Broadband Network Technology 3

Advances in computer technology are driving the requirements for broadband networks. Because of increases in the processing power of computers, there is a need for higher bandwidth networks. Computers are increasingly able to execute ''multimedia'' applications, so it is expected that future networks must be able to carry several kinds of traffic. Broadband networks will lead to applications that are used for a wider range of problems, with more emphasis on image-based communications.

The computer and telecommunications industries have conceived broadband network designs for these requirements. Fiber optic links are a key component of these networks. However, replacing the smaller capacity links in current networks with higher bandwidth fiber optic links is not all that is needed: Improvements in protocol and switch design must also follow. Future switches will have more processing power, in order to keep pace with the faster flow of traffic through the links. They will also be designed in a way that allows them to handle different types of traffic. Today's switching technologies do not have this capability-packet switches only handle text and numerical data efficiently, the telephone network's circuit switches are best suited to voice traffic, and special networks are needed for video. The ''integrated services' concept envisions networks that use the same links and switches for all types of traffic, instead of different technologies for video, data, and voice.

Broadband networks use new switch technologies.

BROADBAND APPLICATIONS

The new high bandwidth integrated services networks would improve the performance of existing applications and enable new applications. Existing applications, such as electronic mail or databases, could be augmented though the use of image files and video clips; higher bandwidth networks would also allow the faster transfer of large files of supercomputer data. Support for real-time high-resolution video would expand possibilities further, allowing videoconferencing or the display of output from a scientific instrument, such as a telescope. More generally, the combination of more powerful computers and integrated services networks will permit wider use of two new categories of applications—multimedia applications and distributed computing.

Multimedia Applications

Multimedia applications take advantage of the capability of high-bandwidth integrated services networks to deliver many different kinds of data-video, image, audio, and text and numerical data. They also take advantage of the processing power of advanced workstations and other devices attached to the network, allowing users to edit, process, and select data arriving from a variety of sources over the network.[1] Multimedia applications have been suggested for a large number of areas[2], including education and health care. There are many different concepts for delivering multimedia services to the home, such as multimedia catalogues for home shopping, information services, entertainment video, and videotelephone services. Many segments of both service and manufacturing industries are increasingly using image-based applications-for example, computers are widely used in the publishing and advertising industries to compose pages using high-resolution images.

Multimedia is also the foundation for a new category of applications that use the combination of computing and communications to create a ''collaborative' work environment in which users at a number of scattered sites are able to work together on the same project.[3] For example, an application might allow several researchers to work on the same set of experimental data at the same time-any processing done by one researcher would automatically be shown on the other researchers' displays. Videoconferencing and collaborative applications might allow closer interaction between researchers in different places. It is expected, for example, that the teams working on the Grand Challenges will include scientists at many locations.

For researchers, "visualization" provides a way to represent large amounts of data in a more understandable form; it uses images and video to show the results of simulations or experiments (box 3-A). For example, the results of a simulation of a city's air quality could be shown as an image, with the concentration of a particular chemical indicated by different colors and color intensity. If a researcher wanted to review the evolution of the air quality over time, a series of images could be used to create a video segment showing the change in pollutant concentration. Other programs running on the workstation could be used to process the data further, perhaps by examining one part of an image more closely or by comparing the simulation data to experimental data.

In education, multimedia could be used in computer-based instructional materials. Multimedia databases would give students and teachers access to image and video data. Videoconferencing and collaborative applications could enable closer interaction between teachers and students at multiple locations. For example, it might

---

[1] Special Issue: Multimedia Communications, *IEEE Communications*, vol. *30, No. 5, h@ 1992.*

[2] Michael L. Dertouzos, Director, MIT Laboratory for Computer Science, testimony at hearings before the Joint Economic Committee, June 12, 1992.

[3] Sara A. Bly et al., "Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment," *Communications of the ACM,* vol. *36, No.* 1, January *1993.*

[4] *Matthew* Arrott and Sara Latta, "Perspectives on Visualization"*IEEE Spectrum,* vol. 29, No. 9, pp. 61-65,

