

The building of the Internet

Implications for the future of broadband networks

Jeffrey A. Hart, Robert R. Reed and François Bar

The Internet, a system of interconnected computer networks primarily in the USA, can be seen as an experiment in the development, deployment and use of high-speed networks, and as such can provide guidance for the shaping of the future national telecommunications infrastructure. Internet's significance lies not only in the technologies it helps develop, but more importantly in the new usage dynamics it helps uncover, the new network management mechanisms it tests and the new policy strategies it explores.

Jeffrey Hart is with the Department of Political Science, Indiana University, Bloomington, IN 47405, USA (Internet: HARTJ@UCS.INDIANA.EDU). Robert R. Reed is also with the Department of Political Science, Indiana University (Internet: RORREED@SILVER.UCS.INDIANA.EDU). François Bar is with the Berkeley Roundtable on the International Economy, University of California, Berkeley, CA 94720, USA (Internet: FBAR@GARNET.BERKELEY.EDU).

¹We will use the capitalized Internet to refer to the international networks that grew up around the TCP/IP-based networks of government agencies and universities, in the USA, initially, and later in many other countries. The lower-case internet refers to any private network that uses the TCP/IP family of protocols.

²The available traffic statistics report only traffic on the NSFNET backbone and not on the Internet as a whole, so traffic that does not flow through the NSFNET backbone is excluded. Yet we know that the
continued on page 667

The Internet¹ is a loosely organized system of interconnected computer networks, which primarily serves the research and education community. Its large and growing community of sophisticated users, the diversity of applications and uses it fosters, and the trials now underway within gigabit testbeds make the Internet one of the boldest real-life experiments in broadband networking today. As a result the unfolding of the Internet's history, the dynamics of its evolution and the policy issues it raises hold useful insights into the future evolution of telecommunications networks. The challenges facing the Internet community today presage the more general challenges that policy makers will have to tackle if they want to foster the emergence of an advanced national network infrastructure. Three facets of the Internet story are especially relevant.

First, data traffic over the Internet is growing explosively. During 1991 traffic on the NSFNET backbone has increased by an average 6% per month. Traffic doubled from May 1990 to May 1991 and again from May 1991 to March 1992. Traffic growth has been accelerating during this period.²

In order to cope with this growth, the NSFNET backbone was upgraded from T1 lines (capable of transferring data at 1.54 megabits per second) to T3 lines (45 megabits per second). At stake in the current policy discussions of the National Research and Education Network (NREN) is the next phase upgrade to gigabit data transfer rates (roughly 50 times the current T3 rate). This next upgrade is planned for the mid-1990s.

Most significantly for our argument, the expansion of the Internet is driven by users. In contrast with the national debates on the deployment of broadband networks, the question facing the Internet community is not 'What will fill the pipe?' but rather 'How to build a pipe big enough to contain the current (or projected) overload?'

The growth in Internet traffic has been fuelled partly by increased use among existing users but also very significantly by the addition of new

users. The Internet is available to users in over 50 countries. It has over 4 million users affiliated with over 5000 organizations.³ In the USA alone there were between 2 and 3 million users of the Internet in 1991. According to 1992 statistics, the number of users is doubling every seven months.⁴ The Internet interconnects roughly 700 000 host computers. It is so important to some computer specialists that they will refuse employment at companies or agencies which do not have the ability to interconnect.

Initially, the ARPANET – the precursor of the Internet – was intended only for the Department of Defense and its contractors. As it became open to civilian research uses, and placed under the responsibility of the National Science Foundation (NSF), the NSFNET's primary users were the academic 'elite' of advanced computer scientists and researchers. More recently two new communities of users have been granted access to the Internet: private corporations and the broader academic community.

Private corporations use the Internet as a wide area network (WAN) to interconnect their local area networks (LANs). The broader academic community, including for example academic libraries as well as K-12 schools,⁵ is using the Internet for electronic mail (e-mail), file transfers and library interconnection. As a result the character of the Internet's user community has changed significantly. The character of network use and the kinds of applications carried over the network are changing accordingly. Interactive applications now constitute the fastest-growing segment among the applications carried by the Internet.

The reason for the widespread enthusiasm is not that the Internet is the optimal high-speed data network. In fact the network's main family of protocols, TCP/IP, is now quite old and probably much less efficient than newer approaches to high-speed data networking, such as Frame Relay or SMDS (switched multimegabit data service). Its success stems from the fact that the Internet is often today the only possible outlet for eager users which offers a standardized and stable interface, along with a deliberate focus on openness and interconnection. While these entail problems, such as vulnerability to worms and viruses,⁶ they make the Internet extremely attractive to very different groups of users in corporations, government agencies and academic institutions.

Further, as a result of its widening use the Internet is becoming a very fertile experimental ground for high-speed networking applications. Here again we find it extremely interesting that the most successful 'broadband' applications now supported by the Internet differ significantly from those most prominent in the public debate. For example, they tend to involve corporate computing (whether by libraries doing cooperative cataloguing, or by corporations linking CAD workstations) rather than video transmission. The open systems approach to network embraced by the Internet favours and stimulates a variety of these experiments. The work underway in the gigabit testbeds reinforces this trend and takes it one step further. Overall, the various Internet user communities are uncovering new facets of, and new potential uses for, broadband networks. Therefore the unfolding results of the Internet experiment deserve attention as we try to understand what the future information infrastructure will look like and what it will be used for.

