CS 107 Fall 2006

Lecture 2: Software Systems

1 During one lifetime

What is a computer? An electronic digital machine that can accept input, process it, and produce useful information. The computer has a controller, a logic and calculation unit, one or more memory areas, and a bus to connect them all together. It has input and output devices. Two kinds of things are stored in the memory – a program (instructions on what to do) and the data (on which to perform the instructions).

The first, limited computer was not fully complete in 1942 when I was born.

> The Atanasoff-Berry Computer (ABC) was the first electronic digital computing device[1]. The machine, conceived in 1937, was capable of solving up to 29 simultaneous linear equations and was successfully tested, though its input/output mechanism was still unreliable in 1942 when its inventors left Iowa State College for World War II assignments. The ABC pioneered important elements of modern computing, including binary arithmetic and electronic switching elements[2], but its special-purpose nature and lack of a changeable, stored program distinguish it from modern computers. [–Wikipedia]

Today. Today computers are involved in almost every aspect of life. None of the things mentioned below existed when my children were born. Consider my family, and the ways we use and rely on them.

1. Publishing: I relied totally on my computers (Commodore, Atari, Next, and now Mac) to write my Ph.D dissertation and publish two textbooks. Words, drawings, and graphs were all done on whatever was my current system. The publisher took my work and used a computer-based system to typeset it.

2. Government: My husband used and uses the computer to communicate (email, website) to organize and operate True Vote Connecticut, an organization that has successfully forced our state government to purchase a cost-effective and more secure voting system instead of the very expensive, fragile, and insecure touch-screen system that was selected 18 months ago.

3. Education: My son Ted uses a website to help run his tutoring business, as my husband and I also do with our classes. And of course, we use spreadsheets for calculating grades.

4. Defense: When Ted was in graduate school, he worked for the Navy one summer. His job was to write a computer program that would fly in an unmanned airplane flying over the ocean, and scanning the water with a video camera. The program was supposed to recognize any ships in the water (any object more than 50 feet long) and radio back its position. Images had to be analyzed as fast as they arrived.

5. Medicine: Ted's wife Teresa does research in the field of public health. She relies heavily on an application package that searches databases and produces statistics. An ultrasound echogram was used early in her pregnancy to verify the size, position, and condition of the fetus. (Paul is 1 week old now, and healthy.)

6. Commerce: I ordered a gift for the baby today, online. It was a hard-to-find item and the internet helped immensely.

7. Entertainment: Ted's 3-year-old son, Richard, can operate the CD player to listen to recordings.

8. Finance: My son Robert works doing arbitrage for a bank on on Wall Street. He analyzes vast numbers of stock sales looking for patterns that can be exploited to make money. This work uses mathematics and databases.

9. Religion: Robert's wife will soon be an Episcopal priest. During her education, she relied on computer applications to help her learn Hebrew and Greek and do her homework in those languages.

10. Business: My son David owns two New Haven businesses. One company designs and builds websites, the other company installs computers, networks, and security systems.

11. Life: Like most teens and 20-somethings, my foster-daughter, Trina, would be lost without her cell phone.

These are just some of the ways that one family uses computers constantly.

2 Computer Systems and Applications

To use a computer, somebody must write (and debug) a program. A program is a set of instructions about how to do a task, step by tiny step. (Writing programs is a science and an art that we will discuss later in the term.) There are several major kinds of programs:

1. Operating systems: These programs form an interface between the computer and the people who use it. Today's popular systems include Windows XP, Linux, and OS-X.

2. General-purpose applications: Programs such as Netscape, Oracle, Excel, Money, and PowerPoint are used by millions of people all over the world.

3. Specialized applications: Schools, libraries, hospitals, businesses, and governments all have specialized programs to handle their records and business processes. Such programs are sometimes created for a company from the ground up, and sometimes built on existing general-purpose applications. SAM and ProGrader, which will be used in this class, are specialized applications.

4. Translators: Programs are written in many different computer languages; each language has its strong and weak points, its defenders and its detractors. Every computer language has a translator that converts the programs written by a programmer in a semi-human language and translates it to a series of bits and bytes a computer (and only a computer) can understand. Every different combination of computer hardware + operating system must have its own specialized translators.