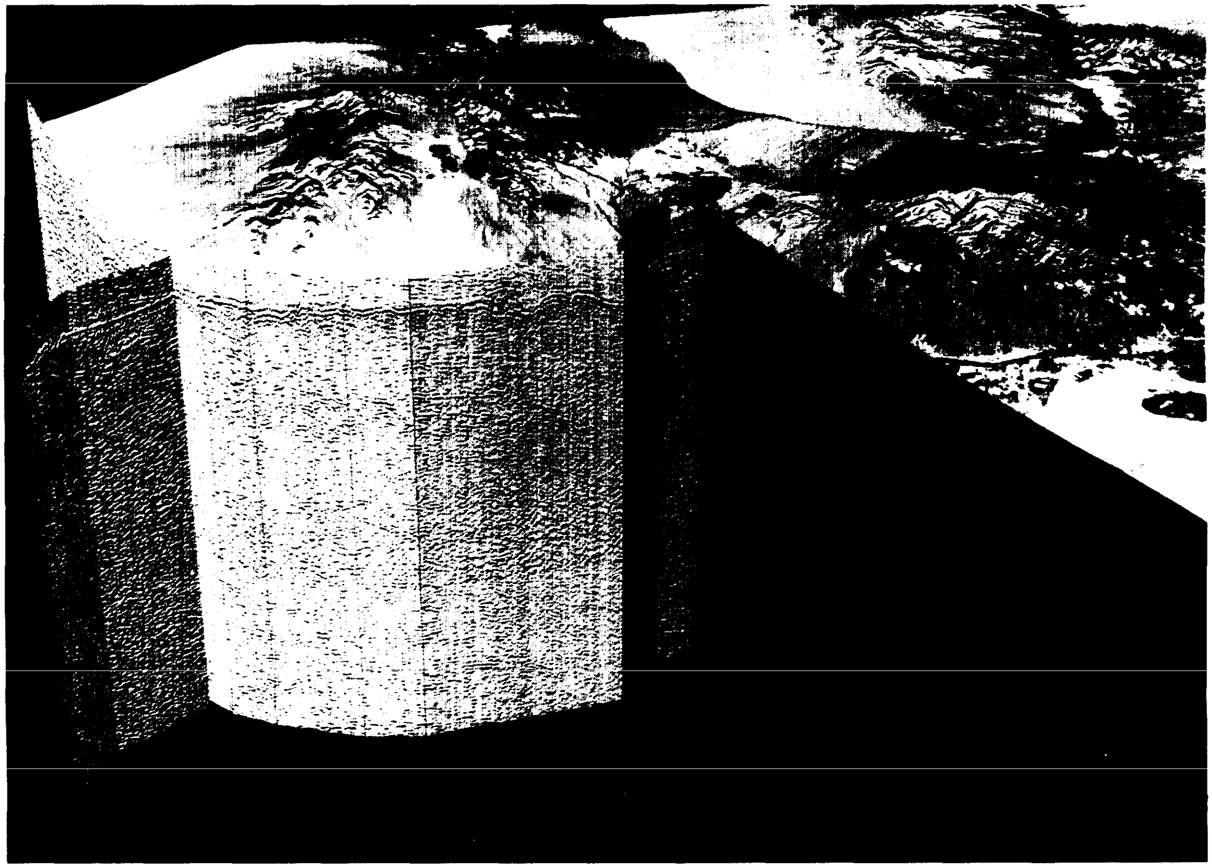

CASA

As part of the CASA testbed research described in chapter 4, a gigabit network will be used to combine data from a variety of sources, such as satellites and digital elevation models, to create three-dimensional views.

be possible to better emulate the classroom environment by allowing more two-way communication than is currently possible. Students might also be able to select a particular view of an experiment being demonstrated by a teacher. In health care, transfer of high-resolution images, such as x-ray and MRI data, combined with videoconferencing and other collaborative applications, could allow doctors to consult with specialists in other areas of the country.[5]

▮ Distributed Computing

Other researchers have begun to consider the relationship between computers and communications in a more general way. "Distributed computing' uses the network to combine the processing power and memory of multiple computers. It is then possible, for example, to combine several low-cost workstations to achieve performance comparable to that of a supercomputer--a very expensive machine to purchase and operate. Computations can also draw on data stored in many different locations. Distributed computing becomes feasible as the network connecting the computers becomes less of a bottleneck, allowing them to work more closely together.

It may also be useful to do distributed supercomputing—using the network to provide proc-

[5] M. Niel Ransom and Dan R. Spears, "Applications of Public Gigabit Networks,' IEEE Network, vol. 6, No. 2, March 1992, p. 31.

Box 3-A-interactive Visualization Using Gigabit Networks

Scientists at the University of Wisconsin's Space Science and Engineering Center have developed a software package that allows "interactive" visualization of data computed by a model. This photograph of a workstation's display shows both a computed forecast of cold fronts moving across the North Atlantic and the "control panel" that the scientist can use to control the images displayed.

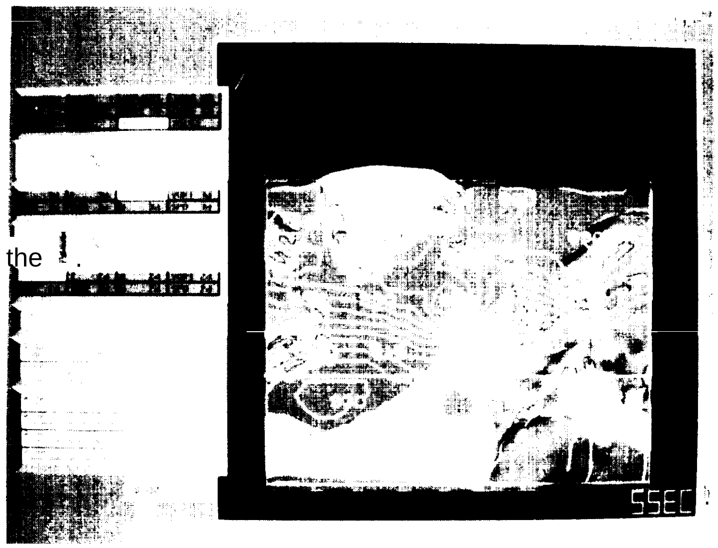

Bill Hibbard, University of Wisconsin.

computers generate new kindss of images, that change in response to their users' needs. If the scientist selects the command "animate," in the upper left hand corner of the control panel, the workstation will display a succession of images that show the evolution of the storm over time. Other commands allow the user to rotate and "zoom" the images, to look at them from any angle.

The time required for a workstation to compute a new image, in response to a user command, can be significant. A supercomputer would be able to reduce the response time and allow interactive exploration of the data computed by the model. However, like many research institutions, the University of Wisconsin does not have a supercomputer.

The Internet could be used to send data to one of the NSF's supercomputercenters. However, the data rate of today's Internet is too low-the advantages of speeding up the computation by using a supercomputer are outweighed by the time needed to transfer the model data to and from the supercomputer. With a gigabit network, the communication time would no longer be a bottleneck.

As part of the BLANCA testbed's applications research (see oh. 4, p. 56), University of Wisconsin scientists will use a gigabit network to support interactive visualization of large data sets. The user's commands would be sent from the workstation through the network to a supercomputer at the National Center for Supercomputing Applications, in Champaign, Illinois, which would do most of the visualization processing and send the image data back to the workstation for display. This testbed research may serve to demonstrate away for the majority of research institutions that do not have supercomputers to do interactive visualization.

SOURCES: William Hibbard, University of Wisconsin--Space Science and Engineering Center, Mar. 11, 1993; William Hibbard, David Santek and Gregory Tripoli, "Interactive Atmospheric Data Access Via High-Speed Networks," Computer Networks and ISDN Systems, vol. 22, pp. 103-109.

essing power that exceeds that of a single supercomputer (see box 3-B). Supercomputer applications are often simulations of real-world phenomena; for example, airplane designers do not need to build a model of an airplane and test it in a wind tunnel, but can simulate the flow of air around the airplane on a supercomputer. Unfortunately, for many interesting problems the process-

[6] Charles E. Catlett, ''In Search of Gigabit Applications, ' *IEEE Communications,* vol. *30, No. 4,* April 1992, pp. 43-45.

