Introduction

1

The networking of computers—a technology more than twenty years old—is substantially impacting individuals, organizations, and business. The global public Internet now interconnects a significant portion of the world's computers, allowing them to communicate freely to realize shared applications. No longer is computing confined to serving isolated individuals or organizations. Networked computing supports groups of users (or organizations or enterprises) jointly participating in applications as well as the citizenry at large. This is a watershed in the history of computing.

Two previous technologies profoundly affected social and cultural institutions: mass transportation and telecommunications. Modern transportation transformed the urban landscape, leading to suburban growth, often at the expense of central cities. The telephone and television irrevocably affected politics and government; the free flow of ideas and a direct window into other societies have impacted political thought. Together, transportation and telecommunications led to a globalization of many communities, organizations, and institutions. They enabled and empowered the multinational corporation and the global superpower nation.

Networked computing is a seminal addition to this technical infrastructure. Like its predecessors, it will have a substantial and lasting impact on individual lives as well as business, social, and cultural institutions. It will facilitate the dissemination of information and knowledge, collaboration among individuals, business processes spanning geographically dispersed organizations, and commerce (among businesses and with consumers). Unlike telecommunications (which has emphasized voice and facsimile media), networked computing supports virtually all forms of information, such as data, images, and video, as well as money, and allows them to be integrated in innovative ways. Beyond information transport, networked computing supports software-defined storage and manipulation of information, and automates many knowledge-intensive tasks.

We have entered an information age, for which the key technological enablers are the network, the computer, and its software applications. This book is a guide to the possibilities of this information age, the technologies, and the supporting industry. It will position you to make best use of these far-reaching new technologies in your own area of expertise. The book will give you an appreciation of the possibilities of networked computing as well as the limitations. It will also position you to work effectively with professional programmers to implement your ideas and equip you with enough understanding of the industry to be an informed customer. Armed with this knowledge, go out and change the world!

1.1 A Historical Perspective

Networked computing is a collection of related technologies that support a broad range of geographically distributed computer applications. (It is distinct from a similar term in current vogue—the network computer, or NC—which describes one computing technology.) The computing portion of networked computing enables the storage, retrieval, and processing of tremendous amounts of information and also serves as an interface to users. A network enables computers to interact and share information, much like the telephone allows people to talk. A computing application is a software program that provides direct and specific value to a user or an organization, and a networked application distributes programs across two or more computers, which then collaborate in realizing the application. Users are the people leveraging the application for their job, to interact or collaborate with other users, or merely to have fun.

Computing technology has changed and expanded over the years, resulting in an expanding range of applications. As originally conceived in the 1930s, the computer performed massive calculations. Due to electronics advances, its computational capability continues to expand rapidly. Later, mass storage media (such as magnetic disks) extended applications to encompass the storage, retrieval, and manipulation of massive quantities of information. The relatively recent addition of networking allows computers to communicate and interact.

1.1.1 Technology View

When running an organization, your first impulse, if you want to become more flexible and responsive, is to decentralize. Then you discover there isn't enough coordination, so you establish hierarchical management structures, hold meetings, and generate memos and reports to improve internal communications. In consonance with technological advances, the computer industry has undergone precisely the same transitions, resulting in the major phases of computing technology shown in Table 1.1.

These phases are not mutually exclusive. Centralized mainframes still flourish—as the repositories of mission-critical corporate information—and are integrated into networked computing applications. Time-sharing still exists, in the sense that departmental-level computers support multiple users at their desktop computers. The isolated computer is a thing of the past; virtually all computers are networked today.

The most important enabling technology for networked computing is the network itself, which builds on data communications media (especially fiber optics) and the advances in electronics (which also made much faster and cheaper computers possible). Networked computing results in the convergence of two industries-computing and communications—irrevocablychanging each.

1.1.2 User and Organization View

Each phase of computing brought expanded opportunities and challenges. Until the early 1990s, mainframes formed the information core of major organizations' work processes. However, centralized and time-shared computing—in part because their applications

Table 1.1 Four major phases of computing.

Phase	Characteristics	Typical applications
Centralized computing (a few computers for a whole organization)	A relatively few mainframe computers were physically large and expensive, affordable only to large organizations.	Automate major business functions, such as payroll and accounting, and manage enterprise information resources, such as customer lists or inventories.
Time-shared computing (a few computers shared by a large number of users)	Terminals were added to allow a large number of workers to directly access applications on a centralized computer.	Workers could directly interact with the centralized applications to input data, initiate transactions, and extract information. Users could participate in shared applications.
Decentralized computing (a computer for every department and every user)	Less expensive centralized computers could be deployed at the departmental level. Inexpensive personal computers could be dedicated to a single user, supplementing the centralized computers.	Personal productivity was enhanced by word processing, spreadsheet, and small data management applications. Home users and students could benefit from similar applications, as well as others dedicated to personal finances, education, or entertainment.
Networked computing (computers can communicate and interact with one another)	All computers are connected by networks, allowing them to participate in geographically distributed applications on a global basis. Desktop computers can provide access to information in centralized computers. Networked computing is distinct from the "network computer (NC)" discussed in Chapter 3.	Networked applications not only benefit from more sophisticated user interfaces (graphics, pointing device, etc.) supported by the personal computers, but also access the processing power and massive data residing on servers and mainframes. Groups of users can participate in applications that enable collective communication and collaboration.

