12

THE INTERNET

IN THE EARLY 1990s the Internet—the system that links millions of computers around the world—became big news. In the fall of 1990 there were just 313,000 computers attached to the Internet; five years later the number was approaching 10 million, and by the end of 2000 the number had exceeded 100 million. Although computer technology is at the heart of the Internet, its importance is economic and social: the Internet gives computer users the ability to communicate, to gain access to information sources, and to conduct business.

I. From the World Brain to the World Wide Web

The Internet sprang from a confluence of three desires, two that emerged in the 1960s and one that originated much further back in time. First, there was the rather utilitarian desire for an efficient, fault-tolerant networking technology, suitable for military communications, that would never break down. Second, there was a wish to unite the world's computer networks into a single system. Just as the telephone would never have become the dominant person-to-person communications medium if users had been restricted to the network of their particular provider, so the world's isolated computer networks would be far more useful if they were joined together. But the most romantic ideal—perhaps dating as far back as the Library of Alexandria in the ancient world—was to make readily available the world's store of knowledge.

FROM THE ENCYCLOPEDIA TO THE MEMEX

The idea of making the world's store of knowledge available to the ordinary person is a very old dream. It was the idea, for example, that drove the French philosopher

Denis Diderot to create the first great encyclopedia in the eighteenth century. The multivolume *Encyclopédie* was one of the central projects of the Age of Enlightenment, which tried to bring about radical and revolutionary reform by giving knowledge and therefore power to the people. The *Encyclopédie* was in part a political act; similarly, the Internet has a political dimension. The first English-language encyclopedia, the *Encyclopedia Britannica,* appeared in 1768 and was modeled directly on the *Encyclopédie.* Of course, neither the *Encyclopédie* nor the *Encyclopedia Britannica* could contain *all* of the world's knowledge. But they both contained a significant fraction of it and—at least as important—they brought order to the universe of knowledge, giving people a sense of what there was to know.

The nineteenth century saw an explosion in the production of human knowledge. In the early decades of the century, it was possible for a person of learning to be comfortable with the whole spectrum of human knowledge, in both the arts and the sciences. For example, Peter Mark Roget—now famous only as the compiler of the thesaurus—earned his living as a physician, but he was also an amateur scientist and a member of the Royal Society of London, an educationalist, and a founder of the University of London. And Charles Babbage, besides being famous for his calculating engines, was an important economist; he also wrote works on mathematics and statistics, geology and natural history, and even theology and politics.

By the twentieth century, however, the enormous increase in the world's knowledge had brought about an age of specialization during which it was very unusual for a person to be equally versed in the arts and the sciences, and virtually impossible for anyone to have a deep knowledge of more than a very narrow field of learning. It was said, for example, that by 1900 no mathematician could be familiar even with all the different subdisciplines of mathematics.

In the years between the two world wars, a number of leading thinkers began to wonder whether it might be possible to arrest this trend toward specialization by organizing the world's knowledge systematically so that, at the very least, people could once again know what there was to know. The most prominent member of this movement was the British socialist, novelist, and science writer H. G. Wells, best known in the United States as the author of *The War of the Worlds* and *The Time Machine.* During his own lifetime, Wells had seen the world's store of knowledge double and double again. He was convinced that narrow specialization—such that even educated people were familiar with no more than a tiny fraction of the world's knowledge—was causing the world to descend into barbarism in which people of learning were being "pushed aside by men like Hitler." During the 1930s he wrote pamphlets and articles and gave speeches about his project for a World Encyclopedia that would do for the twentieth century what Diderot had done for the eighteenth. Wells failed to interest publishers owing to the enormous cost of

such a project and, in the fall of 1937, embarked on a US lecture tour hoping to raise funds.

He covered five cities, and his last venue, in New York, was also broadcast on the radio. In his talk, titled "The Brain Organization of the Modern World," Wells explained that his World Encyclopedia would not be an encyclopedia in the ordinary sense:

> A World Encyclopedia no longer presents itself to a modern imagination as a row of volumes printed and published once for all, but as a sort of mental clearing house for the mind, a depot where knowledge and ideas are received, sorted, summarized, digested, clarified and compared. . . . This Encyclopedic organization need not be concentrated now in one place; it might have the form of a network [that] would constitute the material beginning of a real World Brain.

Wells never explained how his "network" for the World Brain would be achieved, beyond supposing that it would be possible to physically store all the data on microfilm. All the information in the world would do no good, however, if it were not properly organized. He thus envisaged that "[a] great number of workers would be engaged perpetually in perfecting this index of human knowledge and keeping it up to date."

During his tour, Wells was invited to lunch as the distinguished guest of President Roosevelt. Wells lost no time in trying to interest his host in the World Brain project, but, perhaps not surprisingly, Roosevelt had more pressing problems. Wells left the lunch a disappointed man. Time was running out for the World Brain; as the world drifted into World War II, he was forced to abandon the project and fell into a depression from which he never emerged. He lived to see the end of the war but died the following year, at the age of eighty.

Wells's dream did not die with him. In the postwar period the idea resurfaced and was given new vigor by Vannevar Bush, the scientist and inventor who had developed analog computers and risen to become chief scientific adviser to the president and head of the Office of Scientific Research and Development, where he directed much of the United States' scientific war effort.

In fact, Bush had first proposed an information storage machine several years before the war. This was to be a desk-like device that could hold the contents of a university library "in a couple of cubic feet." With the start of the war Bush had to put these ideas to one side, but in its final months he turned his mind to what scientists might do in the postwar era. For him, one problem stood out above all others: coping with the information explosion. In July 1945 he set his ideas down in a popular article, "As We May Think," for the *Atlantic Monthly.* A few weeks later, when it was reprinted in *Life* magazine, it reached a wider readership. This article established a dream that had been pursued by information scientists for decades.

During the war, as Wells had predicted, microfilm technology advanced to a point where the problem of storing information was essentially solved. Through the use of microfilm technology, Bush noted, it would be possible to contain the *Encyclopedia Britannica* in a space the size of a matchbox, and "a library of a million volumes could be compressed into one end of a desk." Thus the problem was not so much *containing* the information explosion as being able to make use of it. Bush envisaged a personal-information machine that he called the memex:

> A memex is 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 speed and flexibility. It is an enlarged intimate supplement to his memory.
>
> It consists of a desk, and while it can presumably be operated from a distance, it is primarily the piece of furniture at which he works. On the top are slanting translucent screens, on which material can be projected for convenient reading. There is a keyboard, and sets of buttons and levers. Otherwise it looks like an ordinary desk.

The memex would allow the user to browse through information:

> If the user wishes to consult a certain book, he taps its code on the keyboard, and the title page of the book promptly appears before him, projected onto one of his viewing positions. Frequently used codes are mnemonic, so that he seldom consults his code book; but when he does, a single tap of a key projects it for his use. Moreover, he has supplemental levers. On deflecting one of these levers to the right he runs through the book before him, each page in turn being projected at a speed which just allows a recognizing glance at each. If he deflects it further to the right, he steps through the book 10 pages at a time; still further at 100 pages at a time. Deflection to the left gives him the same control backwards.
>
> A special button transfers him immediately to the first page of the index. Any given book of his library can thus be called up and consulted with far greater facility than if it were taken from a shelf. As he has several projection positions, he can leave one item in position while he calls up another. He can add marginal notes and comments . . . just as though he had the physical page before him.

Bush did not quite know how, but he was sure that the new technology of computers would be instrumental in realizing the memex. He was one of very few people in 1945 who realized that computers would one day be used for something more than rapid arithmetic.

In the twenty years that followed Bush's article, there were a number of attempts at building a memex using microfilm readers and simple electronic controls,

but the technology was too crude and expensive to make much headway. When Bush revisited the memex idea in 1967 in his book *Science Is Not Enough,* his dream of a personal-information machine was still "in the future, but not so far."

At the time he wrote those words, the intellectual problems involved in constructing a memex-type information system using computer technology had, in principle, been largely solved. J.C.R. Licklider, the head of ARPA's Information Processing Techniques Office, for one, was working as early as 1962 on a project he called the Libraries of the Future, and he dedicated the book he published with that title: "however unworthy it may be, to Dr. Bush." In the mid-1960s, Ted Nelson coined the term *hypertext* and Douglas Engelbart was working on the practical realization of similar ideas at the Stanford Research Institute. Both Nelson and Engelbart claim to have been directly influenced by Bush. Engelbart later recalled that, as a lowly electronics technician in the Philippines during World War II, he "found this article in *Life* magazine about his [Bush's] memex, and it just thrilled the hell out of me that people were thinking about something like that. . . . I wish I could have met him, but by the time I caught onto the work, he was already in a nursing home and wasn't available."