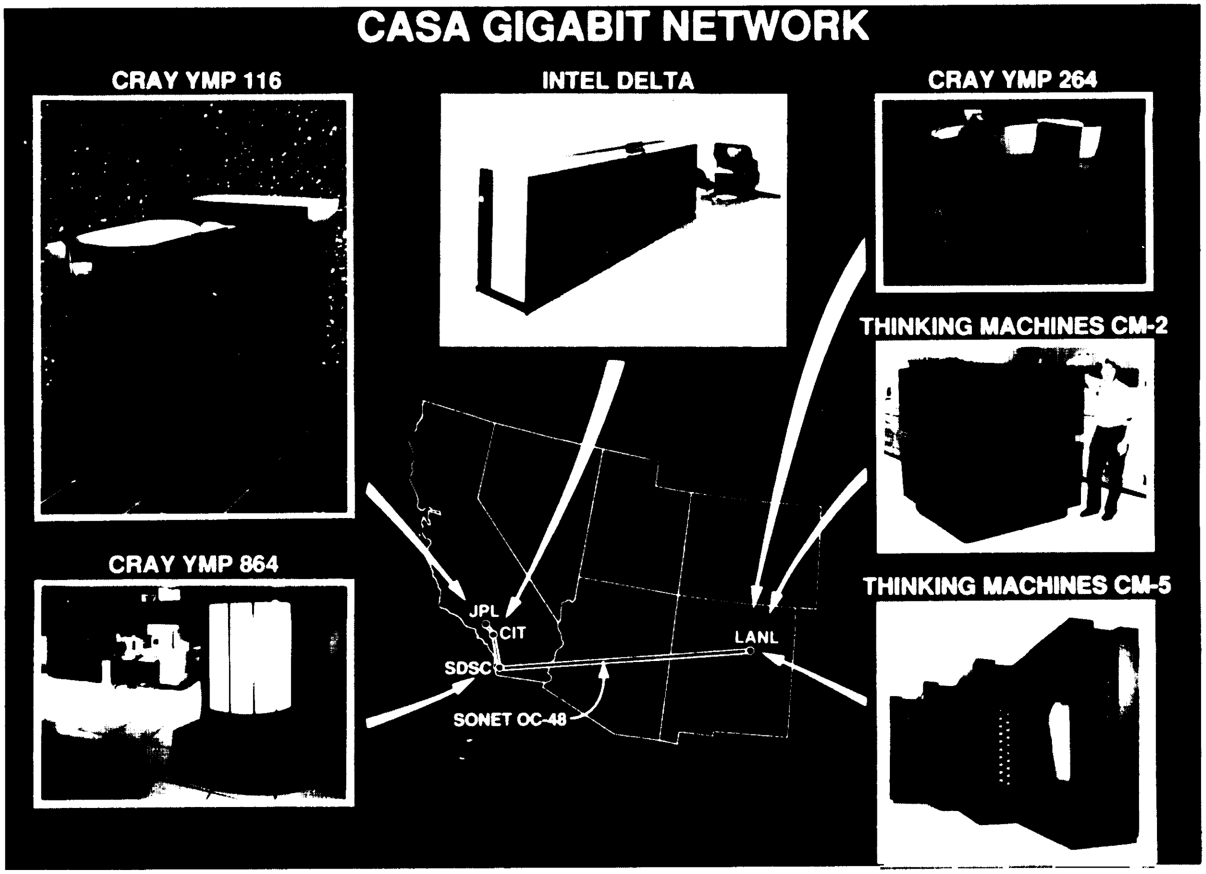

CASA.

The CASA testbed will use a gigabit network to link supercomputers at the California Institute of Technology (CIT), Jet Propulsion Laboratory (JPL), Los Alamos National Laboratory (LANL), and the San Diego Supercomputer Center (SDSC).

ing time required with even the fastest supercomputer can be significant.[7] Researchers hope to reduce this time by using multiple computers in parallel and linking them through a network. The network could also connect the supercomputers to scientific instruments or massive remote databases that would provide data to be used in the calculations.[8]

■ Applications and Network Requirements

Two requirements will be placed on future networks. First, they will need to have much higher bandwidth than today's networks, in order to keep pace with advances in computer technology and support bandwidth-intensive video-based and distributed computing applications. Distributed supercomputing applications would require even greater increases in network bandwidth. Second, the networks will have to be more flexible than today's networks-they will be supporting a more diverse range of services, with a wider range of bandwidth requirements.

HIGHER BANDWIDTH

The bandwidth requirement for each type of application depends on a number of factors.

[7] Office of Science Technology Policy, "Grand Challenges 1993: High Performance Computing and Communications, ' p. 15.

[8] Catlett, op. cit., footnote 6, pp. 46-49.

Box W-Distributed Supercomputing

Supercomputer-based simulations are becoming essential tools for science and engineering. Often, scientists are able to study problems that would be difficult or impossible to study theoretically or experimentally. For example, a number of researchers are developing climate models that can be used to predict the evolution of the Earth's climate over the coming decades. Computational science is becoming more widely used as the increasing power of new supercomputers brings more problems within reach.

One of the goals of the HPCC program is the development of computer technology that will allow scientists to tackle problems that are beyond the capabilities of today's machines. Some problems simply take too long to compute-some current models of the ozone depletion process take 10 hours of supercomputer time to compute the complex chemical reactions that take place in everyday of "real" time. Other problems cannot be studied at a useful level of detail-computer power might limit a climate model to tracking the evolution of temperatures at only a limited number of points on the globe.

Greater processing power is also required for "interactivity," completing a computation in time to allow a user to take some action to control an instrument, change the parameters of a simulation, or "browse" other data sets in a database (see box 3-A). This requires that the computation of the model and the visualization processing be done in a fraction of a second. Today, images can take a considerable amount of time to compute, and are generally processed "off line" after the computation has been completed. Because of the time required to process newly computed or collected data much of the data often goes unused.

The testbed program is investigating the use of gigabit networks to help address difficult computational science problems. High-speed networks may enable increased processing power, by linking several computers through the network. For example, a model could be computed on a supercomputer and then sent to a special graphics processor for the visualization processing, or a model could be split into two parts, with two supercomputers working in parallel to solve the problem. Networks also allow data from multiple sources to be used in a computation--large databases and scientific instruments, for example.