were provided by a centralized information technology (IT) organization—were comparatively unresponsive to the needs of departments and workers. Decentralized computing empowered users by allowing them to add their own applications and process data in personalized ways, but the downside was the organizational chaos resulting from inconsistent solutions. Looking at the transition to decentralized computing in retrospect offers the following insights:

- The users benefited from greater innovation by independent application suppliers. Users had greater computing power at their disposal, enabling, for example, user interfaces based on graphical user interfaces.
- Organizations had to adjust their information technology organizations to a new technological reality. Many organizational problems are still with us, such as high cost of administering a diversity of computers and fragmentation of the organization's information assets.
- For computer suppliers, the old strategy of proprietary turnkey applications became obsolete. The premier firms of the centralized computer industry didn't adapt fast enough (such as IBM, DEC, and Data General). Today, they have recovered, but many companies at the pinnacle of the industry were formed in response to decentralized computing.
- Desktop computers unleashed computing into every walk of life and have impacted the everyday lives of many citizens. Today, yearly sales of home computers exceeds television sets in the United States, and computing as a part of education, every occupation and job, and many avocations is increasingly accepted.

This gives an optimistic view, but these trends also have negative ramifications. There is an increased administrative overhead associated with decentralized computing, including the greater involvement of individual users in maintaining and upgrading their desktop computers.

While desktop computing had widespread ramifications, the impact of networking is even greater. In the mid-1990s, networked computing founded on networked microprocessor-based computers made significant inroads into applications previously the domain of mainframes, as the networked computers became sufficiently reliable and offered a path to new and reorganized work processes and greater customer satisfaction. Some of the implications of networked computing include

The transition from centralized applications that performed hidden "back-office" functions, to the expectation that most computing applications will intimately involve users, is complete.

- Networked applications empower computer-mediated interaction and collaboration.
- Applications can span organizational boundaries, opening up many opportunities in commerce (the buying and selling of goods and services, and their coordination).

Networked computing is a basic infrastructure for the society. As evident from the increasing attention from legislative, regulatory, and judicial authorities, networked computing has substantive impact on society and its citizens.

From an individual's perspective, computing has followed a trajectory from an invisible back-office function (centralized computing), to a tool for enhancing personal productivity and entertainment (decentralized computing), to an expanded role in accessing vast global information resources, and in interacting and collaborating with others.

1.1.3 Unrelenting Change

Progress in the underlying technologies of computing (processing, storage, data communications) has been dramatic and unrelenting for several decades. Likewise, driven by the rapidly decreasing costs of these technologies and better ideas of how to exploit them, the applications of computing have seen dramatic expansion and change. The computing industry is distinctly different from most others in this dramatic rate of change.

Applications implemented in the technology of yesterday are called legacy applications. One result of rapid change in the industry is an expanding set of legacy applications. Unfortunately, any application using today's most modern technology will be tomorrow's legacy application. History has seen—and this is likely to remain so for the foreseeable future—a continual evolution from today's to tomorrow's technology, while having to maintain and integrate legacy applications. One implication is the importance of a forward-looking rather than retrospective view of technology when conceptualizing applications—a view, I hope, that will be aided by the understanding imparted in this book.

1.2 Computing in the Future

The history of computing (and technology more generally) illustrates that progress is far from steady. Technology makes big leaps—from centralized to decentralized, from decentralized to networked—that take a while to be assimilated within the industry and the applications. Can future leaps be anticipated? Probably some can—those listed in Table 1.2—because they are well under way. A quick summary of the table is that computing will be anywhere, everywhere, and within. If you find this alarming, take heart in the observation that computing will put on a much different face. Replacing the big boxes that clutter our homes and offices will be smaller, less obvious, and less obtrusive computers [Wei93]

Computing is not the first technology to follow a similar evolution. A century ago, electrification (analogous to networking) had as its major applications light and mechanical power. Both light and power followed a clear evolution to mobility (battery-operated flashlights and power tools), ubiquity (electric outlets and appliances in every room), and embeddedness (small electric motors within many products). The greatest lesson from electrification-and also one evident in computing—is that the first step is retrofitting a new technology into the existing ways of doing things. Eventually, however, people figure out new ways to use the unique capabilities of the technology, and only then substantial gains become evident (see the sidebar "Electrification: Lessons from an Earlier Technological Advance" for a historical observation).