The policy environment surrounding the evolution of the Internet is dramatically different from the traditional telecommunications policy environment. The policy goals being discussed, the policy mechanisms

continued from page 666

number of foreign networks connected to the Internet is growing rapidly (see Figure 4 below). So total Internet traffic may be growing faster than traffic on the NSFNET backbone.

³The figures are inexact because no one has figured out how to keep an accurate count of the rapidly growing number of users.

⁴Written comments on the first draft of this paper by Vint Cerf, March 1992.

⁵We will use K-12 to stand for education from kindergarten through the 12th grade. This is a standard acronym in the US education policy literature.

⁶Clifford Stoll, *The Cuckoo's Egg*, Doubleday, New York, NY, 1989; Clifford Stoll, 'Stalking the wily hacker', in Peter J. Denning, ed, *Computers Under Attack*, ACM Press, New York, NY, 1990. It should be noted that dealing with security problems is one of the high priority items on the agenda of the Internet Activities Board.

which have permitted the growth of the Internet, and which are now envisioned to guide its future, are actually quite foreign to the telecommunications debate. Although they are not necessarily discussed in those terms, Internet-related policies reflect industrial policy considerations, for example, when they have allowed a military-funded network to be spun off for broader civilian use, or when they viewed the development of the NREN as a way to further US competitiveness.

Cross-subsidies of many kinds – including several that would probably no longer be tolerated within the public telecommunications networks – pervade the Internet. These include the initial military subsidies which supported the Internet's basic technology development and the public subsidies channelled by DARPA, the NSF and the Department of Energy into the exploration of high-speed networks and applications. Current discussions about the conditions under which for-profit use of the Internet by private companies should be allowed – the so-called 'privatization of the Internet' – in effect explore new possibilities for cross-subsidizing Internet use between various categories of users.

The significance of these mechanisms goes beyond simple transfers of funds. Underlying them is a larger effort to assemble a coalition of users who will share a common network infrastructure that is beneficial to all. At stake is the deployment of a network infrastructure able to support joint experimentation and learning among various Internet users and to serve as a conduit for the diffusion of network-related innovations across various user communities. Joint experimentation is becoming an important feature of the bargains established between Internet access provider firms and for-profit users. An example is an agreement between Hewlett-Packard and PSInet to give Hewlett-Packard access to the Internet in exchange for H-P's willingness to share innovations that result. Another example is the use of the Internet as a distribution channel for new software releases, such as the latest version of X-Windows.

The Internet is a fascinating experiment in the development, deployment and use of high-speed networks. This experiment can provide some guidance for national telecommunications policy as the latter faces the task of shaping the future of the nation's network infrastructure. The significance of the Internet as an experiment lies not simply in the technologies it helps develop, but more importantly in the new usage dynamics it helps uncover, the new network management mechanisms it tests and the new policy strategies it explores. Ultimately the NREN may come to represent much more than simply a large-scale experiment. It could become a critical piece of the USA's information infrastructure.

The building of the ARPANET

The ARPANET's origins can be traced to the appointment in 1961 of J.C.R. Licklider, at that time a professor of mathematics at Massachusetts Institute of Technology (MIT), as the first director of computing office of the Advanced Research Projects Agency (ARPA) at the Department of Defense. While Licklider was not a computer scientist, he was a visionary who believed in time-sharing and interactive computing. Licklider's interest in interactive computing led him to support efforts to create radically new communication systems. He was succeeded at ARPA by Ivan Sutherland, who was to hire the people

who eventually designed the first packet-switching network. Sutherland's successor, Robert Taylor, also believed strongly in the idea of networking. Taylor was responsible for selecting the team that built the first packet-switched network. Licklider's, Sutherland's and Taylor's efforts at ARPA helped to accelerate the development of time-sharing computers and computer networks.⁷

Even before Licklider, Sutherland and Taylor were funding early work on time-sharing and networks, Paul Baran at the Rand Corporation was thinking about networking computers to create a robust communication system to survive a nuclear first strike. In reports published in 1964 Baran 'proposed that messages be broken into units of equal size and that the network route these message units along a functioning path to their destination where they would be reassembled into coherent wholes'.⁸ Donald Davies at the National Physical Laboratory in the UK first used the word 'packet' in late 1965 to describe the units of equal size mentioned in Baran's work. He did not see Baran's work until after circulating his own work on packetizing data for storage and forwarding.⁹

While Baran and Davies independently came up with the idea of using packets for the storage and transmission of data on computer networks, it was not until 1967, when ARPA was preparing its 1968 request for proposals (RFP) for a system to reliably link computers in academic, industrial and government research laboratories, that a packet-switching network was actually designed. ARPA's Taylor had hired Larry Roberts away from MIT's Lincoln Laboratory in 1967 to write the RFP and to decide on the sites. In writing the plan for the ARPANET, which was published in June 1967, Roberts essentially had to design a packet-switching network. Roberts had not read the works of either Paul Baran or Donald Davies. Nevertheless he saw quickly the value of using a 'packet-switching' architecture for networks.