In the testbeds described in chapter 4, distributed supercomputing is used to increase processing power to study long-term weather models (part of the CASA testbed research), molecular dynamics (NECTAR), and chemical modeling (CASA). The use of networked computers to speed up the visualization process in an interactive fashion is being explored as part of applications for medical treatment planning (VISTAnet) and radio astronomy (BIANCA). Navigation of multiple large databases and associated visualization are used for terrain visualization (CASA), atmospheric sciences (BLANCA), and terrain navigation (MAGIC).

The testbeds are also working on the systems software and "tools" that will support these applications. Today, implementing distributed applications requires detailed knowledge of the behavior of the network and the characteristics of different computers. Distributed supercomputing will only be widely used by scientists if they can be freed of the need to learn these details, and can concentrate on the science aspects of their simulations. The testbeds are developing software modules that implement commonly used functions, and programs that automate parts of the applications development process. In the long run, the objective is to create software support and distributed operating systems for a "metacomputer," which would hide the complexity of networked computers and appear to the programmer as a single computer.

SOURCES: Gary Stix, **"Gigabit Connection," Scientific American, October 1990, p. 118; Matthew Arrott and Sara Latta, "Perspectives on Visualization,"** *IEEE Specftum, vet. 29, No. 9,* **September 1992, pp. 61-65; Larry** Smarr and Charles E. Catiett, "Metacomputing," *Communications of the ACM, vol.* 35, No. 6, **June 1992, pp. 46-52; office of Science and Technology Policy, "Grand Challenges 1993: High Performance Computing and Communications," p. 54.**

Because of advances in ''compression" technology, it now appears that relatively modest increases in bandwidth can accommodate many simple video and multimedia applications. There are many ways to convert a video signal to a digital stream of bits; new compression algorithms are able to squeeze the information content into fewer bits without significantly affecting picture quality.[9] These improvements have resulted from a better understanding of the mathematics of signal processing and also from research on how people perceive images .10 In addition, increased processing power due to advances in microelectronics has allowed sending and receiving computers to do more complex signal processing.

Advances in compression technology have been dramatic. While it was once believed that a 155 Mb/s fiber optic link could carry only a single high-definition television (HDTV) signal, it is now believed that such a link can carry five or six HDTV signals.[11] In addition, it now appears that many simple video and multimedia applications will not require broadband fiber access to the network. New compression techniques are able to compress VCR-quality video to a few megabits per second, bandwidths that can be supported by new schemes for converting the telephone companies' existing copper local loops to digital service[12] There are a number of emerging viola and image compression standards-the most prominent of these will be the HDTV standard to be chosen by the Federal Communications Commission in 1993.[13]

However, there are still many possible applications that would more fully use the capacity of fiber.[14] These are the kinds of applications that are being investigated by the testbeds described in chapter 4. One possibility is distributed supercomputing-the use of high-bandwidth links to combine the processing power of multiple supercomputers. There are also applications that require images or video of a quality that can be only supported by fiber, despite advances in compression technology. In some cases, such as some medical applications, compression cannot be used because it destroys vital data. Other applications may demand very high bandwidths because many medium bandwidth streams of data are delivered to the user at once, allowing the user to select, combine, or process the streams at the workstation.[15] "Telepresence" or "virtual reality' applications require the delivery of large amounts of data in order to create the illusion of a user being in a distant location.

FLEXIBILITY

The second requirement the envisioned applications place on advanced networks is flexibility. First, new network technologies should be sufficiently flexible to carry all kinds of traffic. The integrated services concept envisions a network in which the same links and switches are used, to the extent possible, for all types of traffic. Integrated services networks may be more efficient than separate networks, and would also match advances in computer technology that allow computers to run multimedia applications. Today, different network technologies are used for voice, video, and ordinary data traffic. As new services were required, new types of networks were constructed. The telephone network was

9 Mark Robichaux, ''Need More TV? TCI May Offer 500 Channels,'' *Wall Street Journal,* Dec. 3, 1992, p. B1.

10 P. H. Ang et al., ''Video Compression Makes Big Gains," *IEEE Spectrum*, October 1991, pp. 16-19.

11 See, for a table of bandwidth requirements for compressed signals, J. Bryan Lyles and Daniel C. Swinehart, ''The Emerging Gigabit Environment and the Role of Local ATM,'' *IEEE Communications,* vol. 30, No. 4, p. 54.

12 Charles F. Mason,''Bell Atlantic Stretches Copper for Video Trial,'' *Telephony,* Oct. 26, 1992, p. 10.

13 Mark Lewyn, ''Sweating Out the HDTV Contest,' *Business Week, No. 3306,* Feb. 22, 1993, pp. 92-94.

14 Ransom and Spears, op. cit., footnote 5, pp. 30-40.

15 Lyles and Swinehart, op. cit., footnote 11, p. 55.

augmented first by packet switched networks for data and then by a variety of specialized networks for video communication[16] and distribution, such as cable television networks. Separate networks were required in part because no switching technologies worked equally well with all services.

Broadband networks should also be able to accommodate a range of applications bandwidths, from the very small amount of bandwidth required for ordinary electronic mail to the gigabit rates needed for distributed supercomputing. Some kinds of switching technologies are more flexible than others in accommodating different bandwidths in the same network. Circuit switches, the type used in the current telephone network, limit applications to a small number of predetermined bandwidths, while packet switches are more flexible.

Flexibility is also important from a network planning standpoint. While there are some general ideas about the ways in which broadband networks will be used, there is no real operational experience. Ideally, the network technology that is deployed would be able to accommodate a range of different scenarios, and its effectiveness would not depend on network planners knowing the exact mix of future applications in advance. In addition, future networks will have to support a more diverse range of users, each with different bandwidth and service requirements. Network operators would like to deploy network technology that could provide services to a diverse range of users with a minimum amount of customization.

FAST PACKET NETWORKS

A number of new concepts for network design may meet the requirements for flexible broadband integrated services networks.[17] There is general agreement that these networks will rely on fiber optic transmission, which has sufficient bandwidth to carry video and other types of bandwidth-intensive services. There is also general agreement that future networks will use a concept called "fast packet switching," which provides the necessary processing power to keep up with increases in link bandwidth and the necessary flexibility to support different kinds of services and a range of bandwidth requirements.

