NSF Opens High-Speed Computer Network

Reaching from Seattle to South Florida, the NSFnet is establishing itself as the world's premier computer network; but problems of billing and security still must be solved

COMPUTER WIZARDS may remember 1988 as the year of the worm—the time when a hacker nearly brought down the entire network that carries their electronic data across the continent. But it was a good year, too. It marked the first time a public system known as NSFnet reached a high-speed transmission rate of 1.5 million bits per second, fast enough to unclog old traffic tie-ups and give scientists at remote locations a feeling of immediate access to U.S. supercomputer centers.

The NSFnet was launched by the National Science Foundation (NSF) in 1985 as means of opening up the supercomputer centers to distant users. It has grown into something far more ambitious. Today it is compared to a superhighway, to be used by any and all comers for whatever kind of research they like. In time, it is supposed to reach farther, run faster, cost less, and be more accessible than any other network in the world.

At present, however, the NSFnet is still under construction. One of the biggest challenges the builders will face in the near future is finding a way to charge for service. At the moment, maintenance is supported by a few private grants and a subsidy from NSF—justified as an aid to NSF researchers. But if the network succeeds in its goals, demand will grow, and users will have to come to grips with steadily rising budget demands. As yet there is no way to meter usage, which, no doubt, is one reason why the network is so popular.

The big news is that the "backbone" lines—crossing the United States in two redundant paths—are up and running. Modest connections have been made to France and West Germany, and others to the United Kingdom and Scandinavia are being negotiated. On the West Coast, an arm is reaching toward Asia. A couple of slow links to Japan (9600 bits per second) are already in place.

According to Stephen Wolff, director of the program, the NSF has committed itself to spending $14 million on the network through 1990. The project has received donations of hardware and software from the IBM Corporation, fiber optic connections from the MCI Communications Corporation, and $5 million from the state of Michigan (where the operations center is located). The combined value of all this is $40 to $50 million, Wolff says. The NSFnet is destined to serve as the high-speed lane in a complex, interlocking system called the Internet, or the National Research Network.

The remarkable thing, some officials say, is that this project has won the cooperation of five sibling and often quarrelsome U.S. agencies. They have gone beyond pro forma attendance at meetings and actually agreed to share routing machines, satellite links, and fiber optic cable.

Overseeing the joint effort is a small group called the Federal Research Internet Coordinating Committee (FRICC), chaired by William Bostwick of the Los Alamos National Laboratory. The agencies involved are NSF, the Defense Advanced Research Projects Agency (DARPA), the Department of Energy, the National Aeronautics and Space Administration (NASA), and the Department of Health and Human Services. By its own claim, FRICC has already improved efficiency by cutting back on duplicate satellite channels to Europe, each of which costs $120,000 a year.

Following the lead of two high-level reports on the use of computers in research,* FRICC is now planning a bigger empire. NSF officials say that another $400 million in federal support will be needed through 1995, and that the speed will have to be raised 1000-fold. Senator Albert Gore (D–TN) introduced a bill to finance the expansion last October and plans to reintroduce it in 1989. In time, NSFnet may become the trunk carrier for the entire tangle of university and government lab connections known as the Internet, offspring of ARPAnet, the granddaddy of them all.

ARPAnet was launched in 1969 as a military experiment in "packet-switching." It used an ingenious method of sending data over a network without first establishing a connection. Because the system is in a sense intelligent, it can send packets of data to their destination even if parts of the circuit are busy or broken. The exact path is not clear and does not have to be known before the packets are launched on their journey. ARPAnet became popular among researchers as an electronic mail delivery service and soon took on a life of its own. In 1983, it had become so loaded with traffic that Defense officials decided to split off operational military traffic into a separate system called Milnet. Now, according to Mark Pullen of DARPA, the plan is to let ARPAnet fade away and yield its role as a general purpose system to NSFnet. A new Defense Research Internet will be created for experimental work.

Although the NSFnet was initially promoted as a link to the supercomputers, it is rapidly growing into something bigger. Wolff says he knew from the start that this would happen. The services it provides—electronic mail, rapid file transfer, access to exotic computing machines—are valued by all kinds of researchers. The goal is to link every research computer to every other research computer, in a kind of global electronic brain.

