The Internet 2

The gigabit National Research and Education Network (NREN) is to develop from the current Internet, a ‘‘network of networks” that connects users in all parts of the United States and around the world. The Internet allows users to communicate using electronic mail, to retrieve data stored in databases, and to access distant computers. The network began as an Advanced Research Projects Agency research project to investigate computer networking technology, and in slightly over 20 years has grown into an essential infrastructure for research and education. The NREN initiative and associated research programs are intended to support the further evolution of research and education networking, broadening access to the network and enabling new applications through the deployment of advanced technologies.

Federal support to further the development of networks that support research and education communications is directed primarily at upgrading the Federal "backbone" networks that have formed the core of the Internet.[1] These networks include the National Science Foundation's NSFNET backbone, the NASA Science Internet (NSI) (figure 2-l), the Department of Energy's Energy Sciences Network (ESnet), and the Department of Defense's DARTnet and Terrestrial Wideband Network (TWBnet). The NASA and DOE networks are primarily intended for traffic related to the mission of the supporting agency, while the current NSFNET backbone serves users in a broader range of disciplines in universities, supercomputer centers, and industry research laboratories. The DOD networks support research and development of new communications technologies. The Federal

Federal agency networks will form the core of the gigabit NREN.

[1] Office of Science and Technology Policy (OSTP), "Grand *Challenges* 1993: High Performance Computing and Communications, ' p. 18.

gure 2- —Federal Agency Backbone Network

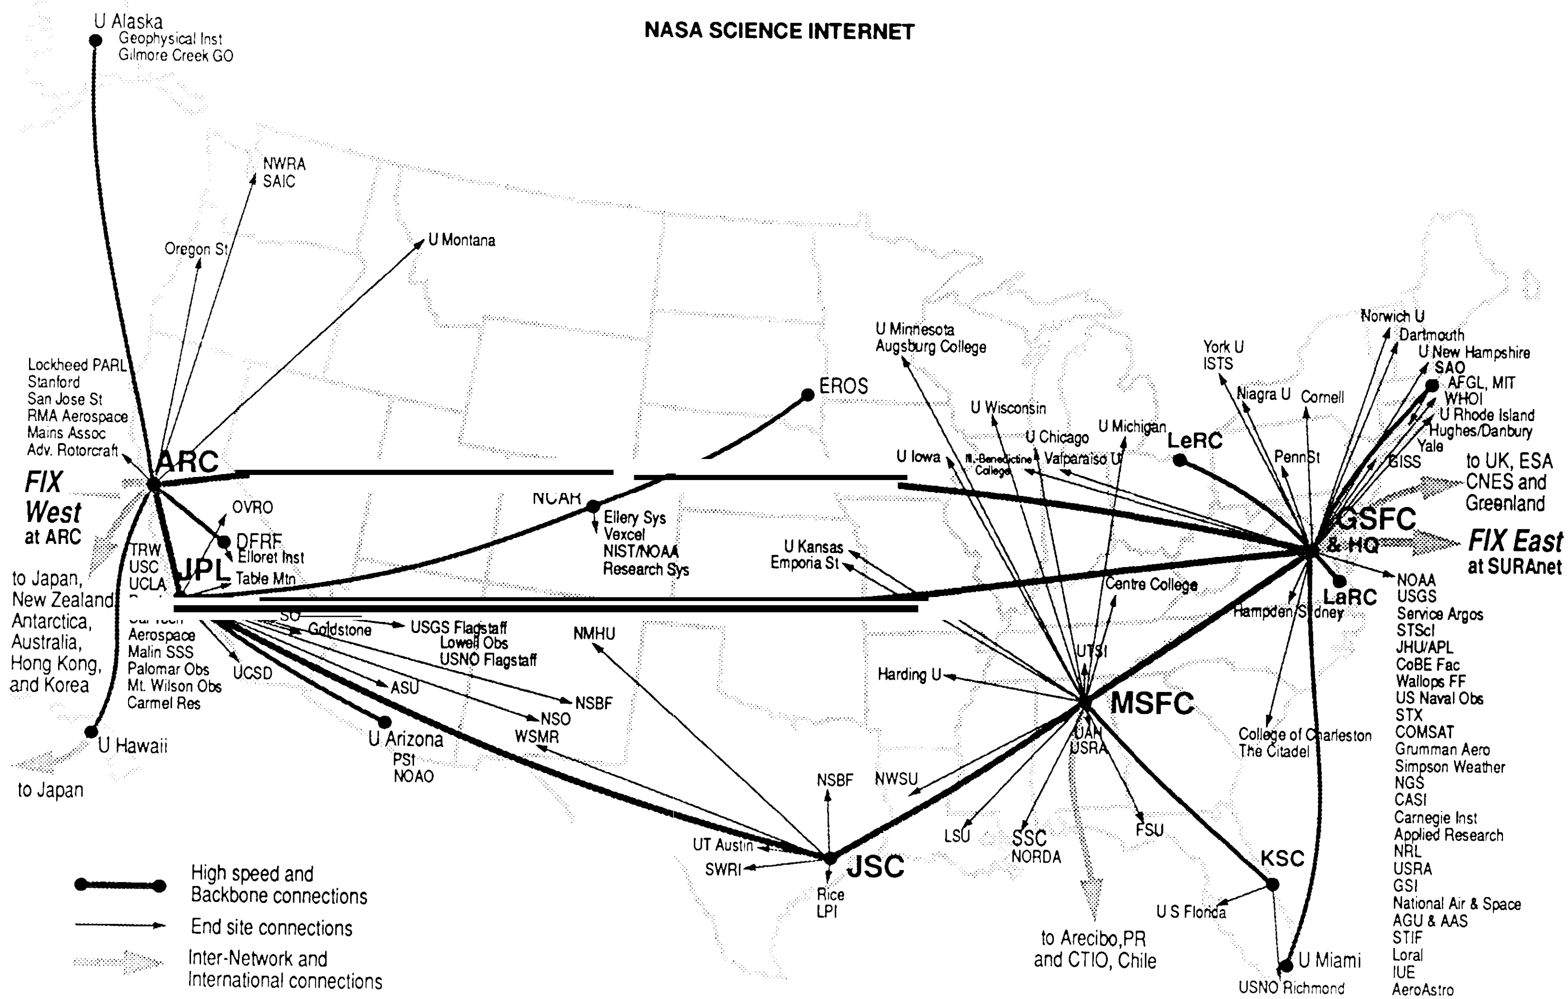

Important sites on the NASA Science Internet (NSI) include Ames Research Center (shown as "ARC"), Jet Propulsion Laboratory (JPL), Johnson Space Center (JSC), Marshall Space Flight Center (MSFC), and Goddard Space Flight Center (GSFC). NSI is connected to other Federal networks at the Federal Internet Exchanges (FIXes) East and West.

SOURCE: National Aeronautics and Space Administration (NASA).

