A Concise History of Computers, Smartphones and the Internet

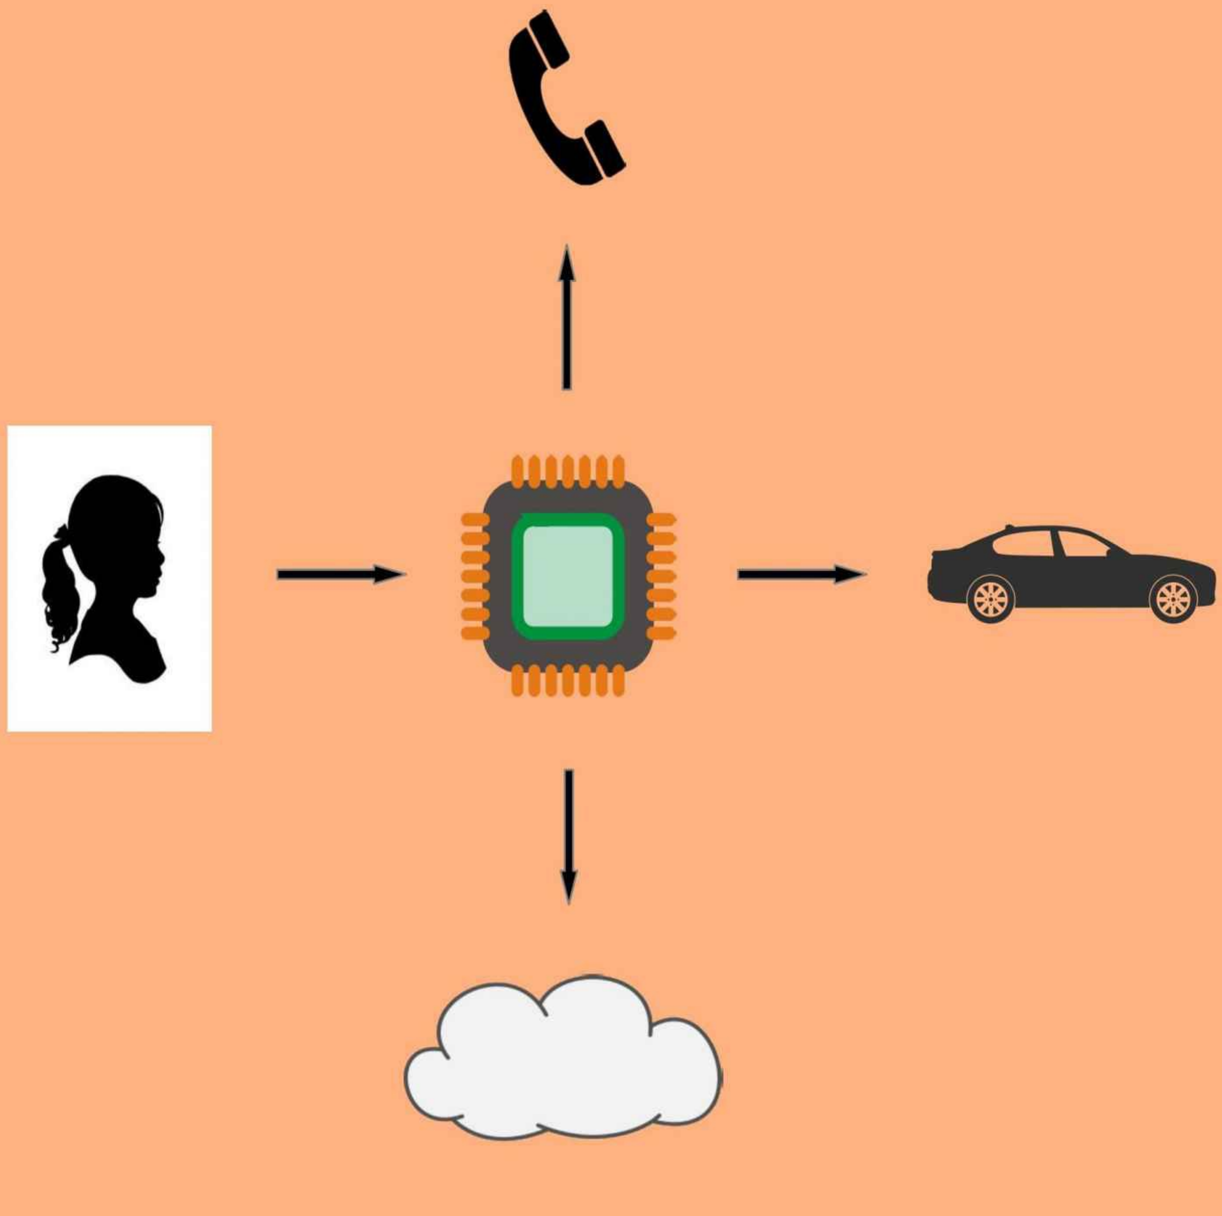

Who is this book for?

The very first electronic computers were invented at the end of World War II. They were very large machines that could only be used in special air conditioned rooms. Today, almost everybody carries a computer in their pocket, in their mobile phone. How did all this come about in only 70 years?

This book is for people who would like to know the answer to this question. It tells this exciting story, with a lot of pictures. It is not a complete history — that would be a very large and detailed book. Rather it is a concise history that in 40 pages covers the most important people, companies and inventions that led to where we are today.

The first chapter covers the evolution of computer hardware — the physical machine. The second chapter focuses on the software — the programs that provide the instructions that tell the hardware what to do. The third chapter covers the most important data networks that were developed so that computers could communicate with each other, ending with the Internet which only became the dominant computer network after 1995. The last chapter on Smartphones traces the history of mobile phones from the discovery of radio waves in the late 19th century to the Apple iPhone.

This book does not require a lot of technical knowledge about computers. People who are interested in learning more about how computers actually work can read the book "Understanding Computers Smartphones and the Internet", by Ernie Dainow.

Ernie Dainow
edainow@gmail.com
Toronto, Canada
February 2017

ISBN: 978-0-9952144-1-5

Table of Contents

About the Author

1. History of Computers

The invention of the computer cannot be credited to any one person. Many people contributed to the evolution of the computer over many years. As with any new technology, there was a burst of inventions and new companies that make up the history of computers, many of which were important for a period of time and then were overtaken by a newer technology. In this book, only the most important innovations that led to where we are today are mentioned. This is only a small number of the total number of companies and people who are part of the complete history.

There has been a major revolution that introduced a new computer era on average every 14 years. The following is an outline of this evolution that this book follows.

1936 Early Computers — a simple theoretical machine was proposed by Alan Turing. While various calculators existed, this was the theoretical groundwork for a general purpose computing machine. The first working computers were developed during World War II and the years following in government research labs and in universities.

1951 First Commercial Computers — The first general purpose computers became available from many different companies in the U.S., Europe and Japan.

1964 IBM Mainframe Era — the IBM System/360 family of compatible computers allowed customers to start with a small system and upgrade to larger computers as their computing needs grew. The success of System/360 drove many computer vendors out of the market and established IBM as a dominating force in the computer industry for the next 25 years.

1981 Microcomputers — a revolution in widespread use of microcomputers started with the IBM Personal Computer (PC) that was released in 1981, even though earlier microcomputers had been popular, such as the Apple I in 1976.

1995 Internet — the Internet was first built in 1969 as a research project under a grant from the U.S. Department of Defense, but widespread use of the Internet did not occur until after the World Wide Web (www) was developed in 1991 and Microsoft Windows 95 made it easy for non-technical users to access the Internet with simple point and click.

2007 Smartphones — the first smartphone was developed by IBM in 1994 and was followed by the BlackBerry in 2002. But the general Smartphone era really began with the revolutionary Apple iPhone in 2007 which led to widespread mobile computing in the general consumer market.

1936 Early Computers

A key early influential idea was published in 1936. Two years before completing his PhD, Alan Turing wrote a paper on an important theoretical question in mathematics, showing that it was not possible to decide if a statement was true or not given a set of axioms and rules of logic. This was the “decision problem”, posed by David Hilbert, one of the great mathematicians of the time. To establish his proof, Turing described a theoretical machine that could manipulate symbols according to a set of instructions. This machine could move a tape forward or backward and write 0 or 1 on the tape depending on what it read on the tape. Turing proved that such a machine could compute the solution to any problem, except those that did not have a solution. This theoretical machine became known as a “Turing machine”.

There were many mechanical calculators (such as adding machines) in use by the end of the 19^{th} century and the idea that it was possible to build a general purpose calculator that would work for more than one type of calculation and was as simple as a Turing machine was revolutionary.

Other important ideas came from Claude Shannon. In 1938 he published his master’s thesis which showed that electrical switches could be used to implement logic and arithmetic. The paper included a circuit design that could add binary numbers and a circuit design that could be used to make a decision based on comparing two values. This became the foundation of digital circuit design. Shannon went on to a distinguished career at Bell Labs and MIT. In 1948, Shannon published a paper “A Mathematical Theory of Communication” which became the basis of information theory.

Turing’s interest in computation led to him being recruited by the British Secret Service during World War II. He worked on a team tasked with deciphering the Enigma code used by the Germans to encrypt their radio transmissions. This dramatic story has been recounted numerous times such as in the play “Breaking the Code” shown by the BBC in 1996 and the film “The Imitation Game” in 2015.

Turing was instrumental in designing and building a machine that sped up the computations that allowed them to break Enigma. However, this machine was not really a computer but more of a specialized calculator that simulated an Enigma coding machine.

There was another Secret Service team in Britain working on breaking encryption of other German devices. They did design and build a machine that could properly be called a computer in 1944. This machine, the Colossus, was one of the first computers ever built.

In 1945 John von Neumann published a detailed design for a stored program computer. Von Neumann was a mathematician who worked out many important calculations for the Manhattan project, the U.S. development of the atomic bomb. He was greatly influenced by the Turing machine concept but he needed a detailed and practical blueprint to develop a working computer that could be built with the technology available at the time. The fundamental design he laid out became the model for all computers.

Computers today are generally called “von Neumann” architecture machines. There have been ideas for other architectures, such as optical computers and quantum computers and research is being done in these areas, but it will be many years before von Neumann machines are supplanted.

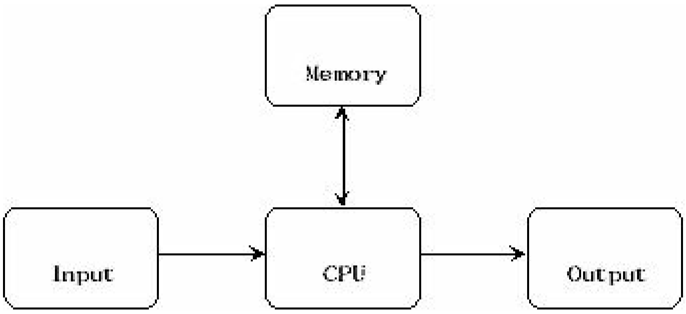

Von Neumann Computer Model

Programs and data are read from an Input device (such as a card reader) into the Central Processing Unit (CPU). Program instructions are processed line by line in sequence. The instructions are able to move data between the CPU and Memory, do computations on the data and send the results to an Output device (such as a printer).

After the war, work continued largely at universities and government projects to build larger and faster computers. At the Institute for Advanced Study (IAS) in Princeton, New Jersey, von Neumann designed and built the IAS computer. It became operational in 1951 and influenced the design of many other computers all over the world.

1951 First Commercial Computers

Many people saw the potential for computers to be used as general scientific and business machines. In 1951, the first commercial computers were offered for sale — the Ferranti Mark 1 based on computers built at the University of Manchester and the Univac I based on the ENIAC built at the University of Pennsylvania.

Univac I Computer, 1951

In 1952, IBM followed with its first electronic computer, the IBM 701, which was heavily influenced by consultations with von Neumann and his IAS design.

This was a growing and competitive market. Many large companies entered the computer business, such as General Electric, RCA and Honeywell in the U.S., ICL in the U.K., Olivetti in Italy, Groupe Bull in France and Fujitsu in Japan.