NSFnet managers came a step closer to this goal in July when they installed seven new gateways to the backbone and boosted speed significantly. There are now 13 access points to the system, and the transmission rate (or band width) has been raised along most of the lines from 56,000 bits per second to 1.5 million bits (megabit) per second. These improvements will make it accessible to 400 universities, research centers, and lesser outposts. The potential clientele is huge: there are an estimated 60,000 "host" systems in the Internet.

In the computer business, capacity and demand always chase one another in overlapping waves. Whenever capacity is improved, new users come pelting in and old users find more demanding ways to use the equipment. The result is exponential growth in traffic. The NSFnet office at Merit Inc., in Ann Arbor, Michigan, reports that traffic on the backbone this year has quadrupled.

Despite this boom in use, the higher speed has made a difference to regular users like Paul Bash, a 36-year-old postdoc in Harvard's Chemistry Department. Bash's field is "computational molecular biology," and he says he has a voracious appetite for computer time. It took 1000 hours on a supercomputer to produce the results for a paper in *Science* (1 May 1987, p. 564). The

*"Toward a National Research Network" (National Research Network Review Committee, National Academy Press, Washington, D.C., 1988) and "A Report to the Office of Science and Technology Policy on Computer Networks to Support Research in the United States" (Federal Coordinating Council on Science, Engineering, and Technology, Washington, D.C., November 1987).

task might have taken a decade to complete on an ordinary machine. His paper describes a model that predicts the behavior of proteins in solution and calculates changes in their free energy levels. Bash says the device could be used to zero in on promising avenues for new drug development.

Bash often roams from supercomputer to supercomputer, looking for clear time in which to conjure up molecules. Sometimes he goes to the supercomputer center at San Diego. More often he stops near home at the Pittsburgh center. He regularly haunts the Illinois supercomputer, where he is a "key user." He also visits the Ames Research Center near Palo Alto, California, the John von Neumann Center at Princeton, and the Minnesota Supercomputer Institute.

This lab-hopping would be stupefying if it meant traveling to these places or shipping computer tapes back and forth in the mail. Instead, Bash travels by cable. Sitting at his desk in Cambridge, he types some commands into his terminal and connects in seconds to Illinois, Minnesota, or California. He ships large data files over the network, logs into a distant supercomputer to put the job in order, and sets the machine running. "Then I can go out for a beer," he says. Later he checks in to collect the results.

"It's very nice when it works right," Bash says. "Six months ago the network was frustrating and unreliable. I would get logged out from Pittsburgh all the time." Since the summer, "it's gotten much better." But the supercomputers themselves are "chock full." Bash says the system will bog down again unless it is expanded.

Another unique use of the network will be to hook into special data collections. Wolff gives as an example an astrophysics database at Strasbourg, France, called SIMBAD. The NSF, in a joint project with NASA and the French National Institute for Information and Automation, has hired a satellite link connecting to SIMBAD. It will enable researchers on the Internet to tie directly into the database and summon up information published since 1950. The beauty of SIMBAD, says Joyce Rey-Watson of the Smithsonian Astrophysical Observatory, is that it is organized not by author or title, but by celestial object. It recognizes any of the profuse, multiple names astronomers have used over the years. (Some objects have as many as 30 different names.) Rey-Watson, who will be the registrar of U.S. users, says the connection to France is not yet reliable enough to open the gates. She hopes to do so early this year.

By far the biggest use of the network is for electronic mail, a way of sending person-to-person messages almost instantaneously. Scientific collaborators separated by thousands of miles use it to keep in touch on the progress of research, sending results and comments back and forth on a daily basis. Although the network at present has no way of distinguishing mail packets from, say, processed data packets, officials guess that 80% or 90% of the traffic is mail.