Figure 2-2—Regional Network

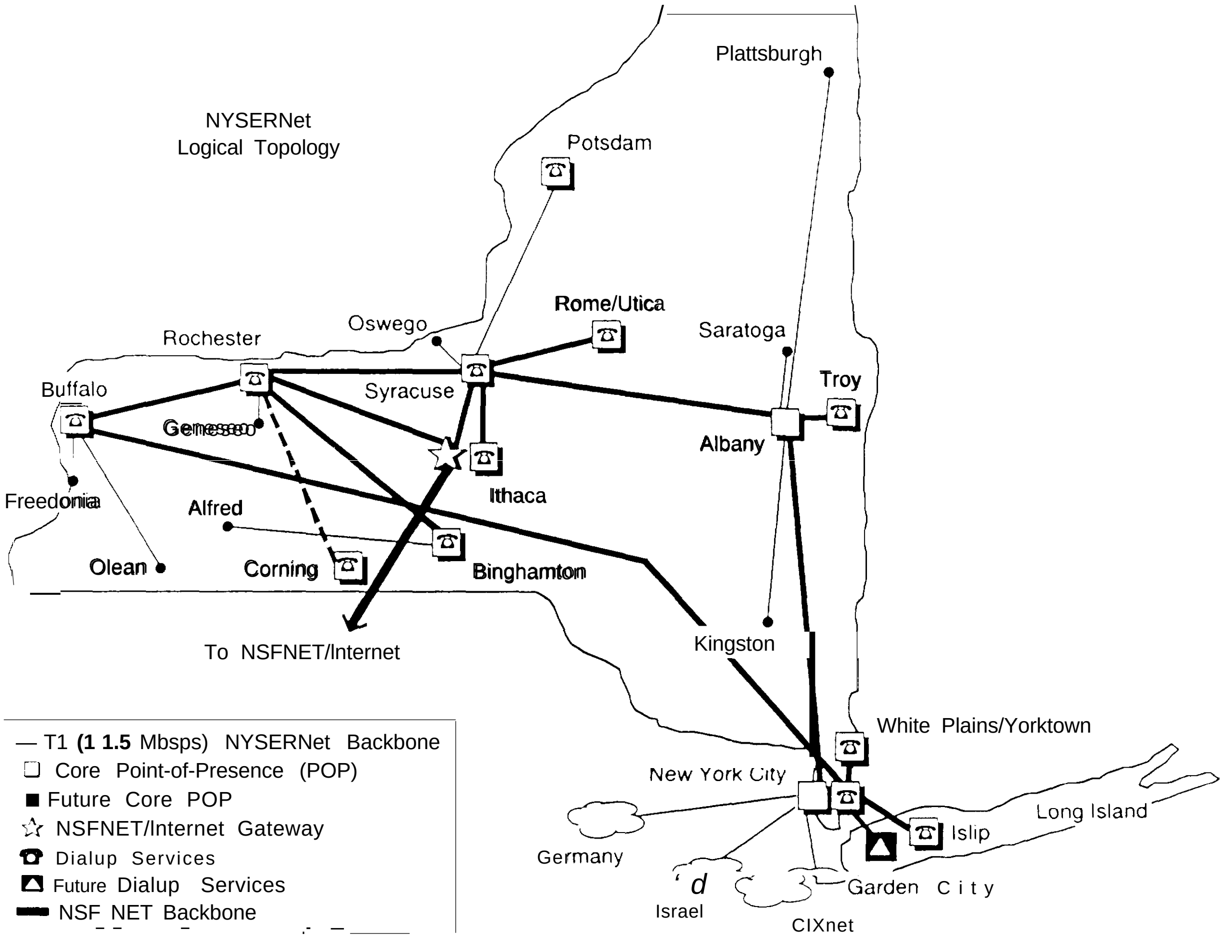

SOURCE: NYSERNet.

networks are interconnected at FIXes (Federal Internet Exchanges) at NASA's Ames Research Center in California and at the University of Maryland.

Upgrading the agency-supported backbones is not the only thing needed to improve research and education networking. The majority of users in universities, schools, and libraries do not have direct access to one of the backbone networks. These users rely on thousands of other networks that, together with the Federal agency backbones, form the Internet. These networks are interconnected, and information typically travels through several networks on its way from one user to another. In order to provide good performance end-to-end, all of the Internet's networks will need to evolve in a coordinated fashion, matched in capability and performance.

Most of the Internet's networks are "campus" or 'corporate networks, connecting users within a university or a company. Campus and corporate networks may in turn be interconnected by "regional" networks. For example, NYSERNet (New York State Education and Research Network) connects campuses and industrial customers in New York State (figure 2-2) and BARRNET (Bay Area Regional Research Network) does the same in northern California.

Figure 2-3—The NSFNET Hierarchical Structure

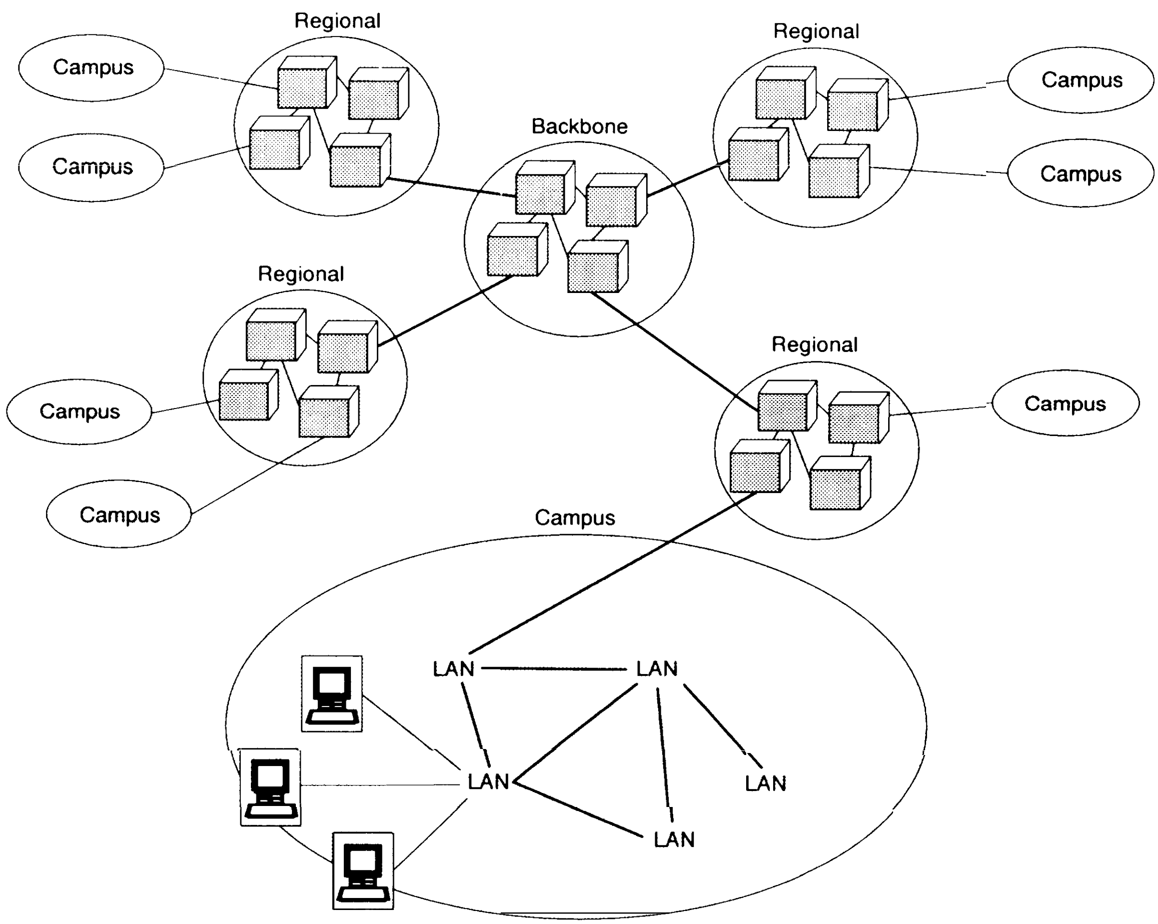

SOURCE: National Science Foundation (NSF).

Regional networks also provide a connection between campus networks and the national NSFNET backbone that carries traffic to other regions.[2] The regional networks, and the resulting three-tier structure of campus, regional, and backbone networks (figure 2-3), evolved with support from the National Science Foundation.[3]

The Internet also includes several networks that provide service on a for-profit basis.[4] The government investment in developing and demonstrating Internet technology during the 1970s and 1980s has created opportunities for the private sector to sell Internet services. The effectiveness of the Internet technology has been proven, and a

2 NASA and DOE sites are connected directly to the agency networks. However, NASA and DOE rely on the regional networks and the NSFNET backbone to connect to university researchers participating in NASA and DOE projects.

3 For a description of evolution of the regional networks and the three-tier structure, see Richard A. Mandelbaum and Paulette A. Mandelbaum, "The Strategic Future of the Mid-Level Networks," Brian Kahin (cd.), *Building Information Infrastructure* (New York, *NY:* McGraw Hill Primis, 1992).

4 Eric Arnum, "The Internet Dilemma: Freeway or Tollway," *Business Communications Review,* December 1992, vol. 22, No. 12, p. 31.

growing number of companies are now using the Internet to conduct business. Even though the NREN program continues government funding for the agency backbone networks, in order to upgrade them to gigabit speeds, government support is becoming less central to the Internet as a whole. New commercial providers of nationwide Internet services have emerged. In addition, NSF has been reducing subsidies to the regional networks, which are increasingly being asked to recover costs from users.

The availability of commercial services is leading to a change in the makeup of the users of the Internet. Until recently, corporate use of the Internet was restricted to scientists and engineers in research laboratories or engineering departments. In part, this was due to the history of the Internet as an experimental network. The limited use of the Internet by the private sector was also due to an "Acceptable Use" policy that reserved the federally supported backbones for research and education traffic.[5] The new commercial providers have no traffic restrictions, allowing the Internet to serve a wider range of users. Today's Internet users can have different security requirements[6] their technical sophistication varies, and the demands they place on the network's capacity differs.

One of the goals of the NREN program is to continue the trend towards provision of Internet services on a commercial basis, rather than solely as the result of a government subsidy.[7] The NREN program continues government support for networking, but the emergence of commercial providers is leading to changes in the mechanisms by which this support is provided. NSI and ESnet will continue to support agency missions, but the next-generation NSFNET backbone will be considerably different from the current NSFNET backbone. As part of its NREN plans, NSF has decided that much of the trafic that is currently carried by its NSFNET backbone will in the future be handled by commercial providers, encouraging the further development of this segment of the Internet.