Paul Baran, in a recent article, quotes a concise explanation of the concepts behind packet switching:

The basic idea is to allocate some of all the system capacity (along some path between subscribers) to one customer at a time; but only for a very short period of time. Customers are required to divide their messages into small units (packets) to be transmitted one at a time. Each packet is accompanied by the identity of its intended recipient. In packet-switched networks each packet is passed from one packet switch to another until it arrives at one connected to that recipient, whereupon it is delivered. Packets arriving at a switch may need to be held temporarily until the transmission line that they need is free. The resulting queues require that packets be stored in the switches and it is not unusual that all packet buffers are occupied in a given switch. Thus both the switch capacity (processing and storage) and transmission capacity between switches is statistically multiplexed by subscribers. The designers of such packet networks are faced with the problem of choosing line capacities and topologies that will result in relatively high utilization without excessive congestion.¹⁰

Packet switching is different from circuit switching in that, in a packet-switching network, packets have no previously determined routes or paths. Each packet travels separately by the best route possible at any given time. The separate packets do not have to take the same route. Once the packets arrive at the destination they are reassembled into the proper sequence.¹¹

Robert Kahn, then a Professor of Mathematics at MIT, took a one-year leave in 1968 to work at a government-funded private

⁷Interview with Robert Kahn, Corporation for National Research Initiatives, Washington, DC, 19 August 1991; and letter to Jeffrey Hart from Lawrence G. Roberts, 24 April 1992. See also Peter J. Denning, 'The ARPANET after twenty years', in Denning, *op cit*, Ref 6; and L.G. Roberts, 'The evolution of packet switching', *Proceedings of the IEEE*, Vol 66, November 1978, pp 1307-1313.

⁸Denning, *op cit*, Ref 7, p 12. Paul Baran, 'On distributed communications', Vols I-XI, Rand memoranda, Rand Corporation, Santa Monica, CA, August 1964; and Paul Baran, 'On distributed communications networks', *IEEE Transactions on Communications Systems*, Vol CS-12, 1964, pp 1-9.

⁹Roberts, 1978, *op cit*, Ref 7, p 1308.

¹⁰Paul Baran, 'Packet switching', in John C. McDonald, *Fundamentals of Digital Switching*, 2nd ed, Plenum, New York, NY, 1990, p 193.

¹¹Defense Data Network, *DDN .New User's Guide*, February 1991, p 5.

Cambridge, MA, thinktank called Bolt Beranek and Newman (BBN), so that he would have more in common with his colleagues at MIT. Many of his colleagues did applied work and Kahn felt that his own work was overly theoretical.¹² Kahn was assigned the problem of responding to ARPA's RFP for computer networks. Frank Heart at BBN had experience building computer hardware and knew how to take Kahn's mathematical ideas and put them into practice. With the help of Heart and Severo Ornstein, Kahn wrote BBN's proposal for the ARPA contract. The contract was awarded to BBN in January 1969.¹³

BBN built a specialized computer for the ARPA contract called an interface message processor (IMP).¹⁴ The IMP was a packet switch that was connected directly to a host computer and could transfer packets to other IMPs via 56 kbps leased telephone lines. In the autumn of 1969 the first IMP was installed at UCLA, which became the first ARPANET node. Kahn and another BBN employee, David Walden, went to UCLA to test and debug the IMP. By December 1969 the network had been expanded to four nodes. There were 23 host computers on the ARPANET in April 1971, 62 in June 1974 and 111 in March 1977.¹⁵

In 1972 Kahn joined DARPA to become Director of the Information Processing Techniques Office where he started the Strategic Computing Program. At this point the ARPANET consisted of around 30 host computers connected with each other by IMPs linked together through leased telephone lines. A new device called a Terminal IMP (TIP) was added to the network in 1972 to allow users to dial up the network from remote terminals over the public switched telephone network (PSTN). In October 1972 Kahn installed a complete ARPANET node at the first International Conference on Computer Communications in Washington, DC, thus making possible the first public demonstration of a packet-switched network.¹⁶

The origins and role of TCP/IP

The ARPANET was immediately useful to the computer science and military community.¹⁷ Now it was possible to link up many different kinds of computers with differing data transfer rates, and data could be transferred reliably. Access to the ARPANET was limited to defence agencies and defence contractors. Within that group the heaviest use of the ARPANET was by the computer scientists. As of 1971 the two most widely used applications on the ARPANET were electronic mail and remote login services.

Its users came to see the ARPANET is an invaluable tool, and later put pressure on the Department of Defense to provide broader access in order to realize its full potential for the scientific community. By 1983 the ARPANET had expanded to over 100 nodes (from four in 1969), but access was still limited to defence agencies and defence contractors. Two new special purpose networks were built in the early 1980s on the model of ARPANET – CSNET, funded by the National Science Foundation, and BITNET, funded by IBM¹⁸ – to give access to electronic mail capabilities to the non-defence-contracting computer science and academic communities respectively. Access to the CSNET gave computer scientists access to all the nodes on the ARPANET.¹⁹ BITNET connected only those local networks that were connected to an IBM mainframe.²⁰

¹²John Markoff, 'Robert Kahn's vision of a national network of information begins to take hold', *New York Times*, 2 September 1990, p F1.

¹³Kahn, *op cit*, Ref 7; and Roberts, 1978, *op cit*, Ref 7, p 1308.

¹⁴IMPs are now called Packet Switched Nodes (PSNs). The original IMPs were Honeywell DDP-516 minicomputers. See Andrew S. Tannenbaum, *Computer Networks*, 2nd ed, Prentice Hall, Englewood Cliffs, NJ, p 35. BBN rebuilt the original IMPs to do packet switching with virtual circuits (instead of leased lines) in 1972, and then provided packet-switched data transfer services through its Telenet subsidiary starting in 1975. Telenet was made available to users nationally by dialling local telephone numbers. Telenet is currently owned by US Sprint.