Although the intellectual problems of designing a memex-type machine were quickly solved, establishing the physical technology—a network of computers—would take a long time.

THE ARPANET

The first concrete proposal for establishing a geographically distributed network of computers was made by Licklider in his 1960 "Man-Computer Symbiosis" paper:

> It seems reasonable to envision, for a time 10 or 15 years hence, a "thinking center" that will incorporate the functions of present-day libraries together with anticipated advances in information storage and retrieval. . . . The picture readily enlarges itself into a network of such centers, connected to one another by wide band communication lines and to individual users by leased-wire services. In such a system, the speed of the computers would be balanced, and the cost of the gigantic memories and the sophisticated programs would be divided by the number of users.

In 1963, when the first ARPA-sponsored time-shared computer systems became operational, Licklider set into action a program he privately called his Intergalactic Computer Network, publicly known as the Arpanet.

The stated motive for Arpanet was an economic one. By networking ARPA's computer systems together, the users of each computer would be able to use the facilities of any other computer on the network; specialized facilities would thus be

available to all, and it would be possible to spread the computing load over many geographically separated sites. For example, because the working day for East Coast computer users started several hours ahead of the West Coast's, they could make use of idle West Coast facilities in the morning; and when it was evening on the East Coast, their positions would be reversed. Arpanet would behave rather like a power grid in which lots of power plants worked in harmony to balance the load.

In July 1964, when Licklider finished his two-year tenure as head of the Information Processing Techniques Office (IPTO), he had a powerful say in the appointment of his successors, who shared his vision and carried it forward. His immediate successor was Ivan Sutherland, an MIT-trained graphics expert, then at the University of Utah, who held the office for two years. He was followed by Robert Taylor, another MIT alumnus, who later became head of the Xerox PARC computer science program.

Between 1963 and 1966 ARPA funded a number of small-scale research projects to explore the emerging technology of computer networking. ARPA was not the only organization sponsoring such activity; there were several other groups interested in computer networking in the United States, England (especially at the National Physical Laboratory), and France. By 1966 the computer networking idea was ready for practical exploitation, and IPTO's head, Robert Taylor, decided to develop a simple experimental network. He invited a rising star in the networking community, Larry Roberts, to head up the project.

Larry Roberts had received his PhD from MIT in 1963. He was first turned on to the subject of networking by a discussion with Licklider at a conference in November 1964, and, as he later put it, "his enthusiasm infected me." When he was approached by Taylor to head IPTO's networking program, Roberts was working on an IPTO-sponsored networking project at MIT's Lincoln Laboratory. Roberts initially did not want to trade his role as a Lincoln Lab scientist for this important bureaucratic task, but Taylor went over his head and reminded Roberts' boss that the lab received more than half its funding from ARPA. When he took over the Arpanet project in 1966, Roberts had to find solutions to three major technical problems before work could begin on a practical network. The first problem was how to physically connect all the time-sharing systems together. The difficulty here was that, if every computer system had to be connected to every other, the number of communications lines would grow geometrically. Networking the seventeen ARPA computers that were then in existence would require a total of 136 (that is, $17 \times 16 / 2$) communications lines. The second problem was how to make economic use of the expensive high-speed communications lines connecting computers. Experience with commercial time-sharing computers had already shown that less than 2 percent of the communications capacity of a telephone line was productively used because most of a user's time was spent thinking, during which the line was idle. This drawback did not matter very much where local phone lines

were concerned, but it would be insupportable on high-speed long-distance lines. The third problem Roberts faced was how to link together all the computer systems, which came from different manufacturers and used many varieties of operating software that had taken several years to develop. Enough was known about the software crisis at this stage to want to avoid the extensive rewriting of operating systems.

Unknown to Roberts, a solution to the first two problems had already been invented. Known as "store-and-forward packet switching," the idea was first put forward by Paul Baran of the RAND Corporation in 1961 and was independently reinvented in 1965 at the National Physical Laboratory in England by Donald Davies, who coined the term *packet switching.* Davies recognized the packet-switching concept to be similar to an older telegraph technology.

In telegraph networks, engineers had already solved the problem of how to avoid having every city connected to every other. Connectivity was achieved by using a number of switching centers located in major cities. Thus if a telegram was sent from, say, New York to San Francisco, the message might pass through intermediate switching centers in Chicago and Los Angeles before arriving in San Francisco. In the early days of the telegraph, during the late nineteenth century, at each switching center an incoming telegram would be received on a Morse sounder and written out by a telegraph clerk. It would then be retransmitted by a second telegraph clerk to the next switching center. This process would be repeated at each switching center, as the telegram was relayed across the country to its final destination. An incidental advantage of creating a written copy of the telegram was that it could act as a storage system, so that if there was a buildup of traffic, or if the onward switching center was too busy to receive the message, it could be held back until the lines became quieter. This was known as the store-and-forward principle. In the 1930s these manual switching centers were mechanized in "torn-tape offices," where incoming messages were automatically recorded on perforated paper tape and then retransmitted mechanically. In the 1960s the same functions were being computerized using disk stores instead of paper tape as the storage medium.

Store-and-forward packet switching was a simple elaboration of these old telegraph ideas. Instead of having every computer connected to every other, store-and-forward technology would be used to route messages through the network; there would be a single "backbone" communications line that connected the computers together, with other connections being added as the need arose. Packet-switching technology addressed the problem of making economic use of the high-speed communications lines. So that a single user did not monopolize a line, data would be shuttled around the network in packets. A packet was rather like a short telegram, with each packet having the address of the destination. A long message would be broken up into a stream of packets, which would be sent as individual items into the network. This would enable a single communications line

to carry many simultaneous human-computer interactions by transmitting packets in quick succession. The computers that acted as the switching centers—called *nodes* in the Arpanet—would simply receive packets and pass them along to the next node on the route toward the destination. The computer at the destination would be responsible for reconstituting the original message from the packets. In effect, by enabling many users to share a communications line simultaneously, packet switching did for telecommunications what time-sharing had done for computing.

All of this was unknown to Roberts until he attended an international meeting of computer network researchers in Gatlinburg, Tennessee, in October 1967. There he learned of the packet-switching concept from one of Donald Davies's English colleagues. He later described this as a kind of revelation: "Suddenly I learned how to route packets."

The final problem that remained for Roberts was how to avoid the horrendous software problems of getting the different computers to handle the network traffic. Fortunately, just as Roberts was confronting this problem, the first minicomputers had started to enter the market and the solution came to him in a eureka moment in a taxicab ride: the interface message processor (IMP). Instead of the existing software having to be modified at each computer center, a separate inexpensive minicomputer, the IMP, would be provided at every node to handle all the data communications traffic. The software on each computer system—called a host—would need only a relatively simple modification to collect and deliver information between itself and the IMP. Thus there was only one piece of software to worry about—a single software system to be used in all the IMPs in the network.

The Arpanet project—little more than a gleam in IPTO's eye when Roberts took it over in 1966—had become concrete. In the summer of 1968 he succeeded Taylor as director of IPTO, and the Arpanet project went full steam ahead. Work began on a $2.5 million pilot scheme to network together four computer centers—at the University of California at Los Angeles, the University of California at Santa Barbara, the University of Utah, and the Stanford Research Institute. The software for the IMPs was contracted out to Bolt, Beranek and Newman (BBN), while on the university campuses a motley collection of graduate students and computer center programmers worked on connecting up their host computers to the IMPs. The participation of so many graduate students in developing Arpanet, typically working on a part-time basis as they completed their computer science graduate studies, created a distinctive and somewhat anarchic culture in the network community. This culture was as strong in its way as the computer-amateur culture of the personal-computer industry in the 1970s. But unlike the personal-computer world, this culture was much more persistent, and it accounted for the unstructured and anarchic state of the early Internet.

By 1970 the four-node network was fully operational and working reliably. The other ARPA-funded computer centers were soon joined up to the network, so that by the spring of 1971 there were twenty-three hosts networked together. Although impressive as a technical accomplishment, Arpanet was of little significance to the tens of thousands of ordinary mainframe computers in the world beyond. If the networking idea was ever to move beyond the ARPA community, Roberts realized that he had to become not merely a project manager but an evangelist. He proselytized Arpanet to the technical community at academic conferences, but found he was preaching to the converted. He therefore decided to organize a public demonstration of the Arpanet, one that simply no one could ignore, at the first International Conference on Computer Communications (ICCC) in Washington in the fall of 1972.

The conference was attended by more than a thousand delegates, and from the Arpanet demonstration area, equipped with forty terminals, people were able to directly use any one of more than a dozen computers—ranging from Project MAC at MIT to the University of California computer center. A hookup was even arranged with a computer in Paris. Users were able to undertake serious computing tasks such as accessing databases, running meteorological models, and exploring interactive graphics; or they could simply amuse themselves with an air-traffic simulator or a chess game. The demonstration lasted three days and left an indelible impression on those who attended.

