10 Licklider's Intergalactic Computer Network

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 gigantic memories and the sophisticated programs would be divided by the number of users.

J. C. R. Licklider[1]

The network is the computer

Today, with the Internet and World Wide Web, it seems very obvious that computers become much more powerful in all sorts of ways if they are connected together. In the 1970s this result was not so obvious. This chapter is about how the Internet of today came about. As we can see from Licklider's (B.10.1) quotation beginning this chapter, in addition to arguing for the importance of interactive computing in his 1960 paper on "Man-Computer Symbiosis," Lick also envisaged linking computers together, a practice we now call *computer networking*. Larry Roberts, Bob Taylor's hand-picked successor at the Department of Defense's Advanced Research Projects Agency (ARPA), was the person responsible for funding and overseeing the construction of the ARPANET, the first North American *wide area network* (WAN). A WAN links together computers over a large geographic area, such as a state or country, enabling the linked computers to share resources and exchange information. As Roberts said later:

> Lick had this concept of the intergalactic network which he believed was everybody could use computers anywhere and get at data anywhere in the world. He didn't envision the number of computers we have today by

B.10.1. Joseph C. R. Licklider (1915–90) was a visionary computer pioneer whose impact is still felt everywhere in computer science. His interests were wide ranging and included psychological aspects of communications and learning, brain studies, computer networks, time-sharing computers, interactive systems, and cooperation between computers and humans. His groundbreaking paper "Man-Computer Symbiosis" investigated the possibility of a closer cooperation between humans and computers with computers being used to augment human intellectual capacity. His 1968 paper with Bob Taylor, "The Computer as a Communications Device," outlined their joint vision of what has become the present-day Internet.

B.10.2. Sun Microsystems was founded in the early 1980s by two Stanford MBAs, Vinod Khosla and Scott McNealy, and a Stanford graduate student, Andy Bechtolsheim, together with another graduate student, Bill Joy, from the University of California, Berkeley. Sun was an acronym for Stanford University Network, where Bechtolsheim's prototypes were already running, connected by the Ethernet.

> any means, but he had the same concept – all of the stuff linked together throughout the world, that you can use a remote computer, get data from a remote computer, or use lots of computers in your job. The vision was really Lick's originally. None of us can really claim to have seen that before him nor [can] anybody in the world. Lick saw this vision in the early sixties. He didn't have a clue how to build it. He didn't have any idea how to make this happen. But he knew it was important, so he sat down with me and really convinced me that it was important and convinced me to move into making it happen.[2]

It was Roberts and a small team of dedicated engineers and graduate students, mostly trained at MIT, who built the ARPANET. The ARPANET was rapidly followed by a worldwide proliferation of similar – but incompatible – networks. In 1974, Bob Kahn and Vint Cerf brought order to this chaos by publishing a paper with the intimidating title "A Protocol for Packet Network Intercommunication." Their paper coined the term *Internet* for what they called the "internetworking of networks."

In 1982, two Stanford MBAs, Vinod Khosla and Scott McNealy, founded Sun Microsystems, together with hardware expert Andy Bechtolsheim, a PhD student in Stanford University's Electrical Engineering Department, and a Unix software expert, Bill Joy, a graduate student from the nearby University of California, Berkeley (B.10.2). They set up the company to develop robust single-user *workstations*, computers with less computing power than minicomputers but more powerful than PCs. From the start, the Sun founders envisioned networking their workstations. As Sun's CEO, Scott McNealy, said,

> The whole concept of "The Network is the Computer" we started at Sun was based on the fact that every computer should be hooked to every other computing device on the planet.[3]

To begin our story, we shall go back in time to the origins of the telegraph (Fig. 10.1) system and briefly describe what has been called "The Victorian Internet" by author and journalist Tom Standage.

Fig. 10.1. An advanced version of a telegraph with a printing receiver and transmitter.

Fig. 10.2. A map of the global communications network of the British Empire in 1872. Telegraph cables reached as far as Hong Kong and Australia.

The Victorian Internet

The preface to Standage's book includes a summary of the impact of the new technology:

> During Queen Victoria's reign, a new communications technology was developed that allowed people to communicate almost instantly across great distances, in effect shrinking the world faster and further than ever before. A worldwide communications network whose cables spanned continents and oceans, it revolutionized business practice, gave rise to new forms of crime, and inundated its users with a deluge of information. Romance blossomed over the wires. Secret codes were devised by some users and cracked by others. The benefits of the network were relentlessly hyped by advocates and dismissed by its skeptics. Governments and regulators tried and failed to control the new medium. Attitudes towards everything from news to diplomacy had to be completely re-thought. Meanwhile, out on the wires, a technological subculture with its own customs and vocabulary was establishing itself.[4]

Although this sounds like a description of the present-day Internet, it is in fact referring to the global telegraph network, which transformed business and personal life in the nineteenth century, over a hundred years earlier than the Internet (Fig. 10.2). Just as the Internet of today is sometimes known as the "Information Superhighway," the nineteenth-century telegraph network was called the "Highway of Thought."

Message routing systems, in which a message is passed from one "station" to the next, have operated since antiquity. The mechanism for passing the messages has changed, however. At first, human runners carried the messages. In 1791, the French engineer Claude Chappe (B.10.3) invented a sophisticated optical signaling system that consisted of a network of towers. An operator in each tower moved two large, jointed arms on a signaling device called a *semaphore* to spell out messages. Each tower also had two telescopes, one pointing backward and the other forward. Because Chappe's semaphore network was an optical system, it required good line-of-sight visibility between the towers. A major drawback of the system was that it only worked well in good weather,

B.10.3. French postage stamp with Claude Chappe (1763–1805), inventor of the optical telegraph.

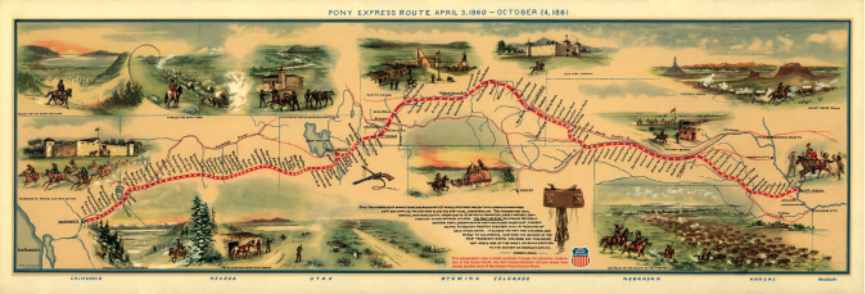

Fig. 10.4. The route of the famous Pony Express in the United States. This mail service showed that fast transcontinental communications were possible, but the service was expensive and was overtaken by telegraph technology along the same route. The Pony Express operated for less than two years, from April 1860 to October 1861.

and then only during the daytime. When Napoleon Bonaparte came to power in 1799, he quickly recognized the military value of such a rapid communication system and arranged for a portable version of Chappe's semaphore signaling equipment to be developed. In searching for a name for his invention, Chappe devised the word *telegraphe* from two Greek words meaning *to write at a distance*. His semaphore system operated with a countrywide network of towers in France for more than fifty years (Fig. 10.3). Finland, Denmark, Sweden, Russia, and the United Kingdom quickly established similar optical telegraph networks. At the height of the network's popularity, nearly a thousand semaphore towers operated across Europe.

For all the success of Chappe's system, it was superseded by a superior technology after little more than fifty years: telegraph technology based on using electric signals that traveled over cables. Before the telegraph network became fully operational, couriers on horseback used to deliver long-distance messages. In the United States, the Pony Express linking Missouri with California operated for only slightly more than a year between 1860 and 1861, but could not compete for long with the telegraph (Fig. 10.4). The electric telegraph had much greater speed and reliability, and operated in all weather conditions. In the United Kingdom, it was an entrepreneur named William Fothergill Cooke who, in an uneasy collaboration with physics professor Charles Wheatstone, worked hardest to raise enthusiasm for building an electric telegraph network to span the country (B.10.4). At about the same time, in the United States, Samuel Morse (B.10.5), a painter and scientist, was also working tirelessly toward the same objective. The electric telegraph sent messages using a code of dots and dashes that we now call Morse code. Although the specific telegraph technology used in the United Kingdom and the United States differed in detail,

Fig. 10.3. One of Claude Chappe's telegraph towers was rebuilt near Saarbrucken, in Germany. The French engineer Chappe succeeded in covering France with a network of 556 stations stretching a total distance of 4,800 kilometers. The network was used for military and national communications from 1792 to the 1850s.

B.10.4. Entrepreneur William Cooke (1806–79) and physicist Charles Wheatstone (1802–75) pioneered the telegraph in the United Kingdom. After a slow start, undersea cables were laid connecting the whole of the British Empire.

Fig. 10.5. Memorial in Telegraph Field, Valentia Island, Ireland, commemorating the site of the European end of the first transatlantic cable in 1858.

the telegraph network expanded rapidly in both countries. By 1850, there were more than two thousand miles of wire in the United Kingdom and more than twelve thousand miles in the United States. In 1852, the first underwater cable linking London and Paris was laid across the English Channel. In the United States, the idea of a transatlantic cable had been suggested by Morse in the 1840s but the idea was thought to be impractical. In 1854, the wealthy businessman Cyrus Field (B.10.6) took up the idea and had a cable laid from New York to St. John's in Newfoundland, preparing to extend the cable across the Atlantic Ocean to Ireland (Fig. 10.5). Field persuaded both the U.S. and British governments to back the project, and the first undersea connection was established in August 1858 (Fig. 10.6). Queen Victoria sent the first transatlantic message to U.S. President James Buchanan.

Unfortunately, the engineer in charge of the project, Edward Whitehouse, had very little understanding of the science of underwater telegraphy. Within a month, the cable had failed. A joint government inquiry was set up to determine the reasons for the failure. It was William Thomson, then a physics professor at Glasgow University, later to become Lord Kelvin, who put submarine telegraphy on a sound scientific foundation. Thomson's understanding of the relevant physics was confirmed by the successful laying of an undersea cable through the Persian Gulf in 1864, connecting Europe to India. One of the key technologies for undersea cables was the use of a rubbery gum called *gutta-percha*, obtained from a tree grown in Southeast Asia, to coat and protect the cables (Fig. 10.7). The London-based Gutta Percha Company suddenly found itself with a virtual monopoly on the production of submarine cables. The Gutta Percha Company eventually became part of Cable & Wireless Worldwide.