The same will be true with computing. You can tell that the technology and its applications have matured when computing is incorporated in everyday life in natural and unobtrusive ways, when computing makes things easier, more efficient, and more pleasurable, and when computing doesn't get in the way. This isn't quite true today, but it is forthcoming.

Electrification: Lessons from an Earlier Technological Advance

Electrification had its greatest impact on productivity and standards of living only after ways were found to exploit its unique characteristics. Computing will follow the same course.

In part, the industrial revolution substituted machinery and water or steam power for human labor. For a single factory, water or steam power from a central source (analogous to centralized computing) was distributed throughout a factory using cumbersome drive belts. The factory was organized around the distribution of power, with compact, multistory factories. Electric power initially had little impact because large electric motors were simply substituted for water or steam sources—nothing else was changed. It took decades to recognize that smaller electric motors powering individual machines (analogous to decentralized computing) enabled the reorganization of the factory around the needs of the work process (linear assembly lines) rather than power distribution, with dramatic improvements in efficiency and quality.

The electric motor also has a parallel to embedded com-

Table 1.2 Future trends beyond networked computing.

Trend	Description	Comments
Mobility (computing anywhere)	Networked computers can be taken anywhere and still benefit from full network services.	Laptop computers and personal digital assistants are the precursors. Mobility requires ubiquitous networking access analogous to the cellular telephone.
Ubiquity (computing everywhere)	Networked computers are unobtrusively sprinkled throughout the physical environment.	Information kiosks, mobile phones with Web browsers, and personal digital assistants are steps in this direction. In the future, as computers gain a similar size and resolution to paper, magazines, and books, computers should become as ubiquitous as the printed word is today.
Embedding (cornputing within)	Networked computers are embedded in most everyday products.	This is already common: Automobiles, consumer electronics, toys, and appliances have computing within. In the future, many more products—even as mundane as light switches and lightbulbs—will not only have computing within but also network connections.

Computing embedded in a plethora of products (such as automobiles and appliances) is "embedded computing in the small," and represents only the tip of the iceberg. Arguably even more important is "embedded networked computing in the large," which focuses on networked computing within many essential systems supporting society and the economy.

For some time networked computing has been embedded within (and crucial to the operation of) much of the infrastructure. It has been the controlling element for telecommunications and electric power networks and many aspects of the transportation system, such as air traffic control and train control. The world's financial systems, including the flow of money and financial markets, depend on a networked computing infrastructure.

Today networked computing is being integrated into a broader range of large systems. Large companies are automating many of their repetitive business processes, controlling the flow of material, finished goods, and money within their organizations. Increasingly, business-to-business relationships with customers and suppliers are automated by electronic commerce applications using networked computing. Even consumers are joining the fray, conducting a small (but rapidly growing) fraction of their financial and commercial transactions over the Internet.

A less tangible (but no less real) impact of networked computing is its application to the interaction and discourse among individual citizens on a global basis, utilizing Internet applications such as email, chatrooms, and newsgroups. This is important because the interactions among citizens (as well as social compacts) most clearly distinguish a society from a collection of individuals.

What is distinctive about computing embedded within these large systems is that computers, networks, and software are only a part of a system, which also includes citizens and workers, procedures, policies, and laws, the flow of material and finished goods, and many other aspects. To be most effective, technology cannot simply replace and automate existing functions (this is also illustrated by the historical example in the sidebar "Electrification: Lessons from an Earlier Technological Advance"). The design of organizational processes should create a holistic combination of workers and technology, determining what is best delegated to technology, and to people, and defining the interface between the two.

puting. Today, the electric motor isn't a separate consurner product, but is embedded in many products. Like computing, electric motors passed through a "personal motor with accessories" phase, until they became small and inexpensive.

1.3 Bits: The Atoms of the Information Economy

The industrial economy is giving way (in relative terms) to the information economy. The former focuses on the manufacturing of physical goods and the latter on the creation, access, and manipulation of information and knowledge.

EXAMPLE: The entertainment and software industries exemplify the information economy. An increasing fraction of the economy is "services," many of which are information based (accounting, law, education, etc.).

"Information" in this context includes text and numbers as well as other media such as art, video entertainment, games, money and financial instruments, and many other goods and services that don't have a physical presence. "Knowledge" means understanding and judgement based on large amounts of information and represents an intangible but essential asset of most organizations. Just as transportation and machinery form the technological foundation of the industrial economy, networks and computers are the technological foundation of the information economy.

