Life After Internet: Making Room for New Applications

Larry L. Smarr and Charles C. Catlett

Over the years there has been tremendous growth on the Internet in the number of users, networks, and the amount of traffic. The kinds of work that can be done has changed from basic remote login, file transfer, and electronic mail, to the creation of animated simulations that can be viewed and manipulated by colleagues in remote locations. Even the user has changed—from primarily computer scientists to scientists, academics, researchers, and others who have no particular expertise in computers (*1*).

The Internet has evolved toward higher performance technology, more users, and more advanced services largely because of the availability of applications that allow for greater productivity. As the network moves toward gigabit/sec technology, advanced applications and new modes of using the network will place additional demands on the system. They will bring change not only to the architecture of the network but also to the policies and management structures which govern its use.

The Evolution of Internet

Since 1980, the Internet has grown from a network of fewer than 50 to over 5,000 sites and from perhaps 1,000 to over 250,000 hosts (*2*). It has brought users access to supercomputers, data, and their colleagues, and has played a major part in bringing computational science to virtually all fields from astronomy and physics to history, art, and economics.

Work that can be done on the Internet has changed from basic remote login, file transfer, and electronic mail, to the creation of animated simulations that can be viewed and manipulated by colleagues in remote locations.

To fully appreciate the changes brought about by the network, it is useful to compare network capabilities in three periods: 1980-85, when general connectivity was available only by modem; 1985-89, after the National Science Foundation Network (NSFNET) was widely available; and 1989-93, when we will be well into the transition to the National Research and Education Network (NREN). Even though these periods span a relatively short time, enormous changes in connectivity, functionality, applications, and technology have occurred and are expected, as shown in Table 1.

1980 - 1985

In the early 1980s, with fewer than 50 sites, the Internet offered electronic mail services to help researchers collaborate, and it was itself a testbed for network and computer science research. The supercomputers were non-UNIX machines then, with typically 64 MBytes memory and fairly primitive software, which the remote user could access through a terminal or a workstation able to emulate a terminal. The functionality offered to supercomputer users was limited to terminal access and file transfer. A user did code development on the supercomputer, including all compiling, debugging, and program execution. Since data were stored at archives at the supercomputer center, there was no widespread need for large, interactive file transfer.

1985 - 1989

When the National Science Foundation supercomputer centers were established in 1985, users accessed the centers through dialup for interactive

login and BITNET if available (or Kermit for asynchronous lines) for file transfer. But when NSFNET connected the centers in 1986 (3), the Internet became the primary method for interactive login and file transfer to supercomputer centers. The functionality provided by BITNET and dialup was now available in a more user-friendly, transparent, and convenient manner.

From 1985 to 1989, the rate of growth of Internet sites, users, and traffic accelerated exponentially as both the number of Internet sites and the amount of Internet traffic at each site increased (4). Typically, university campuses were installing campus networks. The campus network connected departments, and the departmental local area networks (LANs) connected the users' desktops with the campus net. At the same time, networking software such as Massachusetts Institute of Technology's PCIP, the Telnet of the National Center for Supercomputing Applications (NCSA), and similar software from Carnegie Mellon University, Stanford, and Cornell made it possible for PCs and Macintoshes to connect directly to the Internet.

By early 1988, the number of sites had grown to 150 and the backbone, a combination of 56 and 9.6 Kbit/s lines, was significantly overloaded. When the Advanced Research Projects Agency Network (ARPANET) was phased out in spring 1988 (5), its network traffic shifted abruptly onto the already overloaded NSFNET. The resulting network saturation caused such a decline in service that many users stopped using the network. They returned to dialup and to BITNET. Here is an important lesson for the future: users vote with their feet on network services; they have no particular allegiance to the Internet.

By 1989, the widespread use of the UNIX operating system coupled with the increase in Internet throughput and workstation capabilities allowed a more transparent coupling between workstations and supercomputers. Users could edit and store code at their workstations, but compiling and de-

Table 1: Growth of the Internet and its Capabilities for Low-, Medium-, and High-End Users