¹⁵Kahn, *op cit*, Ref 7; and Roberts, 1978, *op cit*, Ref 7, p 1308.

¹⁶Roberts, 1978, *op cit*, Ref 7, p 1309.

¹⁷Robert E. Kahn, 'A national network: today's reality, tomorrow's vision, Part 2', *Educom Bulletin*, Summer/Fall 1988, pp 73–74.

¹⁸BITNET was started by the City University of New York and Yale University, then rapidly grew to 50 or so universities. From the start it used IBM mainframes and e-mail software. Later on IBM provided funding to expand the network in the USA and Europe. The main financial support for BITNET came from the universities themselves. BITNET was governed and managed through an organization called Educom.

¹⁹CSNET, initially funded by the NSF, was developed around 1980. At one time CSNET had almost 200 participating sites and connections to almost 15 countries. CSNET still serves a number of industrial and collegiate sites. See Karen Armstrong McKelvey, Michelle Magolis and Susan Estrada, *Cerfnet User's Guide*, July 1991, Section 6.1.

²⁰See John S. Quarterman and Josiah C. Hoskins, 'Notable computer networks', in Denning, *op cit*, Ref 6. BITNET and CSNET were merged in 1987 to form the Corporation for Research and Educational Networking (CREN). Educom continues to manage the CREN as it did earlier for BITNET.

Vinton Cerf first met Walden and Kahn in 1969 when they went to UCLA to install the first ARPANET node. Kahn wanted Cerf at DARPA because Cerf knew a lot about the early work on network protocols, but Cerf decided instead to join the faculty at Stanford in 1972 (he was not to join the DARPA staff until 1976). That did not keep the two from collaborating on network protocols, however. The Department of Defense wanted to interconnect computers with satellites, and with packet radio systems,²¹ and Kahn needed help from Cerf to rework the older network protocols to handle this difficult problem. In the process of doing this Cerf and Kahn invented the 'gateway' concept, which allowed very different types of networks to be connected, even though they used different-sized data packets and worked at different clock speeds. The IEEE published a paper by Cerf and Kahn in May 1974 which outlined a network interconnection protocol that is now known as TCP/IP.²² Thanks to its robustness, adaptability and relative simplicity TCP/IP has become a *de facto* world standard for interconnecting networks.

²¹In particular, the ALOHA system pioneered by Norman Abrahamson at the University of Hawaii in 1970 and first deployed in 1971. See Baran, *op cit*, Ref 10, pp 211–221; and Tannenbaum, *op cit*, Ref 14, p 182.

²²TCP stands for Transmission Control Protocol and IP for Interconnection Protocol. See Tannenbaum, *op cit*, Ref 14, pp 36–40; John Davidson, *An Introduction to TCP/IP*, Springer Verlag, New York, NY, 1988; and Douglas E. Comer, *Internet-working with TCP/IP: Principles, Protocols, and Architecture*, Prentice Hall, Englewood Cliffs, NJ, 1988. Cerf was later to play a key role, along with Kahn, in building a political coalition to support the construction of a national data network to replace the ARPANET when the ARPANET was decommissioned in 1990. In addition, Cerf chaired the Internet Advisory Board (IAB) from 1989 to 1992. The IAB deals with technical and governance problems in the Internet community.

²³This is somewhat ironic in that the TCP/IP is today challenged by advocates of another 'open systems' approach called Open Systems Interconnection (OSI).

²⁴Mark Hall and John Barry, *Sunburst: The Ascent of Sun Microsystems*, Contemporary Books, Chicago, IL, 1990.

²⁵Written comments on an earlier draft by Vint Cerf, March 1992. See also James Wilson, *Berkeley UNIX: A Simple and Comprehensive Guide*, Wiley, New York, NY, 1991, p 2; and H.M. Dietel, *Operating Systems*, 2nd edn, Addison-Wesley, New York, NY, 1990, pp 571–573. It is interesting to note that the UNIX at Berkeley led not only to the proliferation of TCP/IP systems but also to the formation of Sun Microsystems under the leadership of Berkeley UNIX researchers like Bill Joy.

²⁶Ethernet permits packet switching in local area networks over coaxial cables. The inventors of Ethernet were Robert Metcalfe and David Boggs. See Baran, *op cit*, Ref 10, pp 212–213. ARPANET was easily made compatible with Ethernet LANs because it had already incorporated the idea of using gateways between different types of packet-switched networks in the TCP/IP suite of protocols.

²⁷In the 1970s a variety of regional and limited-purpose national networks joined the ARPANET to offer broad connectivity. Merit, Inc, which operated a wide area network for the state of Michigan, joined the ARPANET early on. BARRNET, linking government and academia in the San Francisco Bay area, also joined the ARPANET. See Joseph Polka, 'Getting together bit by bit', *Science*, Vol 248, 13 April 1990, p 160.

