Introduction and Summary 1

The vision of the Nation's future telecommunications system is that of a broadband network (see box l-A) that can support video, sound, data, and image communications. Toward this end, the High-Performance Computing Act of 1991 called for the Federal computer networks that connect universities and Federal laboratories to be upgraded to "gigabit networks" (see box l-B) by 1996.[1] This background paper reviews technologies that may contribute to achieving this objective, and describes the six prototype gigabit networks or "testbeds" that are being funded as part of the Federal High Performance Computing and Communications Program. These prototype networks are intended to demonstrate new communications technologies, provide experience with the construction of advanced networks, and address some of the unresolved research questions.

FEDERAL SUPPORT FOR GIGABIT NETWORKING

The High Performance Computing and Communications Program (HPCC) is a multiagency program that supports research on advanced supercomputers, software, and networks.[2] In part, these technologies are being developed to attack the "Grand Challenges": science and engineering problems in climate change, chemistry, and other areas that can only be solved with powerful computer systems. Network research is one of four components of the HPCC program, and represents about 15 percent of the program's annual budget of close to $1 billion.[3]

The HPCC program funds the development of new communications technologies.

[1] High-Performance Computing Act of 1991 (HPCA), PL 102-194, Sec. 102(a).

[2] Office of Science and Technology Policy (OSTP), "Grand Challenges 1993: High Performance Computing and Communications, ' 1992.

[3] Ibid., p. 28.

Box I-A–Broadband Networks

Computers and networks handle information as patterns of electronic or optical signals. Text pictures, sound, video, and numerical data can then be stored on floppy disks, used in computations, and sent from computer to computer through a network In digital/computers or networks, the electronic or optical signals that represent information can take on one of two values, such as a high or a low voltage, which are usually thought of either as a “l” or “1” a “O” (figure I-A-l). These 1s and 0s are called bits.

Figure I-A-I—Digital Data

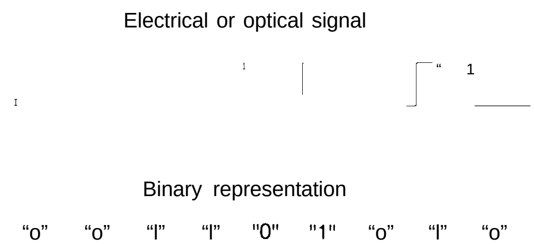

SOURCE: Office of Technology Assessment, 1993.

Different patterns of 1s and Os are used to represent different kinds of data Inmost computers, the letter ”A” is represented by the pattern of electronic signals corresponding to “01 000001 .“ To represent images, different patterns of bits are used to represent different shades (from light to dark) and odors. Sound is represented in much the same way, except that the patterns of bits represent the intensity of sound at points in time.

The number of bits required to represent information depends on a number of factors. One factor is the quality of the representation. A good quality, high-resolution image would require more bits than a low-resolution image. Also, some kinds of information inherently require more bits in order to be represented accurately. A page of a book with only text might contain a few thousand characters, and could be represented with a few tens of thousands of bits. A page of image data on the other hand, could require millions of bits.

Because images and video, which is a sequence of images, require many more bits to be represented accurately, they have strained the capabilities of computers and networks. Images take up too much space in a computer’s memory, and take too long to be sent through a network to be practical. The new high-capacity network technologies described in this background paper have the ability to support two-way digital, image, and video communications in a more efficient manner.

Digital Networks

In the past, networks designed for video or sound used anabg transmission. In the old analog telephone network, for example, the telephone’s microphone converted the spoken sounds into an electrical signal whose

The other three components of the program target supercomputer design, software to solve the Grand Challenges, and research in computer science and mathematics.

The HPCC program is the most visible source of Federal funds for the development of new communications technology. The networking component of the program is divided into two parts: 1) research on gigabit network technology, and 2) developing a National Research and Education Network (NREN). The gigabit research program supports research on advanced network technology and the development of the six testbeds. The NREN program supports the deployment of an advanced network to improve and broaden network access for the research and education community. The High-Performance Computing Act of 1991 specfies that the NREN should operate at gigabit speeds by 1996, if technically possible.[4]

[4] HPCA, op. cit., footnote 1.

strength corresponded to the loudness of the sounds. This signal then traveled through the network's wires until it reached its destination, where it was used to make the telephone's speaker vibrate, recreating the spoken Sounds.