	1980-85	1985-89	1989-93
Connectivity	<50 Internet sites dialup	900 Internet Sites	10,000 Internet sites
Functionality and Performance	*HIgh* 10 Kbit/s, remote login, small file transfer *Medium* 1.2 Kb/s, dialup login *Low* n/a	*High* 200 Kbit/s, large file transfer (1.5 Mbit/s backbone), *Medium* 10 Kbit/s, remote login, small file transfer, (56 Kbit/s mid-level) *Low* 1.2 Kbit/s, dialup login	*High* 20 Mbit/s, large data streams *Medium* 200 Kbit/s, large file transfer *Low* 10 Kbit/s, remote login, small file transfer
Applications	*High* email *Medium* host login *Low* n/a	*High* 2D visualization *Medium* email *Low* host login	*High* 4D visualization, collaboration tools, distributed supercomputing *Medium* 3D visualization, multimedia email *Low* 2D visualization, email
Network Technology	100 packet/s switches 56 Kbit/s backbone TCP/IP	1500 packet/s switches, 1.5 Mbit/s backbone standard OS TCP/IP routing extensions	1 M packet/s switches 1 Gbit/s backbone binary compatible OS stds protocol extensions mgmt/admin extensions
Supercomputer Technology	64 Mbit/s memory 300 MFLOPS throughput Vector Processor Non-Unix OS	1 Gbyte memory 1 GFLOPS throughput Vector processor Unix OS Parallel processors emerging	16 Gbyte memory 50 GFLOPS throughput Parallel Processors Unix OS New OS emerging

bugging were still done on the supercomputer. Users were able to transmit two-dimensional color raster images (8 bits deep) or to download data sets and analyze them on workstations using visualization software.

At high-end Internet sites—e.g., those with 1.5 Mbit/s connections—users could begin to perform three-dimensional visualizations locally. These consisted of either dynamic two-dimensional graphics or three-dimensional images produced by downloading data and processing graphics at the local workstation.

For many applications, however, the Internet throughput, even at high-end sites, was barely adequate to download supercomputer-generated data for local analysis and visualization within reasonable timeframes (i.e., measured in tens of minutes as opposed to hours).

1989 - 1993

From 1989 through 1991 the connectivity both of the Internet itself and of individual Internet sites has continued to grow, and the view through 1993 suggests more growth in the future. Most major universities and research laboratories were connected as of 1990, and an increasing number of corporate research laboratories are now online. In the period from 1985 through 1989, many new Internet sites and mid-level networks received seed funding from the National Science Foundation. In 1990, similar programs assisted undergraduate institutions. In the next few years emphasis will be placed on connecting community colleges, high schools, and elementary schools.

A second wave in connectivity growth may be expected as a result of the many state networks emerging today. They in turn will provide a vehicle for connecting small businesses and research laboratories to the Internet. This continued growth in sites and users per site will be spurred by growth in per-user network requirements and expectations as the new modes of using the higher-bandwidth Internet emerge.



In late 1990 and early 1991 the NSFNET backbone was upgraded from 1.5 Mbit/s to 45 Mbit/s. Sites connected with 1.5 Mbit/s circuits now experience somewhat better throughput, and supercomputing centers are more tightly coupled with one another over the 45 Mbit/s backbone.

In the next few years we can expect to see many mid-level networks upgrade to 45 Mbit/s and all mid-level networks supporting at least a backbone of 1.5 Mbit/s to major institutions. At least three National Science Foundation supercomputing centers are preparing to perform experiments with users at gigabit/sec speeds. Many campuses and some departments within campuses have installed high-performance local area networks such as fiber distributed data interface (FDDI) and systems based on 800 Mbit/s high performance parallel interface (HiPPI). Most of NCSA's industrial partners, all leaders in various areas of the economy, are installing FDDI networks and are already expressing great interest in the interconnection of the NCSA FDDI and HiPPI networks.

By 1993, the notion of transparent distributed computing and animated visualization will be accepted by both high and mid-range users. Those at the high end will be routinely performing interactive steering of remote computations. Distributed client-server software tools will be commonplace on the Internet, with high-end users running applications that are distributed across several supercomputers and using voice mail and video conferencing.

Just as the UNIX operating system provided a basis for more functionality on the Internet, a more transparent coupling of the distributed computational environment will be achieved through adopting a low-level standard such as the IEEE floating point. When workstations and supercomputers can share distributed tasks in a predictable fashion, vendors will have to use this and other low-level standards, as earlier they had to move to UNIX. As a result, the computing environment will cease to feel like a collection of components and will seem to be one single system.

New Applications

Now that the Internet has become an accepted part of the academic and scientific environment, new ways of working in that environment are emerging. The evolving technologies described here—visualization and analysis applications, distributed metacomputers, new ways of accessing growing repositories of information, and innovative modes of collaboration—all increase productivity, but they also put extraordinary demands on network capacity. As such applications proliferate and become easier to use, they will change not only the way scientists work but the function and structure of the Internet itself.

