Gigabit Research

4

As was discussed in chapter 3, the limitations of current networks and advances in computer technology led to new ideas for applications and broadband network design. This in turn led to hardware and software development for switches, computers, and other network components required for advanced networks. This chapter describes some of the research programs that are focusing on the next step-the development of test networks.[1] This task presents a difficult challenge, but it is hoped that the test networks will answer important research questions, provide experience with the construction of high-speed networks, and demonstrate their utility.

Several "testbeds" are being funded as part of the National Research and Education Network (NREN) initiative by the Advanced Research Projects Agency (ARPA) and the National Science Foundation (NSF). The testbed concept was first proposed to NSF in 1987 by the nonprofit Corporation for National Research Initiatives (CNRI). CNRI was then awarded a planning grant, and solicited proposals or white papers" from prospective testbed participants. A subsequent proposal was then reviewed by NSF with a focus on funding levels, research objectives, and the composition of the testbeds. The project, cofunded by ARPA and NSF under a cooperative agreement with CNRI, began in 1990 and originally covered a 3-year research program. The program has now been extended by an additional fifteen months, through the end of 1994. CNRI is coordinat-

The HPCC program's six testbeds will demonstrate gigabit net-working.

---

[1] Corporation for National Research Initiatives, ''A Brief Description of the CNRI Gigabit Testbed Initiative," January 1992; Gary Stix, ''Gigabit Connection,'' *Scientific American,* October 1990, pp. 118-1 19; John Markoff, "Computer Project Would Speed Data,'' *The New York Times,* June *8,* 1990, p. Al; "Gigabit Network Testbeds,'' *IEEE Computer,* vol. *23, No. 9,* September 1990, pp. 77-80.

ing five testbeds; a sixth testbed, funded by ARPA alone, was announced in June of 1992.

The testbeds are investigating gigabit networks, very high-speed broadband networks that represent the limit of what can be achieved today. Most current work on broadband networks is looking at lower bandwidths, such as the 155 Mb/s rate that will be used for the telephone companies' B-ISDN service. Because of the focus on gigabit rates, some aspects of the testbeds' research agenda are unique. In other respects, however, the testbeds are one of a number of research programs whose work will impact the NREN-fast packet switching technologies, for example, are being studied as part of many industry research projects.

RESEARCH OBJECTIVES

In general, the objective of the testbeds is to speed the deployment of advanced network technology, in the NREN and elsewhere. The networks are designed to provide a realistic test environment for the technologies outlined in the previous chapter. The switches and transmission equipment conform to emerging industry standards wherever possible. More speculative concepts such as optical switching are not being investigated by the testbeds—the focus is on the network technologies that are central to near-term industry planning. One purpose of the testbeds is to look at unresolved research questions. However, the most valuable aspect of the testbeds will be to demonstrate the feasibility of these networks and provide experience with their construction.

While much of the research is related to near-term industry plans, the testbeds are also looking into the future. The testbed networks achieve the highest bandwidths possible, given the constraints of emerging industry standards, current technology, and the time horizon of the program. The equipment used in the testbeds had to be such that it could reasonably be expected to be working in time to integrate the components and begin testing the networks by the end of the project. The applications are the most bandwidth-intensive possible, "gigabit applications" that require a full gigabit of bandwidth for each user. For the most part, these are distributed supercomputing applications that use the network to combine the processing power of multiple supercomputers.

The research is also related to the expected use of the network technology in the NREN environment. This emphasizes the use of Internet protocols with the new fast packet switching technologies, because the NREN program is linked to the evolution of the Internet. In addition, supercomputer-based applications of the type being investigated by the testbeds will play an important role in the gigabit NREN. However, not all issues relevant to the future development of the NREN are addressed by the testbeds: because of the emphasis on high-speed applications there is little work being done on applications that will be used outside the supercomputer community. Nor is there significant work being done on topics related to the growing size and complexity of the Internet (see ch. 2, p. 26, and ch. 5, p. 70).

Given the objective of demonstrating the feasibility of the emerging network design concepts, the testbeds are emphasizing the construction of working networks-much of the prior network research used modeling or simulation in "paper studies. Because there is little real experience with broadband networks, these models and simulations are based on assumed traffic patterns that may not be accurate. The testbeds are addressing this problem by building test networks and investigating both network and applications research simultaneously. The applications will provide a source of traffic with which to test the network components and protocols.