Digital/ networks transmit information in digital form, as a series of bits. Digital networks are required for high-speed communications between computers-computers work with digital data. However, digital networks can also transmit real-world information such as sounds and pictures if special digital telephones or video cameras are used to represent the information in digital form. A digital telephone, for example, generates a series of patterns of Is and 0s, corresponding to the loudness of the sounds. At the destination, these 1s and Os are interpreted by the digital telephone and used to recreate the original sounds. Digital networks are quickly replacing analog networks. They are needed to transmit the growing amount of computer data They also transmit voice and video information more cleanly, without interference and distortion. More importantly, digital networks allow a single network to carry all types of information. Today, separate networks are used for voice traffic (the telephone network), computer communications (data networks such as the Internet) and video (broadcast or cable television or other specialized networks). Because these different kinds of information can all be represented in digital form, a single digital network can potentially be used to transmit all types of information. This is not the only requirement, however (see ch. 2 and ch. 3).

Broadband Networks

The Capacity of a digital network is often described in terms of the number of bits that the network can transmit from place to place every second. A digital telephone network can transmit 64,000 bits every second. This is sufficient capacity to carry a telephone conversation with acceptable quality, but is not enough to carry video. Although some videotelephones can use regular telephone lines, users of videoconferencing systems usually prefer to use special services that can transmit at 384,000 bits per second or more. VCR-quality television needs about 1.5 million bits per second, and high-definition television needs about 20 million bits per second--about 300 times the capacity of a digital telephone line.

The Capacity of a network, measured as the number of bits it can transmit every second, is called "bandwidth." Engineers often talk about "narrowband" networks, which are low bandwidth networks, and "broadband" networks, which are high bandwidth networks. The dividing line between the two is not always clear, and changes as technology evolves. Today, any kind of network that transmits at more than 100 million bits per second would definitely be considered a broadband network. Chapter 3 describes fiber optics and other technologies that will be used to build broadband networks.

SOURCE: Office of Technology Assessment, 1993.

Broadband networks such as the NREN will both improve the performance of existing applications and accommodate new types of applications. There will likely be a shift to image- and video-based communications, which are not adequately supported by currently deployed network technology. "Multimedia" applications that use images and video, as well as text and sound, look promising in a number of areas, e.g., education, health care, business, and entertainment. Broadband networks will also allow a closer coupling of the computers on a network; as the network is removed as a bottleneck, the computers will be able to form an integrated system that performs as a single, more powerful, computer.

Broadband networks will require a fundamental rethinking of network design. Several new concepts have been proposed and are being investigated by the testbeds. Fiber is a highly touted technology for constructing broadband networks, but it alone is not sufficient. Switches (see box l-C) and the components that link computers to the network will have to be upgraded at the same time in order to keep pace with

Box I-B-Gigabit Networks

Much of the research described in this background paper is aimed at the development of gigabit networks, broadband networks that can transmit data atone billion bits per second or more (a "gigabit? is one billion bits; "gigabit per second" is abbreviated as Gb/s or Gbps). This represents a 20-fold increase over the most capable links in the networks that currently serve the research and education community. The current National Science Foundation network uses links that transmit data at 45 million bits per second (megabits per second or Mb/s), and even this capacity has not been fully utilized because of bottlenecks in the network's switches. The development of a gigabit network is an ambitious target-most current industry technology planning targets broadband networks with lower bandwidths, in the 150 million bits per second range.

The basic outlines of the technology evolution of the DOD, NASA, DOE, and NSF networks that serve research and education were established in 1987 and 1989 reports issued by the Office of Science and Technology Policy. In the late 1980s, link bandwidths in the Federal networks were 1.5 Mb/s or less. The OSTP reports outlined a three-stage plan for the evolution of these networks to gigabit networks by the mid-to-late 1990s (see figure I-B-I). The gigabit target was also specified by the High-Performance Computing Act of 1991. The OSTP report envisioned that each generation of technology would move from an experimental phase in the Federal networks to commercial service.