Remote Visualization and Steering of Processes

The gigabit/sec testbed projects coordinated by the Corporation for National Research Initiatives (CNRI) are exploring applications that require new modes of network use which challenge network capacity. Some of these applications produce peak loads that are several orders of magnitude greater than the average network load, a degree of variability not seen in the past. Understanding the characteristics of this network traffic and developing

methods for supporting such large variability in throughput are some of the major objectives of research in the gigabit testbeds.

VISTANET is a testbed project that is developing applications to aid physicians in planning radiation treatment (6). Using supercomputers and high-performance visualization workstations, a radiation oncologist is able to send as many as 50 simulated radiation beams through a three-dimensional image of a tumor and the surrounding healthy tissue. The dosage delivered by each beam is calculated by a remote supercomputer which returns volumetric data to the workstation to be rendered and displayed while the physician interactively adjusts the beams for maximum dose to the tumor and minimum dose to the healthy tissue. The volumetric nature of the data being exchanged demands data rates of up to 500 Mbit/s.

One of the BLANCA testbed applications includes a collaboration with the Berkeley-Illinois-Maryland Array (BIMA) astronomical imaging group to interactively process radio telescope data (7). Over several hours, a radio telescope collects a spectral data cube corresponding to a two-dimensional image of the sky at a series of radio frequencies. Depending on the telescope, this data can consist of up to 1,024 two-dimensional spectral images gathered in a single observation, data on the order of tens of MBytes. Supercomputers convert this data into visual images and display it interactively on remote graphics supercomputers, allowing the scientist to "fly" through the data cube (8).

Some supercomputer simulations—for example, large-scale hurricane or global climate models—do not proceed fast enough on current supercomputers to warrant real-time interaction with the simulation. Nevertheless, it is important to monitor the simulation at medium-speed (i.e., at a throughput of several Mbit/s) to ensure that the model is progressing satisfactorily. These slower simulations are often played back for analysis. Their output data, sent through visualization pipelines, generates a sequence of images for study. In some cases, however, a single snapshot of the simulation is so large that the workstation memory is insufficient to store enough of the data to process. In such situations the supercomputer is called upon to re-execute major portions of the visualization pipeline and to send two-dimensional images or graphics commands and volume data to the workstation. Data rates for this type of supercomputer post-process analysis can range from several Mbit/s for sending rough images to over a Gbit/s for sending volumetric data.

The evolving technologies increase productivity, but they also put extraordinary demands on network capacity.

As user interfaces become more intuitive, the potential community of users grows. In the example above, the study of global climate change is limited to those who have the means and the expertise to program supercomputers. But some recent developments indicate that change may lie ahead.

Interactive visualization environments are emerging that provide the user with object-oriented user interfaces through which previously developed visualization modules can be linked together. Data is passed through the pipeline, and the resulting images can be viewed immediately. For instance, a user wishing to view data simulating a thunderstorm can use previously developed isosurface generator and colorizer modules to produce a three-dimensional colored simulation of the surface of a storm cloud.

Further development of systems that take advantage of previously constructed modules is needed. The current visualization environments require significant expertise in scientific visualization techniques. These systems must be able to support distributed pipelines as well as modules which can run on several supercomputers and are therefore themselves distributed. In addition, they must be able to take advantage of new collaboration techniques and the advanced network services and protocols now being developed.

High Definition Television (HDTV) may emerge as a significant part of the user interface for visualization in the 1990s. In the simplest case, a high-performance graphics computer at a supercomputer center would create a series of high-resolution images which would be transmitted to an HDTV screen at the user's workstation, thus reducing the data rate necessary for the transfer to as low as 200 Mbit/s with modest compression and 20 Mbits with significant compression, albeit at the expense of image resolution. More commonly, however, computers in various locations linked by a high-bandwidth data stream would create the HDTV visualization.

Distributed Metacomputers

The largest users of today's supercomputers have run out of capacity on single-vector multiprocessor supercomputers and are active participants in experiments to use the network to give them more power (9). For applications composed primarily of either parallel or vector code, the use of several homogeneous supercomputers fills that need. However, for applications composed of a combination of code suitable for vector processors and code best suited to massively parallel architecture, a distributed environment of vector and massively parallel systems is needed(10).



Storm cloud visualization program at NCSA.