In addition, there is a focus on overall systems performance. The overall performance of a network depends on how well the individual components work together, not solely on the performance of any single component. In the past, researchers have tended to focus on the design of individual components; for example, some have looked mainly at switch design, others at transmission systems, and others at protocol issues. In part, this has been due to the complexity of organizing research programs such as the testbeds that draw on the collaboration among several disciplines.

The five CNRI testbeds are AURORA, BLANCA, CASA, NECTAR, and VISTAnet, and are discussed in more detail in boxes 4-A to 4-E. The sixth testbed, MAGIC, is described in box 4-F.

▮ Testbed Design

Each testbed is building a high-speed network that addresses wide area networking issues. The networks connect three or four sites—industry research laboratories, universities, Federal laboratories, and supercomputer centers-separated by anywhere from about 30 to many hundreds of miles. The focus on wide area networks provides a realistic testbed for the agency backbones and the public switched network. In the past, much of the research done on advanced networks has involved small "local area networks. " These served to demonstrate the basic concepts and could be investigated by a small research group within a laboratory. The development of high-speed wide area networks is much more difficult, both technically and organizationally.

The testbed networks reflect the basic technology trends outlined in the previous chapter. The networks all use optical fiber transmission and fast packet switching. There is major emphasis on the use of the telephone companies' Asynchronous Transfer Mode (ATM) concept—five of the six testbeds use ATM in some fashion. One of the testbeds also uses Packet Transfer Mode (PTM), a second kind of fast packet switching, and is investigating the relationship between ATM and PTM. Industry standard equipment is used wherever possible-the transmission links conform to the current version of the Synchronous Optical Network (SONET) standard, and the switches and other components that process the ATM cells conform as closely as possible to the current versions of the international standards.

In order to focus on the systems issues, an effort was made to draw on component development work that was already underway when the testbed program started in 1990. This would limit the extent to which components had to be specially developed and allow more time to experiment with protocols, applications, and other issues related to the operation of the overall network. Because fiber optic technology is the most advanced part of the system, the testbeds are able to use early production models of SONET transmission equipment, operating at 622 Mb/s or 2.4 Gb/s. The switches, on the other hand, are mainly prototypes, as are the interfaces between the computers and the networks-before the testbed work focused attention on the issue of interconnecting different network elements, network interfaces received less attention than such areas as switch or protocol design.

At each testbed site are computers, switches, and network equipment. Computing resources available on the testbeds include workstations, vector supercomputers, massively parallel supercomputers, and some specialized processors. In some cases this equipment is connected directly to the wide area network; in other cases it is connected through a local area network. The local area networks are using newly emerging gigabit-per-second standards such as the supercomputer community's High Performance Parallel Interface (HIPPI) or pre-standard experimental technologies. A number of different interface devices are being developed to handle the conversion between the local area and wide area network protocols, especially the HIPPI to ATM conversion.

Of particular interest is the investigation of the use of networks to enable collaboration between scientists and bring to bear increased processing power on a scientific simulation. Many of the applications also use the network to support visualization or interactive control of a simulation executing on a distant computer. Scientists and

Box 4-A–AURORA

The AURORA network links four sites in the Northeast: the University of Pennsylvania in Philadelphia; Bell Communications Research (Bellcore) in Morristown, NJ; IBM's T.J. Watson Research Center, in Hawthorne NY; and the Massachusetts Institute of Technology (MIT), in Cambridge, MA (figure 4-A-l). Bellcore is the research arm of the Regional Bell Operating Companies(RBOCS) that provide local telephone service in much of the United States.

The testbed sites are connected by 622 Mb/s SONET channels. The transmission facilities are provided by three different carriers: interexchange links are provided by MCI, local exchange links to IBM and MIT are provided by NYNEX, and local exchange links to the University of Pennsylvania and Bellcore are provided by Bell Atlantic.

Figure 4-A-l—AURORA

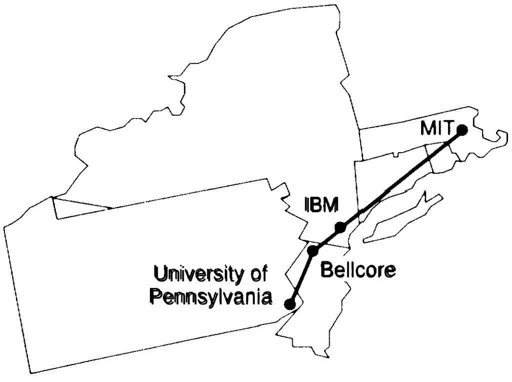

SOURCE: Office of Technology Assessment, 1993.