The ICCC demonstration was a turning point both for the Arpanet and for networking generally. The technology had suddenly become real, and many more research organizations and universities clamored to become connected to Arpanet. A year after the conference, there were 45 hosts on the network; four years later there were 111. It was an improvement, but growth was still slow.

THE POPULARITY OF E-MAIL

It was not, however, the economics of resource sharing, the ability to use remote computers, or even the pleasure of playing computer games that caused the explosion of interest in networking; indeed, most users never made use of any of these facilities. Instead, it was the opportunity for communicating through electronic mail that attracted users.

Electronic mail had never been an important motivation for Arpanet. In fact, no electronic mail facilities were provided initially, even though the idea was by no means unknown. An e-mail system had been provided on MIT's Project MAC in the mid-1960s, for example; but because users could only mail other MIT colleagues, it was never much more than an alternative to the ordinary campus mail.

In July 1971 two BBN programmers developed an experimental mail system for Arpanet. According to one member of the BBN network team:

> When the mail [program] was being developed, nobody thought at the beginning it was going to be the smash hit that it was. People liked it, they thought it was nice, but nobody imagined that it was going to be the explosion of excitement and interest that it became.

Electronic mail soon exceeded all other forms of network traffic on Arpanet, and by 1975 there were over a thousand registered e-mail users. The demand for e-mail facilities was also a major driving force for the first non-ARPA networks. One of the most important of these was Usenet, a network formed in 1978 for colleges that had been excluded from connection to Arpanet. An unplanned spinoff from this network was the Usenet news system. Acting like a giant electronic bulletin board, the system enabled network users to subscribe to news groups where like-minded individuals could exchange views. Initially designed by a student at Duke University and another at the University of North Carolina to exchange news, the idea quickly caught on. At first, news groups exchanged information mainly about computer-related matters, but eventually they included thousands of different topics. By 1991 there were 35,000 nodes on the Usenet system and millions of subscribers to the news network.

However, while most people quickly tired of using news groups, networked computing and e-mail became an integral part of the modern way of doing business, and the existing computer services industry was forced to respond. First, the time-sharing firms began to restructure themselves as network providers. BBN's Telcomp time-sharing service, for example, relaunched itself as the Telnet network in 1975 (with Larry Roberts as its CEO). Telnet initially established nodes in seven cities, and by 1978 there were nodes in 176 US cities and 14 overseas countries. In 1979 Telcomp's rival Tymshare created a network subsidiary, Tymnet. Established communications firms such as Western Union and MCI were also providing e-mail services, and the business press was beginning to talk of "The Great Electronic Mail Shootout." In the public sector, government organizations such as NSF and NASA also developed networks, while in the education sector consortia of colleges developed networks such as Merit, Edunet, and Bitnet, all of which became operational in the first half of the 1980s.

Electronic mail was the driving force behind all these networks. It became popular because it had so many advantages over conventional long-distance communications, and these advantages were consolidated as more and more networks were established: the more people there were on the networks, the more useful e-mail became. Taking just a few minutes to cross the continent, e-mail was much faster than the postal service, soon derisively called "snail mail." Besides being cheaper than a long-distance phone call, e-mail eliminated the need for both parties having to synchronize their activities, freeing them from their desks. E-mail also eliminated some of the problems associated with different time zones. It was not just

e-mail users who liked the new way of doing business; managers, too, had become enthusiastic about its economic benefits, especially for coordinating group activity—it could reduce the number of meetings.

Electronic mail was a completely new communications medium, however, and, as such, it brought with it a range of social issues that fascinated organizational psychologists and social scientists. For example, the speed of communications encouraged kneejerk rather than considered responses, thus increasing rather than decreasing the number of exchanges. Another problem was that the terse "telegraphese" of e-mail exchanges could easily cause offense to the uninitiated, whereas the same message conveyed over the phone would have been softened by the tone of voice and a more leisurely delivery. Gradually an unwritten "netiquette" emerged, enabling more civilized interactions to take place without losing too many of the benefits of e-mail. While some of the new computer networks of the 1970s were based on the technology developed in Arpanet, this was not true of all of them. Computer manufacturers, in particular, developed their own systems such as IBM's Systems Network Architecture (SNA) and Digital Equipment Corporation's DECNET. By reworking a very old marketing strategy, the manufacturers were hoping to keep their customers locked into proprietary networks for which they would supply all the hardware and software. But this was a shortsighted and mistaken strategy, because the real benefits of networking came through *inter*networking—in which the separate networks were connected and literally every computer user could talk to every other.

Fortunately, IPTO had been aware of this problem as early as 1973—not least so that electronic mail would be able to cross network boundaries. Very quickly, what was simply an idea, internetworking, was made concrete as the Internet. The job of IPTO at this stage was to establish "protocols" by which the networks could communicate. (A network protocol is simply the ritual electronic exchange that enables one network to talk to another—a kind of electronic Esperanto.) The system that ARPA devised was known as the Transmission Control Protocol/Internet Protocol, or TCP/IP—a mysterious acronym familiar to most experienced users of the Internet. Although the international communications committees were trying to evolve internetworking standards at the same time, TCP/IP quickly established itself as the de facto standard (which it remains).

Yet it would be a full decade before significant numbers of networks were connected. In 1980 there were fewer than two hundred hosts on the Internet, and as late as 1984 there were still only a thousand. For the most part these networks served research organizations and science and engineering departments in universities—a predominantly technical community, using conventional time-sharing computers. The Internet would become an important economic and social phenomenon only when it also reached the broad community of ordinary personal-computer users.

THE WORLD WIDE WEB

This broad community of users began to exert its influence in the late 1980s. In parallel with the development of the Internet, the personal computer was spreading across the educational and business communities and finding its way into American homes. By the late 1980s most professional computer users (though not home users) in the United States had access to the Internet, and the number of computers on the Internet began an explosive growth.

Once this much broader community of users started to use the Internet, people began to exchange not just e-mail but whole documents. In effect, a new electronic publishing medium had been created. Like e-mail and news groups, this activity was largely unforeseen and unplanned. As a result, there was no way to stop anyone from publishing anything and placing it on the net; soon there were millions of documents but no catalog and no way of finding what was useful. Sifting the grains of gold amid the tons of trash became such a frustrating activity that only the most dedicated computer users had the patience to attempt it. Bringing order to this universe of information became a major research issue.

Numerous parallel developments were under way to create finding aids for information on the Internet. One of the first systems, "archie," was developed at McGill University in 1990. By combing though the Internet, archie created a directory of all the files available for downloading so that a user wanting a file did not need to know on what machine it was actually located. A more impressive system was the Wide Area Information Service (WAIS), developed the following year by the Thinking Machines Corporation of Waltham, Massachusetts. WAIS enabled users to specify documents by using keywords (say, *smallpox* and *vaccine*), at which point it would display all the available documents on the Internet that matched those criteria. The most popular early finding aid was "gopher," developed at the University of Minnesota. (A highly appropriate name, gopher is slang for one who runs errands but also a reference to the university's mascot.) The system, which became operational in 1991, was effectively a catalog of catalogs that enabled a user to drill down and examine the contents of gopher databases maintained by hundreds of different institutions.

All of these systems treated documents as individual entities, rather like books in a library. For example, when a system located a document about smallpox vaccine, say, it might tell the user that its inventor was Edward Jenner; in order to discover more about Jenner, the user would need to search again. The inventors of hypertext—Vannevar Bush in the 1940s and Engelbart and Nelson in the 1960s—had envisaged a system that would enable one to informally skip from document to document. At the press of button, as it were, one could leap from smallpox to Jenner to The Chantry in Gloucestershire, England (the house where Jenner lived and now a museum to his memory). Hypertext was, in fact, a lively computer research

topic throughout the 1980s, but what made it so potent for the Internet—ultimately giving rise to the Word Wide Web—was that it would make it unnecessary to locate documents in centralized directories. Instead, links would be stored in the documents themselves, and they would instantly whisk the reader to related documents. It was all very much as Vannevar Bush had envisioned the memex.