Fast packet networks overcome the main weakness of traditional packet networks by using special control mechanisms to provide the consistent network performance required for video and other real-time services. In traditional packet networks such as the current Internet, the network could become heavily loaded in a way that degraded these services. Researchers are looking at a number of different schemes to either prevent networks from becoming too heavily loaded, or to minimize the effects of a heavily loaded network on traffic such as video that is sensitive to network performance. Fast packet switches can then act as the foundation for integrated services networks.

Both the computer and telecommunications industries are investigating fast-packet approaches. In response to the emerging consensus for these technologies, considerable work has been done on the development of the necessary network components such as switches. However, until recently most experience with these networks had been confined to relatively small-scale experimentation with local area networks, or simulation and mathematical modeling. One of the main purposes of the testbeds described in chapter 4 is to demonstrate these networks in a realistic environment. In addition to the testbeds, a number of other experimental fast packet networks are now being planned or are operational.

Two different kinds of fast packet switching are being studied by the testbeds. The most

16 For example, the CONCERT network in North Carolina uses high-bandwidth microwave links to support videoconferencing and "teleclasses."

17 Nim K. Cheung, "The Infrastructure for Gigabit Computer Networks," *IEEE Communications, vol. 30, No. 4,* April 1992, p. 60.

prominent of the fast packet switching concepts was first championed by the telephone companies and is called Asynchronous Transfer Mode, or ATM.[18] ATM has been chosen by the telecommunications industrys international standards group, the CCITT, as the foundation for the Broadband Integrated Services Digital Network (B-ISDN), a blueprint for the future development of the telephone network.[19] B-ISDN envisions the provision of 155 Mb/s or 622 Mb/s fiberoptic access links to each customer, which would then be used to carry voice, video, and data traffic to support a range of applications.[20]

One of the most significant aspects of ATM is that it has subsequently been adopted by many companies in the computer industry, and by manufacturers of equipment for local area networks and private networks. This convergence with telecommunications industry plans[21] may simplify the task of connecting different kinds of networks—in the past, local and wide area networks have used different technologies. However, technologies other than ATM have also been proposed for local area networks. Most of the testbeds described in the next chapter are using supercomputer industry networking standards that require the construction of modules that convert between the supercomputer network format and ATM.

While most packet networks use packets that can be very long and vary in size depending on the data being carried, ATM networks use short packets called ‘‘cells’ that are always the same length (figure 3-l). If an ATM network is being used to carry Internet traffic, the Internet packets would be broken into a series of cells (figure 3-2(a)). After traveling through the network, the cells would be reassembled into the Internet packet and delivered to the destination computer. The same network could also carry video or sound: as the video or sound was digitized, the computer would load the bits into a cell (figure 3-2(b)-(c)). As soon as the cell was filled, it would be sent into the network and the user would begin filling the next cell. The cells carrying the video and Internet packet data would travel through the network together, sharing the same links (figure 3-3) and being processed by the same switches.

Figure 3-1—ATM Cell

Data (48 bytes)	Header (5 bytes)

International standards specify that ATM cells are 53 “bytes” long, 48 bytes to carry the applications data and 5 bytes for the header, which is used to carry such information as the destination’s address. By contrast, the packets used in traditional packet networks can be several thousand bytes long. A “byte” is the computer science term for eight bits.

SOURCE: Office of Technology Assessment, 1993.

The second approach to fast packet switching being studied by the testbeds is called Packet Transfer Mode or PTM. The version being studied in the testbeds has not been adopted by standards committees. PTM does not use short cells, but more traditional packets that can be longer if necessary. This may simplify the task of carrying long Internet packets, because the computer does not have to break up the packet into many cells. ATM may also encounter problems at very high bandwidths-because the cells are so

18 Robert W. Lucky, Executive Director, Communications Sciences Research Division, AT&T Bell Laboratories, testimony at hearings before the Joint Economic Committee, June 12, 1992.

19 A. Day, “International Standardization of B-ISDN,'' *IEEE LTS*, vol. 2, No. 3, August 1991.

20 peter Broke and Heinrich Armbruster, ''Broadband Services: An Overview,'' *Telecommunications*, vol. 25, No. 12, December 1991, pp. 24-32.

21 John J. Keller, “AT&T Sets Alliance to Make Gear to Provide Multimedia Services, ' *Wall Street Journal,* Jan. 13, 1993, p. B6; John McQuillan, “Who’s Who in ATM,'' *Business Communications Review,* August 1992, p. 10.

Figure 3-2—ATM

(a) Data communications using ATM

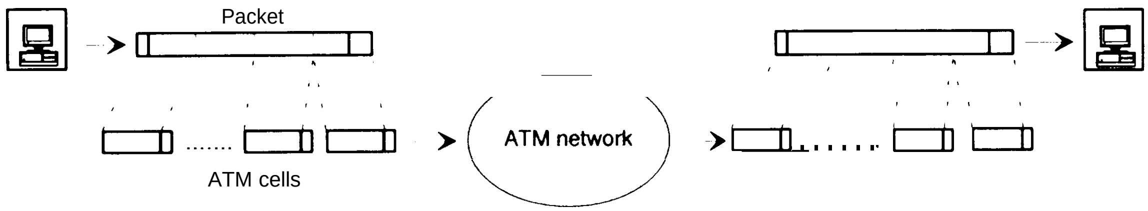

Packets are broken into several ATM cells. After traveling through the network, the cells are reassembled into packets.

(b) Video communications using **ATM**

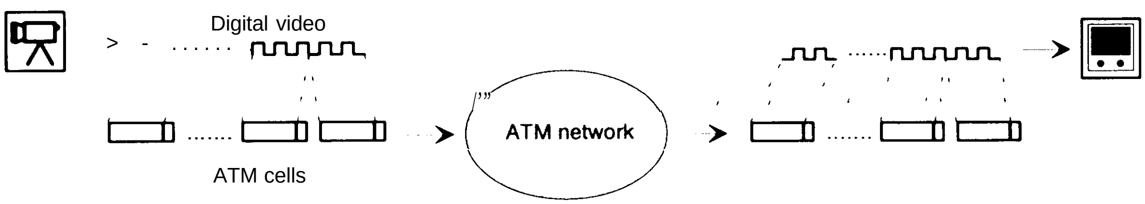

The digital video bits are put in cells and sent through the network. At the destination, the bits are removed from the cells.