Figure I-B-I—Timetable for the National Research and Education Network

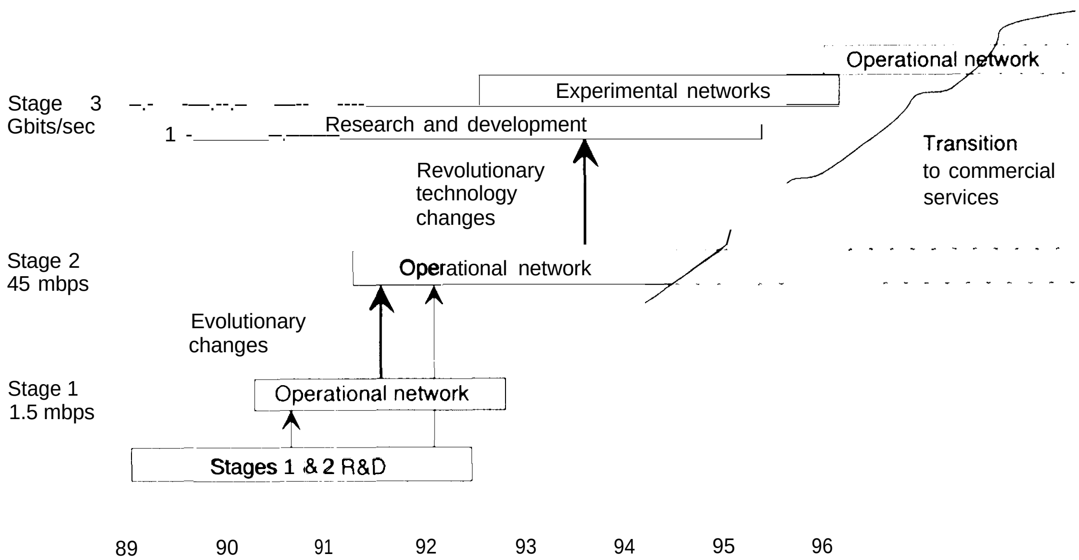

SOURCE: Office of Science and Technology Policy, "The Federal High Performance Computing Program," September 8, 1989.

Currently, the Federal agency networks are in the middle phases of the second stage, the operation of networks with 45 Mb/slinks. At the same time, research and development for the third stage, the deployment of gigabit networks, is underway. In practice, the network capacity will not jump directly from 45 Mb/s to gigabit rates. The next step will be to 155 Mb/s, then to 622 Mb/s, and then to greater than one gigabit per second. The bandwidths used in computer networks (1.5 Mb/s, 45 Mb/s, 155 Mb/s, and 622 Mb/s) correspond to standards chosen by manufacturers of transmission equipment.

SOURCES: Office of Science and Technology Policy (OSTP), "A Research and Development Strategy for High Performance Computing," Nov. 20, 1987; OSTP, "The Federal High-Performance Computing Program," Sept. 8, 1989; High-Performance Computing Act of 1991 (HPCA), Public Law 102-194, Sec. 102(a).

Box I-C-Computer Network Components

A computer network has three main components: computers, links, and switches (figure 1 -C-l). The web of links and switches carry data between the computers. Links are made of copper (either "twisted pair" or "coaxial cable") or fiber optics. Transmission equipment at each end of the fiber or copper generates the electrical or optical signals. There are also satellite and microwave links that send radio waves through the air. Fiber has several advantages over other types of links--most notably its very high bandwidth. The fiberoptic links needed for gigabit networks are already commercially available. However, gigabit networks will not be deployed until research issues in other network components are addressed.

Figure I-C-I—A Simple Computer Network

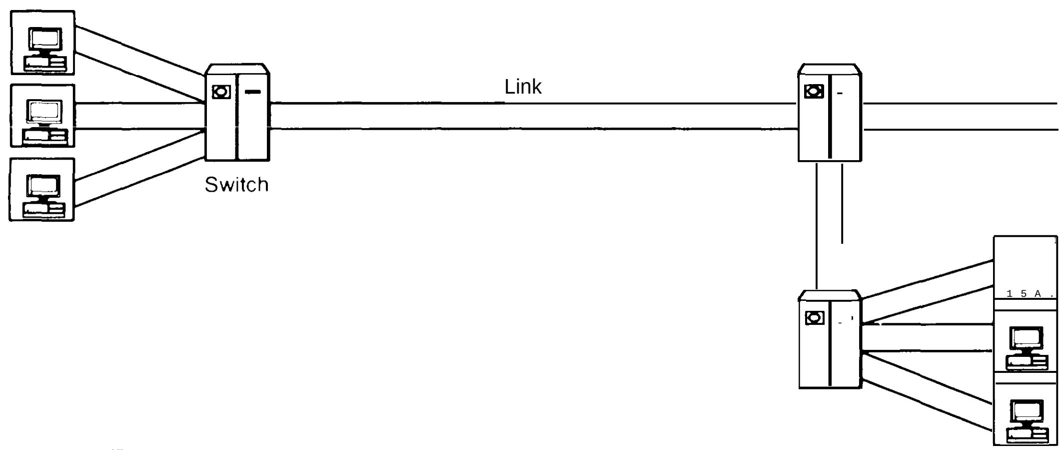

SOURCE: Office of Technology Assessment, 1993.