The World Wide Web was invented by Tim Berners-Lee. Its origins dated back to Berners-Lee's early interest in hypertext in 1980, long before the Internet was widely known. Berners-Lee was born in London in 1955, the son of two mathematicians (who were themselves pioneers of early British computer programming). After graduating in physics from Oxford University in 1976, he worked as a software engineer in the UK before obtaining a six-month consulting post at CERN, the international nuclear physics research laboratory in Geneva. While he was at CERN, Berners-Lee was assigned to develop software for a new particle accelerator, but in his spare time he developed a hobby program for a hypertext system he called Enquire. (The program was named after a famous Victorian household compendium called *Enquire Within Upon Everything* that had long before gone out of print, but that his parents happened to own and he liked to browse as a child.) Despite the similarity to Bush's vision, Berners-Lee had no firsthand knowledge of his work; rather, Bush's ideas had simply become absorbed into the hypertext ideas that were in the air during the 1980s. There was nothing very special about Enquire: it was simply another experimental hypertext system like dozens of others. When Berners-Lee left CERN, it was effectively orphaned.

The personal-computer boom was in full swing when Berners-Lee returned to England, and he found gainful employment in the necessary, if rather mundane, business of developing software for dot-matrix printers. In September 1984, he returned to CERN as a permanent employee. In the years he had been away, computer networking had blossomed. CERN was in the process of linking all its computers together, and he was assigned to help in this activity. But it was not long before he dusted off his Enquire program and revived his interest in hypertext. Berners-Lee has described the World Wide Web as "the marriage of hypertext and the Internet." But it was not a whirlwind affair; rather, it was five years of peering through the fog as the technologies of hypertext and the Internet diffused and intertwined. It was not until 1989 that Berners-Lee and his Belgian collaborator Robert Cailliau got so far as to make a formal project proposal to CERN for the resources necessary for creating what they had grandiosely named the World Wide Web.

There were really two sides to the World Wide Web concept: the server side and the client side. The server would deliver hypertext documents (later know as web pages) to a client computer, typically a personal computer or a workstation, which in turn would display them on the user's screen. By the late 1980s hypertext was a well-established technology that had moved far beyond the academy into consumer

products such as CD-ROM encyclopedias, although it had not yet crossed over to the Internet. By 1991 Berners-Lee and Cailliau were sufficiently confident of their vision to offer a paper on the World Wide Web to the Hypertext '91 conference in December in San Antonio, Texas. The paper was turned down, but they managed to put on a half-hearted demonstration. Theirs was the only project that related to the Internet. Berners-Lee later reflected that when he returned to the conference in 1993, "every project on display would have something to do with the Web."

During that two-year period, the World Wide Web had taken off. It was a classic chicken-and-egg situation. Individuals needed web "browsers" to read web pages on their personal computers and workstations, while organizations needed to set up web servers filled with interesting and relevant information to make the process worthwhile. Several innovative web browsers came from universities, often developed by enthusiastic students. The web browser developed at CERN had been a pedestrian affair that was geared to text-only hypertext documents, not hypermedia documents that were enriched with pictures, sounds, and video clips. The new generation of browsers (such as the University of Helsinki's Erwise and Viola at UC-Berkeley) provided not only the user-friendly point-and-click interface that people were starting to demand but also full-blown hypermedia features. Web servers needed software considerably more complex than a web browser, although this software was largely invisible to the average user. Here again, volunteer efforts, mainly from universities, produced serviceable solutions. The programs evolved rapidly as many individuals, communicating through the Internet, supplied bug fixes, program "patches," and other improvements. The best-known server program to evolve by this process was known as "apache," a pun on "a patchy" server. The Internet enabled collaborative software development (also referred to as the open-source movement) to flourish to an extent previously unknown. And the fact that open-source software was free offered an entirely new model for software development to the conventional for-profit software company.

While all this was happening, Berners-Lee had a useful index of how the web was taking off—namely, the number of "hits" on his original web server at CERN. The number of hits grew from a hundred a day in 1991 to a thousand in 1992 and to ten thousand in 1993. By 1994, there were several hundred publicly available web servers and the World Wide Web was rapidly overtaking the gopher system in popularity—partly owing to the fact that, in the spring of 1993, the University of Minnesota had decided to assert ownership of the intellectual property in gopher software and no doubt would eventually commercialize it. It was perhaps the first time that the commercialization of Internet software had surfaced, and forever after there would be an uneasy tension between the open-source community that promoted free software and the entrepreneurs who saw a business opportunity.

In any event, Berners-Lee was on the side of the angels when it came to commercial exploitation. He persuaded CERN to place web technology in the public domain, so that all could use it free of charge for all time. In the summer of 1994 he moved from CERN to the Laboratory for Computer Science at MIT, where he would head the World Wide Web Consortium (W3C), a nonprofit organization to encourage the creation of web standards through consensus.

II. The Web and Its Consequences

What started the hockey-stick growth of the World Wide Web was the Mosaic browser. The first web browsers mostly came from universities; they were hastily written by students, and it showed. The programs were difficult to install, buggy, and had an unfinished feel. The Mosaic browser from the National Center for Supercomputer Applications (NCSA) at the University of Illinois at Urbana-Champaign was the exception.

Developed by a twenty-two-year-old computer science undergraduate, Marc Andreessen, Mosaic was almost like shrink-wrapped software you could buy in a store. There was a version for the PC, another for the Macintosh, and another for the Unix workstations beloved by computer science departments. Mosaic was made available in November 1993, and immediately thousands—soon hundreds of thousands—of copies were downloaded. Mosaic made it easy for owners of personal computers to get started surfing the web. Of course they had to be enthusiasts, but no deep knowledge was needed.

BROWSER WARS

In the spring of 1994 Andreessen received an invitation to meet with the Californian entrepreneur Jim Clark. In the 1970s Clark had co-founded Silicon Graphics, a highly successful maker of Unix workstations. He had recently sold his share of the company and was on the lookout for a new start-up opportunity. The upshot of the meeting was that on 4 April 1994 the Mosaic Communications Corporation was incorporated to develop browsers and server software for the World Wide Web. The name of the corporation was changed to Netscape Communications a few months later, because the University of Illinois had licensed the Mosaic name and software to another firm, Spyglass Inc.

Andreessen immediately hired some of his programming colleagues from the University of Illinois and got started cranking out code, doing for the second time what they had done once before. The results were very polished. In order to quickly establish the market for their browser, Clark and Andreessen decided to give it

away free to noncommercial users. When they had gotten the lion's share of the market, they could perhaps begin to charge for it. Server software and services were sold to corporations from the beginning, however. This established a common pattern for selling software for the Internet.

In December 1994 Netscape shipped version 1.0 of its browser and complementary server software. In this case "shipping" simply meant making the browser available on the Internet for users—who downloaded it by the millions. From that point on, the rise of the web was unstoppable: by mid-1995 it accounted for a quarter of all Internet traffic, more than any other activity.

In the meantime Microsoft, far and away the dominant force in personal computer software, was seemingly oblivious to the rise of the Internet. Its online service MSN was due to be launched at the same time as the Windows 95 operating system in August 1995. A proprietary network from the pre-Internet world, MSN had passed the point of no return, but it would prove an embarrassment of mistiming. Microsoft covered its bets by licensing the Mosaic software from Spyglass and including a browser dubbed Internet Explorer with Windows 95, but it was a lackluster effort.

Microsoft was not the only organization frozen in the headlights of the Internet juggernaut. A paradigm shift was taking place—a rapid change from one dominant technology to another. It was a transition from closed proprietary networks to the open world of the Internet. Consumers also had a difficult choice to make. They could either subscribe to one of the existing consumer networks—AOL, CompuServe, Prodigy, or MSN—or go with a new type of supplier called an Internet service provider (ISP). The choice was between the mature, user-friendly, safe, and content-rich world of the consumer network or the wild frontier of the World Wide Web. ISPs did not provide much in the way of content—there was plenty of that on the World Wide Web and it was rapidly growing. The ISP was more like a telephone service—it gave you a connection, but then you were on your own. The ISP customer needed to be a little more computer savvy to master the relatively complex software and had to be wary to avoid the less savory aspects of the Internet (or was free to enjoy them).

For the existing consumer networks, responding to the World Wide Web was a huge technical and cultural challenge. AOL licensed Microsoft's Internet Explorer for use in its access software. This gave its subscribers the best of both worlds—the advantages of the existing service plus a window into the World Wide Web. CompuServe did much the same, though it was eventually acquired by AOL in 1998. The other consumer networks did not manage to emulate AOL's masterly segue. Prodigy's owners, Sears and IBM, decided to sell the network in 1996, and it thereafter faded from sight. Microsoft, on the other hand, decided to cut its losses on MSN. In December 1995 Bill Gates announced that Microsoft would "embrace and extend" the Internet; it "was willing to sacrifice one child (MSN) to promote a

more important one (Internet Explorer)." Thereafter, MSN was not much more than an up-market Internet service provider—it was no threat to AOL.