B.10.5. Photograph of the statue of Samuel F. B. Morse (1791–1872), pioneer of the telegraph and inventor of Morse code, in New York's Central Park. The statue was dedicated in 1871.

The telegraph system was an early example of a *store-and-forward* network, in which messages were sent to an intermediate station before they were transmitted to their destination. Telegraph services routed signals through intermediate relay stations because electrical losses from telegraph wires accumulated and degraded a message if it had to travel over long distances. Intermediate stations received messages as dots and dashes of Morse code and recorded them on punched paper tape. Incoming messages were separated from each other by tearing the tape in the appropriate places, which led to the intermediate stations being called "torn-tape relay centers." A telegraph operator at the receiving station read the destination on the message tape

B.10.6. Cyrus Field (1819–92) began work at age fifteen as an office boy for A. T. Stewart & Co., New York City's first department store. By the age of twenty, he was a partner in a paper manufacturing company, and at thirty-three was able to retire as a wealthy man. In 1854, Field became enthusiastic about the possibility of laying an undersea transatlantic telegraph cable from Newfoundland to Ireland. After several failed attempts, in August 1858 Field arranged for Queen Victoria to send the first transatlantic message to President James Buchanan, and there were great celebrations on both sides of the Atlantic. The cable and the project were not successfully completed until 1866. In the portrait, Field is seen touching a globe and holding a length of cable wire.

and took it to the appropriate transmitter for sending to the next relay center, until the message reached its final destination. At the height of the telegraph's popularity, a major relay station had dozens of inbound receiving and outbound transmitting terminals, with scores of operators and thousands of messages queued up during peak hours. Our present-day computer networks also use a store-and-forward system, not for whole messages but for standard-size pieces of information called *packets*.

Fig. 10.6. Isambard Kingdom Brunel's Great Eastern was the largest ship afloat in 1865 and was ideally suited to cable laying. On 13 July 1866, the ship left Valentia Bay in Ireland and arrived in Newfoundland two weeks later.

Nuclear war and packet switching

The idea of a digital communication network using *packet switching*, which breaks up the message into small packets for transmission, occurred almost simultaneously to two researchers on different continents – but for very different reasons. Paul Baran was a researcher at the RAND Corporation in Santa Monica (B.10.7). The RAND Corporation – RAND is an abbreviation of Research ANd Development – was originally part of the Douglas Aircraft Company and was set up as an independent nonprofit research organization in 1948. Baran had previously worked for the Eckert-Mauchly Computer Corporation on the UNIVAC computer and for the Hughes Aircraft Company in Los Angeles on the computerized Semi-Automatic Ground Environment (SAGE) early warning system. Hughes was bidding on a contract for the control system for the Minuteman missiles, and Baran became alarmed, because, as he said later, "You had all these missiles that could go off by anyone's stupidity. The technology was never to be trusted."[5]

Fig. 10.7. The gutta-percha tree grows in Southeast Asia. It produces a rubbery gum that coped better with the underwater environment than traditional rubber and was used to coat the undersea cables.

When Baran joined RAND in 1959, he set out to study the problem of whether U.S. communication systems could survive a nuclear attack by the Russians. The late 1950s were times of great political tension between the United States and the Soviet Union. Two years earlier, the Soviets had managed to put a satellite into Earth orbit, an event that the United States interpreted as a clear threat to national security (Fig. 10.8). Both the United States and the Soviet Union had large stockpiles of nuclear weapons. For these nuclear arsenals to work as a deterrent to war, it had to be evident to each country that the other country would be able to launch a retaliatory attack after a first strike. This was the doctrine of *mutual assured destruction*, usually known by the acronym MAD. If the early warning systems failed and missiles had exploded on U.S. territory, the president needed to have the "minimal essential communications" to launch a counterattack.

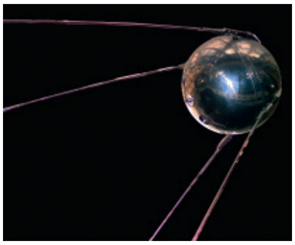

Fig. 10.8. The launch of the Sputnik satellite by the Soviet Union on 4 October 1957 was a wake-up call for the United States. The first Sputnik weighed only 184 pounds and was the size of a basketball. Sputnik II, launched just a month later, weighed half a ton.

Until Baran's arrival at RAND, there had been little progress on how to protect the U.S. communications network in the event of a nuclear attack. Because of his familiarity with digital computers, Baran found that he thought about the problem very differently from most of his colleagues at RAND. He said, "Many of the things I thought possible would tend to sound like utter nonsense, or impractical, depending on the generosity of spirit in those brought up in an earlier world."[6] Signals on a telephone network were *analog* signals, in which the amplitude of the signal varied continuously, and the quality of the signal deteriorates as the number of "hops" through a network increased. Baran realized that by using digitized messages generated by a computer, a signal could be stored, replicated exactly, and retransmitted an unlimited number of times. This technique would allow the signal to be transmitted over long distances through intermediate routing stations without any distortion or loss.

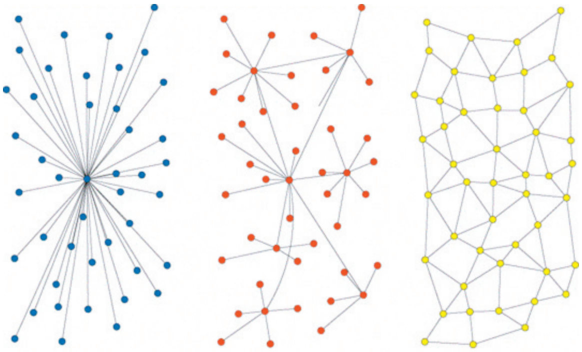

Fig. 10.9. A figure from Paul Baran's original paper on packet switching shows examples of centralized, decentralized, and distributed networks. The distributed network is also called a "fishnet" network, and its most important feature is the lack of centralization so that all nodes have an equal importance. This fishnet architecture also contains a certain level of redundancy because each node is connected to several other nodes. This redundancy allows multiple paths for the message packets to reach their destination. This provides the network with much more robustness and reliability than a centralized network with a single path to the destination node.

In the early 1950s, Warren McCulloch, a neurophysiologist at MIT, had been talking with John von Neumann and others about the brain and its network of *neurons* (nerve cells in the brain). Together with the mathematician Walter Pitts, McCulloch developed a mathematical model of a neuron and linked these model neurons together to form an artificial "neural network" (see Chapter 13). This network was a computer program that operated much like the connected nerve cells in the nervous system. Baran knew of this work and was influenced by McCulloch's ideas of neural networks:

> Warren McCulloch in particular inspired me. He described how he could excise a part of the brain, and the function in that part would move over to another part. . . . McCulloch's version of the brain had the characteristics I felt would be important in designing a really reliable communication system.[7]

At the time, the telephone system depended on a hierarchical system in which switching centers ranked one above the other. Baran had analyzed what would happen to the telephone network after a nuclear attack on U.S. military targets, and his results showed that the system was very vulnerable to even the incidental, "collateral" damage caused by such an attack. In designing a more robust system, Baran introduced what he called "redundancy" – that is, exceeding what is necessary – into the network, by allowing multiple pathways between the different centers. He advocated building a distributed network that looked rather like a fishnet (Fig. 10.9). A network with only one link connecting each center has a redundancy level of one and is obviously extremely vulnerable to disruption. Baran's calculations showed that "just a redundancy level of maybe three or four would permit almost as robust a network as the theoretical limit."[8] This meant that even if some links in the network were destroyed in a nuclear attack, because each center was connected to three or four other centers it would still be possible to find working paths through the network.

B.10.7. Paul Baran (1926–2011) pioneered the concept of distributed networks and packet switching for reliable communications after a nuclear strike. He was never able to persuade AT&T engineers to build even a prototype of his packet-switching system.

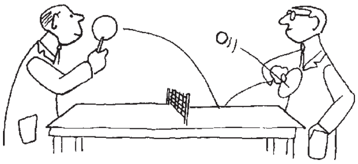

Fig. 10.10. "Interactive Communication consists of short spurts of dialog.[F1]"

Baran's second idea was just as radical: to introduce a type of redundancy in sending messages. He proposed splitting each message into a number of fixed-length "message blocks" – what we now call *packets* – and allowing the blocks to take independent paths through the network. In an interview with the writer Stewart Brand, Baran later described what he called the "wonderful" properties of this method, called *packet switching:*

> Anyway, you send the packet out, and when this station gets it off to the next guy, it sends back, "OK, got it. You can erase the previous one." If the first station doesn't hear back, it sends another copy out in a different direction. The packets can arrive out of order. We just sort them out at the end. Since it didn't have to be synchronous, you didn't have to lock everything all together. It didn't take very long before we started seeing all sorts of wonderful properties in this model. The network would learn where everybody was. You could chop up the network and within half a second of real-world time it would be routing traffic again. Then we had the realization that if there's an overload in one place, traffic will move around it. So it's a lot more efficient than conventional communications. If somebody tries to hog the network, the traffic routes away from them. Packet switching had all these wonderful properties that weren't invented – they were discovered.[9]

In spite of all these advantages, Baran had a hard job trying to convince AT&T – which then operated virtually the entire U.S. telephone network – of the virtues of distributed networks and packet switching. At the time, all telephone networks used what is known as *circuit switching*. When you called someone, the network established a physical connection between the two endpoints and the intermediate links were dedicated to your conversation as long as you stayed on the line. For a phone call, this method makes sense, because there is typically a steady exchange of information during a conversation. By contrast, data communications sent from a user sitting at a computer tend to be "bursty," with large amounts of data sent all at once, followed by pauses with no data being transmitted (Fig. 10.10). Keeping a line reserved for this type of communication is not a very efficient way to use the available bandwidth in the network. *Bandwidth* is the technical term used to define a network's maximum capacity for transmitting information. Using packets that can be routed in multiple ways through a distributed network turns out to be a much more efficient use of the available bandwidth because packets from different messages can use the same links. Because the different packets of the same message can travel over multiple different paths, the packets making up any given message may arrive out of sequence and need to be reassembled in the right order at the receiving center. Each packet must therefore contain a *header* at the beginning of the packet that describes where it is going and to which message it belongs.