(c) Voice communications using ATM

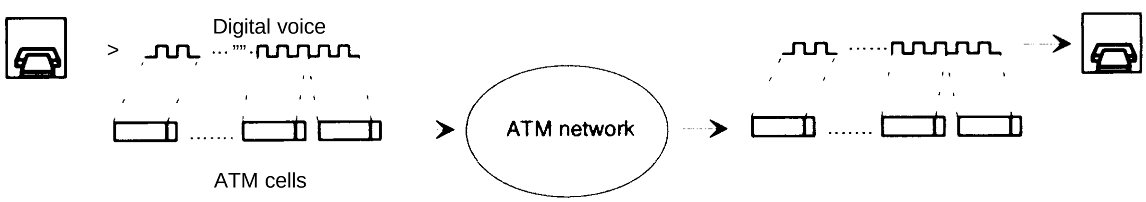

Voice is handled in the same way as video.

SOURCE: Office of Technology Assessment, 1993.

short, there is little time to process each cell before the next one arrives. However, ATM proponents believe that the use of cells makes it easier to develop the control mechanisms that support real-time traffic such as video, and to be better suited to voice traffic. One purpose of the testbed research is to compare the two approaches to fast packet switching with realistic traffic.

NETWORK COMPONENT DEVELOPMENT-CURRENT STATUS

The telecommunications and computer industries have been working intensively to develop components for fast packet networks. The components are in varying stages of development. Fiber optic transmission links are the most advanced in their development—very high bandwidth optical

Figure 3-3-integrated Services Using Fast Packet switching

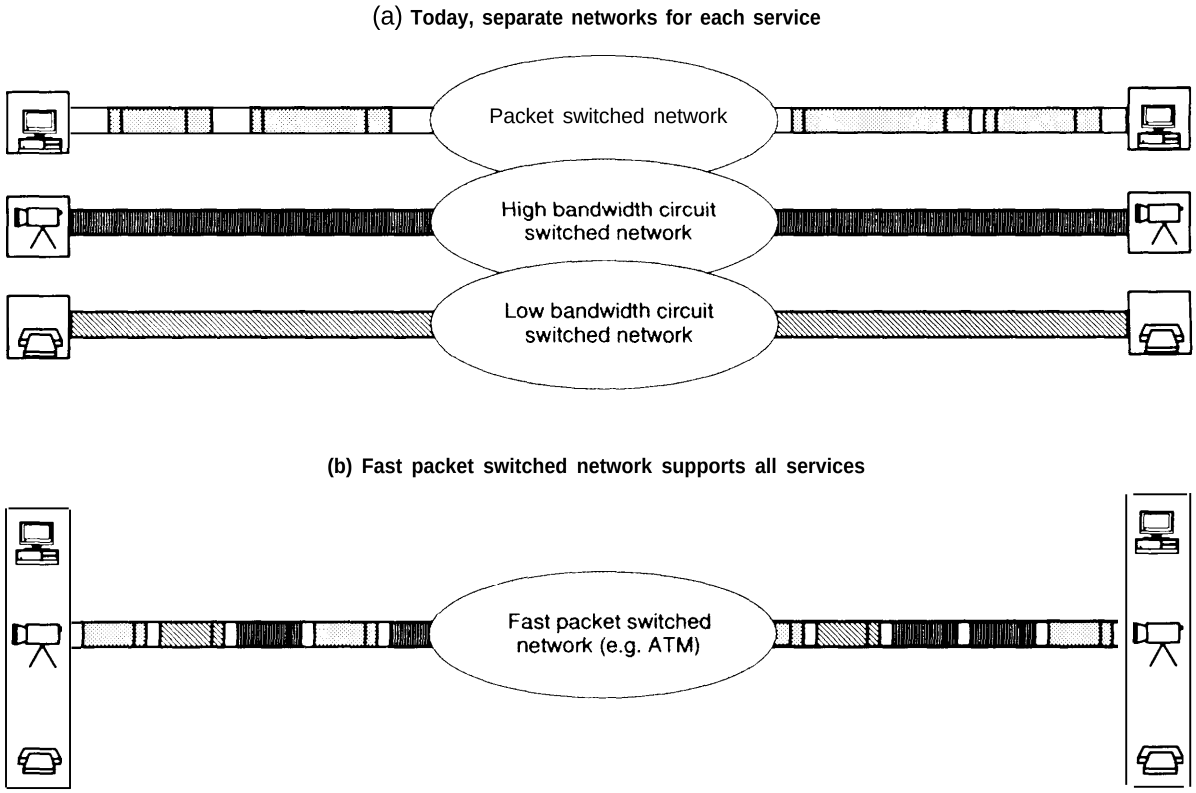

SOURCE: Office of Technology Assessment, 1993.

transmission systems are now commercially available. Fast packet switches are the subject of considerable industry research and development; most of the major telecommunications industry suppliers have had intensive ATM switch development efforts since 1987 or 1988, when it became clear that standards groups were going to adopt ATM. Some fast packet switches are becoming available commercially, but switches are less advanced in their development than the fiber optic links. Important work also remains to be done on the design of the software and hardware that handle the connection between the computer and the network.

Optical Fiber

Optical fiber has clearly emerged as an enabling technology for broadband networks. With increased bandwidth the links will be able to move data more quickly and support the transport of bandwidth-intensive traffic such as video. The development of the transmission equipment that handles gigabit rates is no longer a research issue. Although configured to support voice telephone calls, many fiber optic links in today's telephone network operate at more than one gigabit per second. Furthermore, the fiber cable is already widely deployed in much of the telephone network, especially in the interoffice portions of the network that will provide most of the transmission facilities for the agency backbones and regional networks.

For the telephone company fiber links to support the gigabit NREN and other broadband services, new transmission equipment will have to be deployed. This equipment is expected to conform to a new standard called Synchronous Optical Network (SONET), and is now becoming commercially available. While fiber has been used in the telephone network for a number of years, the link capacity was mainly configured to carry thousands of low bandwidth telephone calls. SONET transmission links, on the other hand, can be configured to support the high-bandwidth channels required for advanced networks. For example, the transmission facilities to be used in the testbeds employ a 2.4 Gb/s SONET link, which can be divided into four 600 Mb/s channels.