Supercomputers are used to model thunderstorms numerically. Temperature, wind, pressure, and other variables are calculated every few seconds at several hundred thousand locations in the area of the developing storm. Mathematical equations are then used to simulate the storm's evolution (Grand Challenges: High Performance Computing and Communications, The FY 1992 U.S. Research and Development Program).

A number of testbed projects are exploring new possibilities in this area. A global climate model composed of both an ocean systems model and an atmospheric model will be distributed across several supercomputers on the CASA testbed (*11*). While the ocean systems model is run on a massively parallel connection machine at the Los Alamos National Laboratory or on the Intel DELTA machine at CalTech, an atmospheric model will be run on a CRAY Y-MP vector supercomputer at San Diego. The NECTAR testbed is investigating these methods with process flowsheeting problems that model the control, design, and economic performance of large-scale chemical plants (*12*). The BLANCA testbed is applying these techniques to the coupling of large- and medium-scale weather simulation systems.

Large-scale systems that need to be adapted to investigate small-scale phenomena are also candidates for distributed processing. Consider, for example, the atmospheric models currently used to simulate thunderstorms or hurricanes. These homogeneous models do not have sufficient spatial resolution to predict small-scale phenomena such as tornados. A typical thunderstorm model requires five hours of dedicated CRAY-2 time, with sustained processor throughput of 370 MFLOPS using a spatial resolution of 1,000 x 1,000 x 500 m zones. A grid of 1.2 million of these zones forms a volume of space measuring 100 x 50 x 15 km. To create a simulation of a tornado would call for a tenfold increase in resolution, requiring delivery of a sustained processor rate of 1 TFLOPS for over 1.2 years. This is clearly untenable, so new methods are being developed to allow for variable grid resolution. Such models allow the researcher to zoom in on portions of the system to be modeled at high resolution while the majority of the large-scale system is modeled at coarse resolution. Some models use artificial intelligence techniques to detect areas that require higher resolution. Systems are also being developed to spread a grid across several networked computers using methods for minimizing the frequency and volume of boundary condition exchange among partitions (*13*).

The development of software that will help noncomputer scientists use these distributed environments is important. Inter-process communications mechanisms such as UNIX sockets or remote procedure calls (RPC) have existed for many years, but to build applications with them requires significant expertise. There must also be tools for debugging these distributed applications and straightforward means of measuring their performance and efficiency, so that scientists can weigh the advantages of distributing their applications by various methods.

Systems that are capable of analyzing an application and running appropriate portions on massively parallel supercomputers and vector supercomputers are necessary in the long term. However, like the visualization analysis applications discussed above, they are highly resouce intensive and will stress the system's capacity.

Accessing Digital Archives

Archiving the flood of new information and making it accessible will chal-

lenge the network of the future. The large information repositories existing on the Internet today are growing at an unprecedented rate. These include online archives of documents and software that are available through public file transfer. At typical supercomputing centers, data archives are growing at 50 to 200 GBytes per month. Major initiatives such as the human genome project, the Hubble Space Telescope, and the Earth Observation System (EOS) will accelerate this growth.

The Internet today can readily support document and software archives, but it is not adequate for interactive access to image and data archives. At the high end, Internet throughput is sufficient to transfer selected images or data sets for local examination but not for interactive browsing. At typical sites in 1991, the transfer of an image in less than one minute is impossible.

The ability to store archives of images and data interactively will offer great advantages. The CASA testbed is developing interactive composite visualization systems for geological applications which will combine Landsat, seismic, and topographic database images. Researchers will be better able to identify fault zones, plate thrusts, and surface erosion, and may even improve their ability to predict the magnitude as a result. These data sets are often much larger than the disk space available locally even on supercomputers and may require higher throughput than many storage media can sustain.

As analysis and visualization software become more widely available, and as the size of the data sets grows beyond available local disk space, access to remote data archives will become increasingly important.

Archives on the Internet currently use diverse indexing schemes, access formats, and storage media. Some data is online, a vast quantity is on secondary or tertiary media such as magnetic tape or optical disk, and much data is not accessible through the Internet at all. In most cases there is no convenient, much less transparent, searching or retrieving capability.The concept of a national distributed digital library exists in the form of public network file archives and electronic bulletin boards. But these archives have virtually no searching, much less browsing, capability. To be useful, the important archives must have well-documented indexing schemes and access protocols and a standard user interface to support searches in multiple archives with varying database query protocols.