In response to criticisms of his ideas from his colleagues, Baran wrote a series of papers that showed in detail how all these problems could be solved. One new factor in understanding his idea was simply the much higher speed of switching that was possible in computer networks. As we have seen, the electromagnetic relay switches of the earliest computers had been replaced by electronic switches capable of much-faster switching rates. To AT&T engineers not familiar with the new digital computer technology, Baran's idea of breaking up a telephone conversation into packets of information seemed truly ridiculous. He recalled:

> The story I tell is of the time I went over to AT&T headquarters – one of many, many times – and there's a group of old graybeards. I start describing how this works. One stops me and says, "Wait a minute, son. Are you trying to tell us that you open the switch up in the middle of the conversation?" I say, "Yes." His eyeballs roll as he looks at his associates and shakes his head. We just weren't on the same wavelength.[10]

By 1965, Baran had managed to persuade the U.S. Air Force to back the creation of a trial distributed switching network. Unfortunately, AT&T declined any involvement in the project and the proposal never got off the ground. Baran reluctantly moved on to study other problems.

At about the same time that Baran was abandoning his research on packet switching, Donald Davies at the United Kingdom's National Physical Laboratory (NPL) was also thinking about packet switching and computer networks (B.10.8). His interest stemmed from reasons very different from Baran's concerns about network survivability in a nuclear war. Davies had joined NPL in 1947 to work on Alan Turing's ambitious project to build the Automatic Computing Engine – the ACE computer. Turing became frustrated by the delays and bureaucracy at NPL and left to join the University of Manchester's computer team, leaving Davies in charge of producing a less ambitious version of ACE. Davies's team delivered the Pilot ACE machine in 1950, and the English Electric Company later successfully marketed a commercial version called the DEUCE computer. After a visit to MIT in 1954, Davies became interested in the problem of sending data communications over a network. While at MIT, he had seen that a significant problem with time-sharing computers was the cost of keeping a phone connection open for each user. Instead of being concerned about the survivability of networks after a nuclear strike, Davies was thinking about the efficient support of online data processing in which users at computer keyboards were generating the data. In 1965, he wrote in an internal NPL note outlining the problem:

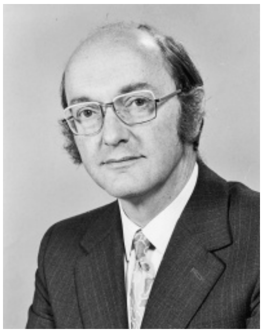

B.10.8. Donald Watts Davies (1924–2000) graduated from Imperial College in London in mathematics and physics. He then joined the United Kingdom's NPL, where he worked with Alan Turing on the ACE computer. Independent of Paul Baran's work in the United States, Davies came to almost the same conclusions as Baran about efficient and reliable communication networks. However Davies's reasons for developing packet-switching networks were related to efficient use of time-shared computers and not for providing network survivability in the event of a nuclear war.

> Starting from the assumption that on-line data processing will increase in importance, and that users will be spread out over the country, it is easily seen that data transmission by a switched network such as the telephone network is not matched to the new communication needs that will be created. The user of an on-line service wishes to be free to push keys sporadically, and at any rate he wishes, without occupying and wasting a communication channel.[11]

Davies also realized that delays in store-and-forward distributed networks could be minimized by using short message blocks he called *packets*. He very deliberately introduced the word *packet* to describe the fixed-length, short blocks of information that made up the message and that traveled separately through the network to their destination. The choice of name made clear that packet switching was fundamentally different from traditional message switching. And it was certainly a much better meme than the phrase "distributed adaptive message block switching" used by Baran!

At the end of 1967, one of Davies's colleagues at NPL, Roger Scantlebury, presented a paper on their packet-switching work at a conference in Gatlinburg, Tennessee. Scantlebury's paper outlined a design for a packet-switching network consisting of relay centers called *nodes* connected by digital links. The nodes handled the transmission of the packets between the nodes, and *interface computers* connected the network of nodes to time-sharing computers and to other subscribers. It was from Scantlebury that Larry Roberts first heard about Baran's work on packet switching (B.10.9). He later said of the revelation he experienced at this conference, "Suddenly I learned how to route packets."[12] When Roberts returned to Washington, he hunted down the RAND reports that Baran had sent to ARPA and found time to meet with him in early 1968. Meanwhile, back in the United Kingdom, Davies and his colleagues set about building a small, packet-switching network at NPL. The Mark I version of the NPL network became operational in 1970, and a Mark II version remained in operation until 1986. Despite the similarity of their ideas, Davies always gave credit to Baran as being the first to publish the idea of packet switching in distributed networks, saying, "The honour for [originating packet switching] must go to Paul Baran."[13]

The ARPANET and the third university

When Licklider went to ARPA, he set out to find and fund the leading computing research centers in the country. From these connections, he established an informal circle of advisers, consisting of about a dozen computer scientists from MIT, Stanford University, and the universities of California at Los Angeles and Berkeley as well as from some computer companies. He lightheartedly called this group his "Intergalactic Computer Network." After he had been at ARPA six months, Lick had seen firsthand the diversity and incompatibility of all the computer hardware and software in the research community. As a result, he wrote a memo to the group that raised the possibility of connecting the different computing systems at each site into a network:

> Consider the situation in which several different centers are netted together, each center being highly individualistic and having its own special language and its own special way of doing things. Is it not desirable or even necessary for all the centers to agree upon some language or, at least, upon some conventions for asking such questions as "What language do you speak."… It seems to me to be important … to develop a capability for integrated network operation.[14]

By this memo, Lick had extended the idea of his Intergalactic Computer Network from just a group of people to a network of interoperating hardware and software.

When Bob Taylor (B.10.10) became director of the computing program at ARPA in 1966, he inherited this vision of Lick's. Obtaining the funding for such a network was only the first hurdle for Taylor. He now needed a program manager who could actually build the network, and he knew exactly the person he wanted. Taylor had just funded an experiment to connect the TX-2 computer at MIT's Lincoln Laboratory to the System Development Corporation's

B.10.9. Larry Roberts was a graduate in electrical engineering from MIT and went to work on research projects at MIT's Lincoln Laboratory. At the age of twenty-nine, he became the project leader and chief architect of the ARPANET. After leaving ARPA, he moved to the communications industry and founded the first commercial packet-switching network, called TELENET.

Q-32 computer, a military mainframe computer in Santa Monica, California. The person in charge of the project was Larry Roberts, an MIT PhD in electrical engineering who, like many others before him, had migrated up the road to the Lincoln Lab. The report on the networking experiment had concluded that although the connection had been successfully established, its reliability and response time to commands were, as later described by Roberts, "just plain lousy." Roberts, with his deep technical knowledge of the problem, was the ideal candidate for the job but, unfortunately for Taylor, Roberts had no interest in leaving Lincoln Lab (Fig. 10.11) to become what he called "just a bureaucrat." In late 1966, nearly a year after he had received the funding for the project, Taylor persuaded ARPA Director Herzfeld to call the director of Lincoln Lab and tell him about ARPA's problem and, at the same time, point out that more than 50 percent of the lab's research funding came from ARPA. Roberts accepted the job two weeks later.

Fig. 10.11. MIT's Lincoln Laboratory was established in 1951 to build the nation's first air defense system. However, its roots date back to the MIT Radiation Laboratory, which was formed out of the Physics Department during World War II to develop radar for the Allied war effort.

When Roberts took over the ARPANET project in late 1966, he saw that he had three major technical challenges to solve, as well as a surprising "sociological" problem involving human behavior and social relationships. The first challenge was how to physically connect all the time-sharing computers at the different sites. In the experiment between Lincoln Lab and Santa Monica, Roberts had shown that a direct telephone line connection between the two computers could work. The problem was that ARPA had funded more than a dozen major time-sharing computers, and to establish direct connections between all these computers would require more than sixty-five long-distance telephone lines, a number that would rapidly increase and become very expensive as the number of computer systems increased. His second challenge was that even if he had these long-distance lines, how could they most efficiently be used? In their book *Computer: A History of the Information Machine*, Martin Campbell-Kelly and William Aspray state that: "Experience with commercial time-sharing systems had shown that less than 2% 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."[15] The last technical problem that Roberts needed to solve was how all the different incompatible computer systems would communicate with one another without each site having to write many different software interfaces. Although unknown to Roberts at this time,

B.10.10. Bob Taylor graduated in 1958 from University of Texas with a degree in psychology and mathematics. In 1965, at the age of thirty-four, he became director of the Information Processing Techniques Office at ARPA. In this role he was responsible for creating the program that led to the creation of the ARPANET. After leaving ARPA, he went on to be the founding director of the Computer Science Laboratory at Xerox's new PARC. It was at PARC where many of the pioneering computing technologies that we see around us today were invented, guided by Taylor's leadership. See Chapter 8 for a more detailed discussion.

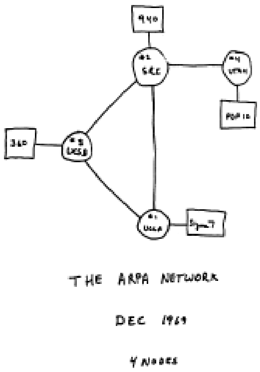

Fig. 10.12. A sketch by Alex McKenzie of the initial four-node configuration of ARPANET connecting the University of California at Santa Barbara, the University of Utah, UCLA, and SRI. The communication backbone of the network was provided by fifty kilobit per second telephone links.

the first two problems had already been solved by Baran and Davies. The use of a store-and-forward network solved the problem of needing to have every computer connected to every other computer. Furthermore, a distributed network, rather than a hierarchical or centralized network, would allow multiple paths between the origin and destination of the messages. Second, the work of Baran and Davies on splitting up the messages into fixed-length packets had shown that a packet-switching network allowed more efficient use of each communication line than a standard message-switching network. Using packets, a single user would not take up the bandwidth of each link, as in a circuit-switched telephone call. Roberts first learned about packet switching and about Baran's work from Roger Scantlebury at the Gatlinburg conference in October 1967.