Each node will have experimental fast packet switches, which can either route traffic to a local area network on the testbed site or to another node. The local area networks will then distribute traffic to workstations, video monitors, and other devices. A number of network interfaces have been built to allow the workstations to connect to the local area networks and SONET transmission links. Bellcore and IBM are also supplying equipment for use in multimedia and videoconferencing applications.

AURORA is unique in two respects. First, it will employ two different switching technologies. Bellcore is contributing an ATM switch, based on the telecommunications industry standard that uses small, fixed length packets called cells. IBM is contributing a switch based on a second fast packet switching technology called Packet Transfer Mode (PTM) (part of IBM's "plaNET" network architecture). The PTM switch was designed to support a network architecture based on variable sized packets; it can, however, also handle ATM cells.

One of the research issues will be to compare the two types of switching technologies and to explore ways in which the two technologies can work together. In the current Internet, networks based on a wide variety of underlying technologies are used. Because both PTM and ATM maybe used in future networks, it is important to gain understanding of how traffic could best be exchanged between these two networks. This work represents an initial step towards gigabit inter networking.

AURORA is also unique in that it is not investigating distributed supercomputing applications. Instead, it emphasizes experimentation with high-speed "multimedia" applications. Because video streams do not in general require a full gigabit of bandwidth, one concept is to deliver a gigabit stream consisting of a large number of medium-bandwidth video signals. For example, the network could be used to support an electronic classroom in which a user could select from different views of a classroom demonstration.

SOURCES: Biereck et al., "Gigabit Networking Research at Bellcore," *IEEE Network, vol. 6, No. 2,* **March 1992, pp. 42-48; Cidon et al., "Bandwidth Management and Congestion Control in plaNET,"** *IEEE Communications, vol. 29, No.* **10, October 1991, pp.** *54-64.*

other researchers are developing applications in a number of areas, such as climate modeling, chemical modeling, and space science. Because, in the long run, scientists will want to develop applications without having to learn all of the details of the network and computers' operation, a number of modules and programs are being developed that simplify the task of applications development in a distributed computing environment.

The protocols generally conform to the existing Internet protocols, the protocols that will be the most widely used in the NREN. The use of well-understood, standard protocols also allows applications researchers to concentrate on applications development. The testbeds will provide a way to test the behavior of the Internet protocols in high-speed networks and to explore their use in a fast-packet-switched environment. However, the testbeds will also be testing a number of experimental protocols that may perform better with new network technologies. This research may serve to test ideas that will be incorporated in the Internet protocols in the future.

▮ Testbed Organization

One of CNRI's key roles has been to assemble the testbed teams. The testbeds draw on researchers in industry, universities, supercomputer centers, and Federal laboratories. Some researchers within the groups have experience with traditional telecommunications issues, while others are more familiar with issues related to the Internet or supercomputer networking. The testbed research is necessarily multidisciplinary. In particular, each research group involves both network and applications researchers. The applications researchers have experience with supercomputers, visualization, graphics, and a variety of scientific disciplines. Network researchers draw on expertise with switches, transmission equipment, protocols, signal processing, and computer architecture.

While regular meetings are held between CNRI and program managers at ARPA and NSF, most of the responsibility for the management of the testbed program lies with CNRI. For example, one of CNRI'S functions was to help develop the specifications for the transmission equipment that would be used in the testbeds. CNRI has also been responsible for maintaining the technical direction of the project, and has held a number of meetings on specific technologies. In addition, there have been annual meetings, which include attendees from a wider group than just the testbed participants, such as workstation manufacturers and government agencies, in an attempt to relate the testbed research to other industry activities and the broader NREN program.

One of CNRI's main contributions has been to ensure the participation of the carriers and other industrial partners. Participation of industry is essential to meeting the research goals of the project. First, the expertise required to develop many of the components required for high-speed network research is only available in industry. These components are complex, and their development involves the fabrication of custom integrated circuits and high-speed circuit design. Second, industry involvement has lowered the cost to the government of the program. The components developed by industry and the transmission capacity between the testbed sites have been contributed at no cost. Because of the contributions of industry, ARPA and NSF's support through the cooperative agreement with CNRI only covers a small part of the total cost of the project.[2]

There are a number of issues associated with the participation of industrial partners in the research venture. Some of these concerns are legal-there are antitrust issues, and further regulatory constraints govern the telecommunications industry. Another factor has been the competitive relationship among the testbed participants-while participating in the same research project, they are also competitors in various lines of business. For example, the wider use of more sophisticated telecommunications industry services may not necessarily be in the interests of companies that have emerged to offer computer networking services.