The next-generation NSFNET backbone will support a narrower range of users and serve fewer sites. Today NSFNET backbone serves many sites nationwide, connecting regional networks and supercomputer centers (figure 2-4). It is a "general purpose" backbone, carrying traffic ranging from ordinary electronic mail to advanced supercomputer applications. In the future, the backbone will primarily be used by the NSF supercomputer centers, in Ithaca, New York, Pittsburgh, Pennsylvania, San Diego, California, and Champaign, Illinois.* Other users, with more routine applications, will use services available from commercial providers. Without the current national backbone, the regional networks will have to make new arrangements for their interconnection (see ch. 5, p. 67).

The next-generation NSFNET backbone *will* continue to contribute to the objective of developing advanced network technology. The new backbone, together with the next-generation NSI and ESNET, will be one of the frost networks to use the technologies studied by the gigabit testbeds described in chapter 4. The Federal networks will provide ''experimental" services, not yet available from commercial providers. They will demonstrate and test new network

---

5 For issues related to NSF'S Acceptable Use Policy, see Hearings before the House Subcommittee on Science, Space, and Technology, Mar. 12, 1992, Serial No. 120.

6 Gary H. Anthes, 'Internet Security Risks, " *Computerworld*, vol. 26, No. 48, Nov. 30, 1992, p. 55.

7 ''[T]he NREN Program has a series of synergistic goals [including] stimulating the availability, at a reasonable cost, of the required services from the private sector,' Office of Science and Technology Policy (OSTP), "The National Research and Education Network Program: A Report to Congress," December 1992, p. 2.

8 For a description of the NSF supercomputer centers, see U.S. Congress, Office of Technology Assessment, *High Performance Computing and Networking for Science*, OTA-BP-CIT-59 *(Washington* DC: U.S. Government Printing Office, September 1989), pp. 9-10.

Figure 2-4--NSFNET Backbone

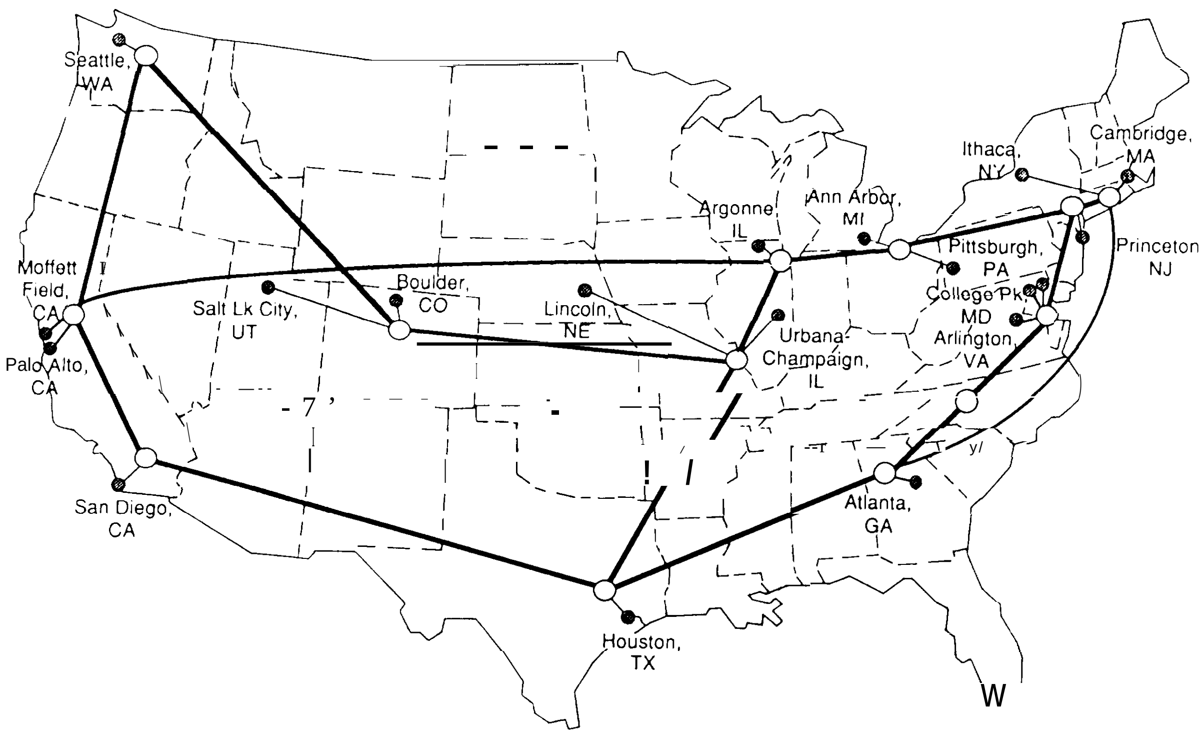

SOURCE: Merit, Inc.

technologies and applications before they are deployed more widely by operators of commercial networks.

Federal agencies may subsidize access to network services for users not at one of the backbone sites.[9] Today, NSF lowers the cost of networking for many users by directly subsidizing a general-purpose backbone and by providing subsidies to the regional networks. This strategy has contributed to broadening network access beyond major universities and supercomputer centers, to include many colleges and schools. In the future, many of these users will no longer be able to use the subsidized NSF network. Instead, Federal agencies may subsidize users' purchases of services from the commercial providers.

The NREN can then be viewed as many interconnected networks, developing from components of the current Internet, Some networks-the agency backbones-will be funded directly by the government. This part of the NREN is sometimes referred to as the 'Interagency Interim NREN" or "NREN proper," and will use advanced network technologies to support high-end users, agency missions, and the science objectives of the I-WCC program. Other Internet networks-such as existing regional networks or new commercial providers-may also carry NREN traffic, from users subsidized by the government, but would carry commercial traffic as well. These networks will likely use less sophisticated network technology than the agency backbone networks.

[9] ''Overtime, NSF will target its funding to those campuses which have **financial impediments** to **connecting** into the U.S. Internet." Robert **Aiken et al.**, "NSFImplementation Plan for **Interagency** Interim **NREN,''** May 1992, p. 4; **''Federal funds.** , will also support users that serve Federal missions whether or not they **access NREN** through the agency networks,' **OSTP,** op. cit., footnote 7, p. 3.

This two-part strategy-agency operation of advanced networks combined with subsidies for Internet access for certain groups of end users—represents a more detailed framework than the general NREN concepts and goals outlined in the High Performance Computing Act of 1991. It is expected to form the basis of NSF's forthcoming solicitation for the operation of its component of the NREN. It is also outlined in recently introduced legislation, the High Performance Computing and High Speed Networking Applications Act of 1993 (H.R. 1757), which would amend the High Performance Computing Act of 1991. However, there is concern in parts of the user community most affected by the change to an environment in which there is no longer a general purpose government operated network about the cost of commercial services and about the timing and management of the transition.

The remainder of this chapter describes the technology used in the current Internet. Chapter 3 provides an overview of emerging concepts that address some of the limitations of current network technology and might be used to construct gigabit networks. Chapter 4 describes the gigabit testbeds, NSF- and ARPA-funded prototype networks that are investigating these new technologies. Chapter 5 outlines NSF, NASA, and DOE plans for the deployment of the testbed technologies in their networks.

APPLICATIONS

From the users' perspective, an "application" is a task that the combination of the computer and the network enables them to perform. For example, a science teacher might use the Internet to locate information that can be used in a class, such as images stored in NASA databases, or databases containing tailored educational materials. Researchers use the Internet to track developments in their field, by exchanging information or drafts of papers and collaborating with other scientists.[10] In the business world, networks are increasingly used to track inventory or manage activities throughout a large company. In the future, networks may be used to help provide medical services to distant locations.

From a network engineering perspective, an ''application '' is a computer program that builds on the basic network service to allow a user to perform tasks. The application program provides interaction with the user; it does not handle the details of moving a message through the network to its destination. These functions are performed by communications software-a second program running on the computer—and specialized hardware that converts the computer's digital data to the format used by the network. When an applications program wants to send information to another computer, it hands the message to the communications software, which then formats the message and sends it over the network.

There are four major Internet applications—electronic mail (e-mail), file transfer, remote login, and news. Electronic mail is used to send messages to other users of the Internet, and for most users it is probably the application they use the most frequently. File transfer (File Transfer Protocol or FTP) is used to retrieve a "file" from another computer; a file could be a computer program, an article, or information from a commercial database. "Remote login" (Telnet) is used to control a distant computer; this is the application used to access a supercomputer or one of the other specialized computing resources on the Internet. "News" is a kind of bulletin board or discussion group-thousands of "newsgroups" address a wide range of different topics.

The current Internet applications are difficult to use. For example, it is difficult to find information resources on the network. First, the user has to know that the information exists somewhere reachable on the network, then where to find it, and, having found the database, how to locate the information in the database. A number of new

10 For an overview of the wide range of uses for the Internet, see Daniel P. Dem, ''Applying the Internet, *Byte*, February 1992, p. 111.

applications assist this process by acting as indexes or catalogues. Second, the user interface for most applications is often difficult to use, requiring a user to recall obscure commands. The difficulty in use is partly due to the Internet's heritage as an experimental network used mainly by scientists and engineers who were comfortable with arcane computer languages.

