Chapter 0
Where We Are – How We Got Here

"Any sufficiently advanced technology is indistinguishable from magic."
— Arthur C. Clarke.

This book will be in the hands of students in the third decade of the 21st century. Although there were early attempts at mechanical computing machines, electronic, programmable, digital computers have been around for only about 80 years.

In the twenty-first century, digital computers are everywhere. You almost certainly have one with you now, although you may call it a "phone." You are more likely to be reading this book on a digital device than on paper.

An electronic digital computer is a machine. It has mechanisms for receiving input of information, for processing and storing information, and for producing and communicating results. You will find computers in all kinds of specialized devices, but the computer itself is one of the most general-purpose machines ever devised.

Figure 0-1
Vannevar Bush
U.S. Library of Congress

The more general a technology, the more and more varied the uses to which it will be put. We use computers for everything from preparing tax returns and processing them to landing airplanes, generating medical images, and displaying pictures of cats.

0.1 Computing in the 21st Century

As we begin the 21st century, those ubiquitous "phones" are the most numerous computers. We use them to communicate with both voice and text, but each one also has more computing power than all of NASA during the first, and so far only, manned moon landing in the summer of 1969. Although phones are the most numerous computers, they aren't the only ones. Many individuals and nearly all businesses also have larger personal-use computers, from tablets like the iPad to laptop computers and desktop computers. People use more powerful personal-use computers called workstations for graphic arts and scientific computing.

Huge data centers house thousands of pizza-box size computers and power enterprises like Google, Facebook, and Amazon along many others. Similar data centers offer "cloud" computing and storage to individuals and to businesses of all sizes from tiny to Fortune 10.

Mainframe computers, machines with very large memories and enormous input and output capabilities, once called "dinosaurs," still power about two thirds of the Fortune 100 businesses. The principal reasons are outstanding ability to perform input and output, which we will cover in more detail in Chapter 4, and the ability to handle hundreds or thousands of transactions per second while maintaining database consistency. Database technology is covered in Chapter 6.

Supercomputers, computing machines vastly more powerful than those used by most people, perform such jobs as weather forecasting and complex simulations. We now call the use of supercomputers *high performance computing*.

Tiny computers embedded in machinery from microwave ovens to automobiles, ships, and airplanes perform control functions formerly handled by special-purpose mechanical or electronic devices. Using programmable controllers in this way makes the functions they perform both more flexible and less expensive than the custom components they replaced.

Not only are computers everywhere, they are nearly all interconnected. Most of the students reading this book cannot remember a time when Internet access was not available to nearly everyone. Today Internet access is available with our phones, in our homes and places of work, and in our cars. The 5G cellular technology being widely deployed as this is written and the 6G technology coming next promise to make speeds in the billions of bits per second available nearly everywhere.

We use computers primarily for communication, from text messages and phone conversations to surfing web sites and streaming sound and video. We also use computers for storage of very rapidly growing amounts of data, for control functions, and for computational tasks like simulations. Making computing machines "intelligent," for some definition of intelligence, has been a goal of computer scientists for nearly as long as there have been computers. The 21st century has seen remarkable progress in machine learning and artificial intelligence. As storage, processing speed, and input and output capabilities have improved, it has become possible to analyze truly huge quantities of data. The ability to process "big data" has spawned new disciplines of data science and analytics.

None of this will seem new to most readers of this book, but the way we got here is a fascinating story. It probably began in pre-history when people first began to count their possessions and to trade with one another. Advances in mathematics from the 17th century onward laid some of the foundations. Although there have been mechanical counting and computing devices since ancient times, the foundation for today's electronic computing machines was laid much later.

0.2 Early Ideas

Computers operate on numbers, and the development of systems of numbers suitable for use as the internal representation of data in a computer is a story on its own. That story appears in the next chapter.

A fundamental characteristic of computers is that they are *programmable*. Different programs allow the computer to do different things; computers are general purpose machines. Possibly the first programmable machine was invented in the years just after 1801 by French weaver Joseph Marie Jacquard. Jacquard invented a weaving loom that could programmed by punching holes in cards. Once a set of cards was prepared, the same fabric design could be woven over and over, error-free. Different cards wove a different design. The Jacquard loom was programmable. In the 1960s and 1970s, punched cards similar to Jacquard's were used as input for electronic computers. Jacquard looms are still used in the 21st century, but the programs are stored in electronic memory, not on punched cards.