5. Libraries: Every computer language has a library consisting of hundreds of functions, large and small, that do commonly-useful things. Putting a function into a library makes it available for other programmers to use. Many of these libraries are specialized, and different languages support different specialties. Often, the presence or absence of the library you need is the biggest difference between one language and another.

3 Operating systems.

An operating system has three main responsibilities:

- To communicate with the human operator. Modern systems use graphic interfaces (desktops) to do that; early systems used coded patterns of lights on the computer. (Remember, there were no monitors for output.)
- To manage the internal parts and systems of the computer in such a way that the human gets the necessary work done.
- To manage the interaction of the computer with the network, including email and browser activity. To protect the local computer from hostile attacks coming from the network.

None of these responsibilities is easy to handle, and it has taken us 50 years since the first systems were written to reach a level where most people can use a computer productively.

Major subsystems and management tasks.

- To make the input and output devices work and bring data to or from them.
- To handle the file system and control access to it. Today most of our files live on a hard disk and we organize them using directories (folders). The file system must keep track of the files themselves and all the data about each file (name, owner, type, size, date). Passwords and file permissions are used to restrict access to the file's owner and people he permits to share the file.
- To start up and run processes (applications), keep track of them and their windows, and shut them down at the right time and in the right way. To change the active application every time a user clicks in a different window. To run some applications in the background so that they continue to make progress while a user is working in another window.

  For example, using Linux or OS-X I can remotely log into any machine on which I have an account. When I log into that machine, a process is created for me and it runs invisibly (in the background) along with the processes of the human that is sitting in front of the computer.

  Each time you browse into a web server, a process is created for you and put in a queue with all the other user's processes. The web server must handle all of the remote requests in turn. When there are too many requests at the same time, the web server becomes sluggish. This has been exploited by some hacker groups to "shut down" websites that they find offensive (such as kiddie porn websites). A hacker breaks into many machines simultaneously (the owners of the machines probably do not know this is happening.) On each machine, the hacker starts up a process that watches the computer's clock. At the specified time and date, all these processes simultaneously attack, that is, access the target website. The target becomes overwhelmed because it cannot handle all those requests, and all outgoing web service ends.
- To allocate and manage memory so that the programs that are running always have enough. Today's computers have a very large memory area, but can still run short. Then the operating system will try to move part of the program or the data from the memory to the hard disk, and bring it back into the memory the next time it is needed. This is called swapping. When the number of demands on the system becomes too great, it spends all its time swapping and does not get any work done! We call this condition "thrashing".

Did you ever notice that response to your mouse movements was getting very sluggish? That was probably caused by thrashing. [demonstrate by launching too many applications.] How do you deal with this? Close some windows.

3.1 So, what are the differences?

There are two common kinds of modern operating systems, servers and clients. A company or university will have at least one server, and probably many. It is common to have a web server, a database server, a file server, etc. Individual users need a client system to interact with these servers.

Windows is a client system. Windows servers do exist, but that is not what you have on your PC. Your system cannot do all of the things a server can do. For example, if your computer is at home, and you want a file that is stored on it, you cannot just log into it remotely and browse through your files.

Windows systems for many years have been produced under market pressure by huge numbers of programmers working without coordination and without good shared libraries. Microsoft prefers to hire smart people, with or without proper professional training, and let them "play". For example, a physics Ph.D. from Yale was hired to design and program while his wife, a Computer Science M.S. from UNH was hired to do a lesser job. The result for many years has been unstable systems that are clumsy to use. They are full of flashy features, and also full of security holes and difficult to administer.

Linux is both a server and a client. When using Linux, you can have one window open on the local system and another window open on a remote system. Moving information back and forth is easy. For example, one year I wrote an exam in Michigan, sent it to my system at home, and had my son print it out and deliver it to school.

Linux is the product of thousands of people worldwide, working as part of the Free Software movement, adding parts and fixing problems. The result is a highly stable system with excellent security features. Well-developed user-friendly GUI desktops have been available for a few years now. All major kinds of general-purpose software are available for free. (This is why Microsoft fears Linux.)

OS-X (Mac OS) is like Linux: both are Unix-based systems, and both supply client and server capabilities. Produced by Apple on the basis of much prior work, OS-X is not free. It is stable, secure, easy to use, and nice to look at. Large amounts of free software are available for it.