Switches

The development of fast packet switches is less advanced than the development of transmission links. However, there has been considerable theoretical work done on switch design, prototypes have been developed, and some early commercial products are becoming available. By the end of 1993 or early 1994, several suppliers should have ATM products on the market. The early products are designed primarily for private networks or carrier networks operating at 45 Mb/s or 155 Mb/s, not gigabit rates. 155 Mb/s is the bandwidth specified by the telecommunications industry's standards group, the CCITT, for the Broadband Integrated Services Digital Network (B-ISDN) service.

There are many different ideas for how to build fast packet switches-the "best" design depends on assumptions about the number of users, the bandwidth of the network, and the mix of traffic. However, all of the proposed switch designs rely on hardware, in order to speed the processing, and are usually based on ''parallel' designs that allow many packets or cells to be processed at once.

If ATM switches do become central to telephone company networks, then there will be demand for large switches to replace the current "central offices' that handle tens of thousands of lines. Most of the ATM switches now becoming available only handle a small number of lines-16 or 64 lines are common configurations. Initially, ATM switches will probably be introduced to support new services, rather than as a replacement for existing central office switches .22 Building an ATM switch that can serve thousands of lines is a difficult task, requiring further research on switch design and device technology and packaging. The move to ATM switching has the potential to change the market positions of telecommunications equipment manufacturers, much as the transition from analog to digital created market opportunities in the late 1970s.[23]

Switches control the flow of packets using considerable software "intelligence." For example, if the network is heavily loaded, a switch may decide to handle video or other performance-sensitive traffic first. Switches may also help prevent the network from becoming too busy—they may prevent a user from sending traffic, or verify that users are not using more than their share of the network capacity. These aspects of the control of the network are still important research issues, however; there are many different proposals for managing fast packet networks. As a result, some of the prototype switches used in the testbeds described in chapter 4 are flexible enough to allow researchers to reprogram the network control mechanisms.

Computers

The use of high-bandwidth links and switches will expose new bottlenecks inside many computers.[24] New computer designs may have to be

22 Richard Karpinski, ''Ameritech Issues End-to-end ATM RFP," Telephony, vol. 224, No. 9, Mar. 1, 1993, p. 14.

23 John J. Keller, ''Telephone Switching Moves Toward Increased Speed,'' *Wall Street Journal, Nov.* 4, 1992, p. B4; Steven Titch, "Northern Keeps ATM Details Slim,'' Telephony, Sept. 28, 1992, p. 13;''NEC Lands ATM Pact for Wiltel,'' *Telephony*, Oct. 26, 1992, p. 8.

24 H.T. Kung, "Gigabit Local Area Networks: A Systems Perspective," *IEEE Communications,* vol. 30, No. 4, April 1992, pp. 86-88.

developed to improve the rate at which data can be transferred from the network, through the computer's internal circuitry, and into memory, where it can be used by the applications software. Both the internal circuitry and the memory of today's computers are limited in the rate at which they can transfer data. In the past, the design and operation of computers has focused on the task of maximizing the processing power once the data is in memory, not the larger problem of maximizing the performance of networked applications.

Computers may also require additional processors or hardware to process protocols. With low speed networks, the computer's main processor was powerful enough to handle the communications functions and still have enough time left to run applications. This may continue to be the case as the processing power of computers continues to increase. However, in some cases it may be necessary to relieve the processor of some of the burden of handling the communications functions. This is likely to be the case with ATM-based networks-because the cells are so short, there are many cells to be processed in a given amount of time. Special ''network interfaces' that speed the protocol processing are being developed for a number of different computers as part of the testbed project discussed in the next chapter.

APPLICATION OF BROADBAND TECHNOLOGIES

The broadband technologies discussed in this chapter will be used in both the Internet and in other networks, such as private networks or the public switched network. The use of broadband technologies in the Internet is linked to their deployment in the public switched network in two respects. First, the Internet will probably continue to rely on the public network's transmission infrastructure. As a result, it is dependent on the rate of deployment and the cost of SONET links. Second, the carriers may use their new ATM-based infrastructure as a way to play a more active role in the computer communications business and offer Internet services.

■ Application to the Internet

The links in the high-speed networks in the core of the Internet are expected to use the SONET-based transmission infrastructure that the telephone companies are planning to deploy. SONET is actually a family of transmission rates-there is 155 Mb/s, 622 Mb/s, and 2.4 Gb/s SONET equipment becoming available now. Users that need access to the Internet at broadband rates will also use SONET for their access links. Large universities and commercial users of the Internet would be able to make special arrangements with their local exchange carrier for the provision of fiber access. The rate at which the carriers will more broadly deploy optical fiber in the local loop depends on the resolution of complex economic and policy issues.

However, many users of the Internet will not require fiber optic access links in the near term. The carriers have proposed several new technologies that would convert existing copper local loops to digital service. These technologies do not support true broadband capabilities, but still represent a significant improvement over the existing analog local loop. They include the Integrated Services Digital Network (ISDN) (see box 3-C for a description of ISDN), High-bit-rate Digital Subscriber Line (HDSL)[25], and Asymmetric Digital Subscriber Line (ADSL) standards. ISDN provides access at 144 kb/s; HDSL provides access at 768 kb/s. These technologies are available on a limited basis from the carriers and are the subject of a number of trials and demonstrations.[26] The pace of their deployment depends

25 Gerald A. Greenen and William R. Murphy, ' 'HDSL: Increasing the Utility of Copper-Based Transmission Networks, *Telecommunications,* August 1992, p. 55.

26 Mason, op. cit., footnote 12.

Box 3-C--ISDN

ISDN (Integrated Services Digital Network) is a telecommunications industry standard for upgrading local loops to digital service. This "last mile" of the network, the wire that connects a telephone network to its customers, is less sophisticated than other parts of the network The core of the telephone network uses high-capacity, digital, fiberoptic links. The local loop, by contrast uses low-capacity, analog, copper wires. This technology is acceptable for ordinary telephone service, but more sophisticated services will require upgraded local loops.

