus and the course wiki.
- There are no required books for this course, but if you do feel like you would benefit from some supplemental learning material, check out the syllabus for our recommendations.
- These first two lectures will cover hardware. In week 3, we'll have our first movie night, a screening of Pirates of Silicon Valley, a docudrama[3] that details the rivalry between Apple and Microsoft from its very beginning. Lectures 3 and 4 will cover the internet and lecture 5 will cover multimedia. Lectures 6 and 7 will cover security and lecture 8 will dive into web development. Lecture 9 will cover the basics of programming. In week 10, we'll have our second movie night with Startup.com, a documentary about a technology startup which grew quickly–too quickly, actually–and eventually busted. Finally, lecture 10 will be the exciting conclusion!
- Among other things, the course assignments will involve editing the course blog and course wiki. A wiki is simply a website whose content is editable by non-technical users. In other words, you can contribute to it without knowing how to program. Each week, David and Dan will post a number of placeholders or key terms like bit and byte which you the students will need to collaboratively define. In a similar vein, you will be tasked with adding articles about current events in technology to the course blog.

4 Hardware (56:00–101:00)

- We've already seen that a magnet's alignment in one direction can be used to represent 1 while its alignment in the polar opposite direction can be used to represent 0. Alternatively, we can use a transistor or a simple electrical switch which in the off position represents 0 and in the on position represents 1.
- When you first turn on a computer, you witness it executing a POST, or power-on self-test, which checks to make sure "everything's okay." It might be checking each bit in your RAM to make sure it's functioning properly. After these tests are passed, the computer begins loading up the bits of data stored on your hard drive.
- The central artery system of your computer is the *motherboard*. Most signals that pass between parts of your computer will travel through the motherboard. The brain of your computer is the CPU, which is attached to the motherboard. The CPU is the device which actually understands the 0's and 1's.

[3] Dramamentary?

- Generally speaking, there are two types of memory in your computer: the hard disk drive (HDD) and the random access memory (RAM). Whereas the HDD may be several hundred gigabytes in size, the RAM will probably only be a few gigabytes in size, if not smaller. RAM is used for temporary storage whereas HDD is used for persistent storage. RAM is generally faster and more expensive, as well. However, in order for you to interact with your data, it must be loaded up in RAM. Even if you have a Microsoft Word document stored on your hard drive, it must be loaded up in RAM in order for you to view it. Ultimately, whatever is loaded up in RAM will be passed to the CPU so that it can be interpreted.
- Actually, to be precise, data will pass through the L2 and L1 caches between the RAM and the CPU. These caches are similar in spirit to RAM, but they are faster, smaller in size, and more expensive.
- CPUs are described by their speed. You might have heard of a 2-gigahertz CPU. What this means is that the CPU can execute approximately 2 billion operations per second thanks to its oscillating crystal. A single operation might be an addition or subtraction of two bits, for example.
- Now we're going to take apart a computer! If you're trying this at home, please make sure to unplug your computer before you do so you don't injure yourself or your computer!
- Down at the bottom of the case, you'll see the motherboard. Perpendicular to that are a number of white expansion slots which allow you to add extra hardware to your computer. If the computer is an older model, you might see a simple audio cable connecting the sound card to the headphone jack. The ribbon-like cables are data cables—you'll probably find one attached to the hard drive. In older computers, this type of connector is called ATA or IDE. As you might expect, each separate device needs a data cable and a power cable connected to it. The power cables generally have red and yellow wires and terminate in pins.
- Knowing a bare minimum about computer hardware can pay dividends: last week, Dan's desktop computer was failing to start and was displaying wavy lines instead of the normal startup screen. With a little bit of investigation, Dan was able to determine that the video card had failed. Thus, he was able to pay a few hundred dollars for a replacement video card rather than a few thousand dollars for an entirely new computer.
- Somewhere close to the CPU you'll find an actual fan. The bulk of the CPU consists of multiple thin metal plates which constitute a heat sink. Both these mechanisms assist in dissipating the large amount of heat which the CPU generates as it oscillates rapidly. Some computers use liquid coolant systems in lieu of a fan and heat sinks.
- We don't recommend removing and replacing your CPU while trying to upgrade your computer. Chances are if you're CPU is outdated, then

most of your other hardware is as well and you'll be better off replacing the whole computer.

- On the motherboard, you'll notice a small circular battery. This is used to power the internal clock which keeps track of the time when you've unplugged the computer.
- Near the CPU are the RAM slots. RAM comes in paired sticks about 5 inches in length. RAM is actually very quick (and relatively inexpensive) to install. It is perhaps the easiest upgrade you can perform on your computer. Do your research, though, to make sure you have the right type. Take a look at crucial or dealram which provide tools to help you find the right RAM for your computer.
- Whenever you're handling computer equipment, you should grasp it by its edges. Make sure you're not on carpet and that you touch one of the metal parts of the computer tower to ground yourself. All of this will prevent static electricity from shorting out your hardware.
- Of course, you don't **have** to open up your computer and tinker with it, but if you find it's no longer meeting your needs, you might save yourself some money if you have the basic skills to install new hardware.
- To contextualize this discussion, let's ask a simple question: why might you want more RAM for your computer? It "speeds" up your computer in the sense that you can run more programs simultaneously without experiencing sluggishness. When your computer runs out of RAM, it starts to put some of your programs in "virtual memory" which is actually space on your hard drive. If you have ever tried switching between programs and noticed that the background program takes an inordinate amount of time to switch to the front, it's probably because it was moved to virtual memory.
- Challenge: try turning off your computer and unplugging all the cables from the back. Can you plug them back in correctly? Some manufacturers are nice enough to color code the connectors. One connector you might be familiar with is USB which is often used to connect cameras, printers, mice, flash drives, and iPods.
- As a teaser for next week, here's a little more about hard drives: hard drives consist of circular platters which spin as data is being written to and read from them. When data is to be stored on the hard drive, it passes from RAM along with software signals that designate how the data is to be stored. Certain signals control how the platters spin and others control the read-write heads. The distance between the heads and platters is less than the width of a human hair, yet the platters spin 5400 RPM or faster. Multiple platters make for greater efficiency than a single platter.