For example, new high-capacity switches are needed to keep pace with the higher bandwidth of fiber optic links. Just as railroad switches direct trains from track to track, the switches in computer networks direct information from link to link. As the information travels through the network, the switches decide which link it will have to traverse next in order to reach its destination. The rules by which the switches and users' computers coordinate the transmission of information through the network are called protocols.

While most computer networks are limited in their ability to carry high-bandwidth signals such as video, cable television networks are widely used to distribute television signals to homes. However, cable networks usually do not have switches. For this reason, they only permit one-way communications: the signal is simply broadcast to everyone on the network. Much of the network research today is devoted to the development of switches that would allow networks to support two-way, high-bandwidth communications.

SOURCE: Office of Technology Assessment, 1993.

the faster flow of data. Broadband networks will be more than simply higher bandwidth versions of today's networks, however. Networks will also be redesigned so that a single type of network can carry video, sound, data, and image services. The existing telephone and data networks do not have sufficient flexibility to carry all types of information efficiently.

■ The NREN

One objective for the NREN is that it serve as an enabling technology for science and engineering researche[5] The gigabit NREN will be able to handle the very large data sets generated by supercomputers. Scientists could use the gigabit NREN to support "visualization," the use of a computer-generated picture to represent data in image form. For example, ocean temperatures computed by a climate model could be represented by different colors superimposed on a map of the world, instead of a list of numbers. Visualization is an essential technique for understanding the results of a simulation. Currently, much of the data collected in experiments and computed by simulations goes unused because of the time needed to compute images on conventional computers. Supercomputers could perform the computations more quickly, but few laboratories have supercomputers. With a high-speed network, a scientist could send the data to a distant supercomputer, which would be able to quickly compute the images and send them back through the network for display on the scientist's computer.

Broadband networks will require a fundamental rethinking of network design.

A second objective for the NREN program is that it demonstrate and test advanced broadband communications technologies before they are deployed in commercial networks. The NREN program will upgrade federally supported networks such as the National Science Foundation's NSFNET, the Department of Energy's Energy Sciences Network (ESnet), and the National Aeronautics and Space Administration's NASA Science Internet (NSI).[6] These networks form the core of the "Internet" a larger collection of interconnected networks that provides electronic mail services and access to databases and supercomputers for users in all parts of the United States and around the world.[7] During 1992, Federal agencies announced plans for upgrading their current networks as part of the NREN program.[8]

The NREN program can be viewed as a continuation and expansion of the Federal support that created the Internet. The Internet's technology evolved from that of the Arpanet, a research project of the Advanced Research Projects Agency. Beginning in 1969, the Arpanet served to demonstrate the then-new technology of 'packet switching." Packet switched networks were able to support computer communications applications that could not be efficiently accommodated by the telephone network's ''circuit switched" technology (see ch. 2, p. 29). Packet switched networks are now widely deployed, Internet services are being offered by the private sector, and the Internet protocols are becoming world standards. In much the same way, the NREN program is intended to catalyze the deployment of a new generation of network technology.

Past government programs have also been successful in broadening access to networks for the larger research and education community. The Internet is increasingly essential to users in the academic community beyond the original core group of users in engineering and computer science. It is now estimated that over 600 colleges

5 For a description of the goals and characteristics of the NREN see HPCA, op. cit., footnote 1, Sec. 102(a)-(c); OSTP, op. cit., footnote 2, p. 18; U.S. Congress, Office of Technology Assessment, *High Performance Computing & Networking for Science,* OTA-BP-CIT-59 *(Washington DC:* U.S. Government Printing Office, September 1989), p. 25.

6 OSTP, op. cit., footnote 2, p. 18; Office of Science and Technology Policy, "The National Research and Education Network Program: A Report to Congress, " December 1992, p. 2.