Most of the early computers were oriented to scientific programming. But IBM's major customers were businesses who used IBM tabulating equipment that could process punched cards, for such purposes as bookkeeping, accounting and printing invoices. In 1953, IBM released the IBM 650 computer which was oriented to business customers. By 1960 the IBM 650 was the most widely sold computer on the market and it established IBM as the leading computer vendor.

These early computers made use of existing components that were available. The instructions in the CPU are built with switches. In electronics at the time, switches were provided by vacuum tubes. Thomas Edison had patented the first vacuum tube design in 1883 for use in his light bulb invention. However it was only later designs by Lee De Forest in 1906 that revolutionized the electronics industry and led to great improvements in telephone systems and radio.

Vacuum Tubes

Use of vacuum tube in computers had many limitations. Even the small computers of this era required many vacuum tubes (over 5000 in the Univac I). Vacuum tubes were quite large, required a lot of electricity and generated a lot of heat. Consequently they burned out quite frequently and the computer was rendered non-operational until the burned out tube was identified and replaced.

There were no existing devices for memory on the earliest computers so various inventions were used on these early computers. These memories were slow and very limited in size. The Univac I had 12 KB (12,000 bytes) of memory based on a mercury delay line, in which a current was passed through a tube of mercury. The time it took the current to go from the beginning to the end of the tube was the memory cycle, after which current had to be applied again.

For input and output, one option was punched cards used by IBM and other companies that made tabulating and sorting machines. Many companies already had such equipment.

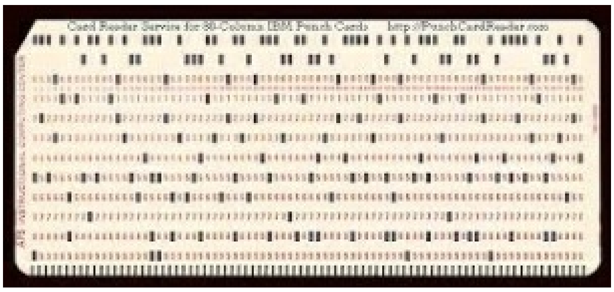

Punched Card

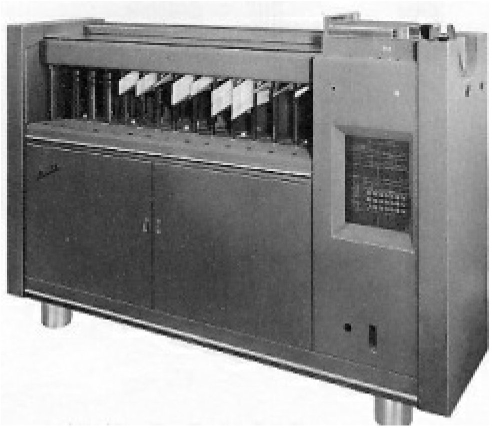

Card Sorting Machine

Another option was paper tape that was used in telegraph equipment, such as a teletypewriter, to send and receive messages by Morse code.

Teletype Model 33

Paper Tape

A more expensive option was magnetic tape to provide greater speed and volume. Magnetic tape had revolutionized the radio broadcasting and gramophone recording industries. For use with a computer, a special tape reader needed to be designed and built (as shown in the background of the picture of the Univac I above).

Another device adapted for computers was the electric typewriter. This provided a basic, but slow, printer for hard copy output (as shown on the right in the Univac I picture).

The larger computers used magnetic tape for input, output and storage. Although tapes could be read at high speed, if a program needed data at the beginning of the tape, there was a delay while the tape was rewound. There were many inventions of "random access" storage devices to solve the rewind problem with tapes. The most successful was magnetic disks with read/write arms that could quickly move to any circular ring on the spinning disk.

IBM 5 MB disk drive 1956

The first commercial disk drive was made by IBM in 1956. It weighed over a ton and stored an “astonishing” 5 MB (5 million bytes) of data. Compare this to a simple USB drive that you can buy today at any electronics store that has 128 GB of storage. This is 25,000 times as much storage and about 500,000 times smaller!

A critical invention in this period was “magnetic core” memory. It consisted of an array of magnetic rings that could be magnetized in one direction for 0 and in the other direction for 1. Sensing wires could read the value of each ring. Core memory was much faster and could be built in much larger sizes than the memory used in earlier computers. This was an important factor that led to rapid advances in computers following the first successful use of core memory in 1953 when it was installed in the Whirlwind computer that was built by MIT for the U.S. Navy.

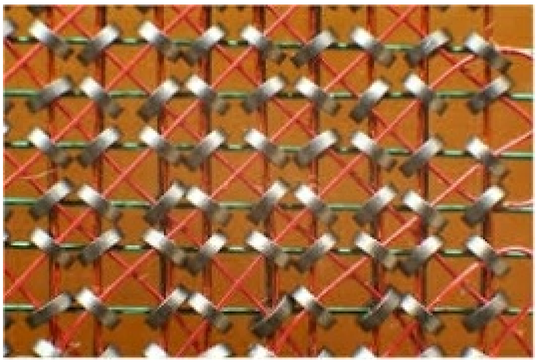

Magnetic Core Memory

Several different inventors worked on various improvements to core memory and received patents. One of these inventors was An Wang who was working at Harvard University's Computation Laboratory but later co-founded Wang Laboratories. Wang Laboratories was a successful company that made early dedicated calculators, then word processing computers and later minicomputers.

1956 The Transistor

In 1947 a group of researchers at Bell Labs in Princeton, NJ demonstrated a “transistor” that could effectively replace a vacuum tube. It was an electronic component that was smaller than vacuum tubes and used less power. Transistors revolutionized the electronic industry and many consider it to be one of the greatest inventions of the 20th century. John Bardeen, Walter Brattain, and William Shockley of Bell Labs won the Nobel Prize in Physics in 1956 for the invention of the transistor.

Early Transistors

It took a number of years for the transistor to be refined enough to be used commercially. One important improvement was changing the material from germanium to silicon. William Shockley left Bell Labs and in 1956 founded the Shockley Semiconductor Laboratory in Mountain View, California. This was the beginning of the electronics industry that became known as "Silicon Valley".

In a semiconductor material like silicon, current does not flow. But if the silicon is combined with other elements, the electrical balance can be changed to a positive or negative bias so that current does flow. This behavior is used to create a transistor which can be used to amplify a signal or to operate as an electronic switch.

The following schematic diagram shows a transistor being used as a switch. In the first case, there is no current through **In** and the switch is off. In the second case, current applied at **In** turns the switch on, and current flows **Out** of the switch.

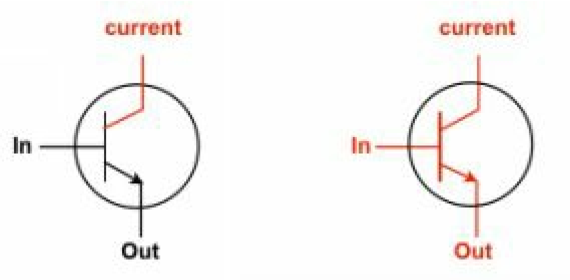

Transistor switches in turn could be used to build logic gates, such as AND, OR, NOT, that Claude Shannon had demonstrated was the basis of circuit design and could be used to build the functions of a computer.

By the mid 1950s transistors were being used in military products, telephone exchanges, portable radios and a few computers. By 1960 it had replaced vacuum tubes in almost all computers. This solved many of the problems of the early computers, bringing greater reliability with reduced size and cost. The use of transistors also led to improvements of many of the peripheral devices needed by computers.

1964 Integrated Circuits

While the transistor greatly reduced the size of electronic components, a circuit board still needed many components to be assembled. The following shows a circuit board containing various electronic components, as well as some black silicon chips.

Circuit Board

In 1958, Jack Kilby at Texas Instruments demonstrated an “integrated circuit” that combined many components on a single chip. He won the 2000 Nobel Prize in Physics for the invention of the integrated circuit.

Six months later, Robert Noyce at Fairchild Semiconductor independently developed an improved design based on silicon. Noyce and Gordon Moore had left Shockley Semiconductor Laboratory and formed Fairchild Semiconductor in 1957. Noyce and Moore later founded Intel in 1968.

Intel 1103 Memory Chip, 1970

The first integrated circuits contained less than ten transistors. With improvements in manufacturing processes, the density of integrated circuits increased dramatically. In 1965, Moore predicted that the density of transistors on integrated circuits would double every two years. This prediction came to be called Moore's Law and it turned out to be quite accurate. In the mid 1970s, thousands of transistors could be built on a chip. In 2015, Intel processor chips had over one billion transistors.

As electronic circuits are made smaller they get faster, since the electrical signals have less distance to travel. In addition more functions can be put on a single chip, making it possible to build smaller and cheaper computers. The "faster, smaller, cheaper" advances in the semiconductor industry have been the engine behind most of the revolutions in the computing industry ever since the invention of the integrated circuit. A laptop computer in 2015 is 130,000 times faster than a mainframe computer was 50 years earlier. On a cost per machine instruction basis, that laptop is about 1 billion times cheaper.

1964 IBM Mainframe Era

With these rapid advances in technology, computer manufacturers designed and built new machines to take advantage of the improvements in cost and reliability. These successive machines were generally incompatible. Programs developed for one machine could not generally be used on other machines without modification.

In 1964, IBM announced System/360. This was a whole new family of computers with compatible peripherals and a common operating system. This allowed customers to start with a small machine and easily upgrade to a larger machine when their needs expanded. This was a revolutionary design. IBM was already the leading computer manufacturer, but after System/360 it completely dominated the industry. In 1970 GE and RCA sold off their computer divisions. The rest of the major computer manufactures were known as the BUNCH, for Burroughs, Univac, NCR, Control Data Corporation and Honeywell. Burroughs and Univac eventually became Unisys who still makes mainframe computers. The rest of the BUNCH gradually left the mainframe computer business.

IBM System/360

Most large organizations from businesses to universities to governments had large data centers running IBM mainframe computers 24 hours a day using computer operations staff working shifts.

The IBM System/360 evolved to compatible System/370 in 1970, System/390 in 1990, zSeries (zero downtime) in 2000 and its current line of mainframes, z Systems in 2005.

IBM also developed a separate line of computers for midrange businesses. These were much smaller than mainframe computers, cost correspondingly less and did not need a large staff for operation and maintenance. The IBM System/3 was introduced in 1969. It was typically used for basic business applications like billing, accounts receivable, inventory control, sales analysis, and payroll. This was followed by System/32, System/34, System/36, System/38 and then AS/400 (Application System/400) in 1988. The AS/400 is still in use today.

However new competition had arisen in the form of “minicomputers”. The first commercial minicomputer was a PDP-8 made by Digital Equipment Corporation (DEC) in 1960.

Digital Equipment PDP-8 Minicomputer

Minicomputers found a ready market in industrial process control, telephone switching and research laboratories in universities and elsewhere. A big market for minicomputers was OEMs (Original Equipment Manufacturers) who bought them to integrate into the proprietary systems that they sold to end customers.

Early minicomputers were followed by larger and more powerful models. These became attractive as departmental computers within larger organizations and as business computers for medium size companies that could not afford mainframes. With such a burgeoning market, a lot of other companies built and sold minicomputers such as Data General, Wang Laboratories, Prime Computer, Varian and Interdata, but Digital Equipment remained the technical and market leader. At its peak in the mid 1980s, Digital Equipment was the second largest computer company in the world after IBM.

Minicomputers declined as microprocessors became cheaper and powerful enough to perform the functions that minicomputers were good at. By 2000 all the major minicomputer companies had been bought or were out of business.

1981 Microprocessors, Microcomputers, Workstations

As the number of transistors that could be built on an integrated circuit increased, two revolutionary breakthroughs came from Intel.

First was the memory chip. In 1970, Intel released the 1103 chip, the first generally available Dynamic Random Access Memory (DRAM) memory chip. It used three transistors as a memory cell to store one bit. It had a total of 125 bytes of memory. DRAM become the standard memory for computers and rapidly replaced magnetic core memory.

Second was the microprocessor chip, a single chip that contained all the logic for a CPU. In 1971 Intel released the 4004 microprocessor chip. It had 2300 transistors.