Roberts had found a solution to his third problem at another conference earlier that year. After a meeting in Ann Arbor, Michigan, in a taxi back to the airport, Wes Clark, another Lincoln Lab veteran, outlined a solution to the problem of each site having to write multiple different versions of interface software for its computer to handle the data communication with the other, different computers. Clark suggested inserting a minicomputer between each mainframe host computer and the network. Each site would then only have to create software to connect its mainframe to the standard network minicomputer. A similar solution had been proposed by the NPL team at the later conference in October. Roberts called the intermediate computers that managed the packet routing and delivery *interface message processors*, or IMPs.

It was at the Ann Arbor meeting that Roberts realized that he also had a sociological challenge on his hands. Many university researchers felt they had worked hard to get funding for "their" computers and were not enthusiastic about having to "waste time" writing software to enable others to use their valuable resources. Roberts noticed a distinct regional bias:

> We actually had more conservatism on the East Coast. When I looked for sites that were willing to start, the four West Coast sites were interested and excited to be involved. And the East Coast sites, like MIT, said "Well, I don't want you to touch my computer." So we went with the ones that were cooperative. . . .[16]

B.10.11. Len Kleinrock ran the Network Measurement Center at UCLA. He is photographed here with the first IMP.

Two of the groups on the West Coast who were enthusiastic about being connected to a network were Doug Engelbart's team at the Stanford Research Institute (SRI) and Len Kleinrock's Network Measurement Center at the University of California, Los Angeles (UCLA) (B.10.11). Engelbart was working on NLS, his oNLine System for interacting with computers, and he saw the ARPANET proposal as an opportunity to extend his system to support distributed collaboration. Kleinrock had been a PhD student at MIT with Roberts, and in his 1962 thesis had used *queueing theory*, the mathematical modeling of queue lengths and waiting times, to simulate the behavior of store-and-forward message-switching networks. His network simulations had used mathematical models for both the rate of message generation and for the distribution of message lengths. The other two groups who responded favorably to networking their sites were the University of California, Santa Barbara, and the University of Utah, both working on ARPA-funded interactive graphics research. The first four nodes of the ARPANET were therefore all on the West Coast of the United States (Fig. 10.12).

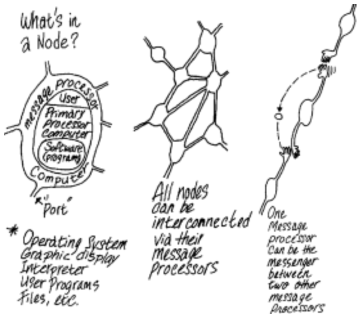

Fig. 10.13. Diagrams of the host nodes and the IMP network: (a) the host nodes; (b) the message processors that allow interconnections through the network; and (c) each message processor can be the messenger between two other message processors.

By early 1968, Roberts had nearly completed a detailed design for the network. There were three basic principles (see Fig. 10.13). The first was that the IMPs should operate as a communication system whose essential task was to transfer packets from anywhere to anywhere else on the network. The IMPs would also take care of route selection and acknowledgment of receipt for the packets. The second requirement was that the network must have a very small message delay time. From his experience with time-sharing systems, Roberts decided that the average transit time through the network must be less than half a second. The third principle was that the IMP system should function independently of whether a host computer had crashed. Network reliability should depend on the IMPs and not the host computers. In addition, the NPL team from the United Kingdom had convinced Roberts that he should specify the use of much faster links than he had specified in his original proposal. By July 1968, Roberts sent out a formal request for quotation to 140 companies giving them information about the project's requirements and inviting them to bid on building the IMPs. One early setback was that two of the major computer companies, IBM and Control Data Corporation (CDC), declined to bid and said that "the network could never be built because there existed no computers small enough to make it cost-effective."[17] Fortunately, Roberts eventually received more than a dozen bids. In December 1968, ARPA announced that it was awarding the contract to build the IMPs to a small consulting firm in Cambridge, Massachusetts, called Bolt, Beranek and Newman, or BBN.

BBN started life as a small consultancy group advising on acoustics in 1948. Richard Bolt and Leo Beranek taught on the MIT faculty, and both were experts on acoustics and, in particular, the acoustics of buildings. Robert Newman, who had been a student of Bolt's and was an architect, joined the consultancy a year later, and BBN was born. The business grew rapidly, and besides its profitable work on buildings, BBN developed unique expertise in the analysis of audiotapes. The company assisted in the analysis of the film of the shooting of President John F. Kennedy and of the shooting deaths at Kent State University. Most famously, in the wake of the Watergate scandal, the White House and the special prosecutor's office called in BBN to examine the 18½-minute gap in the Nixon White House tapes. Dick Bolt headed the investigating committee, which concluded that the erasure was deliberate. By the time that BBN recruited Licklider in 1957, the company had a well-established hiring philosophy summarized by Beranek: "I had the policy that every person we hired had to be better than the previous people."[18] BBN was also well known for its policy of hiring MIT dropouts. The rationale was that if the person had managed to get into MIT, he or she must be smart, and, if the person had subsequently dropped out, all this meant was that he or she could be hired more cheaply than someone who had graduated. Because of the company's recruiting policies and because of the absence of any academic tenure process and any teaching commitments, BBN became a very attractive place for researchers to work. In this way, BBN acquired the informal reputation of being the "third university" in Cambridge, along with Harvard and MIT.

When Licklider joined the company, he convinced Beranek that they should buy him a computer. Beranek later said, "I decided that it was worth the risk to spend $25,000 on an unknown machine for an unknown purpose."[19] The gamble paid off and the expertise that the company developed in computing soon became a major asset. BBN later bought the first PDP-1 from Digital Equipment Corporation, the first relatively inexpensive computer that could be operated by a single person. Licklider and his team used the machine to develop one of the first time-sharing systems capable of supporting four simultaneous users. In 1966, BBN went to the Lincoln Lab to recruit an engineer named Frank Heart to work on a hospital computer project. Heart was a graduate from MIT who had taken MIT's first-ever course in computer programming. He had also earned a master's degree while working on the Whirlwind project, the precursor to the SAGE system, and then followed many other MIT graduates in moving to Lincoln Lab to work on its portfolio of exciting real-time computing projects. By the mid-1960s, when BBN recruited Heart, his Lincoln Lab colleagues were the acknowledged experts in building real-time, interactive computing systems.

B.10.12. Bob Kahn studied at the City College of New York and earned his doctorate at Princeton University in 1964. After working at Bell Labs and then as a professor in electrical engineering at MIT, he took leave from MIT to work at BBN. With Frank Heart, Kahn put together the successful proposal for BBN to build the ARPANET. Kahn is also co-author, along with Vint Cerf, of the TCP/IP protocols that underpin the present-day Internet. The main idea of the protocol is that networks should be connected using gateways that can translate the packets moving between different networks.

ARPA's request for quotation for building the IMP network arrived at BBN in the summer of 1968. Bob Kahn, a professor of electrical engineering was at BBN on leave from MIT (B.10.12). Kahn was an applied mathematician who had been working on communications and information theory and wanted some real-world engineering experience. With his work on the ARPANET, Kahn certainly achieved his goal! Kahn had already sent some of his papers to Bob Taylor at ARPA and the request for proposals landed first on his desk. The original hospital project for which Heart been recruited had not materialized, so he and Kahn were tasked with putting together a team and a proposal to bid for the contract. Heart was fortunate to recruit a strong team of experienced engineers, many of whom had worked with him at Lincoln Lab – Will Crowther, to lead the software team; Severo Ornstein, to lead the hardware effort; and Dave Walden, who had four or five years' experience in programming real-time systems. Other key members of the team were Bernie Cosell, whom BBN called "an ace de-bugger whom every BBN manager had learned to rely on if their projects got into trouble,"[20] and Ben Barker, an engineer from Harvard who played a vital role in getting the actual hardware for the first IMPs debugged and working (B.10.13).

To simplify the design, Heart had insisted on "a clean boundary between the host responsibilities and the network responsibilities."[21] In their book *Where Wizards Stay Up Late*, Katie Hafner and Mathew Lyon summarize the role of the IMP as follows:

B.10.13. A group photo of the IMP team at BBN in 1969. From left to right, Truett Thatch, Bill Bartell (Honeywell), Dave Walden, Jim Geisman, Bob Kahn, Frank Heart, Ben Barker, Marty Thrope (next to Heart), Willy Crowther, and Severo Ornstein. Team member Bernie Cosell is not in the photograph.

> Between Roberts and BBN it was settled: The IMP would be built as a messenger, a sophisticated store-and-forward device, nothing more. Its job would be to carry bits, packets, and messages: To disassemble messages, store packets, check for errors, route the packets, and send acknowledgements for packets arriving error-free; and then reassemble incoming packets into messages and send them up to the host machines – all in a common language.[22]

From his personal experience, Heart knew the value of building into a system as much reliability as possible. He therefore had the team investigate "error-control mechanisms" to cope with random mistakes in data transmission. Due to electrical noise on the lines, a "1" in a message packet could sometimes change to a "0," or vice versa. Fortunately, electrical engineers have developed a number of cunning techniques for not only detecting such errors but also, at some additional cost, correcting them. The basic idea is that of a *checksum*, a small number used to detect whether errors have occurred. The checksum is calculated from the bits of the packet at the source and transmitted along with the packet. On arrival, the checksum is recalculated. If the original and recalculated checksums do not agree, a transmission error due to noise has been detected. There are many types of checksums and error-correction techniques now available. The simplest is a *parity check*, a method of detecting errors by counting the evenness or oddness of the number of bits in the packet. This simple parity check requires only a single extra bit of information to be transmitted but can detect only an odd number of errors. More sophisticated methods such as *Hamming codes*, named for the mathematician Richard Hamming, require additional bits. These techniques can not only detect the precise location of the error but also correct it. For the ARPANET, the BBN team decided on a very simple, pragmatic solution: if an IMP detected an error in a packet it would just discard the packet and not send back any acknowledgment of receipt. The source IMP would wait for an acknowledgment and if it had not received one after a certain time, it would resend the packet.