The gateway concept makes TCP/IP particularly useful for people who want to interconnect computers and networks manufactured by different companies. Thus TCP/IP pioneered what is now called the 'open systems' approach.²³ The UNIX operating system that was developed in 1969 by Ken Thompson and Dennis Ritchie at Bell Laboratories was made available to a number of universities for research purposes through AT&T's liberal licensing policy in 1975.²⁴ DARPA funding, beginning in 1980, made it possible for TCP/IP to be incorporated into the kernel of the BSD 4.1 version of Berkeley UNIX in 1981, which was made freely available to all computing sites with UNIX systems.²⁵

The TCP/IP protocol suite was adopted as the standard for the ARPANET in January 1983. All UNIX systems now contain TCP/IP in the kernel, which includes almost all scientific and engineering workstations. Now even manufacturers of non-UNIX systems, like IBM and Digital Equipment Corporation, support TCP/IP interconnection services as less powerful supplements to their proprietary network protocol suites. The wide availability of TCP/IP systems in the marketplace is a major reason for the rapidly expanding traffic on the Internet.

To summarize, developments in packet switching and network protocol standardization greatly expanded the possibilities for the ARPANET. However, the ARPANET had to face a number of other challenges and opportunities during its lifetime to maintain its viability and utility for the user community. One of these was to adapt to the development of local area networks (LANs), starting from the development of Ethernet by Robert Metcalfe at Xerox Corporation's Palo Alto Research Center (PARC).²⁶ Another was to respond to the building of regional and special-purpose computer networks which eventually needed greater geographic reach.²⁷ The challenge which ARPANET could not handle, and which led to its demise, was to expand access beyond the community of defence agencies and defence contractors.

The building of the NSFNET

In the late 1970s and early 1980s a national programme for supporting research in supercomputing got underway which put very strong pressures on the federal government to build a public network accessible to

all major research facilities, public and private. The Computing for Education and Research Program (CER) was established at the National Science Foundation in the late 1970s. This programme did not include supercomputing initially, but early in 1980 the NSF got congressional approval for the construction of five supercomputing centres. The selection of sites was made in 1983–84 and new supercomputing centres were built in 1985–86 at Cornell University (the Cornell National Supercomputer Facility), Princeton (the John Von Neumann Center), Pittsburgh (the Pittsburgh Supercomputing Center), the University of Illinois at Urbana/Champaign (the National Center for Supercomputing Applications) and the University of California at San Diego (the San Diego Supercomputer Center).²⁸ Four Engineering Centers of Excellence (at the University of Delaware, Purdue University, the University of Washington at Seattle and the University of Minnesota) were included in the NSFNET networking plans in 1986–87.²⁹

There were only five NSF supercomputer centres initially because of the great expense of the new machines. Because scientists who were not based at universities near the centres wanted access to them, the NSF decided to provide access via networks. At a meeting in 1979 at the University of Wisconsin, Kent Curtis of NSF approached Robert Kahn of DARPA to ask whether the ARPANET would be capable of linking the separate supercomputing facilities. Kahn was enthusiastic about the idea, but it was not to be.

The Department of Defense had decided to expand the original ARPANET in the early 1980s. In October 1983 the ARPANET was officially split into two networks: the MILNET and the 'residual' ARPANET. Prior to the split ARPANET had over 100 nodes and combined R&D activities with more strictly military ones. The new MILNET, with 60 nodes, was to be a strictly military network, but there were gateways to connect the MILNET to the 'residual' ARPANET. That left more than 40 nodes for the 'residual' ARPANET which Kahn hoped would become the backbone for the new network linking the NSF supercomputing centres.³⁰

The Pentagon asked Congress for an expansion of the MILNET to 3600 nodes and was authorized to do so, but there were not enough people in the Defense Communications Agency to perform the work necessary for the requested expansion.³¹ So Kahn had the idea that he could create a supercoalition of supporters of both the NSFNET and the MILNET to build the new supercomputer network and the new MILNET by adding nodes to both the MILNET and the residual ARPANET at the same time.

For a while it looked as though this plan would work. However, the labyrinthine acquisition procedures of the Department of Defense and delays in the delivery of circuits from the telephone companies created long lags in the addition of new nodes, so the NSF decided instead to bypass DARPA and to build the new supercomputing network on its own. After an interim period of linking the supercomputing centres with a 'do-it-yourself' network,³² the NSF issued a Request for Proposals. The RFP was awarded in 1988 to a three-company team led by Merit, Inc, which included IBM and MCI.³³ The NSF awarded \$14 million to the Merit-led team to put the NSFNET backbone in place. Merit was responsible for the management and administration of the NSFNET. MCI was responsible for maintaining the network for five years. IBM provided its NetView network management software and switches

²⁸Sidney Karin and Norris Parker Smith, *The Supercomputer Era*, Harcourt Brace Jovanovich, New York, NY, 1987, pp 105–111. Princeton centre was closed after the NSF review in 1989–90. See *High Performance Computing and Networking for Science Advisory Panel: A Background Paper*, Office of Technology Assessment, Washington, DC, 1991.

²⁹Interview with Dr Robert Kahn, Corporation for National Research Initiatives, 29 August 1991. Purdue, Minnesota and Colorado State had supercomputing facilities prior to the NSF initiative in 1984. See *High Performance Computing and Networking*, *op cit*, Ref 28, pp 33–36, for a complete list of supercomputing centres in the USA.

³⁰Defense Data Network, *DDN New User's Guide*, February 1991, p 8.

³¹By this time the DCA had the responsibility of maintaining and expanding the ARPANET: Kahn, *op cit*, Ref 29.