Intel 4004 Microprocessor Chip, 1971

With memory and a CPU available as chips, it became much simpler and cheaper to build computers. By 1975 microcomputers were sold as kits for home hobbyist to assemble. This was soon followed by microcomputers that were sold already assembled. These microcomputers took over this market in 1977 with the introduction of the Apple II, the Tandy TRS-80 and the Commodore PET.

Apple II, 1977

However, the real revolution in microcomputers began when IBM announced the Personal Computer (PC) in 1981.

IBM Personal Computer, 1981

While Apple had been successful selling into the education market, businesses did not widely embrace this new technology until IBM implicitly gave it a stamp of approval with the PC.

The PC was integrated into large data centers as "intelligent" terminals to replace mainframe computer terminals. It also opened the door to medium and small businesses who had not been able to afford the earlier generation mainframe and minicomputers.

IBM did not have much experience developing microcomputer software so they licensed it from a small company founded by written by a Harvard dropout whose claim to fame was developing a version of the BASIC programming language for microcomputers. That programmer was Bill Gates and the small company he founded was Microsoft. With the huge success of the PC, Microsoft became the dominant force in the microcomputer world.

Another revolution based on microprocessors also occurred. Many companies built engineering workstations. These were small computers for high end graphics and computation that rivaled minicomputers at a much lower cost. One of the first workstations was built by Apollo Computer in 1980, later acquired by Hewlett Packard (HP). Sun Microsystems in 1982 was one of the first to use the Unix operating system, which became the foundation for most of the workstation computers, such as those from HP, Silicon Graphics, DEC, and IBM.

Sun Workstation

Most of these early workstations used the Motorola 68000 family of microprocessors, which was a 32 bit chip and more powerful than the Intel 16 bit chips of that time. These Unix workstation computers became widely used in manufacturing design, universities, research labs and as department machines and servers in medium and large size companies. The growth of workstations hastened the end of the minicomputer era.

In turn, as Intel based PCs became more powerful and cheaper, they eventually replaced the Unix workstations. After several years of losses and falling stock prices, Sun Microsystems, the market leader was acquired by Oracle in 2010.

2. History of Software

Software refers to programs that are run on a computer. Software is generally written by using a computer language. A special software component is the Operating System (OS) that controls the computer hardware. A computer generally can't do anything without an OS. When you power on a computer, the operating system software is loaded into memory so that it can run in the CPU. Once the OS is running, you can then run applications (or apps). Some of the familiar operating systems in use today are Microsoft Windows for PCs, Linux for servers, and Apple iOS or Google Android for Smartphones.

The power of software to do many things has revolutionized many industries by providing great cost reductions and/or new ways of doing things. Many start-up companies have been created based on their innovative software. Some of the revolutionary applications of software have been:

- Accounting - one of the first applications to be automated with computers, first large companies using mainframes and then small and medium companies when microcomputers became available.
- Financial - automated teller machines (ATM), online banking, automated stock exchanges, online brokerages.
- Spreadsheets - tabulating information in tables with simple calculations. Advanced users are able to write spreadsheet programs without having to become computer programmers.
- Publishing/Desktop Publishing - word processing for books and articles; layout for newspapers, magazines and newsletters.
- Communications - email, computer based instant messaging/chat.
- Manufacturing
  Computer-Aided Design (CAD) is used to create and analyze designs and prepare data for manufacturing.
  Computer-Aided Manufacturing (CAM) is used to control machine tools in manufacturing.
- Digital photography/video - software to create and manage digital media files made digital cameras and camcorders commercially feasible.
- Computer graphics - has had a significant impact on many types of media. It has revolutionized film animation, movies, advertising and graphic design.
- Games - are a huge market that appeals to many age groups. There are many game platforms: consoles that connect to a TV (Sony PlayStation, Microsoft Xbox, Nintendo), games you can play on your computer and your mobile phone, and online games you can play over the Internet. Improvements in the video quality of games depends a lot on hardware advances. Leading edge game software has driven a lot of the computer graphics industry and has led to the development of special graphics

chips.

- Education - many universities and other educational organizations provide online courses and training over the Internet.
- Research - especially the sciences like physics, chemistry, astronomy, biology, medicine. For example, the Human Genome Project that identified and mapped all the human DNA in 2003 would not have been possible without software.
- Social Networking - is based on providing a service where members can set up a network of friends and communicate with them by posting messages, photos and videos. When the Internet became widely available and made it easy for people to join such services, social networking grew very rapidly. Many people have found this a better way to stay in touch and communicate with friends than using older technologies like telephone and email. Although there are quite a number of different social network systems, Facebook has become the dominant force, with over 1.5 billion members. Twitter is a more public social network that allows people to "follow" anyone else or topics on Twitter. Twitter has become an important source for breaking news which often appear on Twitter before the traditional news media or other web sites.

However, this chapter focuses on the evolution of operating system software and computer languages. Together they provide the foundation for all other software. The history of other important software applications can be found elsewhere, such as in Wikipedia.

1956 Early Operating Systems

The earliest computers had no operating system at all. Machine code programs consisting of a series of numbers had to be keyed in at the computer console each time they were used. It was realized quite early that a program using symbolic names was much easier to write and enter than machine code. Most computers had an "Assembler" to translate the symbolic assembler program to machine code. Here is a simple example of computer instructions to multiply two numbers:

Machine code:

```
0905125000 090528
```

Assembler code:

```
MOV R1, 125000
MUL R1, 28
```

The first assembler instruction moves the number 125,000 into the "register" named R1 . Registers are special memory locations in the CPU where all mathematical and other operations are carried out.

The second instruction simply multiplies the number in R1 by 28.

As computer memories became larger, it became feasible to write simple operating systems that could read in programs ahead of time and store basic libraries that all programs could use. These batch processing systems ran one program at a time in sequence. They were generally developed by customers for their specific needs. One of the first operating systems was developed by General Motors in 1956 for its IBM 704.

But by 1962, most computer manufacturers provided operating systems for their machines.

On a batch operating system, a programmer generally submitted his program to a keypunching service to get it punched onto cards. A large program could be several large boxes of cards. For smaller programs or corrections, the programmer would generally use a keypunch machine himself.

To run the program, the card deck was submitted at a service window in the computing center. A computer operator stacked the punched cards into a card reader which read them into the computer. The operating system put the program into one of several job queues, depending on the resources needed by the program, such as the amount of memory. It would take some time for the job to get to the front of the queue and be run. When completed, the operator would tear the output off the large printers in the computer room and put it in a small cubby hole along with the deck of cards for pickup. It could take several hours just to find out that the compiler identified a typing error on one of the punch cards and the program didn't even run.

As more storage became available on mainframe computers, programmers could save programs in files. If the program did not need changes, then they could just submit a deck of cards with the data needed by the program. For large production systems, input to a program might also be from magnetic tape or from data that had been stored in disk files. However, every operation needed to be submitted as a batch job. Even a simple operation like making a copy of a file required a batch job. That required about five punched cards in IBM's Job Control Language (JCL).

1957 Early Programming Languages

The first widely used general purpose programming language was Fortran (Formula Translation), developed at IBM by a team headed by John Backus. It was proposed in 1954 and completed in 1957. A special program called a compiler had to be developed that would translate the Fortran program instructions into machine code instructions. All programming languages, even the most sophisticated computer languages, today require such a compiler or interpreter to translate the source program code into machine code that the computer can run.

Fortran was especially well suited for numerical calculations used in scientific and engineering programming. Compared to assembly language, Fortran reduced the number of statements in a typical program by a factor of 20 and so it quickly gained acceptance. The popularity of Fortran spurred competing computer manufacturers to provide Fortran on their machines. Fortran has continued to be widely used. Over the years, it was influenced by new ideas of programming in other languages and a succession of new Fortran standards and versions have continued to update the language.

The following is a small Fortran code segment that calculates U.S. Income Tax when the taxable income is between $100,000 and $186,351.

```
IF (INCOME.GT.100000 .AND. INCOME.LT.186351)
TAX = INCOME * .28 - 6824.25
```

Another important language developed around this time was ALGOL (Algorithmic

Language). It was standardized by a committee in 1958. Here is the ALGOL version of the Fortran code above.

```
if income>100000 and income<186351 then
   tax := income * .28 - 6824.25
```

ALGOL was not as widely used as Fortran except in Europe, but it had a greater influence on the development of later languages. It was considered to have much cleaner syntax and logical structure through the use of Begin/End blocks. Some of the important languages today that were influenced by ALGOL are Pascal, C, Java, Javascript, Python, and PHP.

Around the same time, Grace Hopper at Remington Rand found that business data processing customers were uncomfortable with mathematical notation when writing programs. She proposed a programming language called FLOW-MATIC that used English keywords and completed a version for the Univac I computer in 1958.

In 1959, an industry committee was formed to create a programming language for business. A specification for COBOL (COmmon Business-Oriented Language) was produced in 1960. It was heavily influenced by FLOW-MATIC. COBOL was quite different from Fortran. It had very simple calculations but much more capability to format data for input and output and was well suited to reading in large amounts of data from magnetic tape and printing out reports.

Here is an illustrative example of calculating the final cost of an item by adding the sales tax to the price.

```
Fortran:  COST = PRICE + TAX
COBOL:  ADD Price Tax GIVING Cost.
```

Many of the major computer manufacturers committed to provide COBOL on their machines. In 1962 IBM announced that COBOL would be their primary data processing language. COBOL became the most widely used programming language on mainframe computers, whose predominant use was to manage large businesses. There are still many large COBOL production systems in use in large companies today and there is an ongoing need for COBOL programmers to maintain these systems.

With the rapid advances in computing, a number of scientists from different fields thought that computers were on a path to becoming “thinking” machines. This area of research became known as “Artificial Intelligence” (AI).

Quite a number of computer programs produced dramatic results. In 1955, Herbert A. Simon, a professor of political science at Carnegie Mellon collaborated with Allen Newell, a researcher at the RAND Corporation, and created the “Logic Theorist”. It demonstrated human problem solving and proved many of the theorems in Russell and Whitehead’s Principia Mathematica. In 1962, a checkers playing program written by Arthur Samuel at IBM was able to beat a human checkers champion.

A number of special purpose high level languages were developed for AI research. Lisp (List Processor) was developed by John McCarthy at MIT in 1958. It made it much easier to work with lists than in other languages. A list could be a list of symbols like words or a list of numbers. Lisp remains today one of the most widely used languages in AI.

Artificial Intelligence went through a number of cycles. After its initial success, research funding dried up for long periods. Another reason for slow progress was that AI problems typically require a lot of computing power to solve anything more than simple cases. As computers became more powerful, more progress was possible. But it wasn't until almost 40 years after the checkers playing program that a chess-playing program (Deep Blue developed at IBM) was able to beat a chess champion in 1996.

AI has recently had a resurgence as some of its solutions are critical to the success in application areas such as data mining, speech recognition, robotics and self-driving cars.

1961 Time-sharing Operating Systems

Some organizations found the batch processing operating systems very limited. In particular, universities wanted a computer that many students and faculty could use at the same time from a terminal. Terminals at the time were devices from other industries that were modified to work with computers, such as an electric typewriter or a teletypewriter (see picture in Chapter 1).

There were numerous efforts to develop what became known as "time-sharing" systems. The operating system would divide each second up into small time slices. Each terminal connection program and batch job program that was running and would get a limited number of time slices to run before the operating system would suspend the program and switch to another program. This occurred so quickly that a user typing on a terminal would not notice any interruption.

The first successful time-sharing systems were developed in the early 1960s at MIT, the University of Illinois and Dartmouth College.

Time-sharing also supported interactive computing. Interactive computing allows users to type in a command and have the computer run it right away. This provides the ability to test things very quickly and feedback on errors that can be immediately corrected. However interactive computing was only available with a few computer languages. One of the first was Dartmouth BASIC (Beginner's All-purpose Symbolic Instruction Code) in 1964. It was based on Fortran but was designed to be simple enough for non-technical students to learn easily. It was used at Dartmouth and some other schools and colleges for many years but not in commercial mainframe computer applications.

Following the success of time-sharing systems developed at universities and continuing demand for them, the computer vendors added time-sharing to their batch processing systems. In 1972 IBM added time-sharing to System/360. Computers were now powerful enough and had enough memory to support several batch job queues and many concurrent time-sharing users via terminals.