The existing Internet applications programs are beginning to be replaced by more sophisticated versions.[11] Today, for example, the Internet file transfer program, FTP, is used to retrieve a file from a distant computer, but a different program is used to retrieve a file stored on the ''home' computer. Newer versions of these applications are ''transparent, ' so that the user will not know whether a file is located on a distant computer, or that a program is executed on a different machine. These new applications are the beginnings of a foundation for 'distributed computing,' in which the computers on a network form an integrated system that performs as a single computer.

Applications and Network Technology

Some limitations of current applications are due to the applications software itself, but other limitations are due to the underlying network technology. One problem with current network technology is a shortage of bandwidth. Bandwidth is a measure of the amount of data that can be moved through the network in a given period of time, and is typically specified in terms of 'bits per second." Because of the limited capacity of today's network, it is often impractical to move large amounts of data across the network—examples of large files are images (see box 2-A) and the data sets used in supercomputer applications.

A second limitation of current Internet technology is that it is best suited for applications that handle text or numerical data. The Internet is less effective when supporting applications that make use of ''real-time' media such as video and sound.[12] In the case of video, this is due in part to the bandwidth limitation- high-quality video needs to move large amounts of data, and the necessary bandwidth is not available throughout the Internet. Support for video and sound is also limited because the performance of the Internet is highly variable. Because video creates the illusion of motion by sending a "stream" of pictures at regular intervals, a longer delay in the time it takes one of the pictures to get through the network interrupts the video information that is being displayed on the user's computer.[13] A new technology called "fast packet switching," discussed in detail in chapter 3, may provide the more consistent network performance that video applications need. Digital transmission and high bandwidth alone are not always sufficient to enable a network to carry video.

The limited capacity of the current Internet and the variability of its performance also constrain the use of sophisticated 'distributed computing" applications. In distributed computing, one is able to treat the computers on a network as a single, more powerful computer. For example, two computers, exchanging data through the network as necessary, might be able to complete a computation in half the time needed by one computer working alone. If data takes too long to travel between the computers, however, the advantages of dividing a computation among several computers are lost. In the current Internet, the local area network (LAN) technology used in campus networks often performs better than wide area network (WAN technology used in the

[11] For example, "distributed file systems" are beginning to replace the traditional File Transfer Protocol (FTP) application.

[12] Jeffrey Schwartz, "A Push for Packet Video," *CommunicationsWeek*, Aug. 3, 1992, p. 1.

[13] This problem is being attacked in a number of ways. New network architectures, described in chapter 3, try to reduce the degree of variation in network performancee. Other researchers are investigating mechanisms that would compensate for the variable performance. For example, the receiving computer could "even out" some of the variation before the data is displayed to the user.

Box 2-A-images and Video

Images

The screen of a computer's display is made up of many individual picture elements or "pixels," like the little dots that can be seen on television screens. By displaying each pixel with a different shade and different color, the computer forms an image on the screen. The greater the density of pixels, the higher the "resolution" of the image, The displays used for ordinary desktop computers usually have a few hundred pixels in both the horizontal and vertical directions, while a high-definition television display would have about 1,000 pixels vertically and about 2,000 horizontally. Even higher resolution displays are being developed for specialized medical, publishing, and defense-related applications.

The use of high-resolution images places considerable demands on computers and networks. Typically, each pixel on a screen is represented by 24 bits. A high-resolution display with 2,000 pixels horizontally and 2,000 pixels vertically has 4 million pixels (2,000x2,000=4,000,000). This means that 96 million bits are needed to represent the image (4 million x 24 = 96 million).

In the telephone network, voice conversations are sent through links that transmit 64 thousand bits per second. Using these links, an image represented by 96 million bits would take 25 minutes to send through the network. By contrast, it would take less than one-tenth of a second to send the same image through a gigabit network.

Video

Video is a series of images, sent many times a second at regular intervals in order to create the illusion of motion. Typically, 30 or 60 images are sent every second. In a low bandwidth network, in order to send this many images every second, the images have to be of very low resolution.

Two strategies have been adopted for accommodating image and video transport in networks. The first is to use compression techniques that reduce the number of bits needed for each image. Often, some parts of a scene do not have to be shown in great detail. Compression schemes for videotelephones sometimes rely on the fact that users are only interested in the "talking head," not the background. Sometimes little changes from one image to the next (if there is no movement in the scene), in which case the image data does not need to be sent again.

These techniques are being applied to the new high-definition television systems that are being studied by the Federal Communications Commission for selection as a U.S. standard. An uncompressed high-definition television signal that sends 30 images or "frames" every second, with a resolution of 1,000 pixels vertically and 2,000 pixels horizontally, needs about 1.5 gigabits per second. By contrast, new compression algorithms support high-definition television at bandwidths of 30 Mb/s or less, one-fiftieth the bandwidth required for the uncompressed signal.

The second strategy for accommodating video or images is to increase network capacity. Fiber optic technology can transport many more bits every second than the "twisted pair" copper wires that are used for today's telephone service. This background paper outlines some of the research being done on very high-capacity networks that can carry high-resolution video and images. However, even a "gigabit network" is not sufficient for certain kinds of very high-resolution video, and compression techniques might still be used.

SOURCES: Peng H. Ang et al., "Video Compression Makes Big Gains," *IEEE Spectrum,* **vol. 28, No. 10, October 1991,** pp. 1519; **Bernard Cole, "The Technology Framework,"** *IEEE Spectrum, vol. 30, No. 3,* **March 1993, pp. 32-39; J, Bryan Lyies and Daniel C. Swinehart, "The Emerging Gigabit Environment and the Role** of Local ATM," *IEEE Communications, vol. 30, No,* **4, April 1992, pp. 52-58.**

Figure 2-5-Growth in NSFNET Traffic

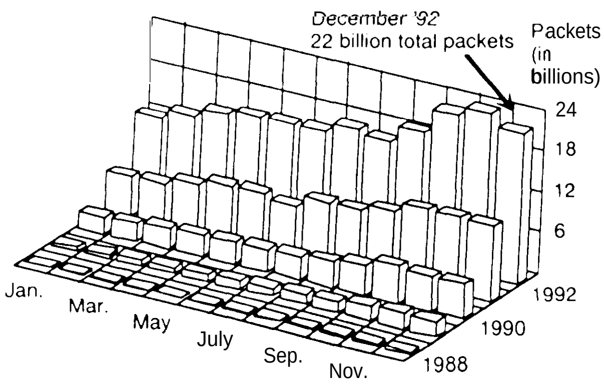

SOURCE: Merit. Inc.

regional and national backbone networks and thus distributed computing applications are used more widely in the local environment.

Limitations in network performance are becoming more apparent as computer technology advances. First, advances in computer power have resulted in demands for more bandwidth (figure 2-5). The size of the files that users would like to send through the network is increasing as a result of greater processing power and larger memories.[14] Some of the new "massively parallel" computers being studied as part of the HPCC Program may accelerate this trend (see box 2-B). Furthermore, the declining cost of computing power has allowed more users to connect to the network, creating more demand for the limited amount of capacity .15 Second, computers are increasingly equipped with display technology that supports video-based applications. As video and sound begin to be processed by computers, there will be greater demand for networks that support this stream-type traffic. Today's networks were designed for an environment in which computers were restricted to working with ordinary text and numerical data.

In response to the limitations of today's networks and the trends in computer design, there is now a general vision of the type of services that future computer networks will have to support—larger, possibly image-oriented files, greater use of stream-type services such as video and sound, and more distributed computing. However, there are a number of issues that must be solved, and researchers are trying to learn more about the applications that users will need in the future. Because most network technologies support some types of applications better than others, arguments in the technical community about the best way to build broadband networks can often be traced to different assumptions about the expected mix of applications. One of the objectives of the NSF/ARPA gigabit testbeds discussed in chapter 4 is to learn about applications for advanced networks by encouraging collaboration among applications developers and network engineers.

PROTOCOLS

The Internet is a "packet-switched' network—a very different design from that used by the telephone network.[16] Data travels through the network as a "packet," a block of digital data consisting of the application's data and some extra information added by the communications software and hardware.[17] This information is sent either before the applications data in a' 'header, ' or after the data, in a "trailer," and tells the network the packet's destination address or instructs the receiving computer as to what to do with the applications data in the packet (figure 2-6). For example, the sending computer could

14 Tim Studt, ''can High-Perforrnan ce Networks Meet Future R&D Needs?' *R&D Magazine, October* 1992, pp. 30-34.

15 Traffic on the NSFNET backbone is growing at a rate of 10% per month. OSTP, op.cit., footnote 7, p. 31.

16 Vinton G. Cerf, ''Networks,'' *Scientific American*, vol. 265, No. 3, September 1991, pp. 72-81.

17 Often, packets are compared t. the envelopes used in ordinary mail service—the extra information in the packet performs much the same function as the address on an envelope.