7 Robert E. Calem, "The Network of AU Networks," *The New York Times, Dec. 6,* 1992, p, F12.

8 National Science Foundation "Public Draft: Network Access Point Manager/Routing Authority and Very High Speed Backbone Network Services Provider for NSFNET and the NREN Program," June 12, 1992; James F. Leighton, Manger of Networking and Engineering, National Energy Research Supercomputer Center, Lawrence Livermore National Laboratory, 'ESnet Fast-Packet Services Requirements Specification Document,'' Feb. 20, 1992.

and universities and an estimated 1,000 high schools are connected to the Internet.[9] As the Internet user community becomes more diverse, there is a growing need for simplifying the applications and their user interfaces.

This background paper primarily describes gigabit NREN applications and network technologies. There are, however, several controversial policy issues related to the NREN program.[10] First, the scope of the NREN is uncertain. As a key component of the HPCC program, a clear role of the NREN is to serve scientists and engineers at Federal laboratories, supercomputer centers, and major research universities. This objective will be met primarily by upgrading the networks operated by the National Science Foundation (NSF), Department of Energy (DOE), and the National Aeronautics and Space Administration (NASA). However, there are several different visions of the extent to which the NREN program should also serve a broader academic community, such as libraries and schools.

A second major issue concerns the "commercialization" of the NREN. The NREN will develop from the current Internet, which is increasingly used by government and businesses, not only by the research and education community. Several new commercial providers have emerged to offer Internet services to this market, which is not served by Federal agency networks. One of the goals of the NREN program is to continue this commercialization process, while at the same time achieving the science and network research goals of the NREN program. There has been considerable uncertainty about the mechanisms by which this objective is to be achieved.

The High-Performance Computing Act does not clearly specify the scope of the NREN or the mechanism for commercialization. NSF has had to address these issues in the course of developing a plan for the development of its network, which will be a central component of the NREN. These debates have slowed considerably the process by which NSF will select the companies that will operate its network. NSF's original plan, released in the summer of 1992, is undergoing significant revisions (see box 5-A). As of May, 1993, a new plan had not been issued. It is increasingly unlikely that NSF will be able to deploy its next-generation network by the Spring of 1994, as was originally planned.

In addition, the growing commercial importance of networking is leading to greater scrutiny of the agencies' choices of contractors to operate their NREN networks. DOE selected a contractor for its component of the NREN in thes summer of 1992, planning to deploy the new network in mid-1993. However, a losing bidder protested DOE's selection to the General Accounting Office (GAO). In March, 1993, GAO overturned DOE's choice of contractor and recommended that DOE revise its solicitation, conduct discussions with potential contractors, and allow contractors a new opportunity to bid.[11] DOE has

---

9 Darleen Fisher, Associate Program Manager, National Science Foundation, personal communication, Feb. 11, 1993.

10 For issues related to the NREN program, see Hearings before the House Subcommittee on Science, Mar. 12, 1992, Serial No. 120.

11 The dispute concerned the parties' interpretation of certain provisions in DOE's Request for Proposals (RFP). AT&T protested DOE's selection of Sprint to be the contractor for the DOE network, arguing successfully that the RFP had specified more fully-developed switches than had been proposed by Sprint as part of its bid. GAO ruled that the switches that Sprint planned to use did not comply with a provision in the RFP that proposals had to "conclusively demonstrate current availability of the required end-to-end operational capability," DOE, by contrast, was satisfied that the switches had been developed to the level envisioned by the RFP and were appropriate to a program designed to explore leading-edge technology.

DOE's RFP had specified the use of "cell relay' technology, which is the basis for both synchronous Transfer Mode (ATM) and Switched Multimegabit Data Service (SMDS) services. ATM is expected to play an important role in the future development of computer networking and the telecommunications industry, while SMDS is viewed primarily as an intermediate step towards ATM. DOE selected Sprint in large part because Sprint proposed to begin ATM services immediately, while AT&T bid a service based on SMDS and evolving to ATM only in 1994. Early deployment of ATM would have provided a valuable opportunity to evaluate and demonstrate a key telecommunications industry technology. Comptroller General of the United States, Decision in the Matter of AT&~ File B-250516.3, March 30, 1993.

asked GAO to reconsider its decision. The DOE example raises questions about the effect of government procurement procedures on the ability of federal agencies to act as pioneers of leading-edge network technology. The additional time that would be required to comply with GAO's recommendations, added to the seven-month GAO process, would delay deployment of DOE's network by over a year.