Support for time-sharing and the emergence of video display terminals in 1974 marked the end of the punched card era. Programmers could edit their code on-screen and save the modified programs to disk files. They could submit programs directly from their terminal for batch processing and retrieve the output to their terminal when the job completed. This model of coding is still predominant today. The only difference is that programmers use PCs to edit source code and run it on their own machines, instead of using a terminal to connect to another computer.

1968 Database Management Systems

As disk storage systems became larger and cheaper, companies started to store more and more information online for more rapid retrieval than was possible with magnetic tapes. The data was stored in a large special file or series of connected files called a database. Databases had a complicated structure to make storage and retrieval as fast as possible. A large software package called a Database Management System was needed to manage the database and provide an interface to make it easy for application programs to use the database. IBM provided one of the earliest database systems, but the market was so large and so lucrative that many third party vendors developed competing database systems.

There were several different models of how to structure the data in a database. Eventually the “relational” model proved to be the most flexible and efficient. In a relational database, data is structured as a table with the columns representing different attributes of the data and the rows are individual records. The following credit card statement is a good example of tabular data that can easily be stored in a relational database.

Transaction Date	Description	Amount
27/02/2016	AA INFLIGHT MC FACET 3 PHOENIX AZ	$5.56
27/02/2016	GREAT AMERICAN GRILL T TUCSON AZ	$55.07
27/02/2016	AMERICAN 000102726190000 PHOENIX ON	$52.50
03/03/2016	KAISER GRILLE PALM SPR PALM SPRINGS CA	$94.46
07/03/2016	PARIS LE VILLAGE BUFFE LAS VEGAS NV	$92.13

IBM released a relational database called DB2 in 1981, which existed along side its other database systems. DB2 came to dominate the mainframe market and drove out most of the earlier database products.

However, the first commercially available relational database was Oracle in 1978. Oracle was first available on minicomputers and then on IBM mainframes, Unix machines and PCs. Oracle became one of the leading database systems and a very large company. Its co-founder and CEO Larry Ellison, became one of the wealthiest people in the world and a high profile personality in the technology industry.

SQL (Structured Query Language), initially developed at IBM, was a language designed to make it easy for application programmers to access data in relational databases. The following example shows the SQL command to read the data shown in the table above:

```
SELECT TransactionDate, Description, Amount FROM Transactions;
```

SQL was standardized in 1986. However most database systems incorporated extensions to SQL or diverged from the standard somewhat. So switching to a different database system usually requires some changes to all the application programs that use the database.

Database software was quite expensive. When Linux became widely used as an inexpensive server, open source database software became popular. MySQL, originally created in 1995, has become the second most widely used database system after Oracle.

1975 Unix

AT&T Bell Labs was involved in a project with MIT and General Electric to develop a time-sharing system called Multics (Multiplexed Information and Computing Service). Dissatisfied with the project's progress, Bell Labs withdrew in 1969. Instead Ken Thompson and Dennis Ritchie worked on a scaled down time-sharing system at Bell Labs that would run on a minicomputer. It was named Unix as a pun on Multics.

Unix was a revolutionary operating system in many ways. It was a multitasking, multiuser system with a powerful set of tools for software developers. Unlike most operating systems at the time which were written in assembler, Unix was largely written in a high level language called C that was developed in conjunction with Unix. This made Unix easily portable to a different hardware platform once it had a C compiler, which was much easier to write than an operating system.

This technique of writing most of the operating system in a high level language instead of assembler had actually been done earlier by Burroughs. In 1961 it released the B5000 mainframe computer using ALGOL for most of the operating system.

In 1975 AT&T licensed Unix for use outside of the company. Unix became the foundation for the high end engineering workstations built with microprocessors beginning in 1980. This became a large market with many companies from start-ups that grew into large companies, like Sun Microsystems, to computer manufacturers that were already large, such as HP and somewhat later IBM. Unix was also the system of choice for web servers in the early days of the Internet. Unix was eventually supplanted by Linux. But Linux is essentially just a different version of Unix.

1981 Microsoft DOS

When IBM entered the microcomputer market with the Personal Computer (PC) in 1981, it came with a simple operating system called DOS (Disk Operating System) that was supplied by Microsoft. Like all operating systems at the time, it had a command line interface in which the user had to type instructions. For example, here is the famous C: prompt followed by a command instructing DOS to make a copy of the file named db.dat

```
C:\>
C:\> copy db.dat db.bak
```

As software, DOS was not an important advancement in operating systems, but it was an important factor leading to the tremendous growth of PCs and the evolution of computers.

With the rapidly growing microcomputer market, IBM felt it did not have time to develop custom hardware as it did for all its existing computer products. So the PC was designed and built with off the shelf electronics. This "open architecture" made it possible for third parties to easily build PCs. For the software, they were able to license DOS from Microsoft. With the huge success of the PC, many small companies and established computer manufacturers jumped on the bandwagon and produced PC compatible "clones". Many innovative features, such as the first laptop computers, originated as PC clones. Descendants of the PC clones comprise the majority of PCs today.

Many other third party vendors developed expansion boards for hardware improvements, like sound and video cards. Many more developed software for applications not available from IBM or Microsoft. Applications for the business market provided the justification for businesses to buy PCs. A few “killer” apps were largely responsible for the explosive growth of PCs. These were Wordstar for word processing, Visicalc followed by Lotus 1-2-3 for spreadsheets and dBase, a database management system.

Programming languages for microcomputers took a different direction than what was being used on mainframes. BASIC was the first widely used language on microcomputers. It was high-level enough to be usable without training and small enough to fit into early microcomputers. But “professional” developers of third party application packages turned to more powerful software tools. C, which was widely used on Unix, became the dominant language. One important reason at the time was that it produced code that ran faster than code produced by other language compilers.

Following is an example of some C code that calculates U.S. Income Tax when the taxable income is between $100,000 and $186,351. This code is very similar to many other computer languages such as Pascal, C, Java, Javascript, Python, and PHP and to the earlier Fortran and ALGOL code. The differences are in the punctuation (brackets and semi-colons) and use of the word *and* instead of the symbol *&&*.

```
if (Income > 100000 && Income < 186351)
     Tax = Income * .28 - 6,824.25;
else
   ... (code for other tax brackets)
```

1984 Apple Macintosh

In 1984 the Macintosh arrived with a revolutionary operating system. It was the first commercial computer with a graphical user interface (GUI) and a mouse.

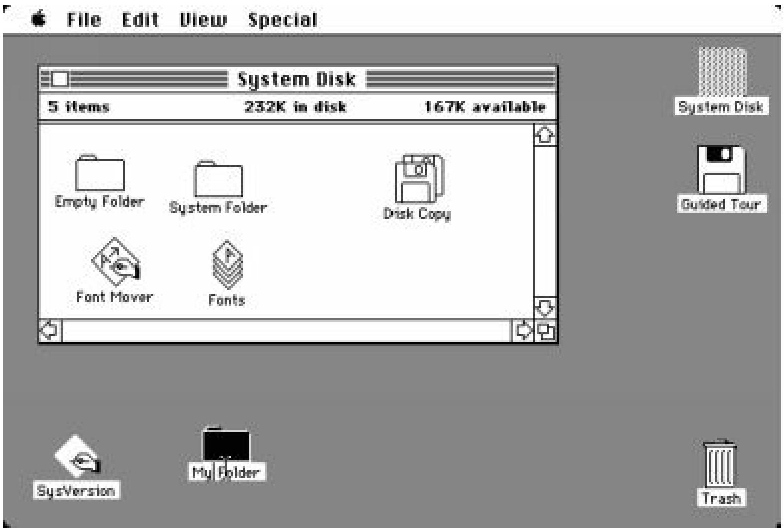

Apple Macintosh

The Macintosh displayed a picture representing the top of a desk that contained small pictures (icons) of folders and programs. When the user clicked on an icon, a rectangular area on the screen, known as a window would open and show what was inside. A folder could contain other folder and program icons.

An application program generally had a window with its own design containing fields for input to the program and display areas to show the results. The user could arrange the desktop by changing window sizes and moving them around. They could quickly switch to another application by clicking with the mouse on any of the other windows on the desktop. It was also easy to move data between applications in different windows by a simple mouse or keyboard "cut and paste" operation. This model of how a user should control a computer has influenced all operating systems since.

The Macintosh was a success for Apple, but it could never really challenge the PC. The very competitive PC market with its many clones was less expensive and had many more third party applications than the Macintosh. The Macintosh share of the microcomputer market peaked at 12% in 1993, except for some specialized areas where it dominated, such as graphic design and desktop publishing.

The Mac OS was used until 2001 when it was replaced by OS X to support Apple's new line of computers. OS X is still in use today (after numerous version updates). OS X is based on a Linux core but has a similar GUI to the original Macintosh.

1985 Object Oriented Programming

As software projects got larger and more expensive to develop and maintain, computer scientists and practitioners in the field explored ways to improve productivity. One of the ways to achieve this was "modular programming". The concept was to break up large pieces of code into functional units or modules. This would make the code easier to understand and provide code modules that could be re-used on other projects.

The problem with modular programming was that it depended a lot on the discipline of the individual programmer. Many programmers found it easier and faster to cut and paste sections of code from one part of a program to another instead of designing a single module and putting it in a library so that it could be called from multiple locations. Cut and paste may have worked well for the original developer, but not for maintenance. Most programs needed to be changed over time to fix bugs or add new features. Often the original programmer had moved on to other projects or another company and the maintenance work fell to someone who did not know the code. If a change was required in code that had been duplicated in several parts of the program, it took more time to find. Often the programmer would not find all of those locations and this resulted in more bugs.

The idea of structuring a programming language to make code more modular gained support. This could be done by using "objects". An object contains all the data, code and methods (operations) to make the object self-contained and useful for use in other parts of the application or in other applications.

To use an object, the programmer needs to code a line to create the object. Then the programmer can call the methods in the object to do the calculations.

Here is an example of using a "CalculateTax" object to calculate the income tax for a given taxable income.

```
TaxObject = CalculateTax( )
IncomeTax = TaxObject.ComputeTax( 125000 )
```

The first line creates an object called TaxObject that has the code provided by the library object named CalculateTax.

The second line uses the object's ComputeTax method to figure out the tax bracket for a taxable income of 125,000, get the income tax rate and calculate the income tax.

While an object oriented language called Simula had been developed as far back as 1965 at the Norwegian Computing Center, it was used for special purpose simulation problems and was not in very widespread use. Bjarne Stroustrup, a Danish computer scientist working at Bell Labs had used Simula but found it too slow for large software development. Stroustrup enhanced the C language with object oriented features and the C++ Programming Language was released in 1985. C was very widely used and programmers found it easy to upgrade to C++, as they could continue to use C and add object oriented code.

This established object oriented programming as a major new paradigm. Object oriented features were added to many existing languages and most new languages developed after this point in time were object oriented.

1990 Microsoft Windows

Microsoft's first two versions of Windows released in 1985 and 1987 were not successful for several reasons. The performance was sluggish, largely due to limitations of the 16 bit Intel chips at the time. By comparison, the Macintosh was built with more powerful 32 bit Motorola microprocessor chips. Many third party developers delayed making the large development effort required to modify their DOS software to run in Windows until they had confidence that Windows would be a successful operating system. So many businesses and users were reluctant to switch from DOS where they had lots of productivity applications.

Windows 3.0 released in 1990 was more successful. The Intel hardware caught up with the demands of the operating system and graphics on PC machines had improved with new video cards. This was the first Microsoft operating system that could compare to the Macintosh.