Moreover, some aspects of the research do not reflect industry priorities. Because of the cost of true gigabit access, it has been estimated that it

---

[2] **Stix, op. cit., footnote 1, p. 118.**

Box 4-B–BLANCA

Figure 4-B-1 —BLANCA

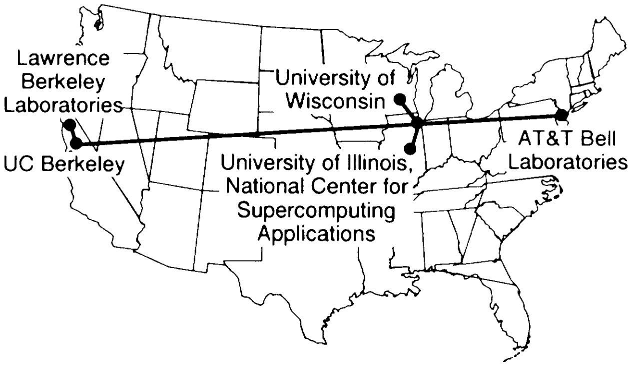

SOURCE: Office of Technology Assessment, 1993.

The sites on the BLANCA network are more widely separated than those of the other testbeds. The network links AT&T Bell Laboratories in New Jersey, the University of Wisconsin and the University of Illinois, and the University of California-Berkeley and Lawrence Berkeley Laboratories (figure 4-B-l). Because of the cost of gigabit transmission facilities, high-speed links will initially be used only for some parts of the network. The cross-country segments of the network will use 45 Mb/s T3 links. While this bandwidth is not sufficient for distributed Supercomputing applications, the BLANCA network will still provide an environment for researching the behavior of new protocols in a large network

BLANCA is an ATM-based network. The experimental ATM switches and other hardware are being supplied by AT&T Bell Labs, the main industrial partner for BLANCA. BLANCA builds on preexisting research relationships between Bell Labs and the University of Wisconsin, University of Illinois, and UC-Berkeley. The switches are designed in such away as to allow researchers at the universities to "take over" the network, to control the switches with computer programs that implement their experimental protocols. The network research interests are similar to those of others looking at ATM-based networks, such as congestion control and the behavior of internet protocols in an ATM-based network, and is being carried out primarily at UC-Berkeley, the University of Illinois, and the University of Wisconsin.

BLANCA emphasizes distributed supercomputing applications, as do most of the other testbeds. The applications work is being done at the National Center for Supercomputing Applications (NCSA), the University

would not be generally available to commercial customers until about 2005.[3] Much of the research agenda focuses on higher bandwidths and more specialized applications than are expected to have near-term commercial significance for the telecommunications industry. Industry planning is oriented more towards medium-bandwidth multimedia applications-applications that require more bandwidth than can be supported by current networks, but significantly less than the gigabit/second rates required by the supercomputer community. For example, the telecommunications industry's ATM-based Broadband Integrated Services Digital Network (B-ISDN) standard envisions 155 Mb/s channels to each customer in the near term. Furthermore, many of the interesting issues related to the operation of fast packet networks can be studied with lower bandwidth networks, although a few issues may only become apparent at gigabit/second speeds.[4]

TESTBED PROGRESS

The major research results of the testbeds are still to come. Most of the networks are not

---

[3] M. Niel Ransom and Dan R. Spears, "Applications of Public Gigabit Networks, ' *IEEE Network,* **vol.** 6, No. 2, March 1992, p. 30.

[4] Leonard Kleinrock, "The Latency/Bandwidth Tradeoff in Gigabit Networks, " *IEEE Communications, vol. 30, No. 4,* **April 1992, pp. 36-40.**

of Wisconsin, and Lawrence Berkeley Laboratories. A significant part of the work involves the development of software packages and modules that make it easier for scientists to use distributed supercomputing applications. For example, NCSA has been developing modules that handle many of the networking functions; these free scientists of the need to learn all the details of t he network's operation-they can simply incorporate the modules in their applications. Another project is developing a digital library that allows the user to control the retrieval and processing of data--one of the programs that can be accessed by this digital library handles visualization processing, for example.

The applications under development as part of BLANCA could be viewed as prototypes for the Grand Challenge problems to be investigated under the HPCC program. One important aspect of these problems is that they will require collaboration between geographically dispersed researchers. The network and computing environment could support this collaboration by providng facilities for videoconferencing. On a more sophisticated level, researchers at NCSA have developed a program that permits collaborative investigation of data. It permits a researcher to highlight a feature in the data displayed on a workstation screen; researchers at other sites would then see the same feature highlighted on their displays.