Box 2-B—Massively Parallel Computers

The conventional computers found on most desktops use a single processor. Programs for these computers consist of a list of instructions, to be executed one after another by the processor. Parallel computers are based on the idea that a computer with several processors can solve a problem more quickly than a computer with a single processor. Much of the HPCC Program's supercomputer design research focuses on the development of "massively parallel" computers with thousands of processors.

Supercomputers are expensive, high-performance machines that have been used mainly for numerical simulations in science and engineering. The first commercially important supercomputer, the CRAY-1, was first sold in 1976. It used a single processor, and achieved its high performance by careful attention to processor design and the use of specialized electronics. Over the next decade, supercomputer designers followed this basic model, trying to achieve the highest possible performance with a single processor.

By the mid-1980s, however, it became increasingly difficult to squeeze better performance out of traditional supercomputer designs, even as more exotic technologies were applied to the task. As a result, supercomputer designers began trying a different route to improved performance-the use of several processors. One approach involved a relatively small number of traditional high-performance supercomputer processors. For example, in 1983, Cray shipped a supercomputer that used four processors to speed up performance.

By contrast, the massively parallel approach to supercomputer design uses hundreds or thousands of low-cost microprocessors (processors that fit on a single semiconductor chip). The greater the number of processors, the more powerful the computer. In many cases, the microprocessors are the same as those used in high-end workstations. The performance of microprocessors increases every year, creating the potential for even more powerful massively parallel supercomputers.

Supercomputer centers and Federal laboratories have purchased several massively parallel supercomputers and are exploring their use in a number of applications. A major challenge for users of massively parallel supercomputers lies in the area of software. Massively parallel computers have to be programmed in new ways, because programs can no longer be thought of as a simple list of instructions. New algorithms, efficient ways of solving numerical problems, will have to be developed. Research on algorithms and software tools that take advantage of the potential of massively parallel supercomputers is one focus of the HPCC program.

SOURCES: Glenn Zorpette, cd., "Special Report: Supercomputers," *IEEE Spectrum, vol. 29, No. 9,* **September 1992; pp. 26-41; Office of Science and Technology Policy, "Grand Challenges 1993: High Performance Computing and Communications," 1992, pp. 13-17; Carl S. Ledbetter, "A Historical Perspective of Scientific Computing in Japan and the United States,"** *Supecomputing Review, vol.* 3, No. 12, **December 1990, pp. 48-58.**

put a short code in the header to tell the receiving computer that the data belongs to an electronic mail message—this allows the receiving computer to process the data appropriately after receiving the packet.

Once the packets have been formatted they are sent out of the computer and through the network's web of links and switches. Switches receive packets coming in on one link and send them out on the next link in the path to their destination (figure 2-7). When the packet arrives at a switch, the switch scans the destination address and determines which link the packet should transit next. The Internet packet switches or 'routers' are special computers that have been provided with connections to a number of links and programmed to carry out the switching functions.

The software in the routers and the users' computers implement 'protocols, ' the rules that determine the format of the packets and the actions taken by the routers and networked computers. The Internet protocols are often referred to as TCP/IP (the acronyms refer to the two

330-073 0 - 93 - 3 QL 3

Figure 2-6-Packet

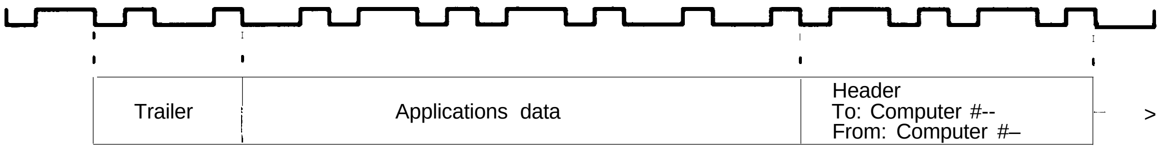

A packet is a block of digital data, consisting of data from the user's application and extra informationused by the network or receiving computer to process the packet. For example, the "header" might contain the "address" of the destination computer. A real packet would be several thousand bits long.

SOURCE: Office of Technology Assessment, 1993.

most important Internet protocols, the Transmission Control Protocol and the Internet Protocol.) Special protocols called "routing protocols" are used by the routers to keep a current map of the Internet and to determine the best path to a destination computer—for example, to choose a path that avoids heavily loaded networks.

One of the most important characteristics of the Internet is that the thousands of linked networks are independently operated; there is no central control of the Internet. However, by sharing the Internet protocols, the networks are able to exchange traffic. One of the functions of the Internet protocols is to mask differences in the technology used by the networks that makeup the Internet. The campus networks' local area network technology differs from the wide area network technology used in the regional and national backbone networks, and there are many different local area network standards. The term "Internet" is short for "internetworking," the practice of linking technologically different and independently operated networks.

The future of the current Internet protocols is the subject of considerable debate in the Internet community. The most significant problem is that today's routing technologies are being strained by rapid growth in the number of connected networks and users.[18] The management of a complex and growing network has been one of the major challenges faced by the current NSFNET. A number of different proposals that would simplify the routers' task of finding paths through today's more complex Internet are being considered. The effect of increases in bandwidth on TCP/IP has also been debated in the technical community, and new protocols have been proposed. Many now believe that TCP/IP can continue to provide good service over gigabit networks, but internetworking in high bandwidth networks is a research topic in itself.

NETWORK COMPONENTS

A network is a complex system, consisting of many computer programs and hardware components such as links, computers, and switches. The overall performance of the network depends on how well these components work together. There are a number of potential bottlenecks-the rate at which data can be transferred from the computer's memory to the network, the rate at which data can be transmitted through the links, and the amount of time the switches need to decide where to send data next. Simply removing one of these bottlenecks does not guarantee that the overall performance of the network will improve. The emergence of fiber optics has removed the links as a bottleneck for the foreseeable future; the research projects described in chapter 4 show that this has exposed research issues in other parts of the network.