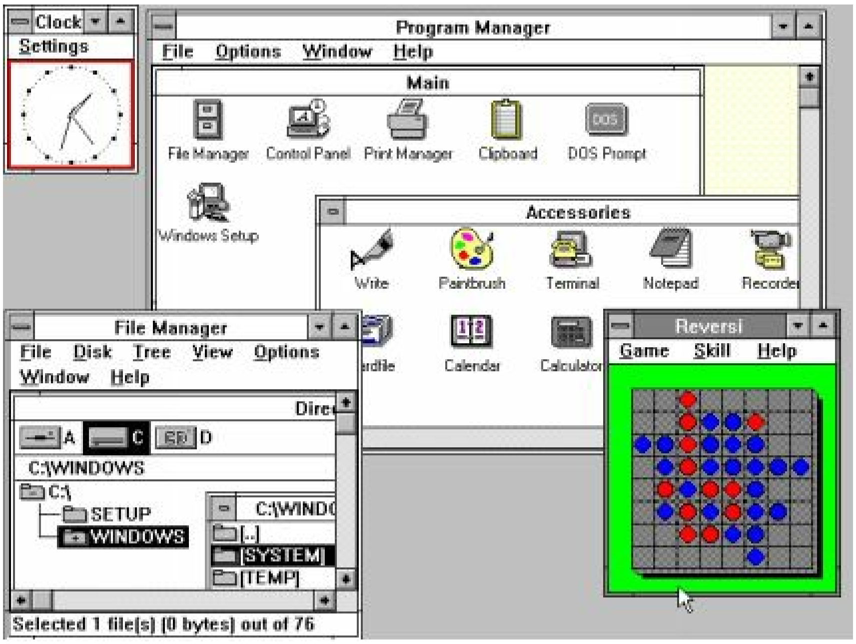

Microsoft Windows 3.0

Many third party applications still did not have Windows versions, in particular the leading word processor WordPerfect and the leading spreadsheet Lotus 1-2-3. But Microsoft had been working on their own word processor Word and their spreadsheet Excel and they were bundled into Microsoft Office. By the time WordPerfect and Lotus had Windows versions of their software ready, Microsoft Office had become a success, and has continued to dominate this application area ever since.

Windows 3.0 was followed by a long line of new versions every few years. The versions targeted for home use were Windows 95, Windows 98, Windows ME, Windows XP, Windows 7, 8 and 10.

Microsoft operating systems from DOS to Windows 3.0 had been plagued by instability and frequent crashes. The user had to “reboot” the system, often losing much of their work. In 1988 Microsoft hired Dave Cutler, one of the chief architects of VMS, a solid operating system that ran VAX computers made by Digital Equipment Corporation (DEC). Cutler brought over a team of engineers to work on the design of a full 32 bit, multiprocessing, multi-user operating system that could run on various processors besides Intel.

Windows NT released in 1993 had a user interface similar to Windows 3.0 but the internal operating system was completely different. Windows NT was a solid operating system and established Microsoft as a vendor of server class software.

Microsoft had plans to merge the underlying operating system of NT with the GUI in the Windows systems but this was not achieved until the release of Windows XP in 2001.

Windows NT continued to be enhanced with successive versions. In 2003, the name was changed to Windows Server. Windows Server is a major Microsoft product. It is widely used in corporate data centers and is second to Linux as a web server.

To support software development in Windows, Microsoft released the programming language Visual Basic in 1991. It was an enhanced version of BASIC that provided an easy way to build applications with a graphical user interface (GUI). The programmer could create input forms visually, using drag and drop to move windows elements like list boxes around on a design screen. Visual Basic became quite popular especially among end users. It was very easy to build Windows applications that provided a GUI interface to a database. Microsoft continued to enhance Visual Basic and it became part of the .NET framework which is Microsoft's current platform for developing software for Windows. However, the most widely used language in .NET is C#, a language developed by Microsoft with similarities to C++ and Java.

1994 Linux

As PCs became more powerful with hardware advances of the Intel chips, a few people in the Unix world saw that inexpensive PCs would eventually compete with the capability of Unix machines at a much lower price. The Intel 80386 released in 1986 and incorporated into PCs was deemed to have the necessary hardware capabilities to run Unix. The GNU project to write a free version of Unix to avoid its licensing restrictions had been started in 1983 by Richard Stallman. Stallman is a fervent believer in free software and founded the Free Software Foundation in 1985. He left the MIT AI lab to work full time on the GNU project, the FSF and many other activities promoting his ideas on free software.

All major components in the GNU project were completed by 1992, except for the critical Unix kernel, the core portion of the Unix operating system. The kernel was first based on code from MIT that proved to be inadequate. It was replaced with code from Carnegie Mellon University. This in turn was delayed for several years due to licensing issues and then proved to be too complicated.

In 1991 Linus Torvalds, a university student in Finland, announced the Linux project to collaborate on writing a Unix kernel for the Intel 80386. Over 100 hundred programmers wrote code and version 1.0 was completed in 1994. When combined with the GNU code, it provided a complete Unix system free of cost that could run on a PC.

The availability of Linux coincided with the explosion in the use of the Internet which began after the World Wide Web was invented in 1991. Organizations needed web servers to mount their web sites. They didn't need big expensive Unix machines. The use of Linux took off. It also started to be used for applications which had been the domain of Unix machines. As well as Internet servers, Linux continued to expand into data centers for general purpose servers and has become the base for cloud computing.

Their was even demand for Linux on enterprise class hardware. In 2009 IBM announced a mainframe-based Enterprise Linux Server. However IBM's proprietary operating systems that evolved from System/360 and provide backwards compatibility for the huge software investments made by their customers continue to be the major part of their software line.

While Linux has graphical user interfaces, many Linux users were system administrators and software developers who relied on the command line interface with its powerful set of commands.

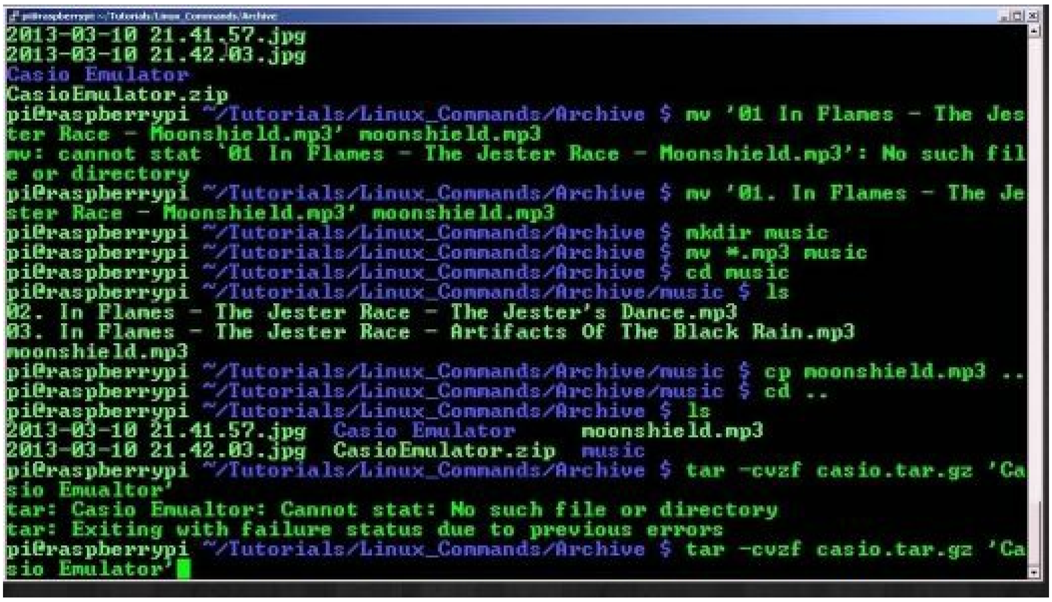

Linux Command Line Interface

The ease with which Linux can be modified to support specialized hardware has made it popular in many computer based devices small and large, from cell phones to robotics to supercomputers. The top 200 of the world's fastest supercomputers run some variant of Linux.

There is only one area where Linux has not made big inroads, desktop computers for end users. This has remained the domain of PCs running Microsoft Windows, and to a lesser extent Apple Mac computers.

3. History of Computer Networking and the Internet

While the history of the Internet goes back quite far, there were many other networks that were much more important than the Internet for a long time. It is important to understand the evolution of computer networking before the Internet came to be the dominant network in the world.

Networking in the Mainframe Era

Early mainframe computers were standalone batch machines without any communication capabilities to send or receive information over a network. In the 1960s, as use of computers within organizations grew and time-sharing became available, there was a desire to connect terminals to computers from locations other than the same building as the computing center.

Early terminals were based on using existing equipment from other technologies. One of the most widely used computer terminals was the teletypewriter (see picture in Chapter 1). This was a specialized typewriter developed for the telegraph industry. It went through various inventions starting as early as 1855. It allowed an operator to transmit telegraph messages by typing on a keyboard. Received messages were printed on paper.

Teletypewriters used a simple transmission scheme, sending one bit at a time over a cable made up of several wires. A bit was either a low voltage electric pulse representing 1 or high voltage pulse representing 0. Each character on the keyboard had a unique code, made up of 7 bits. In addition there was a start bit and a stop bit to mark the beginning and end of the character. This form of transmission is called asynchronous serial communication and has been the basis for a many data communication devices and networks. For many years serial ports conforming to the RS-232 standard were common on most computers and network equipment. This was superceded in 1998 by the Universal Serial Bus (USB) standard which provides higher speeds by transmitting data in blocks rather than one character at a time and supports many different types of devices.

Another early terminal was an electric typewriter that was modified to send electric signals over a wire to a computer. A widely used computer terminal was the IBM 2741, introduced in 1965. Compared to the teletypewriter, the 2741 offered higher speed, higher quality printing, quieter operation, interchangeable type fonts, and both upper and lower case letters.

For short distances within their own premises, companies could run cables from terminals in different offices or buildings to the computer. From more distant locations, companies could use a phone line. However, data signals cannot be sent over the telephone networks so a dial-up “modem” was needed at each end. It “modulated” the asynchronous serial data signals to analog telephone signals and “demodulated” the signal at the other end from telephone analog to data.

1968 Acoustic Coupler Modem

One of the first widely used commercial modems was the Bell 103A, produced by AT&T in 1962. It could transmit 30 characters per second. Modem speeds gradually increased with many technology improvements up to 5,600 characters per second in 1998. In the years following this, with the huge increase in demand for Internet access, higher speed broadband technology using cable and fiber optic networks began to replace dial-up telephone network connections.

Some locations had a high volume of traffic and needed higher transmission rates than what was available over dial-up modems. For these cases, data centers would lease a line from the telephone company for a flat monthly rate. This was not an actual telephone line but a dedicated private circuit, so various types of data communications equipment could be attached at each end. Larger companies with many locations could build large private networks utilizing leased lines. Many companies today continue to use leased lines in a private network.

Besides the teletypewriter, another technology with obvious application for use as a computer terminal was the CRT (cathode ray tube). They were commercially available as early as 1922 and were widely used in television sets in the 1950s. However, it was not until 1971 that mainframe computers had enough memory and fast enough processing speed to manage many of these devices. From this point on, most computer vendors made video display terminals using CRTs and there were also many other manufacturers that got into the market such as Lear-Siegler, ADDS, Hazeltine, Heath/Zenith, Televideo, and Wyse.

One of the most widely used terminals was the IBM 3270 display terminal. It was designed to minimize the amount of data transmitted and the frequency of interrupts to the computer. To do this, multiple display terminals were attached to a controller and it was the controller that communicated with the computer. The user did local editing and when the Enter key was pressed all the changes on the screen were transmitted as a block over a high speed coaxial cable. This was much faster and more efficient than typical asynchronous serial terminals.

IBM 3270 Display Terminal

With the growth in the number of time-sharing users, computer vendors started to provide basic networking capabilities in their operating systems to support large numbers of terminals. IBM System/360 provided various telecommunications access methods to support this.

In 1974 IBM released Systems Network Architecture (SNA), a more sophisticated networking solution than what was in System/360. It comprised hardware and software that allowed computers as well as terminals to interconnect. It improved the efficiency of the use of the data communications lines and expanded the number of terminals that could be supported. An IBM mainframe could handle thousands of 3270 terminals. This was important in large transaction processing applications such as order entry systems. Many large private networks using SNA were built by large companies and continue to be used today. At its peak around 1990, most of the world's networking traffic was SNA.