Celestial navigation – navigation using the positions of the sun, moon, planets, and stars – was important to the British Navy, merchant shipping, and military forces everywhere until the development of radio navigation systems in the 20th century. Celestial navigation used a sighting instrument called a sextant to measure angles to celestial objects and a book of tables called an ephemeris to determine the position of that celestial object at that time. Preparation of the tables was time-consuming, and calculation, transcription, or typesetting errors could result in lost ships.

About 1822 Charles Babbage, a British mathematician and philosopher, proposed to automate the production of such tables, including the automation of printing plates, to eliminate sources of error. With funding by the British government, Babbage designed and began construction of a mechanical computing machine called the Difference Engine. He and the craftsman hired to build the Difference Engine disagreed over costs and the project was never completed. Babbage designed an improved version but was not able to get funding to complete it.

The Difference Engine was a special-purpose mechanical computer. Confident that his design was correct,[1] Babbage went on to design a general-purpose mechanical computing machine, the Analytical Engine, that was remarkably similar to a modern computer. The part Babbage called the mill corresponds to the central processing unit of a modern computer, and Babbage's store corresponds to the memory.

Lord Byron's [2] daughter, Augusta Ada King, Countess of Lovelace, had corresponded with Babbage because of her interest in mathematics. She was fascinated by the Analytical Engine, and in her *Notes* included programs to be run on the Analytical Engine, possibly making Lady Ada the first computer programmer.

---

1 It was. Near the end of the 20th century, the Science Museum in London built two Difference Engines from Babbage's plans and to the manufacturing tolerances of the 19th century. The Difference Engine worked as intended.

2 The same Lord Byron you met or will meet in English classes.

In 1854, George Boole, another British mathematician and philosopher, published *An Investigation of the Laws of Thought* in which he expanded on his earlier idea of an algebra over finite sets of discrete values. We now call that Boolean algebra. It was not until 1937 that Claude Shannon, an American mathematician and electrical engineer, wrote in his master's thesis that, if the discrete values of Boolean algebra were zero and one, or on and off, they could describe switching systems, and that Boolean algebra could be used to optimize switching systems. We'll meet Boole and Shannon again in Chapter Three.

0.3 The World War II Era

During the World War II era, several important inventions were made at about the same time by different people or groups in different places. Simultaneous invention occurred at other times in history, too. Some historians say these inventions were possible "because it's time." That is, the necessary knowledge and technology had been developed to allow invention to occur. By the 1930s electricity was widely available in cities and towns, electrical engineering was an established discipline, and electronic components like vacuum tubes were being mass-produced. Babbage's idea of a programmable computing device was over a hundred years old, as was George Boole's algebra. The Boolean algebra was applied to switching by Shannon in 1937. Alan Turing's 1936 doctoral dissertation described what Alonzo Church later called a Turing machine.

The World War itself provided incentive to complete quickly anything related to the war effort and provided liberal funding for research and development.

Howard Aiken of Harvard conceived the Harvard Mark I computing machine, which IBM called the Automatic Sequence Controlled Calculator. Aiken presented the concept to IBM in 1937 and IBM funded and built it. The Mark I included and improved upon many of the features of Babbage's Analytical Engine. One of the uses of the Mark I was the production of mathematical tables, Babbage's original intention over a hundred years earlier.

Also in 1937, John Vincent Atanasoff of Iowa State College and graduate student Clifford Berry conceived and began work on the Atanasoff-Berry Computer, or ABC. The ABC used binary numbers internally and had an electronic arithmetic and logic unit. World War II interrupted Atanasoff's work on the ABC and it fell into obscurity.

Artillery in World War II was aimed using firing tables much like the tables of an ephemeris. Each kind of artillery piece needed its own firing tables, and the tables included such factors as kind of shell, kind of propellant, distance to the target, and wind direction and velocity. They were computed by hand by women whose job title was "computer." The process was slow, and new artillery pieces were being developed faster than firing tables could be prepared for them. Army Lieutenant Herman Goldstine convinced his superiors that electronic computation of firing tables would get them into the field faster. That was the beginning of the ENIAC[3] project. J. Presper Eckert, John Mauchly, and many others were involved in the design and construction of ENIAC at the University of Pennsylvania. ENIAC began solving operational problem in December 1945. The first real problem was a calculation for the Los Alamos National Laboratory related to the development of the hydrogen bomb, World War II having ended in September.