³²The 'do-it-yourself' network used devices called Fuzzballs to switch packets. It was never meant to be more than an interim solution to the access problem, and it served the programme well until increases in traffic began to produce frequent network crashes.

³³Eric Aupperle, 'Building Merit', *Information Technology Quarterly*, Summer–Fall 1989, pp 7–9; and Kahn, *op cit*, Ref 29.

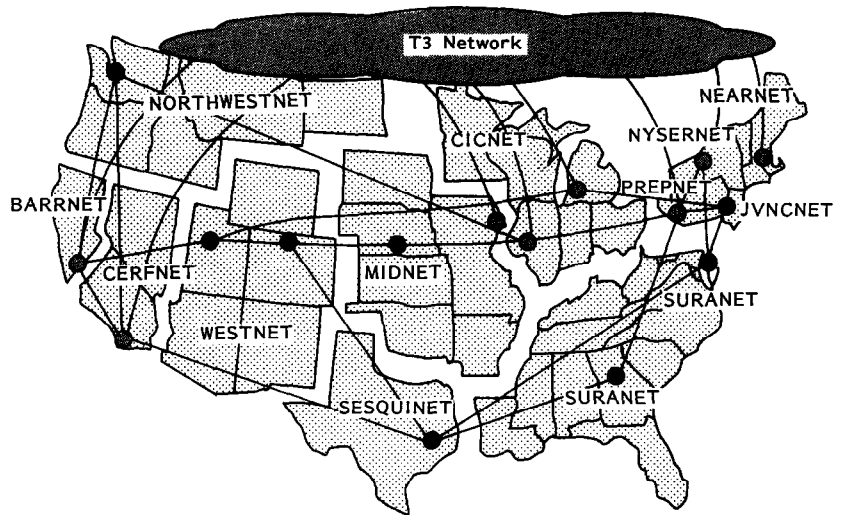


Figure 1. The NSFNET T1 backbone.

Source: Jordan Becker, Vice President for Network Services, ANS CO+RE Systems Inc, presentation at a conference at the Center for Communications and Information Science, Columbia University, New York, NY, 5 September 1991.

based on IBM computers. By July 1988 the new backbone was in place (see Figure 1).³⁴

There were 13 nodes on the NSFNET backbone with the capability of transferring data at 1.5 Mbps, a rate considerably higher than the speed of ARPANET (56 kbps in the early 1980s).³⁵ Despite NSF's decision not to build the NSFNET on the foundation of the ARPANET, the NSFNET shared ARPANET's decision to use the TCP/IP protocol suite. The builders of the NSFNET considered Open Systems Interconnection (OSI) protocols, but opted for TCP/IP because they believed the OSI protocols were not ready. Several new protocols were added to the TCP/IP family to provide new services. The original TCP/IP protocols continued to operate satisfactorily under the higher data transfer rates of the NSFNET.

Rapid traffic growth on the NSFNET

Traffic on the NSFNET grew very rapidly. In May 1989 traffic was approximately 1 billion packets per month. By May 1991 traffic had increased to 7.56 billion packets per month, a 140% increase over the 3.15 billion packets transmitted in May 1990. By March 1992 traffic had almost doubled to 14.9 billion packets (see Figure 2).

Packet use of NSFNET by application was as follows in March 1992: 21% for networked mail applications, 29% for file exchange, 2% for non-TCP/UDP services, 27% for other TCP/UDP services, 7% for domain name look-up, and a remaining 14% for interactive applications (see Figure 3). The category of applications that has been growing most rapidly is other TCP/UDP services.³⁶ Thus one can argue that interactive and X-Windows applications have contributed proportionally more than others to the growth of Internet traffic in the last two years.³⁷

In 1989 around 200 universities were on Internet.³⁸ The total number of networks on the system in May 1989 was 516. Of those, 95 were foreign. By March 1992 there were 4976 networks on the system, including the MILNET networks configured for the NSFNET backbone (added in the summer of 1990). The number of foreign networks on the NSFNET was 1697 in March 1992 (see Figure 4).

The residual ARPANET was decommissioned in June 1990 and all the old civilian ARPANET nodes were taken off the network. DAR-

³⁴Willie Schatz and Mary Jo Foley, 'Users welcome new NSF network with glee and caution', *Datamation*, 1 September 1988.

³⁵See Frederick Williams, *The New Telecommunications: Infrastructure for the Information Age*, Free Press, New York, NY, 1991, p 88.

³⁶TCP is Transmission Control Protocol; UDP is User Data Protocol. Other TCP/UDP services include irc, talk, X-Windows and Appletalk services. X-Windows in particular is a highly network-intensive service and is used extensively by scientists and engineers with UNIX workstations: Internet e-mail from Eric Auopperle to Rob Reed, 19 May 1992. Another source indicates that over 90% of TCP/IP traffic on some Internet nodes is accounted for by six applications: FTP (file transfer protocol), SMTP (simple mail transfer protocol), NNTP (network news transfer protocol), VMNET (an IBM mail exchange application), TELNET (used for remote logins) and RLOGIN (also used for remote logins). See Peter B. Danzig and Sugih Jamin, 'TCP/IP: a library of TCP internetwork characteristics', Computer Science Department, University of Southern California, Los Angeles, USC-CS-91-495 (available from the authors at the following internet address: traffic@excalibur.usc.edu).

³⁷This was confirmed in a recent presentation by Jordan Becker, Vice President of Network Services, ANS CO+RE Systems, Inc, at the Center for Communications and Information Science and Policy Conference, 5 September 1991.