Public networks were also developed and made it possible for organizations to extend their private networks to remote locations. This was spearheaded by the telephone companies. They developed a standard protocol X.25. This was a packet switched network. By 1979 a global service was available, operated by large phone companies in North America, Europe and Japan, and later extended to other parts of the world. It was very cost effective for providing access from a location with a small number of terminals, compared to the cost of a leased line or a telephone line and modems. X.25 evolved into a more efficient network, Frame Relay, which is still used in the data networks of most phone companies. Other X.25 applications gradually migrated to the Internet but some continue to be used in some parts of the world and by some legacy applications.

Packet Switching

Asynchronous serial communication, where a terminal transmitted one character at a time, made very inefficient use of a phone line. Packet switching was a design that transmitted a block of characters at a time through a shared network interface that allowed multiple terminals to use the same communication line.

The first design for a packet switched network was a study conducted for the Air Force by Paul Baran at the Rand Corporation. The Air Force wanted a secure data network that would survive a nuclear attack. Baran published a Rand report in 1962 with three key ideas:

1. A decentralized network with multiple paths between any two points.
2. Messages divided into packets.
3. Delivery of the packets by store and forward switching.

However, the Air Force did not continue the project and did not build a packet switched network.

1969 ARPANET (Internet)

J. C. R. Licklider was a psychologist at MIT (Massachusetts Institute of Technology) who became interested the human/computer interface. In 1962 he became the first director of the Information Processing Techniques Office of the Advanced Research Projects Agency (ARPA) in the U.S. Department of Defense. Under his direction, ARPA sponsored the development of interactive time-sharing computer systems at MIT and other universities and research labs. One of Licklider's strong interests was to link these computers together through a computer network.

A pilot packet switched network named ARPANET was planned and put out for contract. Bolt Beranek and Newman (BBN), a consulting firm with close ties to MIT was selected. They developed software on Honeywell minicomputers for the network nodes, similar to routers, that were inter-connected by leased lines. By the end of 1969, ARPANET connected four universities. By 1981, the number of connected computers was more than 200.

Research into performance and improved network designs continued through the publication of documents titled "Request for Comments". These RFCs ultimately became the standards of the Internet, as managed by the Internet Engineering Task Force (IETF).

In 1973 Robert E. Kahn at ARPA and Vinton Cerf at Stanford University started working on an improved network protocol. The first version of the ARPANET had incorporated flow control to ensure that the sender did not send packets faster than the receiver could handle. This required routers to communicate with each other using a flow control protocol. This added complication to the network. Kahn had worked on satellite packet networks and radio packet networks and was interested in a protocol that would work across these types of networks as well as networks built with electrical communication lines. The solution for these issues was to move the complicated flow control out of the network itself.

The result of Kahn and Cerf's work was to split the network protocol into two separate layers, TCP and IP, commonly referred to as TCP/IP. IP is a simple protocol (Internet Protocol) in which routers simply forward packets based on the IP addresses in the packet. Routers do not use flow control and do not make any attempt to determine if packets have been successfully delivered to the destination. In fact, if a router is congested with more traffic than it can handle, it is permissible to simply discard packets. Ensuring that packets get to their destination is handled by TCP (Transmission Control Protocol), which runs at the end points, in the sending and receiving computers. TCP provides information so that the sending computer can determine if any packets have been lost and retransmit the missing packets.

The robustness of TCP/IP in carrying all kinds of traffic was a key factor that led to the success of the Internet. In recognition of his contribution, Vint Cerf became known as a "Father of the Internet".

The main uses of the ARPANET were for mail, login to remote computers (Telnet), transfer files between computers (FTP) and remote job entry (RJE). This required that each connected computer have the appropriate software. ARPA contracted with the Computer Science Research Group at the University of California, Berkeley to implement the new TCP/IP network protocol designs and add them to Unix. In 1983, the network protocols of ARPANET were replaced by the more flexible and powerful family of TCP/IP protocols, marking the start of the modern Internet.

While the early computers that were connected to the ARPANET were from many different vendors, following this, Unix machines became the preferred computers for using the Internet.

ARPA was a research organization and never intended to operate a network. In 1984 a separate network was split off from ARPANET for the military and run by another agency in the Department of Defense. In 1990 the ARPANET was formally terminated.

Many universities and research organizations could not be connected to ARPANET due to funding or authorization limitations. In 1981, the U.S. National Science Foundation (NSF) agreed to fund a network for academic and research use. The Computer Science Network (CSNET) eventually connected more than 180 institutions, from Australia to Europe and Japan.

Another project of the NSF was support for supercomputing centers. To provide access to these facilities by researchers, the National Science Foundation created NSFNET in 1986. NSFNET became a general-purpose research network connecting the supercomputing centers to each other and to the regional research and education networks that in turn connected campus networks. The combined networks became generally known as the Internet and grew to more than 160,000 host computers by 1989.

Other Networks

Many other networks were also built in this period. There were national research networks similar to ARPANET in several European countries. In the commercial world, there were private networks of large organizations based on proprietary technology such as IBM's SNA and other computer vendors, public X.25 networks provided by many regional phone companies, semi-private networks built by companies in the time-sharing business to make it easy and cost effective for their customers to access to their data centers and networks built for special industries, such as the airline systems for sharing traffic data and buying airline tickets.

1982 Microcomputer Era

As PCs became more widely used in offices, there was a desire to share printers and hard disks among several users as these were quite expensive at the time. The initial PCs had a serial port so it could be used as a terminal but there was no easy way to connect many PCs on a network. In 1976 an influential paper describing a "local area network" (LAN) was published by Robert Metcalfe and David Boggs from Xerox. They had developed a network called Ethernet where all machines shared a single coaxial cable. This was approved as a standard by the IEEE in 1982 and several vendors made Ethernet cards for PCs. There were other competing LAN technologies, especially IBM's token ring, but Ethernet came to dominate the hardware for local area networks. Development on Ethernet continued with speeds increasing from 10 Megabits per second to 10 Gigabit Ethernet (standardized in 2002 but not widely deployed until 2008). Ethernet became the network in large data centers and is now replacing legacy data transmission systems in the world's telecommunications networks.

There were many different networking software systems for connecting computers on a LAN. Microsoft provided networking for printer and file sharing in Windows 3.0 and there were several suppliers of TCP/IP software. However, it was Novell NetWare that came to dominate the LAN. It could run over networking hardware from different vendors, was reliable and fast and provided flexible file sharing.

1995 The Internet

So in the early 1990s there were a lot of different networks in use. Most offices were using Novell NetWare over an Ethernet LAN while mainframe data centers had private networks using IBM's SNA over leased lines. There was a large global public X.25 network provided by telephone companies in many regions of the world. And then there was the Internet, used by university and research organizations.

In a few short years, this all changed dramatically and the Internet came to be the dominant network. How did this happen?

There were several events that converged at this time.

In 1991, restrictions on commercial use of NSFNET were relaxed as more regional commercial networks and ISPs were providing service and wanted to interconnect their networks by using NSFNET. A transition to a commercial Internet backbone was planned and contracts were awarded to large telecommunications carriers for all of the regional NSFNET networks. The final restrictions on commercial traffic ended in 1995 when the transition of NSFNET to private networks was completed and NSFNET was no longer needed.

Internet use in 1991 was largely mail, news, file transfer, remote login and a few search applications. To use the Internet you generally needed an account on a Unix machine and had to use a command line interface. There was no web browser or anything else with simple point and click functionality.

Tim Berners-Lee, a software engineer working at the CERN nuclear research facility in Geneva, felt there must be an easier way to get information that was on other computers. He proposed an idea for a "World Wide Web" of hypertext documents. Hypertext is text that contains links to other text, which can be on another computer. By selecting a hypertext link, you can retrieve those pages. So with a web browser and the Internet, you have access to a world wide web of information. While hypertext had been an idea many others had explored, combining it with the Internet provided a potential beyond earlier hypertext systems.

Berners-Lee developed HTML (Hyper Text Markup Language) a language for creating web pages and HTTP (Hyper Text Transfer Protocol) a protocol that defined how messages and web pages are transmitted over the Internet. In 1991 he demonstrated a working web browser and web server.

Many people were excited about this new way to use the Internet and web browsers were written for many different computers and operating systems. By 1993, users of Unix and Macintosh computers had web browsers which provided easy point and click access to the World Wide Web. This was not easily available to PC users.

The third important event was in 1995 when Microsoft released Windows 95. Windows

95 included a web browser, Internet Explorer, and had support for TCP/IP networking which was necessary to access the Internet. Users did not have to install any software. They only needed to get an Internet account at an ISP. This easy availability of web access in the huge computer base of PC users was a major factor in the explosive growth of the Internet that has continued to today.

The number of people who had Internet use increased from about 15 million in 1995 to 3 billion in 2015. In terms of world population, this was an increase from less than 1% to 40%.

[http://www.internetlivestats.com/internet-users]

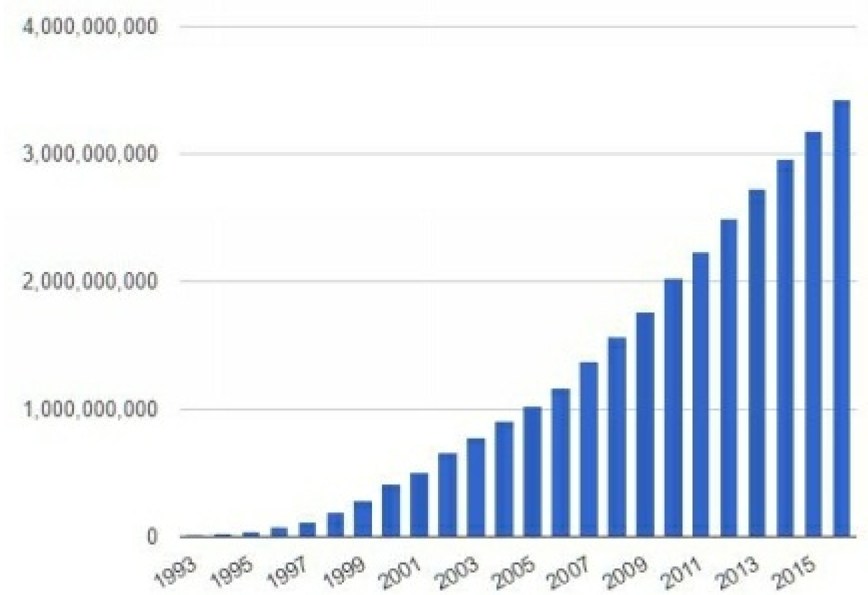

With the huge growth came decreasing costs in network equipment for IP networks. In the space of 20 years, IP networks have largely replaced most other networks — Local Area Networks, data center networks, private networks. The telephone companies have started to build all IP networks that can carry voice traffic as well as data and will eventually replace the legacy circuit switched phone networks.

Internet Software

Organizations large and small wanted to have a presence on the web. There was a burst of web development to create web sites. Initially, web sites were fairly static. Web pages contained HTML which is not a programming language. It is basically text containing markup tags that instruct a web browser how to format and display text.

For example, if you want your web page to display the following,

> **Ice cream** is a frozen food usually made from dairy products.

You would embed the following tags in your text

> Ice cream is a frozen food usually made from dairy products.
>
> The tags ... turn bold on and off.
> The tags <a ... contain a link to another web page.

HTML was very limiting for web designers who wanted ways to make web pages more dynamic. This led to a whole new direction for software development as new tools and languages were created to create richer web sites.

Java

In 1991, a small team of engineers at Sun Microsystems led by James Gosling created a new programming language designed to interface with digital consumer devices. It was demonstrated to the cable TV industry with a handheld interactive controller but it was ahead of its time. Instead, the team saw application to enriching web applications on the Internet and in 1995 worked with Netscape Communications to make Java available in their web browser. The Netscape browser later became Mozilla Firefox.

A Java applet could be downloaded by a web browser from the web server and run on the client machine. Java was a full featured programming language. It was able to control the screen and interact with the user beyond what was possible with HTML.

Java was a revolutionary programming language in many ways. Although the syntax was based on C++, it was designed from the beginning as an object oriented language.

Instead of being compiled to machine code, Java is compiled to an intermediate representation called Java bytecode that is independent of the hardware where it will be run. The target computer must have a Java virtual machine (JVM) which then translates bytecode to actual machine instructions. Sun Microsystems provided JVMs for all major operating systems without charge. This design made Java a highly portable language. Many developers were attracted to its promise of “Write Once, Run Anywhere”. There were a lot incompatibilities between web browsers at the time and Java applets were a seen as a way to avoid having to write different versions of code for different browsers.