The Testbeds

The HPCC program's six gigabit testbeds (table l-l) are intended to demonstrate emerging high-speed network technologies and address unresolved research questions. While each testbed involves a different research team and is emphasizing different topics, there is similarity in their approach. The testbeds typically consist of a high-speed network connecting three or four sites -universities, industry laboratories, supercomputer centers, and Federal laboratories-with high-bandwidth optical fiber. Located at each of the testbed sites are computers, prototype switches, and other network components. Each research group has both network and applications researchers-the applications will be used to test different approaches to network design.

Significant progress has been made toward the development of gigabit technology.

The testbed program is administered by NSF and the Advanced Research Projects Agency[12] (ARPA). Five of the testbeds are jointly funded for 3 years by NSF and ARPA under a cooperative agreement with the Corporation for National Research Initiatives (CNRI). The principals of CNRI, a nonprofit organization, played significant roles in the development of both the Arpanet and the Internet.[13] CNRI is responsible for organizing the testbeds and coordinating their progress. Funding for the testbeds is modest, when compared to their visibility and the overall HPCC budget. The cooperative agreement with CNRI is for $15.8 million over 3 years. Most of the cost of building the networks has been borne by industry, in the form of contributions of transmission capacity, prototype switches, and research personnel.

The testbeds are investigating the use of advanced network technology to match the needs of the NREN. There is an emphasis on delivering the highest bandwidths possible to the users and demonstrating the range of applications that would be used by leading-edge users of the NREN. Most of these applications are supercomputer-based. For example, some applications use the network to link several supercomputers, allowing their combined processing power to compute complex simulations more rapidly. Many of the applications being investigated also use the network to enable visualization of the results of simulations or experiments.

Initially, only a few users would have computers powerful enough to need a gigabit network. However, the processing power of lower cost workstations and ordinary desktop computers is likely to continue to increase rapidly, as a result of advances in microprocessor technology. Gigabit networks and the lessons learned from the testbeds will then be used more widely.

SUMMARY

Progress

Significant progress has been made toward the development of gigabit network technology since 1987, when the Office of Science and Technology Policy (OSTP) noted that considerable research would be needed to determine the design of

12 Formerly the Defense Advanced Research Projects Agency (DARPA).

13 Dr. Robert E. Kahn is President of CNRI; Dr. Vinton G. Cerf is Vice President.

Table I-l-Gigabit Testbed Participants

Testbed	Location	Industry	Federal laboratories	Supercomputer centers	Universities & other
AURORA	Northeast	IBM Bellcore Bell Atlantic NYNEX MCI			MIT University of Pennsylvania
BLANCA	Nationwide	AT&T	Lawrence Berkeley Laboratory	National Center for Supercomputing Applications	University of Illinois University of Wisconsin University of California-Berkeley
CASA	Southwest	MCI Pacific Bell U.S. West	Jet Propulsion Laboratory Los Alamos National Laboratory	San Diego Supercomputer Center	California Institute of Technology
NECTAR	Pittsburgh	Bellcore Bell Atlantic		Pittsburgh Supercomputer Center	Carnegie Mellon University
VISTAnet	North Carolina	Bell South GTE		North Carolina Supercomputer Center (at MCNC)	University of North Carolina-Chapel Hill North Carolina State University MCNC
MAGIC	South Dakota Kansas Minnesota	Sprint MITRE Digital Equipment Corp. Southwestern Bell Northern Telecom Split Rock Telecom SRI International	U.S. Army Future Battle Laboratory U.S. Army High-Performance Computing Research Center U.S. Geological Survey Lawrence Berkeley Laboratory	Minnesota Supercomputer Center	University of Kansas

SOURCE: Corporation for National Research Initiatives (CNRI), Advanced Research Projects Agency (ARPA).

gigabit networks.[14] There has been growing consensus within the technical community on many issues, and the development of the optical fiber links, switches, and other network components is underway. The testbeds represent the next step in the research-integrating the hardware and software components into a working network system and testing it with applications.

The basic characteristics of the design of broadband networks began to emerge in the mid-1980s, supported by the results of simulations and small-scale experiments. Researchers' objective was to develop networks that could support high bandwidths and were also sufficiently flexible to support a range of services. One characteristic of these networks is the use of optical fiber links, which have the necessary capacity to support many new services, including bandwidth-intensive video- and image-based applications. The second major characteristic of the proposed designs for advanced networks is the use of ''fast packet switches, ' a new type of switch that has both the processing power to keep up with increases in link bandwidth and the flexibility to support several kinds of services.