The personal computing and desktop publishing markets exploded as a result of systems that are easier to use. A similar growth will take place as digital archives become more accessible.

Multitiered storage will provide unlimited potential for growth and at the same time support very high transfer rates. This is accomplished by using slower, cheaper media (tape and inexpensive disk) for older files and faster, more expensive media such as parallel disk arrays for interactive work, reserving the best media for frequently accessed data. Support for data input directly from supercomputer simulations will require very high data rates, as will support for interactive visualization and analysis of stored data. Though most high-speed storage mechanisms require the use of very large data blocks to achieve maximum throughput, the user access software should support transparent access at the byte level, particularly as data sets continue to grow to sizes too large for local disk space.

Retrieving information stored on secondary media such as cartridge tape or CD ROM should not result in delays. For data in high demand, robotic systems will be necessary so that the data is available at all times with low, predictable retrieval time. Storage techniques are now being used that keep the first block of all files online for immediate transfer while the remaining data is being retrieved. These techniques and algorithms need to be further deveoped as archives grow.

The personal computing and desktop publishing markets have exploded as a result of systems that are easier to use, have more functionality, and are simply faster. A similar growth will take place as digital archives become more accessible, have more functionality such as browsing, and deliver information faster, with a resulting impact on the Internet.

Technologies for Collaboration

To date the Internet has offered reasonable access to the many resources available through supercomputers and document archives, but it has provided only rudimentary support for collaboration through electronic mail and bulletin board systems. It is very difficult to collaborate with remote colleagues in real-time today using scientific visualization techniques. More often, a videotape of the final product is sent to colleagues for review. Ideally, scientists would wish not only to discuss the model or theory or jointly analyze the data during development but also to discuss the visualization techniques themselves.

Software is becoming available today to allow multiple users to work on documents and applications interactively. These tools might be used by collaborators either to edit a document jointly or to demonstrate applications to one another. Technology such as this must be developed to allow multiple users to cooperate in

steering supercomputer simulations and analyzing volumetric data. Medical applications such as the techniques for planning radiation treatment discussed above coupled with such a system would support real-time second and third opinions, resulting in safer treatments.

As with visualization systems, collaboration systems should provide intuitive, object-oriented interfaces that allow users to construct systems out of existing modules. The systems should be able to support users regardless of their location on the network or the brand of hardware or software on their desktops. Such systems must also have simple mechanisms for setting up conferences, adding or deleting participants, and determining who is in control and who is participating at any particular time. A mechanism for alerting a list of collaborators that a conference is taking place might be included; or a scientist might be able to search the network for existing conferences and request permission to join in as a participant or as an observer.

A real-time resource registration system that allows users to locate existing conferences the way users now locate networked computers would be a convenience. The Internet Domain Name System (DNS) has all but replaced the central management and distribution of Internet host tables, allowing host information to be updated within minutes rather than days and allowing distributed control. Similarly, a directory system is needed to keep track of conference connectivity in real-time (who is attached, what distributed applications modules are being used, and so on). A rational naming system that can accommodate growth will be needed; such a system does not currently exist for networked resources other than computers.

The collaboration systems must also support connections to digital libraries, distributed applications (e.g., distributed supercomputer simulations, and instruments. The result would be a distributed, network-based laboratory as described in great detail at the "Collaboratory" workshop held at the Rockefeller University in March 1989 (*14*).

The collaboration systems must support multiple communication channels to allow users to combine the benefits of the home office environment and those of the face-to-face meeting. Video windows showing the person who is speaking, with perhaps a shared electronic workspace, should be provided; because the telephone would not be well integrated with the network, audio must be added as well.

Will these systems downgrade to the lowest common denominator, leaving high-end participants stranded, or will there be a shared common ground?

Finally, these interactive collaboration systems must be able to function in a distributed environment in which characteristics such as delay, reliability, and throughput may vary widely. One user might be at the end of a 45 Mbit/s line, another at the end of a 600 Mbit/s line, a third might be at the end of a satellite link with 250 msec delay. While the underlying network protocols must support this, the user interface must also take into account the fact that browsing rates and interactive control delays will vary among the participants. Will these systems downgrade to the lowest common denominator, leaving high-end participants stranded, or will there be a shared common ground?

Toward a New Architecture and Policy

As the advanced network applications outlined here become easier to use, more users will take advantage of them and per-user bandwidth requirements will grow. At the same time, new users will continue to join the networks. This growth in both high-end and typical, or low-end users is leading to a somewhat bimodal distribution of user needs. Both the architecture of the system and the policies governing its use must evolve in ways that take this pattern of diverging demands into account.