Physical goods, broken down into their most fundamental and indivisible elements, are composed of atoms. Similarly, the fundamental and indivisible elements that represent information (as it is represented in networks and computers) are bits. Atoms are the building blocks of the physical world, and bits are the building blocks of computer-mediated information [Neg96]. Not all information is in the form of bits—for example, when you listen to the radio, the sound reaching your ears is not bits—but it is possible to represent any information (including sound) by bits (see the sidebar "Any Information Can Be Represented by Bits").

Bits are an appropriate building block for the information economy because they can easily be stored, communicated, and manipulated. The simplicity and universality of bits make possible computers, equipment, and software that can flexibly manipulate different types of information and also combine different types of information as necessary. This last point is important because older technologies (books or newspapers) lack this capability. In effect, bits are a universal alphabet for representing and manipulating all forms of information.

In the remaining chapters, analogies between the physical world and the information world will frequently be used to enhance your understanding of both.

1.4 Road Map to the Book

This book covers three main topics. The first is applications of networked computing; the second is the industry that supplies computing, networking, and software products and services; and the

Any Information Can Be Represented by Bits

A bit (short for "binary digit") is a number that assumes one of two values: "0" (zero) or "1" (one). Text such as you're reading now can be represented as bits by associating each character in the English language with a set of seven bits. By representation, it is meant that the original text can be recovered from the bits.

EXAMPLE: Unique sequences of seven bits can be assigned to each of the letters in the English alphabet, as for example:

a ↔ 0000000.
b ↔ 0000001,
c ↔ 0000010, etc

Altogether, 128 characters could be represented this way, more than enough for the 26 letters of the alphabet and punctuation marks.

In a similar manner, bits can be used to represent all written languages, including those based on ideographs (such as Chinese). Because written language is formed of characters (in English or Arabic) or ideographs (in Asian languages), it is discrete. This means that it is represented by characters drawn from a finite alphabet. The alphabet third is the basic concepts and terminology that underlie the infrastructure and technology of networked computing, from an application perspective. These topics are covered in roughly that order, although the goal is to emphasize their mutual influences. How do economic, legal, and policy issues impact the nature of the supplier industry and the infrastructure? What characteristics of the infrastructure impact the applications? What characteristics of the technology are ripest for exploitation in innovative new ways? What are the limitations of the technology that restrict the application possibilities? How are performance characteristics necessary for many large-scale applications obtained, and particularly, how do performance requirements impact the design of applications?

The first topic (covered in Chapter 2) is the applications of networked computing. These are divided into specific categories: social applications (that support the activities of groups of users), information management, education, and business. Each application area is broken down and classified—and numerous application examples described—with the goal of understanding the value and functionality the application affords to the user or organization. This is the most important chapter in the book and should not be skipped. It is only by addressing the goals and characteristics of applications that the technological and industry requirements can be inferred.

You doubtless have experience in using personal productivity applications on a desktop computer and thus understand what the monitor, graphical user interfaces, keyboard and mouse, etc., do for you. By comparison, you may know relatively little about the arcane terminology underlying networked computing (for example, messages, concurrency, packets, middleware, distributed objects, and transactions). Therefore, beginning in Chapter 3 and continuing in Chapter 9 (middleware), Chapter 11 (networking), and Chapter 12 (communications), the book describes the infrastructure and technologies that support networked applications. The goal is not an understanding of all the details, but rather an appreciation of the opportunities for and limitations on applications imposed by technology. This is imparted through the most significant underlying

can be small, as in English, or as large as the vocabulary, as in ideographs. Anything discrete can be represented as bits by simply associating an appropriate number of bits with each member of the alphabet.

Less obviously, a sequence of bits can also represent audio (such as a voice or music recording), images (such as pictures taken by a camera), and video. Unlike written language, these media are *analog* (meaning continuous, not discrete, in amplitude and time), but they can be approximated by bits with sufficient accuracy that a human cannot tell the difference.

The representation of information by bits is a special case of a digital representation, meaning a representation in terms of numbers. A bit is just a digit with a base of two, but any other base (such as base ten) would be equally effective.

concepts and terminology, supplemented by numerous topical examples.

Beginning in Chapter 4, and especially in Chapter 5, the economic laws, the structure of the supplier industry, and government involvement are discussed. These are arguably as important as the technology in determining the evolution of products and applications. Also, in Chapter 3 the life cycle of an application is discussed, and Chapter 6 describes modern software technology suitable for developing new applications. Together, the coverage of the industry, underlying technologies, and the application life cycle will position you to work effectively with vendors, contractors, and implementers in bringing new application ideas to fruition.

Further Reading

Books that cover general areas of technology and social issues encompassed within this book are recommended at the end of each chapter. There are two other general classes of books that can be recommended: those that convey computer literacy, of which [Bir96, Oak96] are examples, and those that discuss the societal and business impacts of networking and computing, including [Cai97, Haw96, Mit96, Ros97].