There are still major obstacles to be overcome if the NSFnet and Internet are to live up to promises being made. Three big ones are as follows:

■ **International protocols.** The United States is ahead of other countries in networking. Americans got an early start because the heterogeneity of the U.S. computer population and the lack of a central authority, compelled them to develop a system that could be accessed easily, anywhere, by anyone. (Outsiders see this as anarchy.) All that is required of network users is that they "speak" in a common, publicly owned language developed by DARPA, called TCP/IP (Transmission Control Protocol/Internet Protocol). In principle, the United States and its data-sharing partners have agreed to shift in the 1990s to a new international language called OSI (Open Systems Interconnect). It is supposedly just as tolerant of local eccentricity as TCP/IP. But in reality, national post and telephone agencies in Europe and Japan are finding it difficult to relax their monopoly grip on communications lines. The result is that OSI has a slightly different national flavor in each country, and may require a summit conference to reach commonality. However, Wolff says that the technical solution to this problem is already in hand, and will soon be implemented.

■ **Billing and priority routing.** Before the NSFnet grows much bigger, it must develop a system for metering traffic. At present it is open to all comers. Access to any particular segment is controlled by routing machines, which care only about efficiency. Maintenance costs must be allocated according to usage. As a nonprofit system, NSFnet cannot give a free ride to commercial users. Managers would also like to be able to give some uses, such as graphics processing at a supercomputer, a higher claim to fast links. Adding meters and priority sorting devices will be a big task.

■ **Security.** U.S. military users do not want their data packets crossing "hostile" countries. The addition of Hungary to the Internet, possibly in 1989, may soon raise this issue. In addition, many users were shaken by the worm invasion of November 1988 (*Science*, 11 November 1988, p. 855). An official at the U.S. Patent and Trademark Office, for example, says that plans to attach the patent database to Internet have been put on hold partly because of nervousness about infections from outside. Such worries often arise from a poor understanding of the system, computer experts say. But it may be necessary to create some kind of unclassified security group to protect the system and give advice to newcomers with little knowledge of the risks.

In a few months Bostwick and the FRICC committee will publish a new program agenda laying out plans for the next 5 years, and explaining how they intend to tackle these problems. Bostwick says the initial objectives set out for the National Research Network by the White House report on networking have now been achieved, and that "We expect that for the next 10 years, we will remain ahead of commercial" developments. ■ ELIOT MARSHALL

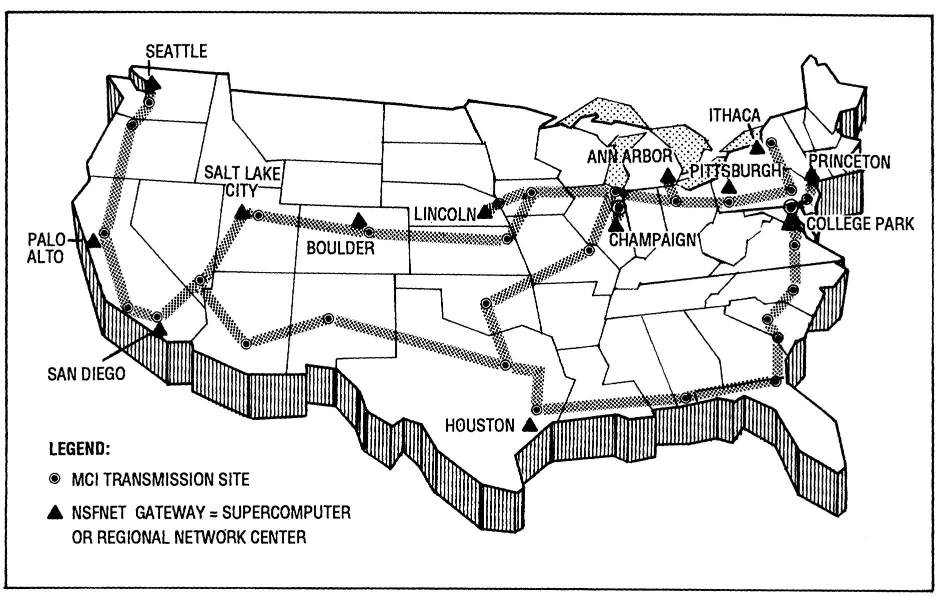

Data superhighway *built by NSF includes 13 gateways (triangles) into regional computer networks and maintenance points (circles) run by MCI Communications, Inc.*