Java applets were initially very popular for web applications. An **<applet>** tag was added to HTML and web browsers added support for downloading Java applets from the web server and passing them to the JVM. However, Java applets took time to download, required a lot of memory and often seemed sluggish. Eventually JavaScript became the preferred language for adding dynamic pages on the client side.

Ironically, Java became very important for developing software that ran on the server. This could be the backend for web applications or for other general purpose applications not related to a web site. Under Sun’s direction, many new features were added to Java and a rich library of Java code became available. In 2007, Java was released as open source. Currently Java is one of the most widely used languages for software development.

JavaScript

A different approach from Java applets for adding dynamic features to web pages was to use a lightweight language that is interpreted and run within the browser. JavaScript was developed at Netscape and released in 1995. JavaScript is not related to Java, it is a completely different language. It has a syntax based on C but is an object oriented language. Some of the things JavaScript is good at is checking user input, providing messages, customizing the appearance of web pages and adding simple animation. Along with HTML and CSS (Cascading Style Sheets), JavaScript has become one of the core technologies of web design and it is supported by all major web browsers.

PHP

Languages that ran in the web browser could make web pages more dynamic, but they could not easily get data from a database. Many web sites needed access to a database, for example when a form is used to allow you to input and edit your user account information (name, address, phone, etc.). What is needed is software that runs on the server. Quite a number of server-side languages have been developed. PHP has become the most widely used server-side language, except on Windows servers where ASP.NET is widely used.

PHP originally stood for "Personal Home Page" and was an ad hoc development that was posted for public use in 1995. It grew organically as a team of developers got involved. PHP 3 was released in 1999 and the name was changed to "PHP Hypertext Preprocessor".

Client-Server Applications

Before the web, the predominant model of time-sharing applications within organizations was client-server, where software was developed for client machines that communicated with software running on central server machines. The most widely used client machines were PCs running Windows and servers were mainframes or the larger minicomputers. A lot of client software was developed for Windows machines and it needed to be distributed and installed on the computers of all users who needed to use the application. With a web browser standard on all user machines, it became preferable to replace these systems with web applications, where the user only needs a web browser to access the server. Custom client software no longer needed to be installed or updated when new versions were released. In addition users had become very familiar with the web user interface, whereas custom client software often had different user interfaces.

So web development has become important in organizations, not only for development of public web sites but also for many internal client-server applications, from expense reports to customer order tracking. In addition, many of the third party software packages purchased by organizations also provide web interfaces for users. The server software needs to be installed and managed on an internal server but the end users only need their web browser.

4. History of Smartphones

Since cell phones are based on radio technology, we need to start about 200 years ago to see how radio waves were discovered and how radio technology evolved through many inventions, business competition and patent wars to the point where they could be used for mobile voice communication.

Discovery of Radio Waves

One of the great problems in physics in the 19th century was understanding light. There were a number of experiments that supported the particle theory of light, which held that small particles of light were sent out from light sources like the sun or fire. But there were also a number of convincing experiments that supported the theory that light was emitted as a wave.

Other great mysteries were electricity and magnetism. Scientists had been doing experiments for centuries. A fair amount was known about their properties but there were a multitude of contradicting theories.

In 1831 Michael Faraday published his discovery of electromagnetic induction. By moving a magnet through a coil of wires he “induced” an electrical current in the wires. He published this and other observations along with a mathematical equation known as Faraday’s law. This was the basis for the electrical industry, such as electric motors and generators. Generators use a physical force to rotate large magnets inside multi-layered coils of wire to create electric power. The physical force may come from waterfalls, wind turbines or steam produced by burning fossil fuels or by nuclear power to boil water. Electric motors work in reverse by passing electricity through wire coils to rotate a magnet. The circular force of the magnet can be transmitted by various mechanical means to many different types of machines and products that use electric motors.

In 1864 James Clerk Maxwell proposed a mathematical model of Faraday’s work to explain how induction travels through space from the magnet to the coil of wires. The equations predicted the existence of waves of oscillating electric and magnetic fields that travel through empty space at a speed that was very close the speed of light. Since this was highly unlikely to be a coincidence, he concluded that light was an electromagnetic wave.

However, the existence of electromagnetic waves was not proved until 1887. Heinrich Hertz did a number of experiments using an electrical spark transmitter to produce waves at a lower frequency than visible light. His measurements showed that the waves obeyed the Maxwell equations and traveled at the speed of light. This proved Maxwell’s theory of the electromagnetic field.

The frequencies of the waves that Hertz’s apparatus produced were what came to be called radio waves. This is how radio was discovered. The unit of frequency, cycle per second, was named the “hertz”.

Radio waves are just one type of electromagnetic radiation. The following figure shows the complete electromagnetic spectrum with the range of all of the frequencies and wavelengths. Note that visible light that humans can see is only a small band of the electromagnetic spectrum and that the frequency determines the color.

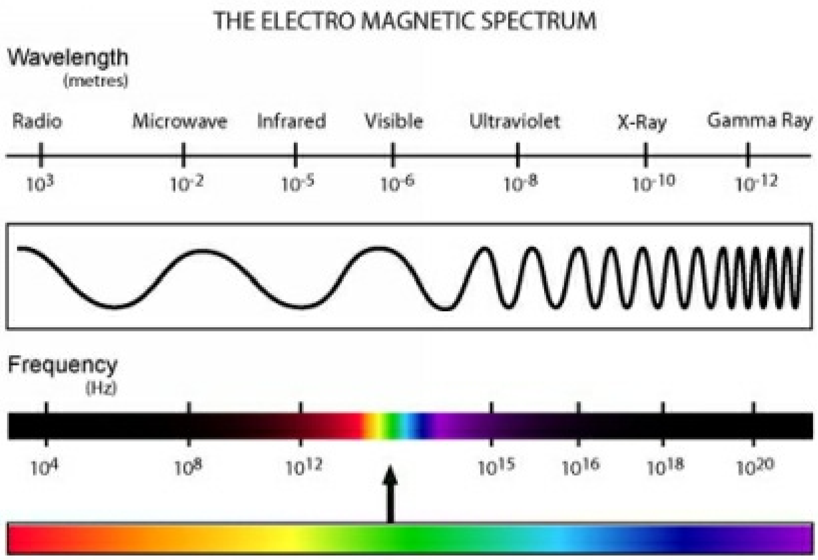

Maxwell’s theory was the second unified theory in physics. It explained what had been though to be separate phenomena of light, electricity and magnetism into a single phenomena, electromagnetic radiation. The first unified theory was of course Newton’s law of universal gravitation that explained both gravity on Earth and the forces that determine planetary motion. Maxwell is considered one of the greatest physicists in history. Albert Einstein described Maxwell’s work as the “most profound and the most fruitful that physics has experienced since the time of Newton”.

Commercialization of Radio

Telegraphy had been invented in 1837. It transmitted a simple electrical signal through wires. Messages were sent using Morse code, where each letter of the alphabet was a unique sequence of dots and dashes. So the international distress signal SOS (Save Our Souls) is sent as . . . - - - . . .

Since telegraphy could only be sent through wires, it could not be used by ships at sea. Many people had tried various systems of wireless telegraphy but had failed to develop a viable method.

After Hertz’s discovery of radio waves, Guglielmo Marconi conducted many experiments using radio and was able to gradually increase the distances over which he could send and receive radio signals. Marconi was granted a British patent in 1896 and built a number of transmitting stations. In 1901 he announced that he had transmitted a radio signal across the Atlantic from England to Newfoundland.

Marconi established companies in Britain, Canada and the U.S. to manufacture wireless telegraph equipment and build a network of ground stations for communication with ships at sea. This was a great improvement — marine communication at the time relied on the use of carrier pigeons. Ships used wireless telegraphy to provide services for passengers who wanted to send and receive telegrams. Ships also received daily news summaries from the shore stations that they printed as newsletters. Cruise ships today still use these services.

The Titanic was outfitted with a wireless telegraph room and operated by two Marconi employees. When it struck an iceberg and sank in 1912, wireless telegraphy was a critical factor in getting aid quickly enough from nearby ships to save many passengers.

The Marconi companies became leading businesses in the radio industry. Guglielmo Marconi was awarded the Nobel Prize in Physics in 1909.

There were many other inventors developing and patenting improvements for wireless telegraphy. Many disputes over patent rights occurred. When Marconi tried to get patents in the U.S. for his wireless telegraphy, they were denied in 1903 on the basis that they infringed on earlier patents granted to Nikola Tesla.

Nikola Tesla was a famous and wealthy inventor, based on his earlier inventions for an electrical motor that could use alternating current (AC). Tesla sold the patents to Westinghouse Electric and worked with them to perfect the technology and make sure it worked with the other components of an AC electrical system. Tesla Motors, the company that makes electric cars, is named after Nikola Tesla.

One of the most important inventions that advanced radio technology was the vacuum tube. While it was first patented by Thomas Edison in 1883 for use in his light bulb, it only functioned as an amplifier after several improvements were patented by Lee De Forest in 1906. By using vacuum tubes to amplify radio signals, much weaker signals could be transmitted and received over longer distances. This also made it possible to transmit audio instead of just the dots and dashes of wireless telegraphy. There were early experiments with radio broadcasting but it did not have good enough sound quality or reliability to make commercialization feasible.

World War I accelerated the development of radio. Countries that entered the war generally shut down all civilian radio stations and directed companies in the radio industry to focus on the war effort. By the end of the war, wireless telegraphy equipment was much smaller and more portable. Marconi in Britain and AT&T (American Telephone & Telegraph) in the U.S. developed air-to-ground voice systems to allow pilots to easily communicate with ground crews instead of having to punch out Morse code. Later in the war, two-way voice radio systems were developed.

Radio and Television Broadcasting

After the war, radio technology was sufficiently advanced to make public radio broadcasting feasible. Starting in 1920, many radio stations popped up all over the world. Radio became an important mass medium. There was a wide variety of programs from news and sports to comedies, dramas, variety shows, game shows and music. It is estimated that by 1929, approximately 40% of American families owned radios, and up to 75% in larger metropolitan areas. The period from the late 1920s until the widespread availability of television in the 1950s was the Golden Age of Radio.

Television broadcasting was a relatively simple extension of audio radio to include video signals. But other technology was needed to record and display the audio/video television signals. This was provided by the cathode ray tube (CRT). Television broadcasting started experimentally in the 1920s and grew very quickly in the 1950s. By

1958 about 83% of homes in the U.S. had a television set. Many other broadcast methods besides radio waves were subsequently used to transmit television signals, such as cable and satellite.

Mobile Car Phones

After World War II, a wide range of mobile telephone services were introduced in different parts of the U.S. and Europe and were used until the advent of cellular phones in the 1980s. These systems used powerful transmitters that covered large areas. Consequently the receivers needed to be powerful, resulting in phones that were large and heavy, so they were generally only used from cars and other vehicles.

Mobile Car Phone

There were not very many channels available, so in a dense urban area you might have to wait before you could make a call. Because of the limited capacity of the system, you might not even be able to get a phone and had to put your name on a waiting list.

Cellular Phone Networks

The limitations of mobile phone technology based on powerful radio transmission from single towers was understood early on. Engineers at Bell Labs proposed using multiple low-power transmitters in a hexagonal grid as early as 1947. But the technology to build such a system did not exist.

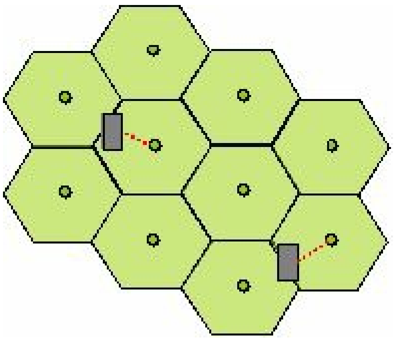

Cellular network showing transmission towers and cell phones

With the advent of transistors, integrated circuits and the microprocessor in 1971, cellular phones became feasible. Another group at Bell Labs expanded the early ideas and published papers that provided the basis for modern cell phone technology. This became the first cellular network standard in the U.S. It became available as a service in 1983 when the FCC granted licenses to an AT&T subsidiary, Advanced Mobile Phone Service (AMPS).