ENIAC was entirely electronic, having no mechanical parts. However, it did not yet have the capability to store programs. A mechanism to allow for stored programs, and eventually an electronic memory, were added later. ENIAC used vacuum tubes as its computing elements. ENIAC remained in operation until 1956.

Those building the ENIAC learned quite a lot about the design of electronic computers, much of it too late to be implement in ENIAC itself. Eckert and Mauchly proposed to the Army's Ballistic Research Laboratory another computer to be called EDVAC.[4] John von Neumann,[5] then at the Los Alamos National Laboratory, was serving as a consultant on the ENIAC project. In June, 1945 while returning to Los Alamos from the Moore School of electrical Engineering in Philadelphia, von Neumann compiled a set of design notes which he called *First*

3 Electronic Numerical Integrator and Computer

4 Electronic Discrete Variable Automatic Computer

5 Pronounced von NOY-man.

Draft of a Report on the EDVAC.[6] Von Neumann mailed the handwritten notes to Herman Goldstine, who was the Army's project manager for the construction of computers at the Moore School. Although von Neumann's notes were a distillation of the work of many people, Goldstine had the notes typed and circulated listing von Neumann as the sole author.

Von Neumann's *First Draft* described what came to be known as the von Neumann architecture and is a description of how computers are still designed. He described a computer with an arithmetic unit, a control unit, a memory, input and output, and storage. Programs and data were to be stored in the same memory, instructions to be executed in sequence, and data represented as binary numbers.

There were significant developments in computing in the United Kingdom during World War II as well, but, because of the British Official Secrets Act, most were not revealed until the 1970s. A key figure was Alan Turing, who designed a code-breaking machine called a Bombe based on an earlier design from Poland. Turing also contributed to the design of the Colossus electronic computers. Probably Turing's most important contribution was the description, in his 1936 doctoral dissertation of what we now call a Turing machine. The Turing machine is a mathematical model of computation that can be used to prove that there exist problems that "look computational" but cannot be solved.

Eckert and Mauchly, having built ENIAC and EDVAC, formed the Eckert–Mauchly Computer Corporation to design, build, and sell electronic computers. Their second computer design was called UNIVAC.[7] It is important as the first of which multiple units were sold commercially.

0.4 Vannevar Bush and the Memex

"…a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding

6 If you are curious about what a description of how computers are still designed, you can find copies of the *First Draft* online or buy a typeset copy of all 47 pages!

7 Universal Automatic Computer

speed and flexibility. It is an enlarged intimate supplement to his memory" (Bush, 1945).

Apart from the somewhat outdated language and the citation date at the end, which is a dead giveaway, you might think this was a description of a 21st century smartphone. This remarkable description is from Vannevar[8] Bush's essay, "As We May Think," published in the *Atlantic Monthly* in July, 1945.

Dr. Bush was describing, not a smartphone, but a hypothetical device he called a **memex**. The size of an office desk, the memex was for information storage and retrieval. Bush was not trying to predict future technology. The memex was based on technology available at the time or likely to be available in the five- to ten-year future. For example, Bush discussed a camera that could take pictures that would be instantly available without development using chemicals in a dark room. The first Polaroid instant camera was sold in 1948; the first digital camera was not made until 1975, thirty years after Bush's article.

Bush's memex addressed the information storage and retrieval needs of that time, 1945. He was concerned by what we now call information overload. An important part of the memex was enough storage to hold all the documents an individual might need and the ability to connect documents with associative links, which we now call hyperlinks. Bush's description included the ability to respond to voice commands and even to take dictation automatically. He envisioned a head-mounted camera that would allow the user to photograph whatever was being seen. That was realized in 2013 as Google Glass.[9]

In short, Bush was describing using technology to manage information in ways that were impossible for an unaided human. In doing so, he foresaw much of the technology which we take for granted in the 21st century.

8 Pronounced va-NEE-var, to rhyme with achiever.

9 For those who may not remember, Google Glass was mounted on a frame like eyeglasses. It could take pictures or video, respond to voice commands, and display information in the wearer's field of vision. Yes, Google Glass is still a thing; Google Glass Enterprise Edition 2 is intended for workers in industrial, medical, or similar settings.