The Grand Challenge problems will also involve very large data sets. Processing the data into image form is computationally intensive, especially when it is necessary to view the data interactively. The University of Wisconsin and NCSA are investigating the use of high-bandwidth connections from a scientist's workstation to a supercomputer to provide the necessary computational resources for visualization processing.

A radio astronomy application being studied as part of the BLANCA project is looking at issues involved in visualizing large data sets. Arrays of radiotelescopes collect the data, which is then sent through the network to a supercomputer. A user at a workstation connected through a high-bandwidth network to the supercomputer can control the processing of the raw data into images, which are then sent through the network to the workstation for display.

SOURCES: BLANCA Annual Report Charles E. Catlett, "In Search of Gigabit Applications," IEEE *Communications, vol. 30, No.* 4, April 1992, pp. 42-51; Larry Smarr and Charles E. Catlett, "Metacomputing," Communications *of the ACM, vol. 35, No. 6,* June 1992, pp. 45-52; Carolyn Duffy Marsan, "Gigabit Network at Siggraph Proves Need, Reveals Limits," Federal *Computer Week,* vol. *6, No. 22, Aug. 3,1992,* p. 1.

expected to be operational until the third quarter of 1993. After the initial planning stage, the testbed work during 1990-92 was mainly devoted to completing hardware development for the switches and interfaces, theoretical and simulation work on protocols, and development of the applications software and tools. The next step will be to integrate these components into a working network; this will occur in stages over the next few months. As the networks become operational, researchers will be able to begin addressing the unresolved research questions.

Work on the testbeds has been proceeding more slowly than expected. It had been hoped that there would be about a year to experiment with functioning networks before the end of the original 3-year program. Because most of the networks were not yet operational, a 15-month extension was granted in order to allow time to look at network-level issues and test the networks with applications. The delay has been due to the late availability of the transmission equipment and problems with the fabrication of switches and other hardware components.

Component Development

During the first 2 years of the testbed project, the participants have been working mainly on the completion of the individual network components. The SONET transmission equipment has taken longer than expected to become available, but is currently being tested and, in some cases,

Box 4-C--CASA

Figure 4-C-1 -CASA

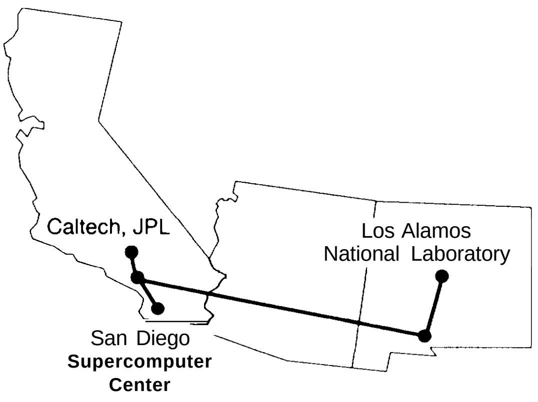

SOURCE: Office of Technology Assessment, 1993.

The CASA network connects four sites--the San Diego Supercomputer Center, Los Alamos National Laboratory, and the Jet Propulsion Laboratory and Caltech in Pasadena (figure 4-C-1). The links between the testbed sites are provided by MCI, Pacific Bell, and US West. Also participating in the project is the UCLA Atmospheric Sciences Department.

The main emphasis of the CASA project is on applications development-using the network to combine the processing power of multiple supercomputers. The three main applications under investigation are modeling of climate change, modeling of chemical reactions, and interactive visualization of data describing the Earth's crust. These applications all require more processing power than is available from a single supercomputer. For example, the CASA climate change model is limited to simulations of a decade or less in the current computing environment One of the research issues concerns the partitioning of a computation among multiple supercomputers. While in theory multiple computers can be combined in order to solve a problem more quickly, the best way to allocate parts of the computation to different computers depends on a number of factors. it maybe necessary, for example, to arrange the computation in such away as to hide the time it takes for data to travel between the computers-even when traveling at speeds dose to the speed of light, data can take a significant amount of time to travel from one computer to another. Efficient implementations would arrange the computation so that the supercomputer would be able to proceed with other calculations while waiting for data to arrive.

Efficient implementations may also be able to take advantage of the strengths and weaknesses of different supercomputer architectures. For example, researchers have determined that the climate modeling application can be split into a number of parts, each of which executes fastest on a particular kind of supercomputer. The part of the simulation that models oceans could be executed on a massively parallel computer, while the atmosphere would be modeled by a more conventional vector supercomputer. The two models would then exchange temperature information and other data at regular intervals. The CASA network provides access to a wide variety of supercomputer architectures, including different types of Cray Y-MPs, and massively parallel machines from Thinking Machines and intel.