First generation cell phone networks were also available in many other countries. Some systems were in service earlier than in the U.S., in Japan in 1979 and the Nordic countries in 1981.

In the meantime, Motorola had been working on a portable, handheld mobile phone. It was demonstrated in 1973, when Martin Cooper, head of the Motorola team, walked down Sixth avenue in New York city and called his rival Joel S. Engel at Bell Labs. Motorola's DynaTAC phone became available to consumers when commercial cellular phone service began in 1983, at a list price of around $4,000.

Martin Cooper with Motorola handheld mobile phone, 1973

Since the first generation (1G) of cellular phone networks, there have been new advances and standards about every ten years. The following picture illustrates the evolution of 4 generations of cellular phone technology, as reflected by typical handsets that were used in each generation.

1G – Analog cellular networks

The first cellular systems used analog radio transmission, by modulating the frequency of the carrier signal, i.e. FM (frequency modulation).

2G – Digital cellular networks

The second generation used digital instead of analog transmission. The sound signal is converted to a digital stream of 0's and 1's. Transmission of a digital signal is more efficient and has less noise loss than analog, so sound quality is better. Data (Internet) service and SMS text messaging were also added. The rise in mobile phone usage exploded as a result of 2G.

3G – Mobile broadband data

The main advance from 2G technology was for data transmission using packet switching over a data network rather than circuit switching over the phone network. The higher connection speeds enabled streaming of media content.

4G – IP networks

The major advance in 4G cellular networks is speed increase for Internet use, potentially approaching the speed of cable broadband connections. The technology is based on using an IP packet-switched network.

Personal Digital Assistants (PDA)

In the 1990s, PDAs were very popular. These were small handheld devices that had a calendar in which you could put appointments, an address book for contacts, a calculator, a note pad and various other functions. They could be connected to your computer to synchronize calendar and contact list information between the PDA and the computer.

In 1993, Apple launched the Newton MessagePad and introduced the term "PDA". It used a pen interface and had handwriting recognition software. It had various technical problems some of which were improved with later releases, but the product could not compete with other PDAs that were much less expensive.

In 1996 the Palm Pilot was released. At a very attractive price of $149, it came to dominate the PDA market.

Palm Pilot PDA, 1996

The other PDAs that were popular were those that licensed Microsoft's Windows CE operating system. Many companies built the hardware and sold PDAs using this strategy, such as Compaq, HP and Toshiba.

Manufacturers of PDAs did make various attempts to integrate cell phone capability, notably the Palm Treo in 2002. But other vendors were ahead in the smartphone market.

Smartphones

The first smartphone was the Simon Personal Communicator, developed by IBM in 1994. It was a cellular phone with a touchscreen that could send and receive emails and faxes and had several other apps. However, at that time, Internet connections over cell phone networks were slow and unreliable so a web browser was not provided. This phone was ahead of its time and IBM decided not to continue with further development.

BlackBerry became an important player in smartphones. Its first products in 1996 were interactive pagers that could also send and receive email. The key BlackBerry feature "always on, always connected" was accomplished by developing a server that pushed mail to the pager as soon as it arrived. The BlackBerry became a must-have device in the business world. In 2002 BlackBerry integrated its pager with a cell phone and established BlackBerry as a leader in the smartphone industry.

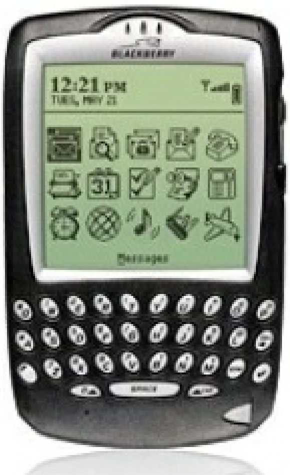

BlackbBerry Smartphone, 2002

BlackBerry concentrated on the business market where it had great success, but this all changed in 2007 when Apple entered the smartphone market with the iPhone.

Apple iPhone I, 1977

The iPhone was a revolutionary product in many ways. It was announced as a merging of the Apple iPod music player, a cell phone and the Internet. A distinctive feature was a touchscreen. While touchscreens had been used as far back as 1972, they had typically not been very precise. Apple perfected the touchscreen so that it worked well with simple finger motions. This allowed the iPhone designers to dispense with a physical keypad and pop up a keyboard on the screen when input was required, resulting in a much larger screen display than typical in mobile phones of the time. It also allowed the phone to use keyboards that the user could configure in languages other than English. The large screen, in combination with a full featured web browser, made using the Internet much more useable than other mobile phones. All these features were integrated into a well designed user interface, an area in which Apple had been an industry leader ever since the Macintosh computer. The result was that the iPhone was easier to use than most other phones.

Unlike the BlackBerry, the iPhone was directed at the consumer. It became a huge hit and propelled Apple to become one of the most valuable companies in the world.

Apple also introduced another revolutionary concept. They encouraged independent software developers to write apps for the iPhone by providing a set of development tools and making it easy for them to distribute their apps on the iTunes store. This was extremely successful and the wealth of apps that became available drove further use and sales of the iPhone. As of 2015 there were 1.5 million apps in Apple's App Store.

Needless to say, many other companies copied the features that made the iPhone so successful. The most successful competitors were the Android phones.

Android was a small start-up company that was developing a generic operating system for smartphones that could be configured for different hardware. In 2005, Google bought the company and promoted the platform to phone manufacturers and telecom carriers. A consortium of major companies in the industry formed the Open Handset Alliance, devoted to advancing open standards for mobile devices and Google released Android as open source code. Android phones started to appear in 2008. The most successful of these has been Samsung. In 2012, Samsung surpassed Apple as the largest seller of smartphones on a worldwide basis. There are also about 1.5 million apps in Google Play, the app store for Android phones.

The Future

Since the first general purpose computers became available 60 years ago, we have gotten used to the rapid changes in computer technology and their impact on many areas of life, from business, science, engineering, entertainment and publishing to information and society in general. While change and innovation have been continuous, as shown in the outline for this book in Chapter 1, there has been a major revolution heralding a new era on average every 14 years.

1936 Early Computers

1951 First Commercial Computers

1964 IBM Mainframe Era

1981 Microcomputers

1995 Internet

2007 Smartphones

Following this 14 year pattern, the next revolution will occur around **2021**. What will it be?

There are a number of emerging technologies that could become the next big thing.

- **Wearable computers** — a number of devices that you can wear are already available, such as fitness monitors, smart watches and virtual reality headsets. After an early trial, Google Glass is being redesigned
- **Robots —** electronic robots have been used in factories since 1961 as special purpose machines that can only do very specific things. Intelligent general purpose robots that can be used by voice commands for many things are an active area of research that is showing much promise.
- **Internet of Things** — the **IoT** is envisioned as a huge extension of the Internet that will interconnect small devices, sensors and appliances as well as computers. According to some estimates there will be 26 billion devices on the Internet of Things by 2020.

But the next revolution, whatever it is, will likely be the last. The reason is the end of Moore's law, which states that density of transistors on integrated circuits doubles every two years. The advances in the semiconductor industry following Moore's law have been the engine behind most of the revolutions in the computing industry.

However, there is a physical limit as to how small the transistors in an integrated circuit can become. In 1995 Moore and others started to predict when we would approach that limit. In 2015 Intel changed its road map and their new chips will lag behind Moore's law for the first time. There is now a consensus in the semiconductor industry that Moore's law will end around 2022.

There is ongoing research to use materials other than silicon for semiconductor manufacturing. In addition, alternative computing techniques such as quantum computing, optical computing and biochemical computing show promise. But it will be a long time before scientists can demonstrate processing at the speed of silicon. It would also take some time for engineers to develop manufacturing processes for new computing methods and scale them up to the size of the semiconductor industry at an acceptable cost.

After 2022, the computer industry will become a "mature industry". Growth and innovation will continue, but it will be characterized by incremental change rather than revolutionary change. This has happened to many great inventions of the 20^{th} century, from automobiles to airplanes to the radio and television industry.

About the Author

After completing a B.Sc. in psychology and mathematics at McGill University in 1967, I became interested in seeing how computers might be used to model human thinking and learning. Based on a lunch time course in Fortran given by the data center at McGill (there was no computer science department at the time), I got hired into a training program at Univac in London, U.K. Univac was the company that had built one of the first general purpose commercial computers. After a stint as the operator of a mainframe computer in a data center, I joined the Univac support team at the National Engineering Laboratories near Glasgow. There I worked on configuring and generating versions of the operating system to test the time-sharing support for Teletype terminals.

I returned to university to study "Artificial Intelligence" and completed a Masters degree in Computer Science at the University of Wisconsin. Following this, I worked in a university computing center in Vancouver, B.C. on a large IBM mainframe computer. There I provided consulting and programming services for faculty research from literature to genetics, using a variety of computer languages: Fortran, COBOL, PL/I, APL, Assembler and a number of statistical analysis packages.

My interest in doing research shifted to an interest in building applied systems. I joined the Toronto branch of a service company that sold computer time to customers. It used Xerox Sigma computers, one of the best timesharing machines then available. I promoted the use of APL ("A Programming Language") and used it to develop a number of financial programs that customers used. APL was a very interesting computer language, unlike any other. It operated on arrays of text or numbers and had a large set of special operators, many of which were Greek letters. APL was originally developed at IBM but it had been implemented on many other computers.

Shortly after the IBM PC was released in 1981, I joined a partner in a start-up company. We were literally "two guys working out of an attic". We developed an integrated financial planning package that achieved moderate success when combined with consulting services, but ultimately it could not compete with the widespread popularity of spreadsheets.

I next worked at a company that had developed a computer terminal that could connect to mainframe computers using different network protocols. These terminals were a precursor to the Bloomberg terminal and were widely used by the Canadian stock brokerage industry to provide stock quotes and financial information. One of the many products I worked on was an IP router that ran in an IBM PC. It forwarded IP packets from a local area network to another IP network over an X.25 connection. At the time, X.25 was a public global network that was more widely used than the Internet. X.25 was a standard that had been developed by the telephone industry and X.25 networks were provided by telephone companies in many countries.

When the Internet started to take off, I joined a start-up that had built one of the first firewalls. A firewall computer is positioned between the Internet and an internal network. It protects the network from hackers on the Internet by blocking unauthorized access to the organization's internal network. Firewalls were a new idea at the time but are now used by every organization with an Internet connection. I worked on several versions of the firewall product using C, Java and Perl in a customized version of Unix that ran on PCs.

I then worked as a software engineer at a company that made network switches for Storage Area Networks (SAN). SANs are high speed networks that allow large computer servers to share very large disk storage systems, such as those manufactured by IBM and EMC. SANs used a network protocol called Fibre Channel but there were new standards being developed at the IETF to support storage traffic over IP networks. Working with the CTO (chief technology officer) and other engineers, I worked on the research and design of a gateway that would transmit storage traffic between IP networks and Fibre Channel networks. This project included numerous trips to Silicon Valley to evaluate start-ups that were developing hardware chips for IP storage. Eventually we bought one of the companies and I worked on the integration of software development between sites. Software development used C/C++ in an embedded operating system that ran on an IBM PowerPC processor chip.

Following this I joined a start-up that was developing an enterprise software security system to protect confidential data on disks and mobile storage. Working in a small company, I took on many roles; I worked to define the product, drafted patent applications, set up and managed the source version control, bug tracking system, developed test plans and supervised the QA team. Code development was for Microsoft Windows using Visual Basic, C and the Windows SDK.

My last position before retiring from a career spanning 48 years was Research & Development for a domain name registrar that manages .org and several other top level domains. I developed prototypes for several research projects using open source code enhanced with custom code written in C, HTML/PHP, Python and Java on Linux using Amazon cloud servers. A "Big Data" project that analyzed DNS traffic was done on a cluster of Linux servers. Other projects were for new Internet standards being developed by the IETF. I participated in the Internationalized Email working group that extended the email standard to support email addresses in any language and another working group that developed a new protocol to replace the Whois domain lookup service.

Ernie Dainow