18 Karen Lynch, "Internet Metamorphosis,' *CommunicationsWeek International,* Aug. 10, 1992, p. 1.

Figure 2-7—Packet Switching

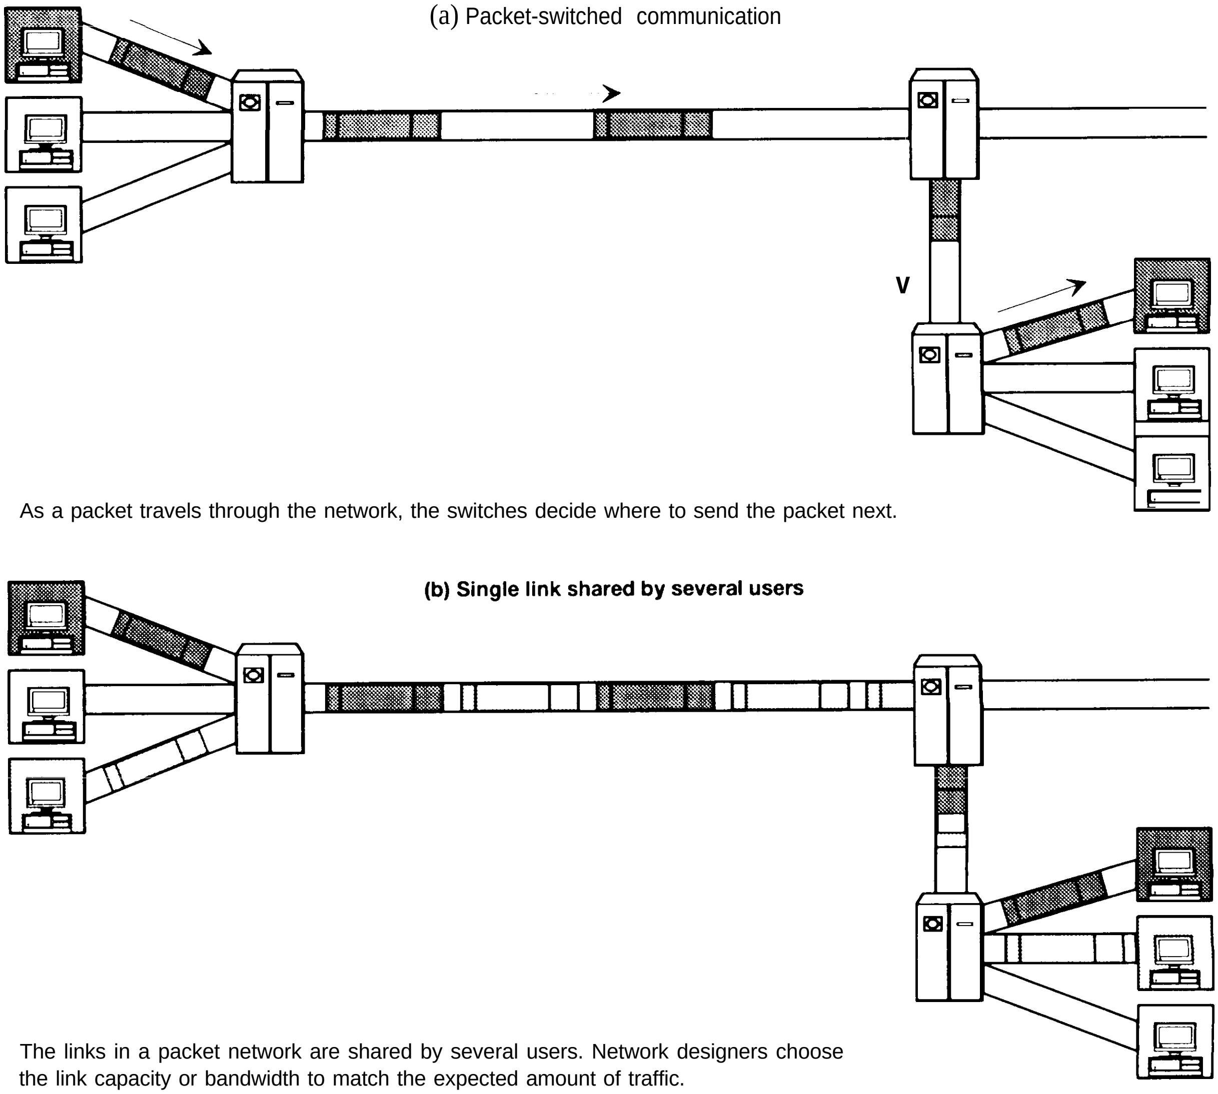

(a) Packet-switched communication

As a packet travels through the network, the switches decide where to send the packet next.

(b) Single link shared by several users

The links in a packet network are shared by several users. Network designers choose the link capacity or bandwidth to match the expected amount of traffic.

SOURCE: Office of Technology Assessment, 1993.

Computers

Many different kinds of computers are attached to the Internet, ranging from desktop personal computers costing a few hundred dollars to supercomputers that cost millions of dollars. Among scientists and engineers, the type of computer that is most widely used is the ‘‘workstation," a powerful desktop computer with enough processing power to support graphical user interfaces and high-resolution displays. For most of today’s applications, almost any computer has enough processing power to attach to the Internet. The low bandwidth of the current Internet places few demands on computers for handling the communications functions, leaving much of the processing power free to run the applications.

One of the reasons for the creation of the NSFNET backbone was to provide access to NSF's four supercomputer centers. Recently, these supercomputer centers have begun to install "massively parallel' supercomputers. This new type of supercomputer attempts to achieve very high processing speeds by combining the processing power of thousands of smaller processors. Other supercomputers use a more traditional design, and are referred to as 'vector' supercomputers. Each design may work best with certain kinds of computations; one of the objectives of the gigabit testbed research is to explore the use of networks to divide up problems in a way that takes advantage of the strengths of both vector and massively parallel supercomputers.[19]

Links

The digital links in computer networks usually use copper or fiber, but satellite and microwave links are also used. At each end of the copper or fiber is the transmission equipment, electronics that convert data into the optical or electrical signals that travel through the network. The capacity of the wires or strands of fiber depends on the characteristics of the material used and on the capabilities of the transmission equipment.

Today's Internet uses both low bandwidth links that operate over copper at a few thousand bits per second (kilobits per second or kb/s), and high bandwidth links that operate over fiber with a data rate of about 45 million bits per second (megabits per second or Mb/s). The test networks described in chapter 4 will use links that operate at a rate of one billion bits per second (a gigabit per second or Gb/s).

Typically, a single wire or strand of fiber carries many links at the same time. Through a process called "multiplexing," several low-bandwidth links can be aggregated into a higher bandwidth link. Gigabit-capacity fiber, for example, can be used either to carry several thousand low-bandwidth links used for telephone calls, or a single high-bandwidth link needed for a gigabit network.

The required link bandwidth depends on both the bandwidth requirement of each user and on the number of users sharing the link One of the main reasons for upgrading the links in the NSFNET backbone from 1.5 Mb/s to 45 Mb/s in 1991 was to accommodate growth in the number of users. Growth in the use of routine applications can also be supported by simply adding more low-bandwidth links. However, new applications that need very large amounts of bandwidth to themselves require the deployment of higher bandwidth links. By increasing the link bandwidth to gigabit rates, the gigabit NREN will be able to support new classes of advanced applications, not just growth in the number of users.

Operators of wide area computer networks, such as the regional networks and the agency backbones, typically lease their links from the telephone companies. The telephone companies have already obtained the rights-of-way and have installed the transmission facilities for use in their core business, voice telephone service. Because of the reliance on telephone company facilities, discussions of computer network link bandwidth often use telecommunications industry designations of link capacity. For example, the current NSFNET backbone is often referredtoasa''T3° network, after the industry designation of 45 Mb/s links. ''Tl links, which operate at 1.5 Mb/s, are used in the current Department of Energy and NASA networks and in the regional networks. As the Federal networks are upgraded to bandwidths above the 45 Mb/s T3 rate, they will use a new family of transmission standards designed for high-capacity fiber optic links, called Synchronous Optical Network (SONET) (see table 2-l).

Universities and corporations install their own links in their buildings for use in local area networks. Local area networks can provide users with higher bandwidth than wide area networks—

[19] ''A Union of Superpowers," *IEEE Spectrum,* vol. *28, No. 6,* June 1991, p. 18.

Figure 2-8-Access Link

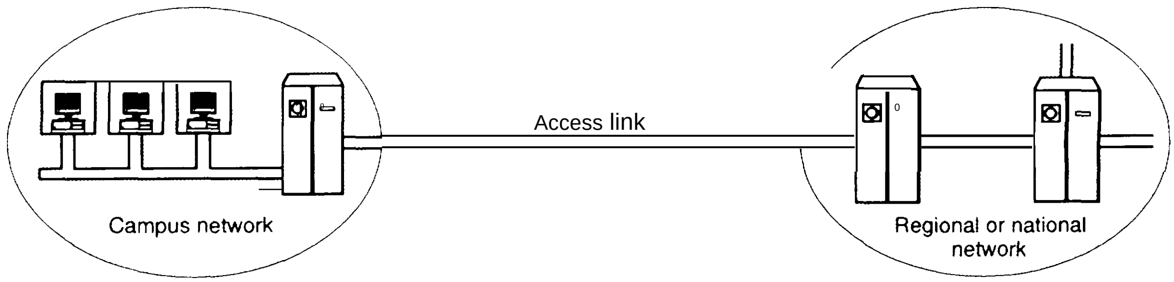

SOURCE: Office of Technology Assessment, 1993.

Table 2-l—Transmission Rates

Industry designation	Transmission rate
DSO	64 kb/S
T1	1.5 Mb/s
T3	45 Mb/s
SONET OC-3	155 Mb/s
SONET OC-12	622 Mb/s
SONET OC-48	2.4 Gb/s

SOURCE: Office of **Technology Assessment, 1993.**

this is due in large part to the high cost of high bandwidth wide area links. Because of the higher bandwidth available on local area networks, they have been used for experimentation with high-bandwidth distributed computing and video applications. In the future, however, users will want wide area networks that match the performance of local area networks; one of the objectives of the testbed project outlined in chapter 4 is to investigate high-speed wide area networking.

When campus networks arrange to be connected to the closest regional or national network, they obtain an ''access' link (figure 2-8). This is usually leased from the telephone company, just as the links inside wide area networks are leased from the telephone company. The cost of the Internet service depends on the access bandwidth; high bandwidth access is extremely expensive. It is common to find local area networks operating at 10 Mb/s or 100 Mb/s, while the access link to the rest of the Internet operates at 56 kb/s or less (some organizations have 1.5 Mb/s access links, but these are considerably more expensive). Most individuals, schools, and small businesses are required to use their ordinary analog telephone line to access Internet services-a device called a ''modem' is needed to send digital computer data over these lines, usually at 14.4 kb/s or less.

Switches

Packet switches in the Internet, also known as routers, direct packets to the next link in the path to their destination. Packet switched networks emerged to handle data communications, services not well supported by the "circuit switches' used for ordinary telephone calls (figure 2-9). Packet networks are more efficient for typical computer communications traffic-short transactions or "bursts" separated by periods of no traffic (box 2-C). In a packet network, several users share the same link-during the period in which one group of users is not using the link, other users can send their packets. In a circuit switched network, by contrast, each communication gets its own link. For this reason, circuit switches are most efficient when a communication involves a relatively long, steady stream of data such as video or voice.