SOURCES: The CASA Gigabit Network Testbed, brochure; Charles E. Catlett, "In Search of Gigabit Applications," IEEECornrnun/cat&rrs, vol. 30, No. 4, April 1992, pp. 43-44; Elizabeth Smith, "NREN-Computer Expressway of the Future," *Gather/Scatter,* p. 3.

installed in the carrier networks. While the development of this equipment did not present any research issues, its availability was subject to factors affecting vendor development schedules. In part, these were hardware and software engineering issues. However, other factors have played a role; for example, the SONET equipment is very expensive and it is "high end' compared to the equipment that vendors expect to be the bulk of early demand. In addition, some aspects of the SONET standard have taken longer to complete than expected.

The development of the switch prototypes had been underway when the testbed work began, but in some cases the testbeds presented a more aggressive research target. The interfaces that

Box 4-D-NECTAR

The NECTAR network consists of a high-speed link that connects two local area networks, at Carnegie Mellon University and the Pittsburgh Supercomputer Center. The fiber links are being supplied by Bell Atlantic, and Bellcore and CMU are collaborating on the hardware design.

One area of research focuses on the interconnection of high-speed local and wide area networks. The NECTAR local area networks conform to anew standard called HIPPI (High Performance Parallel Interface), while the wide area connection between the two sites will use ATM cells over a SONET link. Research on this configuration is important because HIPPI is expected to be widely used by the supercomputer community, and the telephone companies are expected to deploy ATM- and SONET-based networks. A better understanding of the interactions between the two kinds of networks is expected to support future distributed supercomputing applications. As part of the research, a special interface circuit that converts between the HIPPI and ATM/SONET formats is being developed.

Another area of research is the implications of new high-speed networks for computer design. Most of today's computers were not designed in such a way as to optimize the task of moving data to and from the network. Applications may not be able to take advantage of improvements in the network bandwidth without improvements in the internal hardware or systems software of the computer. The NECTAR researchers are investigating different approaches to delivering data from the network to applications at high speed. Part of this effort has involved the construction of special "interface" circuits that free the computer's main processor of some of the protocol processing tasks.

Software development for the applications has been proceeding in parallel with the development of the hardware components. The applications have been tested in the local environment and it is hoped that the applications can be made to work on the gigabit network with a minimum of modifications when it becomes operational. The applications are distributed supercomputing applications that take advantage of the combined power of multiple supercomputers to reduce the time needed to solve a problem, The NECTAR network will connect a number of different computers, including the workstations, the experimental iWarp parallel computer, and a variety of machines at PSC.

SOURCES: NECTAR annual report; H.T. Kung, "Gigabit Local Area Networks: A Systems Perspective," *IEEE Communications, vol.* **30, No. 4, April 1992, p. 79.**

connect the computers to the network, or connect local and wide area networks, were designed specifically for the testbeds. Delays in the development of these components are due to their complexity and the demands of high-speed electronic design. A switch, for example, consists of a number of subsystems, each with a large number of standard and newly designed integrated circuits.[5] At the end of 1992, the custom integrated circuits had been designed, and most of the subsystems tested. The PTM switch to be used in the AURORA testbed has been completed, and the other switches and interfaces should be completed shortly.

To the extent possible, much of the work on protocols has been proceeding in parallel with the hardware development. This is expected to lead to faster research results once the networks become operational. Some of the work on protocols is conceptual and theoretical, and is done by simulation or by mathematically modeling the flow of data through a network. One of the main reasons for building the testbed networks is to test the assumptions that underlie these models and

---

[5] **For a description of the components of one prototype switch see Biersack et. al., "Gigabit Networking Research at Bellcore, "** *IEEE Communican'ens,* **vol.** *6, No. 2,* **March 1992, p. 47.**

Box 4-E-VISTAnet

One of the VISTAnet objectives was to use emerging public network technology and standards wherever possible. The switches and transmission equipment were supplied by the local telephone companies, Bell South and GTE, and are early production models from major equipment vendors. In contrast to the other testbeds, the switches are located not at one of the research sites but in telephone company central offices.

The VISTAnet network research involves Bell South, GTE, MCNC, North Carolina State University, and the computer science department at UNC-Chapel Hill. As with the other testbeds, a number of interface circuits have been developed. One of these boards also has the capability to collect data on traffic patterns, which will be used to develop more accurate traffic models for network research. This technology has been licensed, and may soon be available as a commercial product.