The testbeds have established a useful model for network research.

As these ideas began to emerge, computer and telecommunications companies initiated the development of the network components required for broadband networks. There appear to be no significant technological barriers to the development of the components required for the gigabit NREN. Transmission equipment of the type that would be required for the gigabit NREN is already becoming available commercially and is being used in the testbeds. Some fast packet switches are also becoming commercially available. Versions of these switches that operate at gigabit rates are in prototype form and will be incorporated in the testbeds over the coming year.

The testbeds are looking to the next step in the research-the development of test networks. This is a systems integration task-developing the individual components is only part of the process of building an advanced network. There is often much to be learned about making the components work together and solving unforeseen problems in the implementation. In addition, there are research questions that can only be investigated with a realistic test network. The testbeds will provide a way to test various proposed approaches to network design.

Progress on the testbeds has been slower than expected, due to delays in making the transmision equipment available and in completing work on the switches and other components. Switches are complex systems, requiring the fabrication of numerous electronic circuits. It was originally hoped that the optical fiber links could be deployed and the gigabit switches and other components finished in time to have a year to experiment with the working testbed networks before the end of the program in mid-1993. It now appears that the testbeds will not be operational until the third quarter of 1993. The testbed program has been extended to permit a year's research on the testbed facilities once they become operational.

Testbed Concept

The testbeds have established a useful model for network research. The design and construction of a test network fills a gap between the earlier stages of the network research-small scale experiments and component development—and the deployment of the technology in production

14 Office of Science and Technology Policy, "A Research and Development Strategy for High Performance Computing," Nov. 20, 1987, p. 21.

networks. The testbed networks model the configuration in which the technology is expected to be deployed—the test sites are separated by realistic distances and the networks will be tested with applications of the type expected to be used in the gigabit NREN. In addition, the participants in the testbeds will play important roles when the networks are deployed.

The testbed research contributes in a number of ways to a knowledge base that reduces the risks involved in deploying advanced network technology. First, there are a number of research issues that are difficult to address without a working network that can be used to try different approaches. Second, the systems integration process provides experience that can be applied when the production network is constructed. In many ways the experience gained in the process of getting the testbeds to work will be as valuable as any research done with the operational testbeds. Third, the testbeds serve to demonstrate the utility of the technology, which serves to create interest among potential users and commercial network providers.

The relatively small amount of government money invested has been used primarily to organize and manage the testbeds and to encourage academic involvement. The testbeds have mainly drawn on other government and industry investment. The organization of the testbeds as a collaborative effort of government, academic, and industry groups is essential, because of the many disciplines required to build and test a network. Industry has contributed expertise in a number of areas. For example, it would be too difficult and expensive for academic researchers to develop the high-speed electronics needed for the switches and other components. Academic researchers are involved in the Internet community, and have contributed ideas for new protocols and applications. Other applications work has come from a number of scientific disciplines and the supercomputer community.

One of CNRI's main contributions was to encourage the involvement of the telecommunications carriers in the testbeds. The transmission facilities required for the testbeds are expensive because of the long distances between the testbed sites and the demands for very high bandwidth. Most experimental work in the past was on small scale networks in a laboratory, due to the prohibitive cost of linking distant test sites. However, the carriers are installing the required transmission capacity and making it available to the testbeds at no cost. All three major interexchange carriers (AT&T, MCI, and Sprint), and most of the Regional Bell Operating Companies (RBOCs) are playing a role in the testbeds.

The testbed research overlaps with industry priorities in some areas and not in others. The basic design of the networks—the types of switches and transmission equipment—reflects emerging industry concepts. However, much of the research agenda focuses on higher bandwidths and more specialized applications than will be used with commercial broadband networks in the near term. Only a few users will use the types of supercomputer-based applications being emphasized by the testbeds. Of greater near-term commercial importance to industry are medium bandwidth ''multimedia' applications that require more bandwidth than can be supported by current networks, but significantly less than the gigabit speeds required by the supercomputer community.