Whereas in the 1980s the operating system had been the most keenly fought-for territory in personal computing, in the 1990s it was the web browser. In 1995 Netscape had seemed unstoppable. In the summer—when the firm was just eighteen months old—its initial public offering netted $2.2 billion, making Marc Andreessen extraordinarily wealthy. By year's end its browser had been downloaded 15 million times, and it enjoyed over 70 percent of the market. But Microsoft did not intend to concede the browser market to Netscape.

During the next two years Microsoft and Netscape battled for browser supremacy, each producing browser upgrades every few months. By January 1998, with a reported investment of $100 million a year, version 4.0 of Microsoft's Internet Explorer achieved technical parity with Netscape. Because Internet Explorer was bundled at no extra cost with the new Windows 98 operating system, it became the most commonly used browser on PCs—perhaps not through choice, but by the path of least resistance. A version of Internet Explorer was also made available free of charge for the Macintosh computer.

Distributing Internet Explorer to consumers at no cost completely undermined the business plans of Netscape. How could it sell a browser when Microsoft was giving its away for nothing? Microsoft had fought hard, perhaps too hard, for its browser monopoly. Its alleged practices of tying, bundling, and coercion provoked the US Department of Justice into filing an antitrust suit in May 1998. Like all antitrust suits, this one proceeded slowly—especially for an industry that spoke of "Internet time" in the way that canine enthusiasts spoke of dog years. Microsoft's self-serving and arrogant attitude in court did nothing to endear it to consumers or the press, though this stance did little to arrest its progress. In the five years between 1995, when Microsoft was first wrong-footed by the rise of the Internet, and 2000, when Judge Thomas Penfield Jackson produced his verdict, Microsoft's revenues nearly quadrupled from $6 billion to $23 billion, and it staff numbers doubled from 18,000 to 39,000. Judge Jackson directed that Microsoft be broken up—a classic antitrust remedy for the alleged wrongdoings of a monopolist—although this was set aside on appeal and less drastic remedies were imposed.

THE INTERNET LAND RUSH

During the second half of the 1990s, the Internet became not so much an information revolution as a commercial revolution; users were able to purchase the full range of goods and services that society could offer. In 1990, however, the Internet was largely owned by agencies of the US Government, and the political establishment did not permit the use of publicly owned assets for private profit. Hence privatization of the Internet was an essential precursor for electronic commerce to flourish.

Even before privatization had become an issue, it was necessary to separate the civilian and military functions of the Internet. Since its inception in 1969, Arpanet had been funded by the Advanced Research Projects Agency. In 1983 the military network was hived off as Milnet, and Arpanet became the exclusive domain of ARPA's research community. Once the military constraints were removed, the network flourished—by 1985 some two thousand computers had access to the Internet. To broaden access to the US academic community as a whole, beyond the exclusive circle of ARPA's research community, the National Science Foundation had created another network, NSFnet. In a way that proved characteristic of the development of the Internet, NSFnet rapidly overtook Arpanet in size and importance, eventually becoming the "backbone" of the entire Internet. By 1987 some thirty thousand computers—mostly from US academic and research communities—had access to the Internet. At the same time, other public and private networks attached themselves, such as Usenet, FidoNet (created by amateur "bulletin board" operators), and IBM's Bitnet. Who paid for their access was moot, because NSF did not have any formal charging mechanisms. It was a wonderful example of the occasional importance of turning a blind eye. If any bureaucrat had looked closely at the finances, the whole project for integrating the world's computer networks would likely have been paralyzed. As more and more commercial networks came on board, additional infrastructure was added and they evolved their own labyrinthine mechanisms for allocating the costs. By 1995, the commercially owned parts of the Internet far exceeded those owned by the government. On 30 April 1995, the old NSFnet backbone was shut down, ending altogether US government ownership of the Internet's infrastructure.

The explosive growth of the Internet was in large part due to its informal, decentralized structure; anyone was free to join in. However, the Internet could not function as a commercial entity in a wholly unregulated way—or chaos and lawlessness would ensue. The minimal, light-touch regulation that the Internet pioneers evolved was one of its most impressive features. A good example of this is the domain-name system. Domain names—such as *amazon.com, whitehouse.gov,* and *princeton.edu*—soon became almost as familiar as telephone numbers.

In the mid-1980s the Internet community adopted the domain-name system devised by Paul Mockapetris of the Information Sciences Institute at the University of Southern California. The system would decentralize the allocation of thousands (and eventually millions) of domain names. The process started with the creation of six top-level domains, each denoted by a three-letter suffix: *com* for commercial organizations, *edu* for educational institutions, *net* for network operators, *mil* for military, *gov* for government, and *org* for all other organizations. Six registration authorities were created to allocate names within each of the top-level domains. Once an organization had been given a unique domain name, it was free to subdivide it within the organization by adding prefixes as needed (for example, *cs* for

computer science in *cs.princeton.edu*), without the necessity of permission from any external authority.

Outside the United States, countries were allocated a two-letter suffix—*uk* for the United Kingdom, *ch* for Switzerland, and so on. Each country was then free to create its own second-level domain names. In Britain, for example, *ac* was used for educational (that is, academic) institutions, *co* for commercial organizations, and so on. Thus the Computer Laboratory at Cambridge had the domain name *cl.cam.ac.uk,* while *bbc.co.uk* needs no explanation. The United States is the only country that does not use a country suffix; in this regard it is exercising a privilege rather like that of Great Britain, which invented the postage stamp in 1841 and is the only country whose stamps do not bear the name of the issuing nation.

Early predictions of how the Internet would evolve were seemingly extrapolations of H. G. Wells's World Brain or Vannevar Bush's memex. In 1937 Wells wrote of his World Brain: "The time is close at hand when any student, in any part of the world, will be able to sit with his projector in his own study at his or her own convenience to examine *any* book, *any* document, in an exact replica." In 1996 (when the first edition of the present book went to press) this seemed a reasonable prediction, but one that might have been twenty or thirty years in the making. Although we see no reason to revise the time scale, the progress made has nonetheless been astonishing. Moreover—and this Wells did not predict—the information resources on the Internet go far beyond books and documents, encompassing audio, video, and multimedia as well. A large fraction of the world's current printed output is now available online. Indeed, for university and industrial researchers, visits to libraries are made with ever-decreasing frequency.

Extraordinary as the progress of the Internet as an information resource has been, the rise of electronic commerce—a phenomenon largely unpredicted as late as the mid-1990s—has been even more so.

Some of the early business successes were almost accidental. For example Yahoo! began in the days of the Mosaic browser as a simple listing service ("Jerry's Guide to the World Wide Web") created by David Filo and Jerry Yang, two computer science graduate students at Stanford University. By late 1993, the site listed a modest two hundred websites, though at the time that was a significant fraction of the world's websites. During 1994, however, the web experienced explosive growth. As new sites came online daily, Filo and Yang sifted, sorted, and indexed them. Toward the end of 1994 Yahoo! experienced its first million-hit day—representing perhaps 100,000 users. The following spring Filo and Yang secured venture capital, moved into offices in Mountain View, California, and began to hire staff to surf the web to maintain and expand the index. It was very much as Wells had envisaged for the World Brain when he wrote of "a great number of workers . . . engaged in perfecting this index of human knowledge and keeping it up to date." Yahoo! was not without competition: Lycos, Excite, and a dozen others had come up with the same

concept, and listing and information search services became one of the first established categories of the web. One question remained: How to pay for the service? The choices included subscriptions, sponsorship, commissions, or advertising. As with early broadcasting, advertising was the obvious choice. Another firm focused on helping users find information on the web—Google Inc.—soon demonstrated how lucrative web advertising could be.

Yahoo! was already well established when two other Stanford University doctoral students, Larry Page and Sergey Brin, began work on the Stanford Digital Library Project (funded in part by the National Science Foundation)—research that would not only forever change the process of finding things on the Internet but also, in time, lead to an unprecedentedly successful web advertising model.

Page became interested in a dissertation project on the mathematical properties of the web, and found strong support from his adviser Terry Winograd, a pioneer of artificial intelligence research on natural language processing. Using a "web crawler" to gather back-link data (that is, the websites that linked to a particular site), Page, now teamed up with Brin, created their "PageRank" algorithm based on back-links ranked by importance—the more prominent the linking site, the more influence it would have on the linked site's page rank. They insightfully reasoned that this would provide the basis for more useful web searches than any existing tools and, moreover, that there would be no need to hire a corps of indexing staff. Thus was born their "search engine," Backrub, renamed Google shortly before they launched the URL *google.stanford.edu* in September 1997. The name was a modification of a friend's suggestion of *googol*—a term referring to the number 1 followed by 100 zeros. Brin misspelled the term as *google,* but the Internet address for *googol* was already taken so the catchy misspelling stuck. While just a silly made-up word to most users, the original term was indicative of the complex math behind Page and Brin's creation, as well as of the large numbers (in terms of web indexing and searches) that their search tool would later attain.