Architecture

Remote visualization applications are already in use in many local area networks, though with signficantly reduced functionality. Their use will increase as Internet performance improves. However, if a single remote visualization process were to produce a 100 Mbit/s burst, it would take only a handful of users on the national network to generate a load of over 1 Gbit/s. When these three-dimensional simulations are visualized using time as a fourth dimension, as in volumetric animation, the single-user bursts will increase to several hundred Mbit/s.

The prospect of such demands raises a central issue. Predicting the load on the Internet, difficult though this may be in a packet-switched network in which a single user can generate the peak load, will be the key to providing stable service. In a telephone network, all users have roughly the same requirement—the availability of a voice channel. Variation lies in the frequency and duration of calls. Taking into consideration peak business hours and holidays, network providers can predict load with great accuracy. Should the network become overloaded, new users are blocked out until resources are available. By contrast, the Internet has little informa-

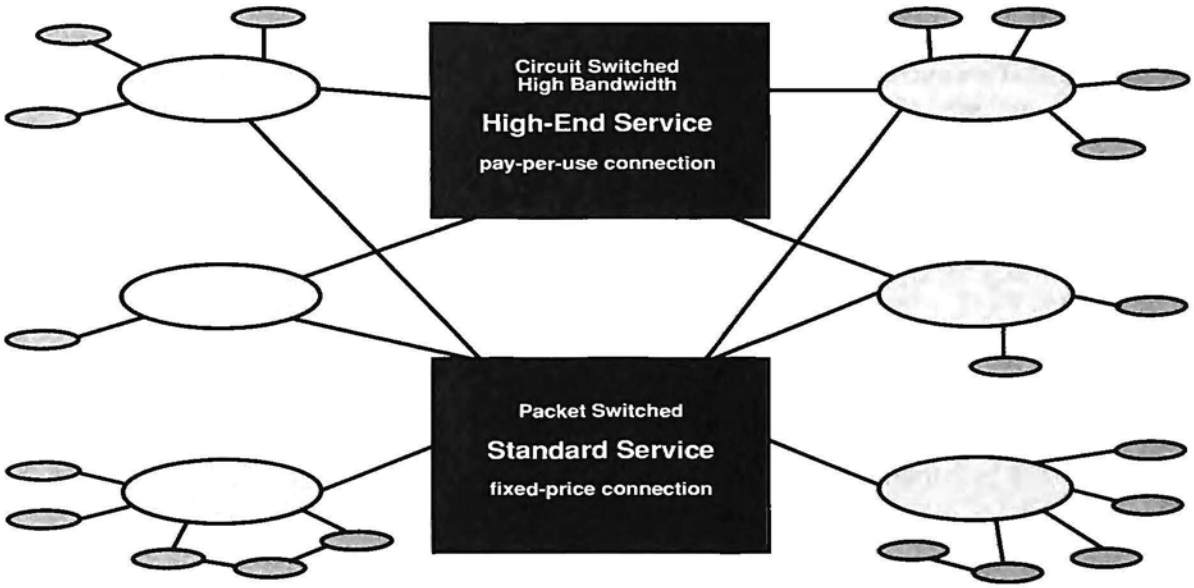


Figure 1: Multi-tiered NREN Services

In the multi-tiered NREN of the future, high-end services would be provided through switched, fixed bandwidth circuits. Standard services would be provided over packet switched lower bandwidth service. Thus, a user might take advantage of the standard packet switched connection when sending email, but could be dynamically switched to higher bandwidths when performing resource intensive activities.

tion as yet about patterns of use, and its packet-switch architecture now provides no mechanism for blocking users to avoid overload.

But consider a possible Internet connection scheme in which each site would have the ability to increase capacity dynamically based on the sustained and peak throughput levels anticipated. Most of the time this connection would support average to low-end applications. Capacity would only temporarily increase to support high-end applications.

In such an architecture (Figure 1), each Internet site might subscribe to a combination of service levels with a fixed rate (for example 1.5, 10, or 45 Mbit/s) and another rate for dynamically allocated additional service (such as 45, 155, 310, or 622 Mbit/s). The network infrastructure of today provides a fixed level of service for each NREN site. To this might be added the capacity to meet predicted high-end service requests. The NSFNET backbone network is currently being upgraded from Phase II 1.5 Mbit/s to Phase III 45 Mbit/s. During early 1991 we have two separate but interconnected backbone networks (the 1.5 Mbit/s network interconnects all 16 backbone nodes and the 45 Mbit/s network interconnects 8 of these nodes). In the future, it is likely that the NSFNET backbone will also employ switched services to reconfigure the network dynamically, enabling it to respond to variations in levels of use.