³⁸Ivars Peterson, 'Highways for information', *Science News*, Vol 133, 18 June 1988, pp 394-395.

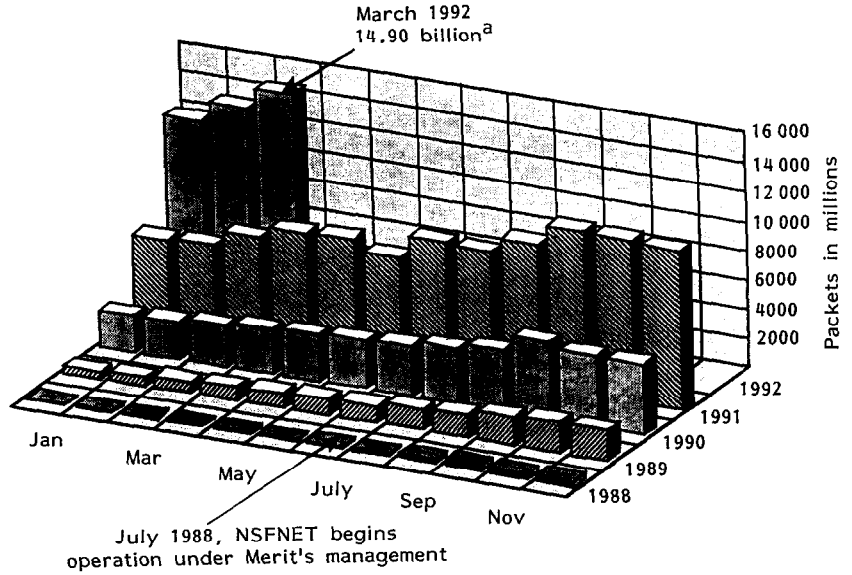


Figure 2. NSFNET packet traffic history.

^a Total packets, T1 and T3 networks.
Source: Merit, Inc, April 1992.




Figure 3. NSFNET applications history by percentage.

Source: Merit, Inc, April 1992.

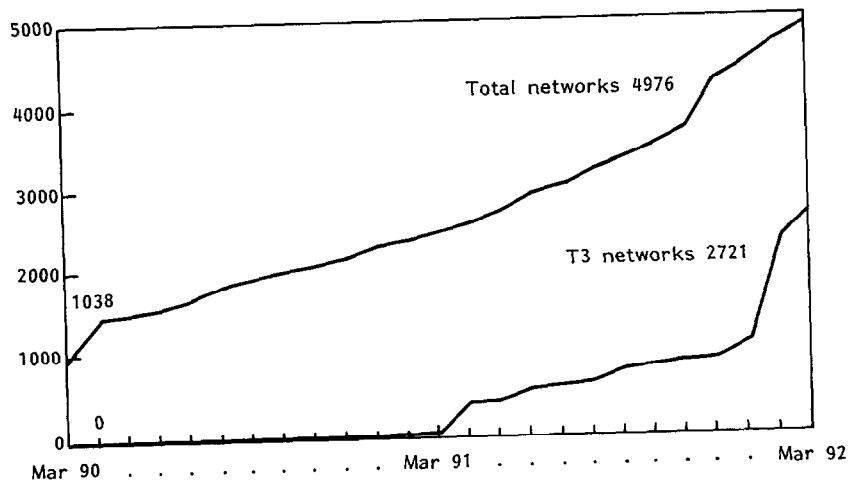


Figure 4. Foreign, regional, state and local networks on the NSFNET, March 1992.

Source: Merit, Inc, April 1992.

Table 1. NSFNET mid-level networks.

BARRNET	North California
CERFnet	Western USA
CICNET	Illinois, Iowa, Michigan, Minnesota, Ohio, Wisconsin
JVNCNET	Eastern USA and international
LOS NETTOS	Los Angeles
MichNet/Merit	Michigan
MIDNET	Arkansas, Iowa, Kansas, Missouri, Nebraska, Oklahoma
MRNET	Minnesota
NSCANET	Illinois, Indiana, Wisconsin
NEARNET	Nevada
NORTHWESTNET	Northwestern USA
NYSERNET	New York
OARNET	Ohio
PREPNET	Eastern USA
PSCNET	Eastern USA
SDSCNET	San Diego Supercomputer Network
SESQUINET	Texas
SURANET	Southeastern USA
THENET	Texas
USAN	National
VERNET	Virginia
WESTNET	Western USA

Source: This list was published in the *Cernet User's Guide*, July 1991, p 17.

PA's cost for maintaining the ARPANET in that last year was around \$11 million, and it decided the money would be better spent on other forms of research and development. Most of the civilian users of the ARPANET had made the transition to the NSFNET, often through regional or mid-level networks. ARPANET users were granted access to the NSFNET under an agreement reached between NSF and DARPA in October 1985.³⁹ The ARPANET had served its purpose well, but was not able to become the backbone for the new NSFNET because access could not be expanded beyond the military and computer science communities. The visions behind the NSFNET and the Internet were more expansive and inclusive. Virtually all academics, most employees of the US government and some employees of private businesses would have access to the Internet. The next step would be to expand access to all employees of commercial businesses able to pay an interconnection fee.