In 1998 Page and Brin launched Google Inc. in a friend's Menlo Park garage. Early the following year, they moved the small company to offices in Palo Alto. By the early 2000s Google had gained a loyal following, and thereafter it rapidly rose to become the leading web search service. Taking $25 million in loans from leading Silicon Valley venture-capital firms to refine the technology, hire more staff, and greatly extend the infrastructure (the ever-expanding number of servers), the two founders were forced to hire a professional CEO, Eric Schmidt, early in 2001. Now with "adult supervision," Google perfected a disciplined business model of sponsored search. The company kept its search page sparse and simple, with just the search box and the Google logo—a stark contrast to websites like Yahoo! whose web pages were cluttered and busy, offering a variety of complementary services and ads. Initially Google's search-page design boosted speed, but its front page also soon garnered praise for its "Zen-like use of white space." Google search results

included both nonsponsored and sponsored links. Companies could purchase sponsored links associated with particular keyword searches; when users clicked on these sponsored links, Google was paid a small sum. Sponsored links met the needs of many users, whose clicks made Google—which rapidly expanded to include many languages spanning the world—the global leader in web advertising revenue.

Another early web success was mail-order selling. Jeff Bezos's Amazon.com established many of web commerce's early practices (including the incorporation of *com* in the firm's name). Bezos, a Princeton graduate in electrical engineering and computer science, had enjoyed a brief career in financial services before deciding to set up a retail operation on the web. After some deliberation, he lighted on bookselling as being the most promising opportunity: there were relatively few major players in bookselling, and a virtual bookstore could offer an inventory far larger than any bricks-and-mortar establishment. Launched in July 1995, Amazon.com would quickly become the leading online bookseller and, in time, the world's top online retailer.

The eBay auction website was another operation that succeeded largely through mastery of the new medium. The auction concept was developed as an experimental, free website—AuctionWeb—by Pierre Omidyar, a Silicon Valley entrepreneur, in late 1995. The site proved very popular and by 1996 Omidyar was able to begin charging commissions, after which "he left his day job and changed the site's name to eBay." In early 1998 professional management was brought in, in the form of Meg Whitman, a senior executive from Hasbro toys. Over the next few years, a worldwide virtual community of buyers and sellers evolved. By 2003, eBay had 30 million users worldwide and $20 billion in sales, and it was said that 150,000 entrepreneurs made their living as eBay traders. Emblematic of the transformative power and unpredictability of the World Wide Web, a global flea market is as unlikely a concept as one could have imagined.

For all of their dramatic growth, profits from the commercial web ventures such as Yahoo!, Amazon.com, and eBay remained elusive in their early years even as their stock prices soared. Inevitably, when Internet euphoria subsided, stock prices plunged; in the spring of 2000 many prominent e-commerce operations lost as much as 80 percent of their value. It had happened before: in the late 1960s software and computer services stocks suffered a similar collapse, from which it took several years to recover. But there was never any doubt that recovery would come, and today the software crash of 1969 is but a distant memory, and only for those who lived through it. The same is proving true for the dot-com crash of 2000. The symbolic return of Silicon Valley to glory came with the success of Google.

In 2004 Google's public offering valued the company at more than $26 billion. By 2007 Google facilitated more searches than all other search and listing services combined. That year Google achieved revenue of $16.6 billion and net income of $4.2 billion. Google continues to dominate the search field with 1.7 trillion annual

searches (in 2011, representing roughly a two-thirds share). While search-based advertising revenue remained its primary source of income, Google successfully moved into e-mail services (Gmail), maps and satellite photos, Internet video (with its 2006 acquisition of YouTube), cloud computing, digitizing books, and other endeavors. More recently, Google has also been an important participant in open-source mobile platforms that are transforming computing.

GOING MOBILE

From shortly after the advent of personal computing, computers have become increasingly mobile. In 1968 Alan Kay first conceptualized the notion of a portable computer, ideas formalized as the "Dynabook" concept at Xerox PARC in 1972. That year, many user-friendly elements of Kay's Dynabook were incorporated into the Xerox Alto workstation, but easy mobility was not one of them. A decade later, in 1982, GRiD Systems Corporation launched the GRiD Compass 1101, arguably the first laptop computer. Though it lacked compatibility with existing platforms and cost roughly $8,000, it succeeded in the price-insensitive military and aerospace markets. In the broader commercial arena, the first "portable" computers were available by the early 1980s, but they were a far cry from the modern laptop. The $1,800 Osborne 1 (released in 1981) lacked a battery, weighed more than twenty-three pounds, had a tiny five-inch screen, and, when folded up for carrying, resembled a mid-sized hard-shell suitcase. Various models of commercial portable computers hit the market as the decade progressed, including the Tandy TRS-80 100, the first popular portable to use an LCD screen. These portables, however, gave up much in screen size, processing power, memory, and compatibility to achieve degrees of mobility. By the early 1990s advances in microprocessors and memory-chip capacity, coupled with improvements in LCD screens, were sufficient for the affordable clamshell-design modern laptop—with models from IBM, Compaq, and other manufacturers becoming thinner, lighter, and more powerful with each passing year. The price-to-performance of laptops (now commonly referred to as notebooks) approached that of desktop computers, and the advantage of mobility led laptop production to surpass that of desktops by 2008.

While carrying laptops on the go, especially for work, became increasingly common, their size and weight prevented these machines from becoming personal accessories. Computers became common everyday and everywhere devices for the masses only with the advent and growing popularity of smartphones—ushering in a transformative new era of mobile computing by altering how people communicate, work, and socialize.

The smartphone, a hand-held computer that includes a wireless phone, is the product of a broad convergence of computing and telecommunications. Smartphones largely evolved from an earlier technology, Personal Digital Assistants

(PDAs), which, like smartphones, were defined by their operating system or platform. Before 2007, there were four primary platforms: Symbian, Blackberry OS, Palm OS, and Windows Mobile. Since then, two new platforms have emerged and gained dominant positions in the smartphone marketplace: Apple's iOS and Google's Android. These two new platforms benefited greatly from third-party application software providers. Exploring the evolution of the various platforms is instructive in understanding the rapidly expanding smartphone market.

Psion, a UK software company formed in 1980, soon diversified into hardware and in 1986 introduced the Psion Organizer II, an operating system–based PDA. Psion's operating system EPOC was successful in its home market but gained minimal traction internationally only after the firm formed a consortium with the world's leading cellphone manufacturer of the time, Nokia, and EPOC evolved into Symbian. Nokia was by far the leading user of Symbian, and its first smartphone hit the market in 2002. At its peak in 2007, two-thirds of smartphones, often called handsets, were based on the Symbian platform.

Competitors' handset features and software applications, however, quickly eroded Symbian's early market lead. Silicon Valley–based Palm, Inc. stumbled along with handwriting-recognition PDAs in the early 1990s before achieving success with its Palm Pilot in 1996. Though Palm's cultivation of a modest ecosystem of third-party applications developers foreshadowed what would become critical in the smartphone market, memory was too small for truly compelling applications. Palm never achieved more than a 3 percent market share and was quickly overshadowed in its primary market—business users—by Research in Motion (RIM), a Canadian specialist in paging, messaging, data capture, and modem equipment that launched the PDA "Blackberry" in 1999. Blackberry benefited in the business and government handset markets from RIM's private data network, user-friendly e-mail, and miniature QWERTY keyboard. Microsoft, which came late to the PDA/smartphone platform business by licensing Windows-based mobile operating systems, had some success in the enterprise market before smartphones became consumer oriented and the touchscreen-based Apple iOS and Android systems rose to dominance.

While Apple's Macintosh was a technical success at its launch in 1984, it helped Microsoft far more than Apple itself (by showing the dominant operating-system company the way to a user-friendly graphics-based operating system). Apple Computer was struggling as a company in the mid-1980s, and co-founder and Macintosh team leader Steve Jobs lost a boardroom battle, was isolated from Apple's management, and elected to resign from the firm. In 1985 Jobs formed NeXT, a computer platform development company focused on the educational and business markets. NeXT acquired the small computer graphics division of Lucasfilms, which it later spun-off as Pixar—the IPO made Jobs a billionaire. Pixar was acquired by Walt Disney, and Jobs became the parent's largest individual stockholder and a

Disney board member. This broad exposure beyond computing was influential in making Jobs aware of new media and consumer electronics opportunities when he was invited to lead Apple Computer in 1997.