VISTAnet differs from the other testbeds in its focus on a single application. MCNC and UNC are collaborating on the development of an application that uses a network of powerful computers to help doctors plan cancer treatments. In radiation therapy, a number of treatment beams are used to deliver radiation to a tumor. Planning the orientation and strength of these beams is essential, because of the need to deliver radiation precisely and avoid the surrounding tissue. Planning involves calculating the distribution of radiation patterns forgiven beam strengths and orientations.

One limitation of existing technology is that the treatment planning is typically done only in two dimensions--doctors are only able to look at the distribution of radiation on a "slice" of the patient's body. The VISTAnet system, on the other hand, would be able to display data in three dimensions, giving doctors a much better view of the distribution of radiation.

Another limitation of current technology is that the planning process is not "interactive." Using today's technologies, interactivity is possible only at low resolution-generating a high-resolution image takes too long and is done "off-line." VISTAnet is developing a system in which doctors can immediately seethe effects of varying a parameter such as beam strength, allowing doctors to examine a greater range of treatment plans.

The interactive display of 3-D images of radiation distribution is computationally intensive. VISTAnet is using a high-bandwidth network to combine the processing power of a supercomputer, medical workstations, and a PixelPlanes machine, a special graphics processor developed at the University of North Carolina. A user at the workstation sends a description of the proposed beam strengths and orientation to a Cray supercomputer at MCNC, which then calculates the corresponding distribution of radiation. Data resulting from this computation is then sent to the PixelPlanes graphics processor at UNC-CH, which generates the image data that shows the radiation distribution as a 3-D image superimposed on the patient's body. The image data is then sent to the workstation for display. Much of the software has been developed and tested on low-speed versions of the VISTAnet network.

The VISTAnet application is a good example of the many different disciplines required to developadistributed computing application. The medical component draws on expertise at the School of Medicine at UNC-CH. The application also requires the development of a user interface that allows doctors to rotate the image or highlight certain features. The graphics algorithms required to interactively generate 3-D volumes are themselves an important research area

SOURCES: Dan Stevenson, ed., "VISTAnet Annual Report" April 1992; B.E. Basch et al., "VISTAnet: A BISDN Field Trial:" *IEEE LTS, vol. 2, No. 3,* **August 1991, pp. 22-30; M. Nisi Ransom and Dan R. Spears, "Applications of Public Gigabit Networks,"** *IEEE Network, vol. 6, No. 2,* **March 1992, p. 32; Daniel S. Stevenson and Julian G. Rosenman, "VISTAnet Gigabit Testbed,"** ***IEEE Journal on Selected Areas in Telecommunications*****, vol. 10, No. 9, December 1992, pp. 1423-1420.**

Box 4-F-MAGIC

The MAGIC testbed is similar in many respects to the five CNRI testbeds, in that a high-speed network is used to provide access to supercomputing resources. As in the CNRI testbeds, there is considerable involvement of industry, the Internet protocols will be used, and the telecommunications services will conform to emerging industry standards like SONET and ATM. The application that will be used to test the network technology is of direct interest to the Department of Defense.

The participants in MAGIC are the Earth Resources Observation Systems Data Center, U.S. Army High-Performance Computing Research Center, the U.S. Army's Future Battle Laboratory, U.S. Geological Survey, Minnesota Supercomputer Center, SRI International, Lawrence Berkeley Laboratory, U.S. Department of Energy, MITRE, Digital Equipment Corp., the University of Kansas, Sprint, Southwestern Bell, Northern Telecom, and Split Rock Telecom.

The MAGIC network will connect four sites, the University of Kansas in Lawrence, Kansas, the U.S. Geological Survey in Sioux Falls, South Dakota, the U.S. Army's Future Battle laboratory in Fort Leavenworth, Kansas, and the Minnesota Supercomputer Center in Minneapolis, Minnesota (figure 4-F-l). In the first phase of the project, the sites will be connected with point to point, 155 Mb/s or 622 Mb/s SONET circuits. In the second phase of the project, the network will use an ATM switch. The SONET and ATM services will be provided by Sprint.