While the Internet networks use telephone company links, the packet switches are usually not operated by the telephone companies. Instead, a second organization plans the network and installs the packet switches at the sites it has chosen—the involvement of the telephone company is usually limited to providing the links between the sites. From the perspective of the

Figure 2-9-Circuit Switching

(a) Telephone network

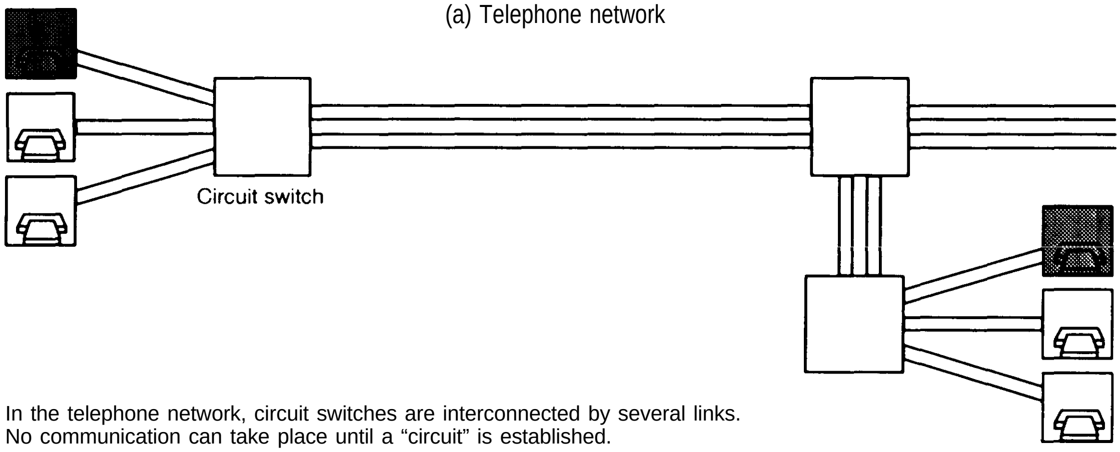

In the telephone network, circuit switches are interconnected by several links. No communication can take place until a "circuit" is established.

(b) Circuit-switched call

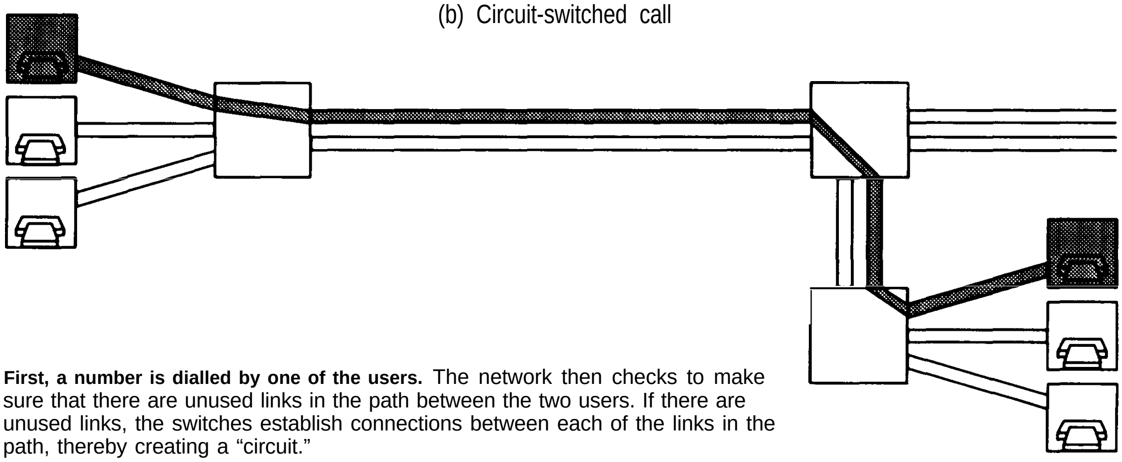

First, a number is dialled by one of the users. The network then checks to make sure that there are unused links in the path between the two users. If there are unused links, the switches establish connections between each of the links in the path, thereby creating a "circuit."

(c) Separate circuits for each call

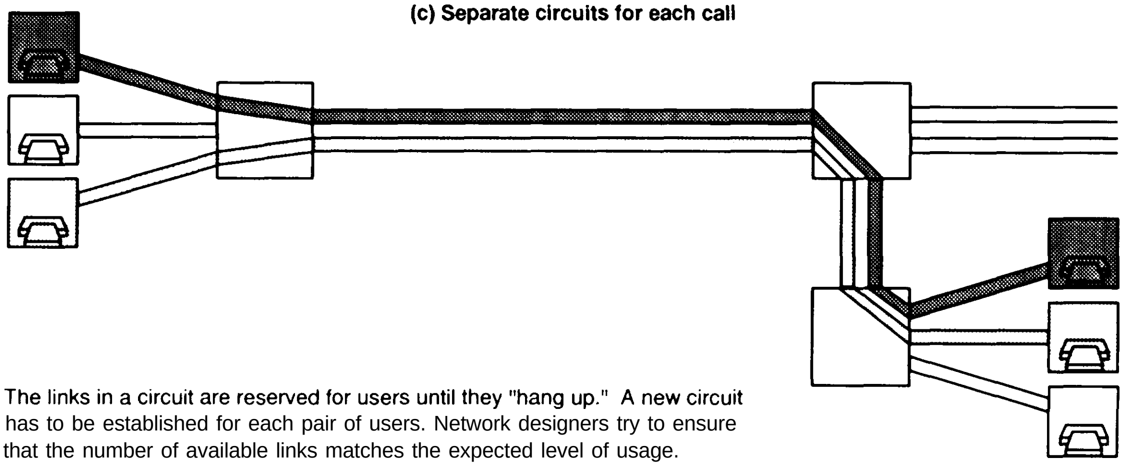

The links in a circuit are reserved for users until they "hang up." A new circuit has to be established for each pair of users. Network designers try to ensure that the number of available links matches the expected level of usage.

SOURCE: Office of Technology Assessment, 1993.

telephone company, the computer network traffic is just "bits" traveling over its links-the telephone company's equipment does not make decisions about where to send the packets. Beginning in the mid- 1970s, the telephone companies began installing some packet switches in their networks in order to support the growing data communications market, but their efforts to enter this market were considered to be unsuccessful.

The processing power required of a packet switch depends on the link bandwidth and the complexity of the network. As the link bandwidth increases, switches must be able to process packets more quickly. The processing power needed will also increase as the network gets larger and more complex, because it becomes more difficult to determine the best path through the network. Currently, the NSFNET backbone's router technology does not allow the use of applications that need more than 22.5 Mb/s, half the potential maximum of a 45 Mb/s T3 network.[20] This shows how the overall performance of the network depends on many different components; increasing the link bandwidth is not the only requirement for an advanced network.

THE INTERNET AND THE PUBLIC SWITCHED NETWORK

In some ways, the Internet and the "public switched network" that is operated by common carrier telephone companies are separate. They differ in the services they provide-the telephone network mainly provides ordinary voice communications services, while the Internet provides data communications services such as electronic mail and access to remote computers. They also differ as to the communities that they serve-almost everyone has a telephone, while the Internet and other computer networks primarily serve users in the academic community or in industry. Finally, they differ in their network technology-the Internet is a packet-switched network, while the telephone network is a circuit-switched network. However, the Internet and the telephone network are related in a number of ways. Any discussion of the evolution of networking has to consider both the traditional telecommunications companies and the Internet community.

First, the Internet and the public switched network are related in that the links in wide-area computer networks are usually supplied by the telephone companies--computer network operators do not usually put their own fiber in the ground. As a result, the availability of new computer network capabilities can depend on the extent to which the telephone companies deploy advanced transmission facilities, and on the cost of leasing the links.

The availability of advanced transmission facilities varies[21], depending on whether a computer network will operate over the telephone network's ''interoffice' or "local loop" segments. Most of the links required for a wide area network such as the NSFNET backbone operate over the interoffice core of the telephone network, which has largely been upgraded to optical fiber and digital transmission. The telephone companies upgraded this part of their networks in part to achieve operational savings, even when delivering existing services.

However, access links, such as those between a campus and a regional network, need to use local loop facilities. For the most part, this segment of the telephone network still consists of copper, analog lines. Large users are able to avoid this bottleneck by making special arrangements with the local exchange carrier for higher bandwidth digital lines. However, individuals, schools, and small businesses generally have to

20 Alan Baratz, IBM, personal communication, Feb. 3, 1993.

21 National Telecommunications and Information Administration Department of Commerce, ''Telecommunications in the Age of Information" October 1991, pp. 97-109.

Box 2-C-Packet Switching and Circuit Switching

Computer networks such as the Internet use packet switches, which direct packets from link to link through a network. Today's telephone network by contrast, uses circuit switches. Each type of switching technology works best with different kinds of communications. Packet switching is more efficient for the transfer of typicai computer communications traffic such as files of text or numerical data (figure 2-C-1). Circuit switching, on the other hand,

Figure 2-C-l—Packet Switching More Efficient for Data

(a) **Data** communications does not use circuits fully

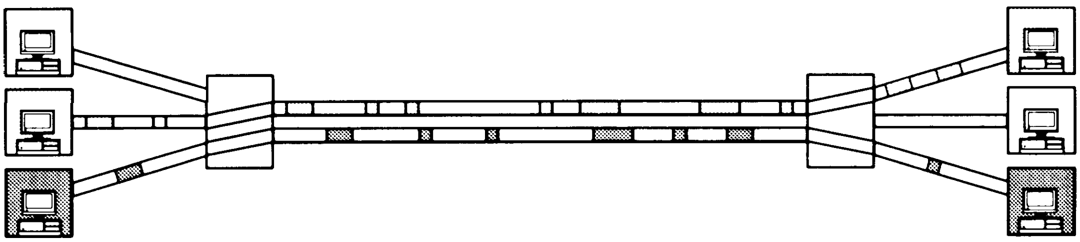

Circuit switching can be used for computer communications. Here, circuits have been set up between two pairs of computers. However, computer communications often have a "bursty" character -- periods in which data IS sent followed by periods of "silence." When no data is sent, the circuit's capacity goes unused. The capacity IS used more efficiently when the communications Involve a steady flow of Information, such as video or voice transmission.