One of Kahn's responsibilities was to specify exactly how the host computers would interact with the IMP machine. His BBN Report 1822 issued in the spring of 1969 told the host sites how to write a piece of software called a *device driver* to implement the host-IMP interface. At UCLA, Kleinrock had put graduate student Steve Crocker in charge of the programming effort to connect the university's Sigma-7 mainframe host computer to the IMP, along with fellow graduate students Vint Cerf, Jon Postel, and Charley Kline. Cerf later recalled: "It was a little funny because we were just graduate students. We kept expecting that professional managers would show up and tell us what to do. But they never did, so we just went on our merry way."[23]

Kahn's report also made clear that the IMP would not contain software for performing host-to-host communication. The responsibility for the host-to-host communications was effectively taken on by an informal collection of graduate students from the first four host sites. They called themselves the Network Working Group and initiated a series of notes called "requests for comments," or RFCs. Because network users needed to sign a collective agreement to enable each system to work with other systems, the group introduced the word *protocol* into the language of networking. In ancient Greece,

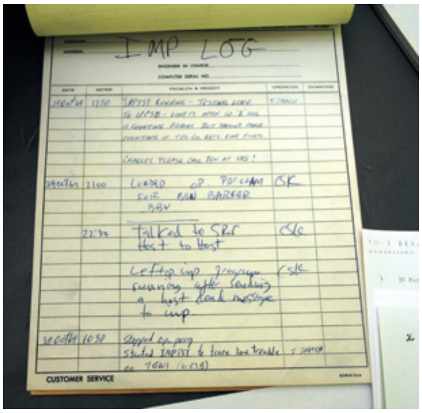

IMP LOG

CUSTOMER SERVICE

29 Oct 69 2100 LOADED OP. PROGRAM FOR BEN BARKER BBN CSK

22:30 Talked to SRI Host to Host CSK

Left op. program running after sending a host dead message to imp. CSK

Fig. 10.14. The first message sent over the ARPANET on 29 October 1969. This was an attempt to log into a remote computer at the SRI from UCLA.

a *protokollon* was the first page of a manuscript that contained a brief summary of the contents, the date, and who authenticated the document. In networking, the meaning of *protocol* resembled the role of the packet header, which specified the destination and information needed to reconstruct the message. A more informal definition was suggested by Cerf: "The other definition of protocol is that it's a handwritten agreement between parties, typically worked out on the back of a lunch bag, which describes pretty accurately how most of the protocol designs were done."[24]

Although the team did not complete a draft of the host-to-host protocol until the summer of 1970, the group had early on adopted a "layered" approach and had also written a couple of key applications. The lowest layer of the protocol specified how to move packets from host to host as a stream of unidentified bits, regardless of what sort of data the bits might represent. The two key applications were one for transferring files and another for *remote login*, software to enable a user at one computer to *log in* – that is, to identify himself or herself at a remote host computer and begin using the site. The final version of the protocol for exchanging files over the network, the File Transfer Protocol, or FTP as it is now usually abbreviated, was not completed until July 1972 when Jon Postel issued the Networking Group's RFC #354. The remote log-in application was called Telnet and was also very widely used. This application, together with the extension of the IMP interface to allow terminals as well as host computers to send data into the network – an interface called the Terminal IMP controller or TIP – later set the scene for the rapid expansion of the network.

The first IMP node was delivered to Len Kleinrock's team at UCLA on schedule in September 1969. The software written by Steve Crocker's team and the hardware interface designed and built by another UCLA graduate student, Mike Wingfield, worked perfectly for the host-to-IMP connection. But it was not until the second IMP was delivered to Engelbart's team at SRI in October that BBN was able to test intercomputer communication between two different host computers, the Sigma-7 at UCLA and the Scientific Data Systems 940 at SRI (Fig. 10.14). The system worked perfectly and the network was soon extended to the sites at Santa Barbara and Utah (see Fig. 10.12 in the preceding text). From the diagram, it is clear that the only link to Utah was through SRI, meaning that this first prototype distributed network was not yet what Baran had called "a robust web of redundant interconnections." By 1973, the ARPANET had expanded to become a much more redundant network (Fig. 10.15).

One of the concerns that Bob Kahn had about BBN's design of the IMP network software was about the flow control of the packets across the network. He said, "I could see things that to me were obvious flaws. The most obvious one was that the network could deadlock."[25] Deadlock is a situation in which the system gets into a state where no action is possible. The scenario that most worried Kahn was one caused by congestion at a destination IMP. If the storage *buffers*, the areas of computer memory used to temporarily store the information while it was being moved from one place to another, became too full at the receiving node, the packets containing the instructions to reassemble the messages would not be received. The destination IMP would be full up with packets unable to be assembled into complete messages. The pragmatic engineer

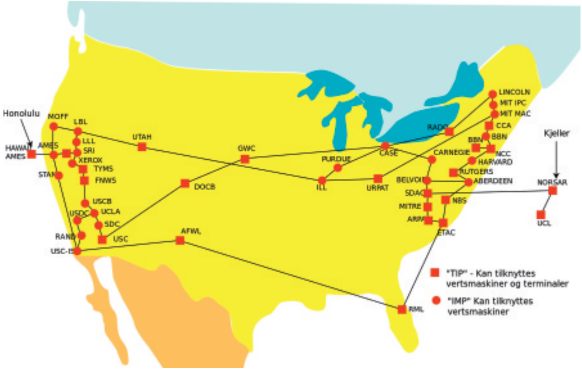

Fig. 10.15. The ARPANET topology in September 1973 including links to London and Hawaii.

Crowther and the theoretician Kahn had debated the problem, and pragmatism had won out. Now that the first four nodes were up and working, Heart gave permission for Kahn to fly out to UCLA with Dave Walden to test his theories. Kahn's first experiment demonstrated that the problem was real within only a few minutes of sending out specific patterns of packets. He later said, "I think we did it in the first twelve packets. The whole thing came to a grinding halt."[26] As a result of Kahn's results, BBN had to redesign the control system so that enough space was always reserved in the IMP memory buffers for reassembly of incoming packets. Under the redesign, the sending IMP would check that there was sufficient space and, if necessary, delay sending the next message.

Email: The ARPANET's killer app

Bob Taylor had proposed the ARPANET to provide interactive access to ARPA-funded computers across the United States and save money for ARPA by sharing resources. Larry Roberts believes that both of Taylor's goals were achieved:

> By 1973, I had cut our computer budget to 30 percent of what it would have been if I hadn't had the network. And saved more money than the network cost. Because I could share computers all across the world and not have to buy computers for every research group that wanted one.[27]

Although it is clear that the ARPANET did enable some sharing of resources, such cooperation was not its primary use. As Licklider and Taylor had pointed out in their paper on "The Computer as a Communication Device," one of the primary uses of a network would be for communication. A 1973 ARPA report showed that three-quarters of all traffic on the network was for what it called "E-mail," or more commonly nowadays called *email*. Time-sharing computers had had email systems linking users of the same computer for some time. BBN engineer Ray Tomlinson realized that he could extend this idea to send emails between different computers using the ARPANET (B.10.14):

B.10.14. Ray Tomlinson was an engineer at BBN who had developed some early email systems for time-sharing computers. When he thought about extending this idea to messages across the ARPANET, he needed a way to separate the name of the recipient in the email address from the machine the recipient was using. He chose the @ sign.

> Once we had the ability to transfer a file from one machine to the other, it became fairly clear that one thing you could do was just write the file across the network and send mail to somebody else. I also happened to be working

> on a piece of software to be used to compose and send mail, called "send message." And it seemed like an interesting hack to tie those two together to use the file-transfer program to send the mail to the other machine. So that's what I did. I spent not a whole lot of time, maybe two or three weeks, putting that together and it worked.[28]

Note Tomlinson's uncritical use of the word *hack:* in these early days it was a term of respect for some technically clever programming exploit. To be called a "hacker" was a compliment and did not have the unfavorable connotations the word now has. Once email between different sites became possible, Kahn said, "It had tremendous benefits: overcoming the obstacles of time zones, messaging multiple recipients, transferring materials with messages, simple collegial and friendly contacts."[29] Use of email grew rapidly and completely transformed the nature of collaboration. In an early experiment, an email sent to 130 people all around the United States at 5 P.M. generated seven responses within ninety minutes and twenty-eight responses in twenty-four hours. Such a response time now seems very slow, but in the 1970s it was revolutionary.

Tomlinson also got to choose a symbol for designating email addresses: his choice has become an icon for the networked world. Tomlinson needed a symbol to separate the name of the user from the machine that the user was using. He said: "The one that was most obvious was the '@' sign, because this person was @ this other computer, or, in some sense, he was @ it. He was in the same room with it anyway. And so it seemed fairly obvious and I just chose it."[30]

Although email usage took off like wildfire, there was much debate about the need to establish a separate email transmission protocol that was independent of FTP. In 1975, the first electronic discussion group called MsgGroup was established. The group had many heated online debates about email headers and requests for comments, and the practice of email *flaming*, sending an angry, critical, or abusive email, became an occasional feature of this and other discussion groups.


B.10.15. Bob Metcalfe and David Boggs invented and built a LAN technology that they called Ethernet to link up the PARC Altos and Gary Starkweather's laser printer. Metcalfe later left PARC to found 3Com, a computer networking company, with the three Cs standing for computers, communication, and compatibility.