0.5 Transistors, Chips and Moore's Law

The ENIAC was an important advance in electronic computing, but it had an important physical limitation: its computational elements were vacuum tubes. Vacuum tubes are made of glass, and so relatively delicate to begin with. If the vacuum is lost, the tub will not function. Vacuum tubes work on the principle of thermionic emission. The cathode (negative element) was heated by a heater similar to the filament in an incandescent light bulb. Like the filaments in incandescent light bulbs, those heaters eventually failed; the bulb burned out. At first several tubes a day failed. By running the tubes at a slightly lower voltage, they reduced the failure rate to one every two or so days. (Randall, 2006) Even doubling the computing power would lead to a crash a day.

The other problem is heat. The ENIAC consumed about 150 kW of power, the equivalent of 1,500 100-watt incandescent bulbs and produced as much heat as 1,500 incandescent bulbs. It required two twelve-horsepower blowers to conduct that heat away (Williams, 1985).

The transistor, invented at Bell Laboratories in 1947 and independently by Herbert Mataré in 1948, could perform the same switching functions as vacuum tubes. That solved both problems. Transistors are solid-state devices, so do not need a vacuum or glass envelope. There is no heater, so transistors consume much less power than vacuum tubes and produce much less heat. A transistor as a discrete electronic component is about the size of a pencil eraser, while a vacuum tube is about the size of a human thumb.

The next step was the invention in 1958 of the integrated circuit "chip" nearly simultaneously by Robert Noyce, Jack Kilby, and Kurt Lehovec. An integrated circuit packages many transistors and other electronic devices on a single semiconductor chip. The first commercial chips had tens of transistors. Modern chips can have billions.

In 1965, Gordon Moore, then director of research and development at Fairchild Semiconductor, wrote in the *Electronics* magazine that the density of devices on integrated circuits doubles about every year, and that he expected this trend to continue at least until 1975 (Moore, 1965). Moore later revised his estimate to doubling about every two years. This came to be known as Moore's Law, although it is an empirical observation, not a physical law.

The significance of Moore's Law is this: the cost of a semiconductor chip is roughly proportional to the area and the computational power is roughly proportional to the number of devices. If the density doubles every 12 to 24 months, in two years or less, we can get the same computational power at half the cost or twice the poser at the same cost.

People have predicted the end of Moore's Law several times in the last decade. In Chapter 3 you will meet Chenming Hu, whose invention of the FinFET transistor probably extends the life of Moore's Law to 2030 or later.

0.6 The Internet and the World Wide Web

Most of the people reading this book are accustomed to send a message to a friend or acquaintance anywhere in the world and having it delivered in seconds or less. An observation called Metcalfe's Law states that the value of a network to a user is proportional to the number of people one can contact. The *total* value is $n \times (n-1) / 2$, the number of connections. (Shapiro & Varian, 1998) [10] According to the Pew Research Center (2021), as of February, 2021, 77% of U.S. households had broadband Internet access at home and another 15% did not have broadband access but use smartphones as their primary means of Internet access at home. An astonishing 92% of Americans have Internet access at home or via a personal device. Of those who don't, poverty or geography explain some. There are even some Americans who choose not to have Internet access.

The Internet is an *Inter*net to emphasize that it is not a simple network, it is a network of networks. Your home or apartment router acts as a gateway to a connection to any other Internet-connected network in the world. It is *The* Internet because there's only one, connecting nearly all the world's networks. Metcalf's Law shows the enormous utility of having a single internetwork to connect all the world's networks.

---

[10] According to Shapiro and Varian, this observation was formulated by George Gilder and attributed to Metcalfe.

The Internet had its' beginnings in 1966 at the U.S. Defense Advanced Research Projects Agency. A design based on the work of J.C. R. Licklider resulted, by 1969, in a packet-switched network of four nodes. Packet switching and the Transmission Control Protocol, TCP, both of which are integral to the Modern Internet, were developed at ARPA. We will visit those fundamental concepts in detail in Chapter 5.

A large number of standards, called RFCs, came out of the research funded by ARPA. A few of those are still in use, and many are direct ancestors of today's standards.

ARPANet was primarily a military network with university and other sites connected under contract to ARPA. Military sites were segregated onto their own network, with work completed in 1984. In 1986 the U.S. National Science Foundation founded a network using packet switching and the TCP protocol. The NSFNet was eventually opened to commercial access and finally operated entirely by commercial organizations since 1995.