(b) Link sharing makes packet networks **more efficient for data**

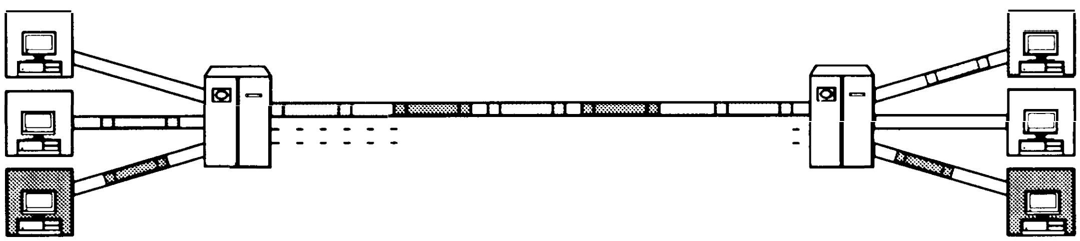

In a packet-switched network, several users' traffic shares the same link. If one user IS not using the link's capacity, it can be used by others. The figure shows bursts of data assembled into packets and travelling through the network on the same link. Here, one link's capacity is sufficient to handle communications between both pairs of users, freeing the second link for other uses.

SOURCE: Office of Technology Assessment, 1993.

use the combination of a modem and their telephone line to access computer networks. The bandwidth of such an arrangement is relatively low--only a few kb/s—and is clearly a bottleneck that limits widespread use of sophisticated services.

The telephone network and computer networks are also related in the sense that the traditional telecommunications providers are beginning to offer data communications services, including Internet services. In the past, efforts by the industry to enter this market have not been successful. This has been attributed to a ''culture clash'—a lack of understanding of computer network technology and of the needs of users of computer networks. However, the telephone comcan provide the consistent performance needed by video or voice traffic (figure 2-C-2). One of the objectives of the research described in chapter 3 and chapter 4 is to develop switches that combine the efficiency and flexibility of packet switching with the consistent performance of circuit switching.

Figure 2-C-2—Circuit Switching Better for Voice or Video

(a) **Variable performance due to packet network link sharing**

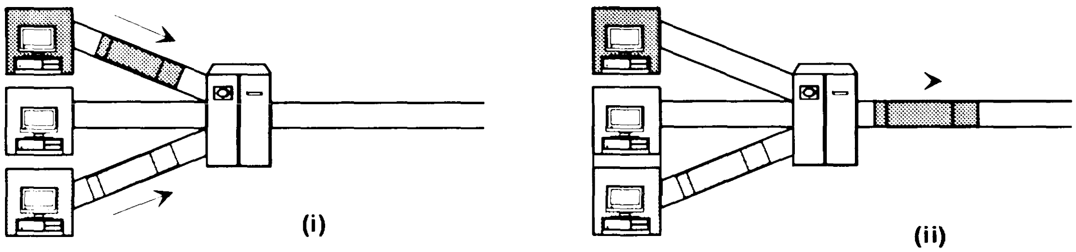

If two packets arrive at a switch at the same time and need to use the same outgoing link (I), one of the packets will have to wait (II). It is difficult for a user to know in advance what the network performance will be. The packet may experience no delay (the dark gray packet), or it may have to wait at each switch (the light gray packet). This variation in delay has limited the use of packet networks for time-sensdive communications such as video or voice.

(b) Circuit switched performance is predictable

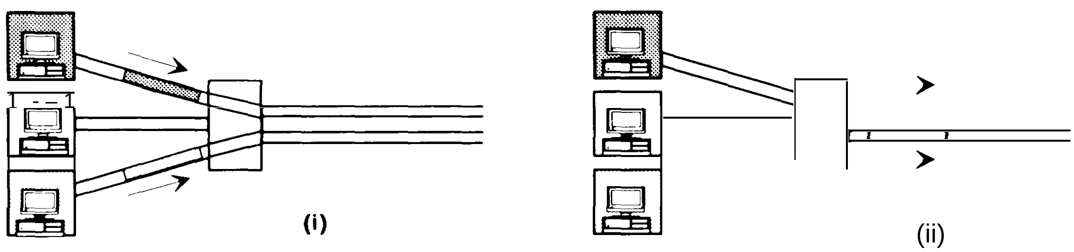

In a circuit-switched network, each communication has Its own circuit. Users' information travels through the network without being affected by the characteristics of other communications (1)–(11). The time needed for information to travel through the network will always be the same.

SOURCE: Office of Technology Assessment, 1993.

panics hope to play a more active role in this market.

The telephone companies have two main competitors in this venture. First, there are already a number of commercial providers of Internet services and other data communications services. These providers lease lines from the telephone companies, install packet switches, and operate their network without any further involvement from the telephone companies, sharing their network's capacity among different groups of users for a fee. The current T3 NSFNET backbone is obtained as a service from one of these commercial Internet providers.

Second, many users choose to operate 'private networks' '—they build a network of their own using leased lines and bypass the public network. Most corporations use this strategy to interconnect local area networks at different sites within their organization. Equipment used in private networks is provided by computer companies and others, who have taken advantage of the telephone companies' lack of success in providing data communications services. United States firms that specialize in the development of routers and other equipment for private networks are world leaders and are among today's fastest growing companies.[22]

The telephone companies have introduced a number of new packet-switched services that are intended to encourage users to abandon their private networks.[23] One of these services is called Switched Multimegabit Data Service (SMDS); another is called Frame Relay.[24] The SMDS and Frame Relay switches do not understand the Internet protocols, but they can still be used to carry Internet traffic. The Internet packets are ''encapsulated, ' or put inside an SMDS or Frame Relay 'envelope,' and sent through the network; at the other end the Internet packet is extracted and delivered to the computer. The carriers view SMDS and Frame Relay as transitional steps to a new technology called Asynchronous Transfer Mode (ATM), described in chapter 3. They can potentially be used to provide data communications services up to 45 Mb/s.

Because of the interrelationship between the Internet and the public switched network, the evolution of the Internet is affected by two different sets of standards committees. The telecommunications industry standards affect mainly low level issues, such as transmission standards, but some of the standards for new telecommunications industry packet switched services may play a role as well. The most important international standards group is the CCITT (International Telegraph and Telephone Consultative Committee). The CCITT is a technical committee of the International Telecommunications Union (ITU), a specialized agency of the United Nations that is headquartered in Geneva.[25] United States telecommunications standards are the responsibility of Committee Tl, which is accredited by the American National Standards Institute (ANSI) and sponsored by the Exchange Carriers Standards Association (ECSA).[26] Telecommunications industry standards setting has often been criticized as excessively bureaucratic and slow.

By contrast, the Internet standards community, which addresses higher level issues related to routing, the TCP/IP protocols, and applications, is more informal. Much of the work is done by electronic mail, and there is a greater emphasis on proving that something works before it is standardized.[27] The two groups responsible for Internet standards are the Internet Engineering Task Force (IETF) and the Internet Activities Board (IAB). The IETF has a number of different working groups, each looking at a different aspect of the Internet's operation.

---

22 G. Pascal Zachary, "U.S. High-Tech Firms Have Begun Staging Little-Noticed Revival," *Wall Street Journal,* Dec. 14, 1992, p. 1; G. Pascal Zachary and Stephen Kreider Yoder, "Computer Industry Divides Into Camps of Winners and Liners," Jan. 27, 1993, p. 1; Alan Deutschman, "America's Fastest-Growing Companies, " *Fortune, vol., 126, No. 7,* Oct. 5, 1992, p. *58.*

23 Another market for these services is the smaller companies that cannot currently justify private networks.

24 Tim Wilson, "Local Carriers Lay Out Data Service Agendas," *CommunicationsWeek,* May 25, 1992.

25 G.A. Codding and A.M. Rutkowski, *The International Telecommunications Union in a Changing World* (Dedham, MA: Artech House, 1982).

26 I.M. Lifchus, "Standards commit@ T1—Telecommunications," *IEEE Communications, vol. 23, No.* 1, January 1985, pp. 34-37.

27 Carl Malamud, "Stacks: Interoperability in Today's Computer Networks," (Englewood Cliffs, NJ: Prentice Hall, 1992), p. 223.