By the time Jobs returned, Apple Computer already had a long history in portable computing. Beyond its continuous stream of new laptop models from the 1989 Mac Portable forward, the firm also launched a PDA, called Newton, in 1993. The Newton was too expensive, too large, and, like the first Palm, Inc. products, too focused on a feature that consumers did not really care for—handwriting recognition—to be commercially successful. Apple Computer dropped Newton in 1998. Early in his second act leading Apple, Jobs envisioned the firm's expansion into consumer electronics. Small digital music players, based on the compressed digital music standard MP3, had been around for several years before Apple launched the iPod, but Apple's player rapidly captivated consumers and gained market leadership. In addition to the iPod's sleek design, Jobs' charisma with product introductions, and Apple's exciting television advertising using bold animation and the music of pop mega group U2, the key to the iPod's immense success was the simultaneous launch of Apple's iTunes store. Seeing the opportunity to negotiate aggressively with executives in the recorded-music industry—which had been battered by online piracy and file-sharing—Apple secured attractive terms to sell music at a profit and, in time, to extend the iTunes store into other media. With the iPod line and iTunes store, Apple gained a leading reputation in marketing popular hand-held devices and in supplying content—both critical to the success of the iPhone.

The iPhone, launched in 2007, quickly overcame being late to the smartphone market with its easy-to-use touchscreen and its successful cultivation of third-party applications (apps) developers. Apps for the iPhone quickly included an abundance of games, as well as thousands of other entertainment, work, and educational offerings. The iPhone App store, resembling the familiar iTunes store, dwarfed the offerings of competitors and drove demand for Apple phones that for years sold at a substantial price premium to other smartphones.

Meanwhile, in 2005, Google acquired Android, Inc., a producer of an open-source, Linux-based mobile operating system. Retaining the Android trademark, Google ensured adherence to standards and licensed the platform on an open-source, royalty-free basis. Soon an extensive third-party Android App ecosystem emerged. With iOS and Android's ascendance, RIM's Blackberry was forced to fight an uphill battle to hold onto the business market and to try to add consumer-oriented models, while Nokia's Symbian platform declined and Palm all but disappeared. Meanwhile, Apple and Android extended their offerings to include popular tablet computers (for Apple, the iPad), which segued the features of a laptop with the mobility of a smartphone. During 2010–2012 Microsoft made a major foray into consumer smartphones and tablet computers. At the time of this writing,

Apple, Samsung (using Android), Microsoft, and lesser brands are engaged in intense competition for smartphone market dominance. The outcome of this commercial battle, however, will be less important than the momentum that mobile computing has given to the phenomenon of social networking.

WEB 2.0

As World Wide Web usage accelerated in the second half of the 1990s, the distinction between producers of content—primarily companies and other organizations—and the many more consumers of content remained relatively clear. To be sure, some early web businesses such as eBay provided platforms for users to add text and images (so their goods could be effectively auctioned) and some computer-savvy individuals set up web logs or "blogs" (to display text of their political opinions, sports commentary, or other writings), but overall those creating content on the web were few in number compared to the millions who browsed the web on a daily basis. In the early 2000s the relatively static web, characterized by a limited number of active producers and many passive consumers, began to rapidly change as platforms for facilitating and encouraging user-generated content and interaction became increasingly common. A prescient industry consultant in 1999 coined the term *Web 2.0* for this nascent trend, and a Web 2.0 industry conference in 2004 solidified this nomenclature. The advent of Web 2.0, among other things, transformed production and uses of encyclopedias, altered the nature of commerce, and established new models for how people commonly socialize.

As discussed in the previous chapter, the *Encyclopedia Britannica* was wounded by the massive reduction in production costs of lesser encyclopedias in digital form—forcing *Britannica*'s new owners to follow suit with an inexpensive CD-ROM version in 1996. In March 2012 *Britannica* announced that it would cease to print encyclopedias in order to focus entirely on its online version. And back in 2008 *Britannica* had announced that it would begin to accept unsolicited user content, which, upon acceptance by *Britannica* editors, would be published on a dedicated portion of the *Britannica* website. Shifting to exist only online and accepting some content from users represented tacit acceptance of the model of what had become the most popular encyclopedia in the world: *Wikipedia.*

Wikipedia (combining the Hawaiian term *wiki,* meaning "quick," with *encyclopedia*) was launched by Jimmy Wales and Larry Sanger in 2001. *Wikipedia* is based almost exclusively on volunteer labor—it is essentially a platform for user writing, reading, and editing. This model facilitated the very low-cost creation of a comprehensive encyclopedia in short order, sharply contrasting with the expensive, multi-decade efforts involved in producing print encyclopedias. In 2005 the British science journal *Nature* carried out a peer review of selected scientific articles of *Wikipedia* and *Britannica* and found that "the difference in accuracy was not

particularly great," lending credence to the value of user-generated content. As of this writing, *Wikipedia* has editions in more than two hundred languages and many millions of articles; "nearly a half billion people read Wikipedia every month" and it is one of the web's most visited sites. Nevertheless, it faces the ongoing challenge of ensuring that participation of quality volunteer editors does not wane over time.

Interactive platforms have become increasingly important to electronic commerce and, in many cases, are now expected by users. Amazon.com was one of the first commercial enterprises to capitalize on user reviews of books, movies, consumer electronics, and virtually all of the many other items it sold. The vast data that Amazon.com has collected on users enables it to offer well-targeted product suggestions based on past purchasing and browsing as well as patterns of preferences among other users. But traditional retailers are catching up. Among the revenue leaders of Internet retail in North America are long-established brick-and-mortar companies such as Staples, Inc. and Walmart.com. Like Amazon.com, these companies encourage users to rate products and benefit from user input to understand customer preferences.

The data these firms collect and present on their websites generally rests on server farms that are set up with redundancy to make disruption of service rare. While most people see Amazon.com as just an online retailer, it has leveraged its expertise with maintaining a massive number of servers to sell a host of web infrastructure and application services to corporate and other organizational clients. Companies and organizations, which in the past tended to maintain their own data locally, increasingly are contracting with specialists such as Amazon.com, Salesforce.com, EMC, IBM, Google, and other leaders in the rapidly growing field of cloud computing—using remote servers for various data and software applications needs. This allows the staff of client firms and organizations to readily access and share data with colleagues, customers, suppliers, and other authorized users—and to benefit from suppliers' economies of scale and expertise in data storage and delivery.

SOCIAL NETWORKING: FACEBOOK AND TWITTER

Cloud computing is also at the heart of online social networking. Web-based social networking generally facilitates users' ability to create a profile on the web and interact with others. For most individuals, social networking has been the most life-changing aspect of Web 2.0. Many people in the developed world, and increasingly the developing world, spend hours each week on social networking websites. For some, especially teens and young adults, web-based social networking has become an everyday activity and a fundamental part of their social lives.

In the early 2000s a host of social networking firms emerged. Among the influential early participants were Friendster (formed in 2002) and MySpace (formed in

2003). These firms, both of which were founded in California and initially focused on the United States, allowed users to create individual public- or semipublic-profile web pages and to connect with others. Friendster and MySpace grew rapidly in their first half-decade and gained millions of users, but in recent years they have been greatly overshadowed by industry-leading Facebook.

Harvard University freshman Mark Zuckerberg founded Facebook—then called Thefacebook—in his shared Kirkland House dorm suite. Frequently occupied with designing and programming computer applications during his first semester, Zuckerberg created two hit programs. The first, Course Match, enabled students to match up classes with others; the second, Facemash, allowed students to compare and choose (based on attractiveness) between two Harvard freshmen portrait photos. Drawing from these dual programs in early January 2004, and also from Friendster (which he belonged to), Zuckerberg registered the domain name *thefacebook.com* and created a platform to allow Harvard University students, faculty, staff, and alumni to create profiles, post a photo of themselves, and invite other members to connect with them as "friends." On 4 February the site went live; within four days, hundreds of Harvard students had registered and created profiles, and in three weeks Thefacebook had more than 6,000 members. From the start Zuckerberg wanted to expand it beyond Harvard to other colleges and universities.

Zuckerberg, and a few friends he partnered with (including his roommate Dustin Moskovitz), expanded Thefacebook to other Ivy League schools and then other universities. By the end of spring 2004 the site had more than 100,000 users. In the summer of 2004, having seen the benefit to operating the firm in Silicon Valley, Zuckerberg moved Thefacebook, by then a corporation, to a house he and his small team rented in Palo Alto, California. One early benefit of this relocation was that Sean Parker, the founder of some early Web 2.0 ventures, reached out to Zuckerberg, the two quickly became friends, and Parker served as the early president of Thefacebook. Though Parker was still young (in his mid-twenties), his experience was particularly helpful in the early years with regard to financing and helping Zuckerberg maintain and maximize ownership control. Subtle and nonintrusive advertising was introduced to help fund the operation.