The Internet of the 1980s was limited to electronic mail and file transfers. The University of Minnesota's Gopher[11] protocol provided a mechanism for indexing and searching Internet resources. There are still Gopher servers running, and Gopher clients are available for Windows, MacOS, Linux and Android.

The invention of the World Wide Web in 1989 by Tim Berners-Lee at CERN[12] changed the way most people access information, goods, and services. The Web protocols alone were not enough. The Mosaic Web browser was the first browser to be able to render pictures, and one of the first graphical browsers. Mosaic was developed at the National Center for Supercomputer Applications at the University of Illinois at Urbana-Champaign and released in 1993. The other missing piece was the search engine. Digital Equipment Corporation's Alta Vista search engine, available to the public at the end of 1995, was among the earliest attempts to index all the World Wide Web.

To put the magnitude of this change into perspective, in 1985 Delta Air Lines had over 20 city ticket offices in the Atlanta area and hundreds throughout the

11 The University of Minnesota's athletic teams were the Golden Gophers.

12 *Conseil Européen pour la Recherche Nucléaire*, or European Council for Nuclear Research

United States. Today there are no city ticket offices in Atlanta, and only one in the entire United States. It's in Manhattan on East 45th Street.

0.7 Convergence

By 1990 the telephone system in the United States was nearly all digital apart from the "first mile" from telephone subscriber to the telephone office or multiplexing point. Invention of the compact disc in 1982 and the digital video disk[13] about 1995 meant that nearly all recorded audio and video were in digital formats by the end of the 20th century. The invention of the World Wide Web in 1989 and the transition of the Internet from a U.S. government funded network to the commercial Internet beginning in 1993 fundamentally changed the way most people get access to information, goods, and services. Nearly every kind of information is now transmitted using the Internet Protocol. These changes led to the realization that all information, not just numbers and symbols, could be received, transformed, stored, and retrieved using computing and communication technology. That realization came to be called technological convergence, digital convergence, or just *convergence*. Convergence broadens the scope of practice of information technology. That means a broader scope of technology to be mastered.

0.8 Systems and Architecture

We talk routinely about systems, computer systems, operating systems, sound systems, alarm systems, and many others. A major part of information technology is *systems integration,* that is, making multiple systems work together to achieve some purpose. Given that *systems* are so important in the profession of information technology, it makes sense to have a definition.

We can define *system* as a collection of components linked together to perform specific functions, and recognizable from the outside as a unit. Everything outside the system is its *environment,* and systems communicate with their environments through *interfaces.*

13 Later called "digital versatile disk," and still later, just DVD.

For example, a sound system might have a CD player, an amplifier, and speakers. The CD player is an interface; recorded music is communicated to the sound system through the interface. The speakers are also an interface, changing electrical signals into acoustic waves that we can hear.

In information technology, systems provide computation or communication, or some other service such a projecting media on a screen. Such systems consist of hardware, at least two layers of software: operating system and application software, and interfaces.

There are several ways to design information technology systems. These design paradigms are called architectures. Common architectures include monolithic, client-server, and multi-tier.

In a monolithic architecture, a single computing subsystem performs all the system's functions. A desktop computer is a simple example of a monolithic architecture. Everything that happens within that desktop computing system happens within a single computing subsystem.

The World Wide Web is an example of a client-server architecture. The web server receives requests from the client, a web browser, and responds by serving web pages. Client and server nearly always run in separate computing subsystems and communicate through a network.

Web mail is an example of a multi-tier architecture. There would be a browser client and web server, but the web server might communicate with both the browser and a "back end" mail server running in a different computing subsystem.

0.9 The Importance of Standards

Every day we use computing and communication equipment made by myriad companies and we expect it all to work together. Usually, it does. No one is surprised when a Samsung smart phone connects to a network operated by AT&T. We aren't even surprised when we place a voice call to a friend on another continent and have a perfect connection. That music CD made in Germany plays flawlessly on a player made in Japan. The pieces all fit together seamlessly because of thousands of standards published by more than a dozen standard-setting organizations.

Absence or incompatibility of standards can range from an annoyance to a serious problem. The three wireless devices on the author's desk have three different chargers with three different connectors. In the early days of the World Wide Web, browser makers competed in part by inventing their own HTML tags. Writing complex HTML that displayed identically across multiple browsers was effectively impossible before the World Wide Web Consortium took over the maintenance of standards. Even after the publication of standards it was many years before browser-makers implemented the official standards uniformly.