From Hawaii to the Ethernet

Taylor's original vision for personal computing at Xerox Palo Alto Research Center (PARC) had always included networking the Alto machines in a local area network, or LAN. The networking technology needed to be cheap – Taylor's goal was it should cost no more than 5 percent of the cost of the computers they were connecting. The technology also needed to be easily expandable, capable of linking hundreds of Altos. Researcher Bob Metcalfe (B.10.15) arrived at PARC in the summer of 1972 still smarting from the indignity of having his doctoral thesis turned down by Harvard as being "insufficiently theoretical."[31] While at Harvard, he had spent much of his time at MIT working on the ARPANET, and he had written up this very practical work for his thesis. At PARC, Metcalfe found several experimental networking projects in progress, but he believed that none of them would satisfy Taylor's requirements. There was also a deadline looming: in his design for the Alto hardware, Chuck Thacker had left space for a yet-to-be-designed *network controller card*, a device to connect the Alto to a computer network, and PARC was nearly ready to roll out the Alto. Metcalfe remembered reading a paper about an ARPA-funded networking project linking computers on the different islands of Hawaii. This Hawaiian network was called ALOHAnet and was designed by University of Hawaii professor Norman Abramson. Instead of sending electrical signals down phone lines, ALOHAnet used a radio network that sent digital signals through the atmosphere. Abramson had come up with a simple way to manage interference caused by two stations trying to transmit a message to the same receiving station simultaneously. If the transmitting stations did not get an acknowledgment of a successful receipt of the message from the receiving station, they were programmed to resend their messages after waiting different, random times. This delay ensured that their messages would not collide a second time. Metcalfe realized that this feature could be very useful for a LAN, in which there could be many computers trying to send messages at the same time. He also proposed joining the Altos together with a physical cable. Metcalfe thought of the Altos as sending their digital messages onto a wire that would merely act as a passive channel, much as the atmosphere served as an inactive medium through which radio signals traveled for ALOHAnet. He likened this use of the connecting wire to the way that physicists before Albert Einstein had assumed the presence of a *luminiferous* (light-bearing) *ether*, a substance through which light signals were supposed to travel. Metcalfe's first memo on his idea in May 1973 was titled "The ETHER network" and described his ideas for linking computers in a LAN.

In collaboration with Stanford graduate student David Boggs, who was then working part-time at PARC, Metcalfe tested out his ideas using a coaxial cable to connect the computers. A *coaxial cable* is a cable with an inner copper conductor surrounded by an insulator and a copper conducting sheath, all encased in plastic. The cable is called *coaxial* because the inner conductor and the outer conducting sheath have the same *axis* (center). The cable acts as a transmission line for radio-frequency signals. The cable connecting the computers was usually silent and so resembled the inert "ether." When a machine wanted to transmit a message, it sent a "wake up" bit onto the cable to alert the other machines. It then sent a packet with a destination address, a sending address, the message, and some checksum bits for error checking. If another machine sent a packet at the same time, resulting in a collision, both machines stopped sending and each waited a different random time before resending the message. Metcalfe and Boggs developed the electronics to make this happen and produced the first *Ethernet card* for the Alto, a device to attach the computer to the Ethernet. Adding new machines proved to be very easy and just required the end of a branch cable to be plugged into the main coaxial cable.

Metcalfe resubmitted his doctoral thesis to Harvard, still about packet-switching but now including a suitably theoretical analysis of ALOHAnet. The university accepted his thesis in June 1973. Boggs took a leave of absence from Stanford and joined PARC full-time. It would be another nine years before he received his Stanford doctorate. The patent for Ethernet was filed in March 1975 under the names of Metcalfe, Boggs, Thacker, and Butler Lampson. Metcalfe left Xerox PARC in 1979 to set up a new company called 3Com to produce Ethernet networking equipment. With the support of the Digital Equipment, Intel, and Xerox corporations, the DIX standard was published in 1980. It carried Ethernet

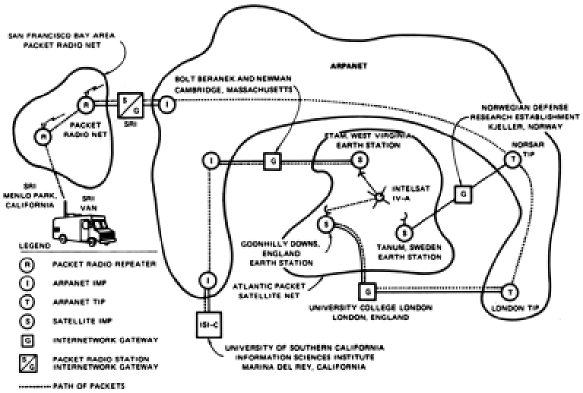

Fig. 10.16. The original conceptual plan suggested by Cerf and Kahn for a multinetwork system. This idea led to the birth of the "network of networks," better known as the Internet.

traffic at a rate of ten megabits per second. Adoption of Ethernet technology rapidly outpaced several competing technologies, and Ethernet became the dominant technology for building LANs.

The TCP/IP protocol and the Internet

Kahn and Cerf had first worked together in 1970 when Kahn was performing his network congestion experiments (B.10.16). By 1973, there were a number of different international packet-switching networks, and members of the network community had discussed what was needed to link them together. Typically the networks would have different interfaces, packet sizes, and transmission rates. Kahn and Cerf came up with the idea of a "gateway" computer. To connect the ARPANET to the gateway machine, the machine must just look like a host to the ARPANET IMPs. But to connect to another network, such as the ALOHAnet radio network, the gateway software must also allow the machine to look like a host on that network. Each network could still send messages to its own machines using its own protocols, but when data needed to be sent to another network, a new and universal networking protocol was needed – a set of rules that any computer connected to the network could follow to send or receive data.

In May 1974, Cerf and Kahn wrote a paper describing such a protocol (Fig. 10.16). They proposed that messages should be encapsulated in *datagrams*, self-contained packets of information with the source and destination addresses written into the header, in much the same way that letters are put into envelopes and sent to their destination. The gateways would read the addresses on the "envelopes," but only the receiving host would read the contents. This set of rules, called the Transmission Control Protocol (TCP), also made the assumption that the packet-switching network was intrinsically unreliable. In the construction of the ARPANET, it had been the IMPs that provided the reliability: now the focus of reliability shifted from the network to the communicating host computers. Cerf and Kahn led the discussion with the ARPANET participants and with the growing international networking community. This community included Donald Davies in the United Kingdom and Louis Pouzin in France, who was building a packet-switching network

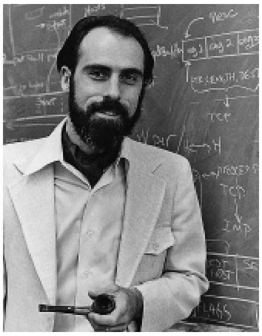

B.10.16. Vint Cerf studied mathematics at Stanford and earned his PhD in computer science from UCLA in 1972. After receiving his doctorate, he became an assistant professor at Stanford and worked with Bob Kahn on the TCP/IP protocol for network interoperability, the basis for the modern Internet.

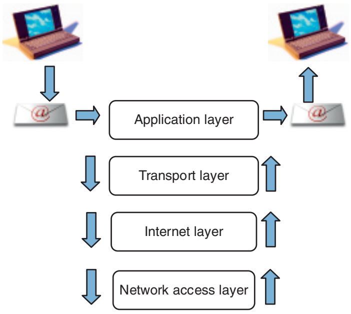

Fig. 10.17. The layered structure of the TCP/IP protocol.

called Cyclades. By 1975, the TCP specification had advanced sufficiently to be implemented concurrently at three sites – BBN, Cerf's group at Stanford, and a group headed by Peter Kirstein at University College London. In October 1977, the effort reached a significant milestone when Cerf and Kahn demonstrated sending messages across three interworking networks: a packet radio network, a packet satellite network called SATNET, and the ARPANET.

The networking protocol took its final form in early 1978 when the part of the transmission protocol that deals with the routing of the packets was separated off as a discrete Internet Protocol (IP). Under the new TCP/IP protocol, TCP was now responsible for breaking up messages into datagrams, reassembling them at the other end, detecting errors, and resending lost messages. The IP was responsible for routing the individual datagrams. Jon Postel (B.10.17) later summarized the guiding principle for what should be in the IP: "I remember having a general guideline about what went into IP versus what was in TCP. The rule was 'Do the gateways need this information in order to move the packet?' If not, then that information does not go into IP."[32]

Fig. 10.18. Matryoshkas are Russian nesting dolls. Here they illustrate the concept of message encapsulation.

In the first chapter we mentioned the layered approach as one of the frequently occurring "mental tools" used in computer science. The TCP/IP protocol is a further illustration of this idea (see Fig. 10.17). TCP/IP consists of four layers. At the top there is the application that sends a message to another application running on a remote computer. At the bottom lies the actual physical wire or fiber where the message is translated into electric or light impulses. If the message is long, then it is sliced up into smaller pieces, each put into a separate envelope. Imagine if we sent an entire book to somebody one page at a time. On the Internet, as the message travels down it is put into a bigger envelope at each layer. We can think of the protocol as resembling *matryoshkas*, or Russian nesting dolls (Fig. 10.18). On the receiving end, the message travels upward, and at each layer the envelope is stripped off. Because the original message was sliced into pieces, the message must also be reassembled on the receiving end.

B.10.17. Jon Postel (1943–98), "the unsung hero of networking," was the coordinator of the RFC discussion forum and later became chairman of the Internet Assigned Numbers Authority, in charge of assigning Internet numbers. Cerf wrote Postel's obituary, which was published as RFC 2468.