Rolling Thefacebook out to students, staff, and alumni one university at a time had numerous benefits that helped it eclipse other competing services (in the early years student members were the primary users). In this way, Zuckerberg created anticipation and pent-up demand, targeted a computer literate group particularly focused on their social life, lessened initial needs for operating capital, and allowed the site to grow in a measured way so that greater attention could be paid to reliability. Friendster's reputation, by contrast, was severely hurt by service problems. Further, by starting only with universities and validating accounts based on university-issued "*edu*" addresses, Facebook (the new company name, with the domain *www.facebook.com*, as of August 2005) validated all users. This prevented the

problem of fake profiles that at times plagued Friendster (sometimes derisively called "fakester") and MySpace. After opening up to most higher-education institutions, Facebook next became available to high schools; by the spring of 2006 more than a million high school students were users. The following fall, Facebook became available to all users aged thirteen or older and soon began to concentrate on international expansion—eventually completing a translation project that allowed it to be used in thirty-five languages by the end of 2008. By that time "70 percent of Facebook's then 145 million active users were already outside the United States." Although the company's expanding beyond just universities and "*edu*" addresses opened the door to some fraudulent accounts (an issue interestingly portrayed in the controversial documentary film *Catfish*), the validation of having a substantial number of Facebook friends encouraged authenticity in users' accounts; indeed, Facebook always had a smaller percentage of fraudulent accounts than other large social networking sites.

As with Facebook's growth, its features were added gradually and methodically, and keeping the site simple and user-friendly was a goal from the start. Other than adding profile text and a photo, or changing one's relationship status (for instance, from "single" to "dating"), one of the few initial functions was the ability to "poke" someone—a somewhat ambiguous gesture, ranging from innocent to sexually suggestive, that, unsurprisingly, was popular with college students. Early in its history, Facebook added a feature for sharing photos. While some specialty sites facilitated photo display and sharing (most notably Flickr)—with Facebook it was part of a broader personal platform and thus, for many people, more useful. In 2006 the firm hit its largest hurdle with its launch of News Feed, which automatically updated a stream of news based on changes and updates to friends' profiles. While it did not provide any information that was not already available, the act of automatically pushing information through a user's friend network led many to find it creepy and "stalker-esque." By that time Facebook allowed the creation of profiles for companies, organizations, or interest groups that users could join (or "like"—the organizational equivalent of "friending"). A "Students Against Facebook News Feed" group quickly formed after the application went live, and within days 700,000 people had joined in the online protest, more than 7 percent of all Facebook users at that time. Zuckerberg, who has long stated an ideological preference for information sharing and transparency, was blindsided by this reaction. He took weeks to respond with an apology and with Facebook's installation of enhanced privacy settings/choices that allowed users to disable News Feed. Though most chose not to disable it, this episode highlighted one of Facebook's greatest vulnerabilities—users' privacy concerns.

A platform dependent on voluntary sharing of personal information is highly dependent on not alienating users. This is true of all social networking sites as well

as other key Internet applications such as search. To an unparalleled degree, both Facebook and Google maintain and use massive collections of personal information on users—data that is invaluable to advertisers in targeting customers. Perceived responsible use of this information is fundamental to Facebook and Google's continued existence. Google, which has long publicized "Don't be evil" as its company motto, is perhaps more vulnerable as switching to another search engine is quick, free, and easy. With social networking sites like Facebook, much of the value to users extends from the substantial time they have already invested in building a network and profile content. As with physical spaces, in cyberspace people often want to be together with their friends.

As this chapter is being written, countless social networking sites exist, but Facebook—with around a billion users—is the largest and dwarfs all others. The company went public in May 2012 with a record-setting valuation of more than $100 billion; but in a matter of months it lost more than half its value. As of early 2013, Facebook has recovered much of its lost capitalization owing to strong earnings and successfully meeting the challenge of generating sizable revenue from mobile device users.

Twitter is Facebook's closest rival, but it has only a fraction of its number of users. A San Francisco–based company founded by Jack Dorsey in 2006, Twitter provides a platform for "microblogs"—short (140 characters or less) text-based messages, known as tweets. Some Twitter users, generally those of a young age, have gotten into the habit of microblogging their entire day—from the mundane (a trip to the grocery store) to the more meaningful (participation in a political rally). For some frequent Facebook or Twitter users, smartphones facilitate wide-ranging access, allowing tweets to be a surrogate for person-to-person communication—letting others know where they are, what they are doing, or when they will be back. Other tweeters focus more on writing brief perspectives on political events, entertainment news, or products that they like or dislike.

POLITICS OF THE INTERNET

As usage of the Internet greatly broadened with the advent of the World Wide Web, many journalists, politicians, and others have presented it as a transformative technology of freedom and democracy. Organizations such as the Electronic Frontier Foundation, formed in 1990 by Lotus Development founder Mitch Kapor and Grateful Dead lyricist John Perry Barlow, have emerged to defend individual Internet rights—with some such advocates from the political left and many others from the libertarian right. The broadening participation in user-created web content—the defining characteristic of Web 2.0—and the newfound mobile computing of smartphones not only fuel such framings of the Internet but also highlight it as a democratizing tool to help battle authoritarian regimes.

Journalists quickly christened the 2009 protests in Iran—against alleged fraud in reelecting president Mahmoud Ahmadinejad—as the "Twitter Revolution." Users' tweets reportedly contributed to spreading the word about the protests and to voicing the opinions of dissenting locals. Closer analysis showed that only a small percentage of Iranians used Twitter (many tweets on the alleged election fraud originated from westerners), and that early coverage of the events instead "revealed intense Western longing for a world where information technology is the liberator rather than the oppressor." Similar early characterizations by journalists and politicians highlighted the important role of Facebook and Twitter in the 2011 Arab Spring (protests throughout the Arab world that led to the overthrow of authoritarian regimes in Tunisia, Libya, and Egypt) and the 2011 international Occupy Movement (sit-in demonstrations against political and economic inequality and concentrated corporate power that began with an Occupy Wall Street protest in lower Manhattan's Zuccotti Park). Although social networking apparently played a considerable role in the Occupy Movement, Internet penetration is still fairly low in many Arab nations and relatively few citizens of these countries use Twitter or Facebook. At least as important to these varied protests was the older technology of cellphone texting, which enabled information to rapidly percolate, from person to person, throughout a community. Some opposition organizers and protesters are justifiably hesitant to go on social networking sites, given likely retribution by authoritarian regimes. This raises the question of whether the Internet is primarily a tool of individual freedom or of government and corporate control. In 2011 the Western community was alarmed when the Egyptian government, working with service providers, shut down Internet and cellphone services in the face of political unrest. Clearly the Egyptian government understood the potential power of social networking and (with the cooperation of corporate service providers) exercised ultimate control.

Such control is not limited to specific times of political unrest such as the Arab Spring. Starting in 2000, the Chinese Information Industry Ministry banned "harmful texts or news likely to jeopardize national security and social stability." Yahoo! agreed to such a "self-discipline pact" in 2002 and used filtering mechanisms to limit speech and prevent "controversial" messages containing words or ideas about topics such as "Taiwan independence" or "human rights" from ever appearing on Chinese discussion forums. Google, too, enforced China's Internet censorship (until March 2010). The Chinese government deploys software to automatically identify political viewpoints, and bloggers in China must register with authorities. Internet cafes in China—though generally believed to offer anonymity and privacy—are actually sites of surveillance where cameras are directed at users and their computer screens.

How much and what types of regulation should be established for the Internet are bitterly contested questions within and between nations. A number of transportation and communications technologies have long required international

cooperation, but perhaps none to the same extent as the largely borderless Internet. The history of radio may offer some guide: there was little regulation at first, but in the 1920s and 1930s substantial national regulation emerged in the United States and Europe, along with some international regulation with regard to short-wave radio. The evolving dominance of a small number of corporate giants is another common thread connecting the history of radio and that of the Internet. But what is truly new is the systematic collection and use of personal data on individuals for commercial gain to the extent now being exercised by Google, Facebook, Amazon.com, Twitter, and other Internet and e-commerce firms.

Whereas some social scientists have presented Web 2.0, social networking sites, and mobile computing as positive forces and valuable tools that enhance our lives (a new "social operating system"), others have highlighted the ways in which such technologies ultimately fail to deliver deep and meaningful social connections and tend to make us "alone together." Attempts to characterize recent technologies' particular roles in social, political, cultural, and psychological change are fraught with pitfalls best avoided by the historian. What is clear at this point is that, with mobile computing and the Internet a part of most peoples' work, economic, and social lives as never before, decisions concerning the Internet and its usage by companies, other organizations, legislators, judges, and ourselves will increasingly shape our future.