Application of Testbed Research

The testbed research is applicable both to the NREN and to other networks. The NREN will serve only the research and education community and is best viewed as only part of the broader national information infrastructure.[15] The scope of the national information infrastructure will in-

[15] For one view of the relationship between the NREN and the "National Information Infrastructure,'' see Michael M. Roberts, "Positioning the National Research and Education Network" *EDUCOM Review*, vol. 26, No. 3, summer 1991, pp. 11-13.

elude both the United States' part of the Internet and a wide array of other services offered by the computer and information industries, the carriers, the cable television industry, and others.

APPLICATION TO NREN

During 1992, DOE, NASA, and NSF published plans for the future development of their networks, a key component in the evolution to the gigabit NREN.[16] Some aspects of these plans are still unclear; for example, NSF has left to prospective bidders the choice of switching technology, from among those being investigated by the testbeds and elsewhere. However, the agency plans appear to be consistent with the target established by the testbeds. Initially, the agency networks will operate at lower bandwidths than the testbed networks, but they will incorporate more of the testbed technology as they evolve over time to meet the goal of the gigabit NREN.

The rate of NREN evolution is less dependent on technology issues than on delays in the selection of service providers.

Today, the highest bandwidth of the agency networks is 45 Mb/s; it appears that they will move to 155 Mb/s in 1994, with 622 Mb/s the highest rate that is realistically achievable by 1996. The rate of evolution is less dependent on technology issues than on delays in the process by which the Federal agencies select suppliers of NREN network services. Because agency choices of technologies and suppliers have broad implications for the Internet and the national information infrastructure in general, there have been several disputes over agency plans (see p. 7). While the NREN program has created a high level of interest in advanced networks, further delays in the deployment of agency networks may reduce the degree to which they will play the role of technology pioneers.

The agency networks' evolution depends in part on the timely deployment of the necessary high bandwidth transmission infrastructure by the telecommunications carriers. Computer networks generally use links supplied by the carriers-the network operators do not normally put their own fiber in the ground. The carriers' networks already have gigabit-capacity fiber installed, but today the capacity is usually divided among thousands of low-bandwidth channels used for telephone calls. New transmission equipment, the electronics at each end of the fiber, is required to allow the fiber's capacity to be divided into the high-bandwidth channels needed by the gigabit NREN. This equipment is being used in the testbeds and is becoming available commercially, but is very expensive.

The testbed applications research helps researchers to understand how the NREN would be used to achieve the science goals of the overall HPCC program. For example, some of the testbed applications show how networks can be used to bring greater computer power to bear on complex simulations such as the Grand Challenge problems. They may also show how networks can be used to help researchers collaborate-the Grand Challenge teams are expected to involve scientists at widely separated locations. In 1992, the NSF supercomputer centers proposed the concept of a ''metacenter, which uses a high-speed network to link the computing power of the four NSF supercomputer centers.

The testbeds do not address all of the technology issues that are key to the future development of the NREN. Because the NREN will develop from the federally funded segment of the current Internet, it is affected by issues related to the growing number of users of the Internet. This growth in the number of users is straining some of the Internet protocols, and their future development is a topic of intensive study and debate

16 NSF, Op. cit., footnote 8; Leighton, op. cit., footnote 8.

within the Internet community. Also, the testbeds are not looking at applications that would be used by a broad range of users in the *near* term, or at issues related to making the Internet applications easier to use.

OTHER NETWORKS

One of the roles of the NREN is to serve as a testbed in itself, demonstrating technology that will then be deployed more broadly in the national information infrastructure. The testbed program will also impact the evolution of the national information infrastructure more directly, bypassing the intermediate stage of deployment in the NREN. This is because the network technology used in the testbeds reflects near-term industry planning. While the testbeds have emphasized higher bandwidths and more specialized applications than are of immediate commercial importance, the testbed networks reflect ideas that figure prominently in industry plans and, wherever possible, use equipment that conforms to emerging standards.

For example, many of the testbeds use a switching technology called Asynchronous Transfer Mode or ATM. This technology has become central to telecommunications industry planning because it is designed to support many different kinds of services-today's telephone network switches are limited mainly to carrying ordinary telephone calls. ATM can support Internet-type services such as will be used in the NREN, and also video, voice, and other data communications services-the carriers plan to use ATM to enter a variety of markets. Although ATM has been widely accepted by the telecommunications industry and progress has been made towards its implementation, there are a number of unresolved research issues. The testbeds are providing a large-scale opportunity to test this technology and possibly provide input to the standards process.