Standards come from two types of organizations: official standards development organizations and standards development interest groups. The defining characteristic of standards development organizations is government authority or recognition. In the United States, it is the standards themselves, and not necessarily the standards development organizations, that receive government recognition. Interest groups have a less formal authority and are often formed as contractual organizations among companies that want to make their products interoperable (Biddle et al, 2010).

In the United States, the American National Standards Institute (ANSI) accredits standards-setting organizations. ANSI is a non-profit organization, not a government agency (Biddle et al, 2012).

The difference between a standards development organization and an interest group is illustrated by the wireless local area communication standards. The standard for short-range wireless Ethernet is established by the Institute for Electrical and Electronics Engineers, IEEE. IEEE acts as a standards development organization. The Wi-Fi Alliance, a group of over 400 manufacturers that provides testing, compliance certification and the "Wi-Fi Certified" logo is an interest group.

You will meet many important standards development organizations and interest groups as you work through this book.

0.10 No Such Thing as a Free Lunch

A recurring theme in Robert A. Heinlein's 1966 novel *The Moon is a Harsh Mistress* is TANSTAAFL, "There ain't no such thing as a free lunch." As you go through this book, and as you go through a career in computing, you will find that TANSTAAFL is a recurring theme in computing as well. The rapid progress described by Moore's Law can sometimes fool us into thinking we're getting something for nothing; technology that was prohibitively expensive even five years ago is suddenly within reach of almost anyone. As you dig deeper, you will find that all of computing is a study in trade-offs.

The design of computing hardware is full of such trade-offs. It's possible to add circuitry to memory that will allow memory errors to be detected and even corrected. Doing so not only increases the cost of memory modules, it also slows the memory down slightly. Most of the slowdown can be avoided with a technique called triple modular redundancy, but avoiding the slowdown increases the cost by nearly three times and also increases the power consumption and the physical space required. Chapter 2 introduces the ripple-carry adder. It is simple to design, easy to understand, but very slow if numbers of more than a few bits are to be added. There are other adder designs, not considered in this book, that are much faster, but also more complex and consequently more expensive. In Chapter 3, you will learn that there are two approaches to designing a control unit for a CPU. One approach is fast but relatively expensive; the other approach is less expensive, but also slower. In the field of software, there are whole classes of algorithms that can be replaced by others that are much faster, but only if sufficient memory is available. This circumstance is so common that it has a name: space-time trade-off.

0.11 Information Technology

The term information technology was coined in 1958 in an article in the *Harvard Business Review*. Although electronic computers were barely ten years old, the authors predicted coming major changes not only in how information was processed but also in how it would be used, with the result of major changes in organizational structures (Leavitt & Whisler, 1958).

We define the *practice* of information technology as the integration, operation, and maintenance of computing hardware, software, communication facilities, services, and infrastructure to receive, transform, store, and retrieve information of all types in an efficient, effective, and secure manner.[14]

The Association for Computing Machinery defines the *academic discipline* of information technology this way: "Information Technology is the study of systemic approaches to select, develop, apply, integrate, and administer secure computing technologies to enable users to accomplish their personal, organizational, and societal goals" (Task Group on Information Technology Curricula, 2017). Information technology is the applied discipline of the field of computing.

A look at the breadth of the ACM definition explains why information technology is such an engaging and fun profession: of the computing disciplines, it is the one requiring the most breadth of knowledge. Sometimes it's enough to know that a particular area of knowledge exists so that you can recognize when you need it and look it up. Sometimes knowledge at the conceptual level is sufficient. In other cases, substantial depth of knowledge is needed. Your information technology curriculum will prepare you with the knowledge you need. This book is a step in that direction.

0.12 Professionalism

This book focuses on the *computing* concepts for information technology, that is, a high-level look at the technology itself, with more detail in some areas. Other courses will provide even more technical detail. Not covered here is the broad area of skills, knowledge and behaviors generally called professionalism.

Information technology is increasingly a strategic asset and a customer-facing function for organizations. The obvious example is Amazon, but consider an airline trying to operate without information technology, and even the pharmacy that lets patients request refills online.