Thus we have the beginnings of a multitiered system: fixed service would be carried by the Phase II network and the Phase III network would be used only when necessary to support high-end applications at loads greater than 1.5 Mbit/s. If resources were unavailable, high-end applications could be denied access temporarily to the Phase III facilities. The user would then choose either to delay running the application or to accept the lower performance and functionality available over the fixed network.

Today switched 45 Mbit/s services have been demonstrated, and in the next several years we will see switched services at 155 Mbit/s and possibly as high as 622 Mbit/s. A two-tiered architecture would help assure that the needs of both the high-end and typical users are met, but better information on Internet users must be collected and analyzed if we are to be able to plan and provide for future needs.There is currently no firm understanding of who and where the users are or what their needs are. If the Internet is to survive the ongoing expansion, much less evolve into a commercial enterprise, such information will be crucial.

Federal, State, and Institutional Roles

Clearly industry has benefited from federally funded supercomputer centers, technology development, software development, and other research. Federal funding has often provided the resources to spark a new technology, and computer networks are no differ-



ent. The development of cutting-edge capabilities must continue to receive support from the Federal government through its support of high performance computing centers to the degree that these centers provide unique services and promote the development of technology and expertise, both of which are transferable to industry.

State funding can help to underwrite the necessary transfer of technology and services from the advanced, high-end community present at national centers to the community of typical users. The establishment of state supercomputing centers is an example of this transfer, in that these centers are providing some of the same services originally offered by the national centers. This passing of the torch allows national centers to focus on the advanced technologies and new services, always advancing the state of the art.

Both federal and state agencies should identify and support those leaders who provide unique services. As far as network services are concerned, the states must increase support for mid-level networks so that the federal agencies can invest in an advanced national infrastructure, including network backbones, advanced network testbeds, and national centers.

Other avenues for state involvement in the Internet would include the use of state telecommunications facilities to support research and education. Many states have a telecommunications infrastructure that is used primarily by state agencies such as law enforcement and those who manage state information. These facilities offer excellent potential for support of state Internet components. For example, some states operate multiple 45 Mbit/s trunks, including some on private facilities. To increase the capacity of these trunks to carry 1.5 and 45 Mbit/s connections to Internet users would be much less expensive than to install new facilities. However, many state networks provide services that are quite different than those desired by Internet users. For example, state networks are often oriented around remote terminal access to mainframes. Such facilities may be high in capacity, but they are aggregations of terminals. For this reason, state governments must rely heavily on the Internet expertise of universities rather than attempt to merge Internet services with state telecommunications networks on their own terms. From the perspective of the state government, it is much more cost-effective to use the Internet to access national resources such as digital libraries, supercomputer centers, and other universities than to attempt to replicate these resources within the state.

Finally, access to advanced network services such as gigabit network testbeds should be allocated in the same way that access to advanced supercomputers is allocated—through peer review and with the aid of federal funds. Ubiquitous, medium-capacity service should be available to the national community; however, a parallel network access program should provide state-of-the-art services to those who can make the case that these services are required to make advances not otherwise possible. If the evolution of the Internet is shaped by the market demands of users, the average requirements will be met but the leaders will be stifled. A national investment in supercomputer centers has enabled leading researchers to make crucial progress on the frontiers of their respective fields. That investment must be nutured and protected and must be enhanced by similar investment in a state-of-the-art national network infrastructure.

Dr. Larry L. Smarr is Director of the National Center for Supercomputer Applications (NCSA), University of Illinois at Urbana-Champaign. Charles Catlett is Manager of Networking Development at NCSA.