TCP/IP emerged as the dominant Internet standard after a long battle with a rival proposal for standardizing networking. The rival standard was called *Open Systems Interconnection* (OSI). It had been developed by the International Organization for Standardization (ISO), a world body that tries to establish uniform sizes and other specifications to ease the international exchange of goods. The U.S. government and European national governments all decreed that OSI was the official networking standard for the Internet. Similarly, major computer manufacturers like IBM, Digital, and Hewlett-Packard also adopted the OSI standard instead of TCP/IP. Amazingly, it was the popularity of the Unix operating system and Ethernet LANs in universities that turned the tide in favor of TCP/IP. Unix, developed at Bell Labs in the 1960s, is a popular operating system because it is powerful and stable, and can be installed on many different kinds of machines. Except for Microsoft Windows, nearly all major operating systems have some kind of Unix at their core. Bill Joy, later one of the founders of Sun Microsystems, received an ARPA grant to write the TCP/IP *stack*, a complete set of networking protocols, into the free Berkeley version of Unix. The first Sun machines were sold with Berkeley Unix including the TCP/IP networking software. With the worldwide popularity of Sun workstations and the increasing availability of Ethernet as a commercial product in the early

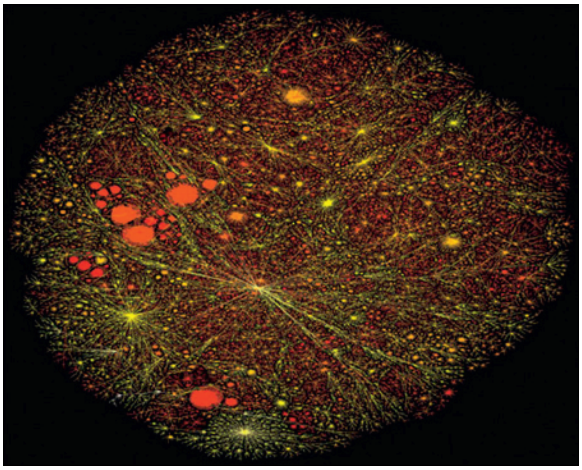

Fig. 10.19. An Internet map in 2005 generated by tracing the packets from source to destination.

1980s, TCP/IP-based computer networking grew rapidly in universities. In the end, TCP/IP won out, showing that an open, grassroots process can sometimes prevail over an imposed official standard.

The Internet soon became much more than a research experiment for academic researchers (B.10.18). By the time the ARPANET was retired in 1990, commercial *Internet service providers* (ISPs) had begun to emerge, allowing business and the general public to connect to the Internet, usually for a monthly fee. For many years, the U.S. government, which funded ARPANET and its successor NSFNET, had tried to limit the Internet to research and educational uses. In 1995, however, the government removed the last restrictions on commercial use of the Internet. Traffic on the public Internet grew at more than 100 percent per year through the late 1990s. As of March 2011, the number of Internet users exceeded two billion, some 30 percent of the world's population (Fig. 10.19).

From copper to glass

The chapter began with an account of Chappe's early optical communication system, and it now ends with a discussion of how lasers and optical fibers

B.10.18. A group photo of Internet pioneers taken in 1994 to mark the twenty-fifth anniversary of the ARPANET. The photo was taken at a Christian Science Church in Boston in front of their world map. Left to right, front row: Bob Taylor, Vint Cerf, and Frank Heart; second row: Larry Roberts, Len Kleinrock, and Bob Kahn; third row: Wes Clark, Doug Engelbart, and Barry Wessler; fourth row: Dave Walden, Severo Ornstein, Truett Thach, Roger Scantlebury, and Charlie Herzfeld; fifth row: Ben Barker, Jon Postel, and Steve Crocker; last row: Bill Naylor and Roland Bryan.

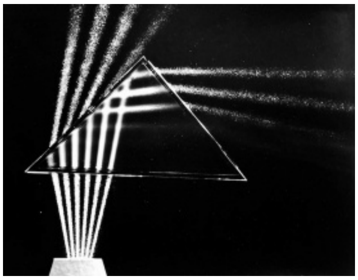

Fig. 10.20. Several rays of light striking a prism at a variety of angles. As can be seen, beyond a certain "critical" angle the light rays are entirely reflected and no light is transmitted through the prism. The ray on the extreme right is entirely reflected while the other rays show both transmitted and reflected beams.

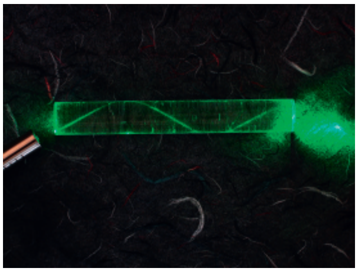

Fig. 10.21. Figure illustrating total internal reflection in an optical fiber.

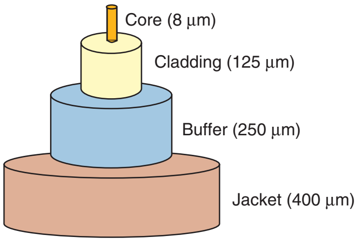

Fig. 10.22. Diagram of a modern fiber-optic cable showing the relative sizes of cladding and core.

have replaced copper wires for high-bandwidth communications and made possible *the broadband* Internet we see today. *Broadband* has various definitions, but the term now generally refers to high-speed data connections to the Internet. Optical fibers are thin "wires" of glass that transmit light using a phenomenon called *total internal reflection*. We can see how such reflection comes about by considering what happens when light travels from air into a block of glass. Because light travels more slowly in glass than air, the light changes direction, bending toward the vertical. This bending of light at a surface is called *refraction*. Now consider light traveling from glass to air, when the light bends away from the vertical. If we increase the angle at which we shine light on the glass-air surface, the transmitted ray emerges at an angle closer and closer to the surface. At an angle called the *critical angle*, the light just grazes the surface. If we increase the angle of incidence beyond this critical angle, all the light will now be reflected and no light will escape into the air (Fig. 10.20). This is the phenomenon of *total internal reflection* (see Fig. 10.21). It is this mechanism that allows light to be transmitted down a glass fiber and to travel around bends in the fiber. *Cladding* or covering (see Fig. 10.22) the core of the fiber with glass of a lower refractive index provides a way to bend or reflect inward light rays that strike its interior surface. This makes it possible for the light to travel long distances with little loss of intensity. Durable optical fibers in medical imaging devices like endoscopes carry light that enable doctors to examine the inside of the stomach and other organs (Fig. 10.23).

Optical fibers were fine for use in applications like endoscopes where the light only needs to be transmitted a few meters. But for transmission of light over the long distances required for telecommunications, scientists believed that the light in the fiber would lose too much intensity for optical fibers to be practical. Engineers use units called *decibels* to measure the energy loss of signals. In 1960, glass fibers had an *attenuation* (reduction in strength) of about one decibel per meter – meaning that about 20 percent of the light entering the fiber was lost in just traveling the width of a table. After traveling a hundred meters through a fiber, only one ten-billionth of the light would remain (see Table 10.1). A loss of ten decibels per kilometer would mean that a tenth of the power remained after a kilometer: a loss of one thousand decibels per kilometer – one decibel per meter – meant that almost no light remained. This gives an idea of the scale of the challenge. It was for this reason that Rudolf Kompfner, head of transmission research at Bell Labs, had dismissed optical fibers as a practical transmission technology in 1961. For optical fibers to become a viable communications technology, the attenuation in the fiber needed to be reduced to about ten decibels per kilometer.

Because of the attenuation problem, Bell Labs and almost all of the global telecommunications industry thought that the future for high-bandwidth communications would be based on *millimeter waveguides*, hollow pipes that could channel millimeter wavelength transmissions. But a few groups in the United Kingdom persisted in their research on optical fibers. At a meeting of the British Association for the Advancement of Science held in Southampton, England, in 1964, Alec Gambling (B.10.19), a professor in the Department of Electronics at the University of Southampton, suggested that glass fibers should be investigated "not because they looked at all promising, but Sherlock Holmes like, they seemed to be the least unlikely possibility to pursue."[33] The real stimulus to optical fiber communication was a classic paper written by Charles Kao and George Hockham in 1966, working at ITT Corporation's Standard Telecommunication Laboratories (STL) in the United Kingdom (B.10.20). They had made a very careful study of the losses in various types of glasses and concluded that the high attenuation rates were due to impurities in the glass. The 1 April 1966, issue of *Laser Focus World* magazine noted:

> At the IEE [Institute of Electrical Engineers] meeting in London last month, Dr. C. K. Kao observed that short-distance runs have shown that the experimental optical waveguide developed by Standard Telecommunications Laboratories has an information-carrying capacity … equivalent to about 200 TV channels or more than 200,000 telephone channels. He described STL's device as consisting of a glass core about three or four microns in diameter, clad with a coaxial layer of another glass having a refractive index about one percent smaller than that of the core. … According to Dr. Kao, the fiber is relatively strong and can be easily supported. Also, the guidance surface is protected from external influences … [and] the waveguide has a mechanical bending radius low enough to make the fiber almost completely flexible. Despite the fact that the best readily available low-loss material

Table 10.1 Decibel table (the values of decibel loss are negative).

Decibel loss	Output/Input signal ratio
1	0.79
2	0.63
10	0.1
20	0.01
30	0.001
40	0.0001
50	0.00001
60	0.000001
70	0.0000001
80	0.00000001
90	0.000000001
100	0.0000000001

Fig. 10.23. Light conducted by a bundle of optical fibers. The core of the fiber is only a few microns wide. A typical human hair is about fifty microns wide, so optical fibers are about a tenth of the diameter of hair.

B.10.19. Alec Gambling joined the Electronics Department at the University of Southampton in 1957. The department at Southampton had been founded in 1947 by Eric Zepler, a German refugee who had previously been head of radio receiver design at Telefunken. Gambling started his research in lasers as a potential source of high-frequency carrier waves for communications. After the classic paper on optical fibers by Charles Kai and George Hockham, Gambling received a research contract from the United Kingdom's Ministry of Defence and started research on manufacturing low-loss fibers. Gambling's group invented the method of Chemical Vapor Decomposition, which is still the most widely used method of optical fiber fabrication worldwide.