---

14 This is the kind of long-winded sentence that gets a D– from English teachers, yet you will encounter sentences like this in another definition immediately below and in all manner of disclaimers, software contracts, and the like. See the endnote of this chapter for advice on making sense of such a sentence.

When information technology is a strategic asset, its practitioners must work within the organization, not performing an isolated back-office function. That means practitioners need excellent communication and teamwork skills. They must also know the business of the organization as well as the details of the technology.

When information is customer-facing, errors and defects are highly visible and not well tolerated. In the batch processing environment of 60 years ago, if there was an error, one could fix it and start over. Now errors and defects cause dissatisfaction, loss of business, and even major news stories. These can be what is euphemistically known as career-limiting events.

"Ethics" means doing the right thing (Baase 2008). Sometimes the "right thing" is obvious, sometimes less so. For example, is it ethical for your spouse or partner to bid on a contract at your organization? Professionals make a concerted effort to understand what is right in varying circumstances, and then to do what is right. The incredible power of information technology makes ethical decisions far more reaching than they were in the past. Professional associations such as the Association of Information Technology Professionals and the Association for Computing Machinery have codes of ethics to which members subscribe and which provide help in discerning what is right in difficult situations.

The practice of information technology is also influenced by local, national, and sometimes international law. Often such laws address the operation of types of organizations and not the information technology directly. Professionals will know the laws that affect their industry and how those laws constrain the operation of the organization.

You will see as you work through this book that the face of information technology changes incredibly rapidly. The nature of the world in which organizations operate also changes, sometimes very rapidly. The practitioner of information technology must stay current with both the technology and with the organization's business. That requires a conscious effort at lifelong learning.

If you are using this book as part of a formal curriculum of study, you will find that the curriculum includes the elements of professionalism. If you are studying on your own, be careful not to neglect the elements of professionalism.

0.13 Summary of Learning Objectives

> *This section of the chapter tells you the things you should know about, but not* ***what*** *you should know about them. To test your knowledge, discuss the following concepts and topics with a study partner or in writing, ideally from memory.*

Be able to state Moore’s Law and explain its significance.

*** to be supplied after edits ***

0.14 References

Baase, S. (2008). *A Gift of Fire* (Third ed.). Upper Saddle River, NJ: Pearson Education.

Biddle, Brad, Frank X. Curci, Timothy Haslach, Gary E. Marchant, Andrew Askland, and Lyn Gaudet (2012). The Expanding Role and Importance of Standards the Information and Communications Technology Industry, 52 *Jurimetrics J*ournal. 177-(2012).

Bush, V. (1945). As We May Think. *Atlantic Monthly, 176*, pp. 101-108.

Leavitt, H. J., & Whisler, T. L. (1958). Management in the 1980's. *Harvard Business Review, 36*(6), 41-48, July, 1945.

Moore, Gordon E. (1965). Cramming more components onto integrated circuits. *Electronics*, (38)8, 114-117, April 19, 1965.

Pew Research Center (2021). Internet/Broadband Fact Sheet. April 7, 2021. https://www.pewresearch.org/internet/fact-sheet/internet-broadband/ Retrieved September 8, 2022.

Randall, Alexander (2006) Q&A: A lost interview with ENIAC co-inventor J. Presper Eckert *Computerworld* 40(8) 18, February 20, 2006.

Shapiro, C., & Varian, H. R. (1998). *Information rules: a strategic guide to the network economy.* Boston: Harvard Business School Press.

Task Group on Information Technology Curricula. (2017). Information Technology Curricula 2017: Curriculum Guidelines for Baccalaureate Degree Programs in Information Technology. New York: Association for Computing Machinery.

Williams, Michael R. (1985). *A History of Computing Technology*. Englewood Cliffs NJ. Prentice-Hall.

0.15 Endnote

This chapter includes the following horrible sentence: "We define the *practice* of information technology as the integration, operation, and maintenance of computing hardware, software, communication facilities, services, and infrastructure to receive, transform, store, and retrieve information of all types in an efficient, effective, and secure manner."

To make sense of such a thing, you have to deconstruct it. Often you can do this mentally, but here is an example in print. Try it on the other definition in this chapter.

We define
=====
the practice of information technology
===== as
the integration,
operation,
and maintenance
===== of
computing hardware,
software,
communication facilities,
services,
and infrastructure
===== to

receive,
transform,
store, and retrieve information
===== of
all types
===== in
an efficient,
effective,
and secure manner.