Figure 4-F-1—MAGIC

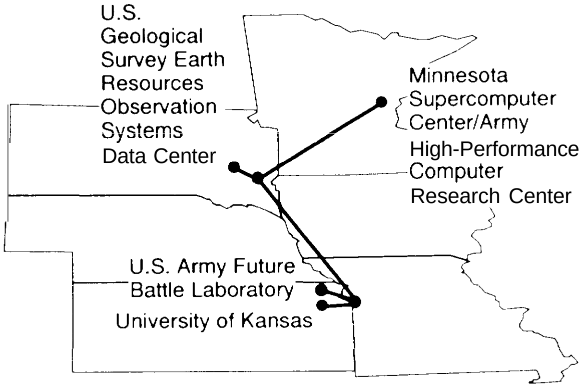

SOURCE: Office of Technology Assessment, 1993.

One of the research issues is the interconnection of different types of gigabit local area networks. Three different types of local area networks will be connected through the ATM wide area network. As part of the research effort, new modules will be built that convert from the local area network technology to ATM, and allow the interconnection of the different networks.

The application will investigate remote visualization of data drawn from a number of different sources. Information from a database at the U.S. Geological Survey will be sent through the network to a massively parallel supercomputer at the Minnesota Supercomputer Center. The supercomputer will compute images based on this data, and send the image data through the network to the Future Battle Laboratory, where it will be displayed on a workstation. The supercomputer provides the necessary processing power to select and view the images interactively (see the discussion of the VISTAnet application in box 4-E).

The test application will allow the simulation of walking or flying through a representation of a landscape. The Army believes "that field officers could benefit from this capability, and that the application could be used for planning and educational activities. The landscape images are created from aerial images, satellite data, and geographic elevation data. Researchers will also study user interfaces to this type of application.

SOURCES: Anita Taff, "Sprint to Provide Services for DARPA Research Net," Network World, vol. 9, No. 27, July 6,1992, p. 9; Tim Wilson, "Group Plots Gigabit Networking," Communications Week International, July 6, 1992, p. 35.

simulations. The protocol research also involves evaluating the behavior of existing networks like the Internet and writing software that will be used to program the switches, computers, and interfaces.

Work on the distributed supercomputing applications has also been proceeding in parallel with the hardware development. Much of the software development for the applications has been completed. In many cases, it has been possible to test these applications to a limited extent using existing high-speed local area networks or low-speed wide area networks like the Internet. Before writing the software, extensive analysis was done of the required computations, to determine how best to divide up the computations among the multiple computers that make up the overall system. Other important software development has involved the development of user interfaces and software tools that would make it easier to program distributed computing applications.

Systems Integration

The next objective of the testbed project will be to combine the network components into an operational network. This will begin once the transmission equipment is in place and work on the switches and other hardware has been completed. The systems integration task will proceed in stages, beginning with the simplest network possible, to minimize the number of sources of possible problems. VISTAnet began the integration process in the fall of 1992; the other testbeds should be in position to start this work by the third quarter of 1993. Over time, the networks will be expanded into more complex configurations.

The issues addressed in the early part of the systems integration phase are the low-level details of making sure that components designed by different groups work together or that a signal arrives in the format expected by a component's designer. These are the kinds of problems that are difficult to find when components are tested individually. For example, when the NSFNET backbone was upgraded from T1 to T3 links during 1990-92, the technical staff of the NSFNET backbone provider found that some components did not behave as expected under certain conditions, or unexpected traffic patterns required changes to the software and hardware.[6] Similar problems will probably be encountered as the testbeds begin to work through this stage with prototype or newly developed network components.

Network Research

One research issue concerns the algorithms used to control fast packet networks. These mechanisms are used to enable fast packet networks to support many different kinds of services using the same links and switches; one of the weaknesses of traditional packet networks was that they could not guarantee the kind of performance required for real-time applications such as video. In a fast packet network, software in the users' computers and in the switches will have to cooperate in managing the flow of traffic through the network in a way that supports all kinds of services. There have been many different mechanisms proposed for accomplishing this objective, but it is regarded as the most difficult problem with fast packet networks. The testbeds will provide an opportunity to test different control algorithms.

Another research issue is related to the development of distributed supercomputing applications. In these applications a computation is divided among multiple supercomputers; the network is then used to exchange data as the computation proceeds. Deciding how to allocate different parts of the computation to different supercomputers is a difficult problem. The best strategy depends in part on the characteristics of the network and the strengths and weaknesses of

---

[6] ''T3 Upgrades to ANS Network Near Completion' *ANS Update*, vol. 1, *No.* 1, 1991, p. 1.

different computers connected to the network—for example, some parts of a computation maybe executed fastest on a massively parallel computer, while other parts may run faster on a vector computer. In order to maximize processing power, computers should not be idle while they are waiting for one of the other computers to finish its task or for data to be sent through the network.