When work began on the ISDN standards in the mid-1970s, it was believed that ISDN would soon be deployed to all of the telephone network's customers. Today, ISDN is used only on a very limited basis, due to delays in completing the standards and several regulatory and economic questions. Because of the delays in deploying ISDN, large business customers found more capable technologies. More importantly, a new vision of the future of the local loop emerged-the wiring of homes with fiberoptic links--and ISDN was no longer viewed as a technology with an important role to play.

There is now renewed interest in ISDN, however, as an "intermediate" step between the current analog local loop and the use of fiber optics. Because of the cost of deploying fiber, it may be many years before significant numbers of homes are connected. ISDN is cheaper than fiber, can be deployed sooner, and, while its capacity is only a fraction of fiber's, represents a significant improvement over the current analog local loop. While ISDN will not become the universal network standard once envisioned, it may play a role in providing better network access to certain groups of users.

For example, one possible application might be telecommuting, which allows employees to work at a desktop computer at home. To connect to the office computers, workers today would need a device called a modem, which lets them send digital computer data over the analog local loop. Common modem standards transmit data at 2,400 or 9,600 bits per second; ISDN, by contrast, provides two 64,000 bits per second (64 kb/s) channels. This would allow videoconferencing of reasonable quality, faster transfer of graphics information, and better quality fax transmission. It would also permit much-improved access to the Internet for home users. Today, good-quality access to the Internet is usually only available to large customers who are able to arrange for special digital access lines to be provided by their local telephone company.

ISDN allows the existing copper local loops to be used for digital service. However, it requires users to buy new equipment for their end of the line, which converts their data to the ISDN format. It also requires that the telephone company's equipment, such as the "central office" switches, be upgraded. Currently, the user equipment is expensive and only one-third of the telephone lines are connected to switches that are "ISDN ready." In addition, ISDN communications are hampered by the fact that different equipment manufacturers have implemented their own versions of ISDN, despite the fact that it was developed to be a standard. In most of the United States, ISDN is not available as a regular service.

However, some progress is being made toward overcoming ISDN'S problems. The industry has a number of initiatives that are intended to encourage the development of ISDN equipment that conforms to a common specification. The Regional Bell Operating Companies, which provide local telephone service in most of the United States, have announced that they are planning to make 56 percent of their lines ISDN-ready by the end of 1994. In addition, the cost of users' ISDN equipment may decline as it becomes more widely used.

Broadband ISDN, which is discussed on p. 46, uses very different technology from "ordinary" or "narrowband" ISDN. Narrowband ISDN is best viewed as a digital upgrade of the telephone network's copper local loop. Broadband ISDN, by contrast, requires fiber optics and Asynchronous Transfer Mode (ATM), a new approach to network design discussed in detail in this background paper. ISDN and Broadband ISDN have little in common other than their names.

SOURCES: James N. Budwey, "h's Time to Get Off the 'POTS'," *Telecfommunications,* August 1992, p. 4; Cindy Skrzycki, " 'Data Highway' Plan Costs May Decline," The Washington Post Jan. 20, 1993, p. G1; Bob Wallace, "Study Raises Concerns About National ISDN," Network World, June 29, 1992, p. 27; Mitchell Kapor, President, Electronic Frontier Foundation, testimony at hearings before the House Subcommittee on Telecommunications and Finance, Jan. 19, 1993; Steve Lohr, "Computer Makers Told to Get Involved In Rules," The New York *Times,* Feb. 23,1993, p. D2.

on resolving standards issues and on business decisions made by the carriers.[27]

Internet traffic may be handled by some of the new fast packet switching systems. As was noted above (p. 43), fast packet networks can carry Internet traffic if the Internet packets are first converted to the appropriate fast packet format—for example, if the Internet packets are broken up into a series of ATM cells. It is likely that other types of switching technologies will also be used. The success of the Internet is due in substantial part to the commonality of protocols that support the technological diversity of the interconnected networks. Some networks will continue to use ''routers," similar to those used in today's backbone networks, while others may employ the new fast packet switching technologies or some of the new data communications services that the carriers may offer, such as Frame Relay or SMDS.

Public Network

In many ways, the most significant aspect of ATM is that it was first championed by the telephone companies and is now a key component of telephone company planning. ATM represents a dramatic change in the design of telecommunications industry networks. Traditionally, the industry has not used packet switches. It used the circuit switches that were ideally suited to carrying voice telephone traffic. The industry standards group, the CCITT, chose ATM because it believed that simply upgrading the existing circuit switched network to higher bandwidths would not provide the necessary flexibility to support future services. ATM is a central component of carrier strategy; they hope to use ATM as the basis for a range of future services, including video, Internet services, and other data communications services such as Frame Relay or SMDS.[28]

ATM's flexibility offers the carriers an opportunity to enter a variety of markets and quickly offer new services with a common infrastructure. However, some believe that ATM's flexibility also means that it is a compromise technology, and that more specialized network technology will continue to play a role.[29] Moreover, there are still important economic considerations for the telephone companies as they determine the best way to evolve from the current network to an ATM-based infrastructure.[30] Both service providers and manufacturers are facing difficult decisions about the timing of their investments and the appropriate migration scenarios.[31] Deployment decisions depend on estimates of future revenues, equipment costs, the viability of competing technologies, and the carriers' investment in their existing networks.

27 Edmund L. Andrews, "'Baby Bells' Rift Threatens An Advanced Phone Service,' *The New York Times, Dee.* 1, 1992, p. D1.

28 Ben Lisowski and Louise Reingold, ''Sprint's Evolution to Broadband ISDN,'' *IEEE Communications,* August 1992, vol. 30, No. 8, pp. 28-30; John Williamson and Steven Titch, "Gazing Toward the Broadband Horizon,'' *Telephony,* Oct. 5, 1992, p. 38.

29 Chin-Tau Lea, ''What Should Be the Goal for ATM?' *IEEE Communications,* September 1992, vol. 6, No. 5, PP. 60-66.

30 Stephen M. Walters, "A New Direction for Broadband ISDN," *IEEE Communications,* vol. 29, No. 9, September 1991, pp. 39-42.

31 Carol Wilson, ''It's Not That You DO It, But HOW You Do It That Counts,' *Telephony,* June 15, 1992, p. 9.