Notes

1. A version of this paper was presented at "Information Infrastructure for the 1990s, A Workshop/ Symposium of the Programs on Science, Technology, and Public Policy and Strategic Computing and Telecommunications in the Public Sector" held at the John F. Kennedy School of Government, Harvard University in January 1990.
2. As of October 1990, there were 2,004 registered Internet sites, according to data from the Merit Computer Network. A site is defined here as a particular network address which corresponds to a particular administrative/geographic unit such as a university campus. A recent search of the Domain Name System by Stanford Research Institute found over 250,000 hosts.
3. This complex included the five supercomputer centers established in 1985-86 (NCSA at Illinois, SDSC in San Diego, JvNC at Princeton, CNSF at Cornell, and PSC in Pittsburgh) and the already established National Center for Atmospheric Research (NCAR).
4. While it is not known where these curves will level off, we are clearly in the period of rapid growth that characterizes the early stages of the technological S-curve. An excellent discussion of the S-curve as related to network growth can be

found in Vijay Gurbaxani, "Diffusion in Computing Networks: The Case of BITNET," *Communications of the ACM* (December 1990) 33:65-75.

5. The Advanced Research Projects Agency has now become the Defense Advanced Research Projects Agency (DARPA). NCSA statistics show a significant increase in dialup access together with a substantial decrease in Internet access during this period.

6. The *VISTANET* testbed participants include Bellsouth, GTE, MCNC (formerly known as the Microelectronics Center of North Carolina), and the University of North Carolina at Chapel Hill.

7. The *BLANCA* testbed participants include AT&T Bell Laboratories, Lawrence Berkeley Laboratory, the National Center for Supercomputing Applications, the University of California at Berkeley, the University of Illinois at Urbana-Champaign, and the University of Wisconsin at Madison. Collaborating telecommunications carriers and computer companies are Ameritech, Astronautics, Bell Atlantic, and Pacific Bell.

8. Another radio astronomy visualization project uses software on a supercomputer to interactively modify parameters such as instrument settings and gain adjustments, refining a 4k x 4k x 24-bit color image (called a "supermap") over many iterations. The scientist must determine at what point the process has converged on a satisfactory image or when it is diverging and must be restarted with different settings. Optimally, the scientist views two to five images per second over a period of several minutes. Using a 1k x 1k x 24-bit frame buffer, this process requires a data rate of between 50 and 125 Mbit/s. Using a 2k x 2k x 24-bit frame buffer such as has recently become available, the data rate increases to between 100 and 500 Mbit/s.

9. See Elizabeth Corcoran, "Calculating Reality," *Scientific American*, (January 1991) 264:108-09, and specifically the item entitled "Modeling on Multiple Machines."

10. A number of efforts are ongoing in this area, including network-based publication and distribution systems such as the Xanadu HYPERserver from Autodesk and digital library systems coupled to community-specific data archives such as those at the Computer and Genomic Systems Laboratory at the University of Arizona (Bruce R. Schatz, "A Digital Library for an Electronic Scientific Community," [September 1989]. Systems such as the Digital Library System being developed by CNRI will combine digital libraries with expert systems, distributed systems, and in some cases programs that migrate semi-autonomously ("Knowbots") over the network from archive to archive, collecting information for the process owner (Robert E. Kahn and Vinton G. Cerf, *The World of Knowbots*, Vol. I, The Digital Library Project [CNRI, 1988]).

11. The CASA testbed participants include Los Alamos National Laboratory, CalTech, the Jet Propulsion Laboratory , and the San Diego Supercomputer Center. The carriers collaborating in the CASA testbed are MCI, Pacific Bell, and US West. CASA is part of the CNRI gigabit testbed effort.

12. The NECTAR testbed participants include Carnegie Mellon University and the Pittsburgh Supercomputing Center. The carrier collaborating in the NECTAR testbed is Bell Atlantic/Bell of Pennsylvania. NECTAR is part of the CNRI gigabit testbed effort.

13. These algorithms must take into account not only the heterogeneity of the processors but also, in the case of coupled simulations, the heterogeneity of the simulations themselves. The progression rates of the coupled simulations can differ by as much as two orders of magnitude. These simulations must be synchronized through allocation of resources or other methods such as a looser coupling..

14. See Joshua Lederberg and Keith Uncapher, "Towards a National Collaboratory," report of a workshop held at The Rockefeller University in March 1989.

15. The AURORA testbed participants include Bellcore, IBM Research, MIT and the University of Pennsylvania. Collaborating telecommunications carriers are Bell-Atlantic, MCI, and NYNEX. AURORA is part of the CNRI gigbit testbed effort.

Additional Reading

Gary Stix, "Gigabit Connection," *Scientific American* (October 1990) 263: 118-120.

Supercomputers: Directions in Technology and Applications (Washington, D.C.: National Academy Press, 1989)

John S. Quarterman, *The Matrix: Computer Networks and Conferencing Systems Worldwide*, Digital Equipment Corporation, 1990).