> has a loss of about 1000 dB/km [decibels per kilometer], STL believes that materials having losses of only tens of decibels per kilometer will eventually be developed.[34]

At this time, the British Post Office operated the British telephone network. After Kao's and Hockham's paper, Frank Roberts at the Post Office Research Station at Dollis Hill in the United Kingdom started a research program to reduce fiber losses. The third group researching fibers in the United Kingdom was Gambling's team at the University of Southampton. One of Gambling's research students, David Payne, built a drawing tower to make optical fibers. A *drawing tower* is the apparatus where a block of molten glass called a *preform* is *pulled* into a long thin fiber with the desired width and thickness. Gambling and Payne were able to reduce losses from thousands of decibels per kilometer to about 140. At about the same time, Charles Kao at STL made careful measurements in different types of glasses and concluded that the high purity of commercially available silica made it a good candidate material for optical communication fibers. After Bell Labs heard about these results from the United Kingdom, work on optical fibers restarted in earnest. It took only four years for the global research community to reach Kao's goal of creating an optical fiber with less than twenty decibels of light loss per kilometer.

Most groups were trying to purify the typical compound glasses used for standard optics, which were easy to melt and draw into fibers. At the Corning Glass Works in Corning, New York, Bob Maurer, Donald Keck, and Peter Schultz started with fused silica, a material that can be made extremely pure but has a high melting point and a low refractive index (B.10.21). By carefully adding controlled amounts of impurities called *dopants* they were able to raise the refractive index of the silica core to be slightly higher than the cladding without increasing the attenuation significantly. In September 1970, the Corning team announced the manufacture of a fiber with attenuation below twenty decibels per kilometer but gave no details of the manufacturing process. By 1972, the Corning team had managed to reduce losses to only four decibels per kilometer. Because the Corning group provided little information about how they had achieved their results, it was difficult for other groups to verify these claims. At a 1974 conference in Brighton, England, Gambling and Payne from Southampton announced their discovery of a new method of fiber fabrication

B.10.20. Charles Kao was awarded the Nobel Prize in physics in 2009 for "groundbreaking achievements concerning the transmission of light in fibers for optical communication."[B1] He was born in Shanghai in 1943 during World War II and, with his family he escaped by boat from mainland China to Hong Kong in 1948. He went to England in 1952 to study electrical engineering at what is now the University of Greenwich, London. After graduating, Kao stayed in the United Kingdom to work for Standard Telephones and Cables. In 1960 he was offered the opportunity to join STL research center and at the same time study for a PhD at University College London. At STL, Kao was asked to explore optical communication technologies and became a passionate advocate for optical fibers. In those days, glass fibers lost far too much energy in the light transmission through the glass. Kao showed that these losses were a result of impurities in the glass and speculated that fibers could be made with loss levels low enough to transmit signals long distances.

B.10.21. Optical fiber pioneers from Corning Glass Works: from left to right, Donald Keck, Bob Maurer, and Peter Schultz. The three developed the first low-loss optical fiber in 1970 after hearing of Charles Kao's goal of a loss of only twenty decibels per kilometer, equivalent to 1 percent of the light remaining after transmission through a one-kilometer glass fiber.

using *chemical vapor deposition* (B.10.22). This is a chemical process that deposits a thin coating on the silica and allows the creation of preforms with a specific refractive index profile. Simultaneous publications from the Southampton group and Bell Labs then gave details of a feasible method of manufacturing doped optical fibers. Using this method, Gambling and Payne were able to make fibers with a loss of only 2.7 decibels per kilometer. Although British Telecom became the first phone company to commit to optical fiber, it was the decision of MCI Communications Corporation to use optical fiber to build a nationwide long-distance communication network in the United States that really opened the floodgates. In 1982, MCI ordered one hundred thousand kilometers of fiber from Corning and took on the battle with the dominant AT&T. By the mid-1980s, the U.S. government had loosened many restrictions that limited competition in the long-distance telephone industry in the United States, creating a very competitive market for fiber networks.

The last piece of the puzzle arrived in 1987 when Payne and a team at the University of Southampton discovered a way to amplify the optical signal in the fiber. Up to then, when the optical signal grew too weak for further transmission, the signal had to be converted from optical to electrical, amplified, and converted back again to an optical signal. By adding a small amount of the element erbium to the core of the fiber, Payne and his colleagues at Southampton demonstrated that the signal could be amplified using a semiconductor laser as a *pump*, a device that boosts or amplifies a signal (Fig. 10.24). The first commercial product was produced in 1990 by Pirelli in the United Kingdom, in collaboration with the Southampton team. The first transatlantic fiber cable had a capacity of eight hundred telephone circuits, and the cost per circuit was $30,000; only ten years later, transatlantic cables with fiber amplifiers had a capacity of six hundred thousand circuits at a cost per circuit of only $500. In this way, fibers and *erbium-doped fiber amplifiers* (EDFAs) have revolutionized telecommunications and computer networks. Kao's original vision for optical fiber communications has come true: "If you really look at it, I was trying to sell a dream.... There was very little I could put in concrete to tell these people it was really real."[35]

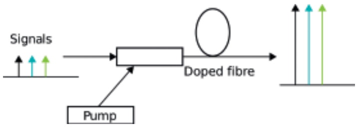

Fig. 10.24. Schematic representation of the working of the erbium-doped fiber amplifier (EDFA).

One postscript to this story is the curious fact that although the United States led the world in developing lasers, AT&T and Bell Labs virtually ignored the idea of fiber-optic communications. Equally surprising is how a small university research group at the University of Southampton with very limited funding not only competed with all the major telecommunication research laboratories but also made a string of fundamental discoveries. Payne believes

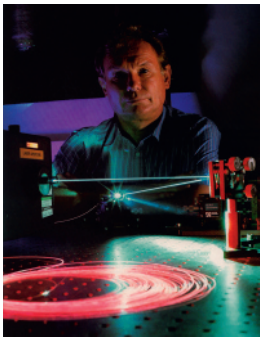

B.10.22. David Payne is director of the Optoelectronics Research Centre at the University of Southampton in the United Kingdom. He was a student of Alec Gambling's and built one of the world's first optical fiber drawing towers during his research for his PhD thesis. Payne later led the team that discovered the erbium-doped fiber amplifier in 1987. With Emmanuel Desurvire of Bell Labs, Payne was awarded the 1998 Benjamin Franklin Medal "for fundamental technical contributions and leadership critical to the successful development of the Erbium-Doped Fiber Amplifier and for the championship of this project."[B2]

that part of the answer is the ability of university research groups to pursue blue-sky research, scientific investigations without any obvious applications in the real world:

> The optical amplifier was developed with no [research] contract. It did not have a proposal to do it, it did not have any milestones, it did not have any deliverables, we did it completely on the site. And I think that you'll find that most of the great developments in all areas of science are like that. So the present move towards focused and managed research might well be the wrong thing to be doing.[36]

Key concepts

- Store-and-forward networks
- Centralized, hierarchical, and distributed networks
- Circuit switching, message switching, and packet switching
- Networking protocols and the ARPANET
- File Transfer Protocol and Telnet
- Wide area networks and email
- Ethernet and local area networks
- The Internet and the TCP/IP protocol
- Optical fibers for communication
- Erbium-doped fiber amplifier

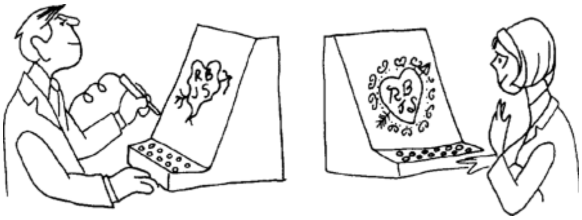

"A communication system should make a positive contribution to the discovery and arousal of interests."[F2]

Women and telegraphy

The telegraph system required a large labor force to operate the transmitting and receiving machines (Fig. 10.25). This need provided new "high tech" jobs for women, many of whom could operate telegraph machines with great dexterity. In February 1846, only two years after Samuel Morse's first successful demonstration of the electric telegraph, the Magnetic Telegraph Company opened an office in Lowell, Massachusetts, and hired a woman named Sarah Bagley as one of the first female telegraphers in the United States. Early in 1847, she was promoted to run the magnetic telegraph office in nearby Springfield, Massachusetts, but was understandably unhappy to learn that she earned only three-quarters as much as the man she replaced. This experience, together with her earlier experiences in the Lowell textile mills, led her to be an early advocate of women's rights.

The *Illustrated London News* described the scene in a telegraph office in 1874:

> It is a cheerful scene of orderly industry, and it is, of course, not the less pleasing because the majority of the persons here are young women, looking brisk and happy, not to say pretty, and certainly quite at home. Each has her own instrument on the desk before her. She is either just now actually busied in working off or in reading some message, or else, for the moment she awaits the signal, from a distant station, to announce a message for her reception. Boys move here and there about the galleries, with the forms of telegrams, which have been received in one part of the instrument-room, and which have to be signaled from another, but which first have to be conveyed, for record, to the nearest check-tables and sorting tables in the centre.[37]

Fig. 10.25. The Central Telegraph Office London in 1874 employed 1,200 telegraphists of whom 740 were female and 270 boy messengers. Each day around 18,000 messages were transmitted. This engraving was produced to accompany an article in the *Illustrated London News*.

ARPANET gets funded

The story of how Bob Taylor got the funding to build the ARPANET is now the stuff of legend. In his office, Taylor had multiple terminals connecting him to different computers funded by ARPA at the various research centers across the country. Each terminal required a different login procedure, and no computer could "talk" to another. Frustrated by these incompatibilities, Taylor decided to act on Licklider's idea. Without even writing a short memo about his plan, he went straight to the office of ARPA Director Charles Herzfeld. Fortunately, Herzfeld had seen the "multiple terminal problem" for himself and had also previously talked with both Licklider and Taylor about their ideas for interactive computing and networking. The gist of Taylor's argument to Herzfeld was that because more and more researchers were requesting funds from ARPA to have their own expensive computers, it would be more cost-effective for ARPA to network the computers at the different sites so that researchers could share both the hardware and access each other's results. Taylor suggested that ARPA should fund a small test network connecting these ARPA computers, initially with only four nodes, but then expanding to a dozen or more if it was successful. After just twenty minutes of discussion, Taylor left Herzfeld's office with an extra million dollars in his budget to build such a network. The original program plan for the ARPANET talks about accessing time-shared computers and does not mention network survivability in case of nuclear attack.