Private enterprise comes to the Internet

To arrange the interconnection of universities to the NSFNET, a system of NSF mid-level networks was established around a number of existing regional networks and some new ones.⁴⁰ Table 1 gives an incomplete list of the NSFNET mid-level networks. Some of these mid-level networks, like CERFnet, provide interconnection services for private businesses for a fee consistent with the NSFNET Acceptable Use Policy. Acceptable use involves 'research and instruction at not-for-profit institutions in the United States'. Commercial (for-profit) uses are permissible only if they are consistent with the overall purposes of the NSFNET. Exceptions must be approved at the NSF Project Office on a case-by-case basis.⁴¹ This seems not to have placed too many limitations on the information firms of California, many of whom access the Internet through CERFnet.

There were enough commercial users of the Internet who did not want to be bound by the NSFNET Acceptable Use Policy that there was considerable demand for purely private Internet interconnection services. For example, Hewlett-Packard set up its own proprietary internet for the purpose of linking its geographically dispersed research opera-

³⁹Quarterman and Hoskins, *op cit*, Ref 20, p 78.

⁴⁰The establishment of the regional networks was to some extent a reaction to the decommissioning of the ARPANET.

⁴¹*Cernet User's Guide*, July 1991, pp 19-20.

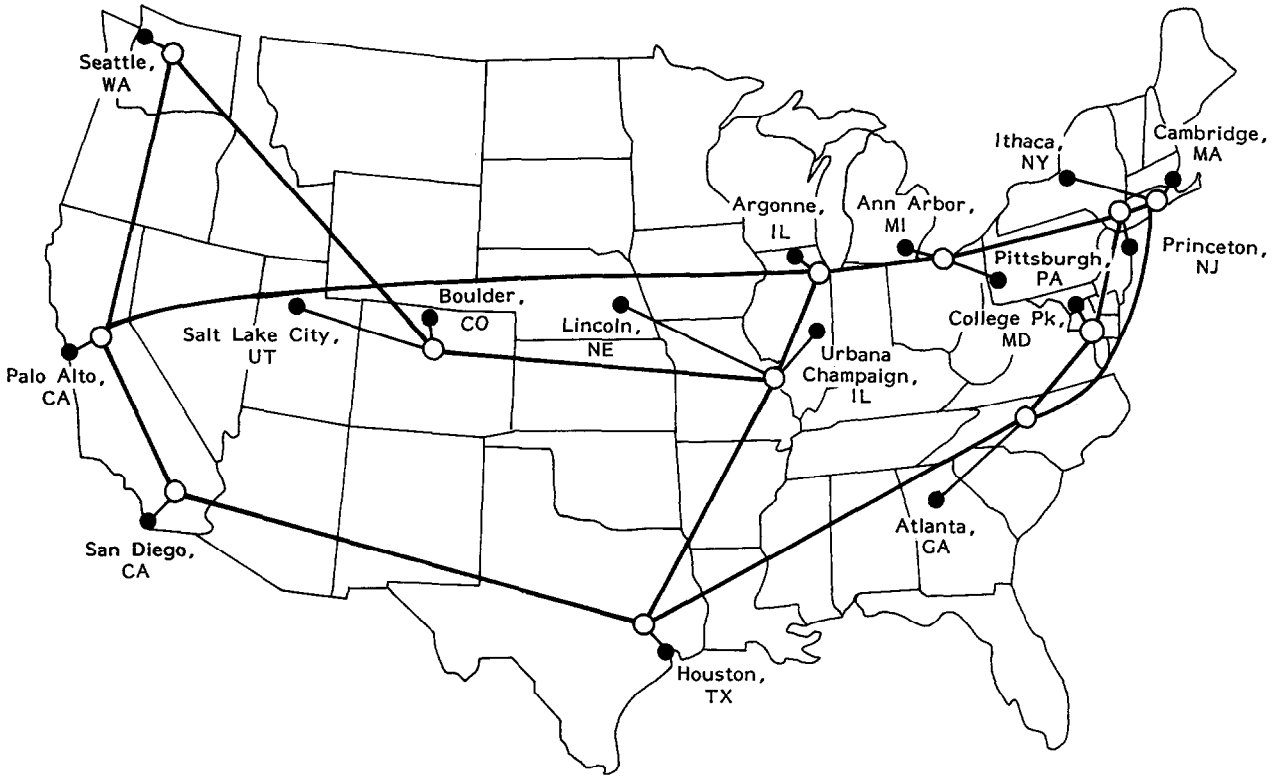


Figure 5. NSFNET T3 backbone, 1992.

Source: Merit, Inc, April 1992.

tions. National Semiconductor Corporation used its private internet to network advanced workstations to conduct simulations of new circuit designs.⁴²

In addition, some of the burden of managing the existing networks was shifted to private sector firms, through service contracts with the mid-level networks. For example, PSInet (run by Performance Systems International of Reston, VA) now provides network management services to NYSERNET that used to be performed by NYSERNET itself. PSInet also provides access to the Internet via NYSERNET or commercial firms in the New York area. Another private firm, UUNET Technologies of Falls Church, VA, now provides Internet connection through its AlterNet. General Atomics of San Diego does the same through its CERFnet operations. CONCERT is a private mid-level network established in 1985 to interconnect universities in North Carolina. Similarly, the state of Michigan funded the building of MICHNET to connect universities, state agencies and private businesses in the state of Michigan.⁴³

The NSFNET has added T3 leased lines to its existing T1 leased lines (bringing trunk transmission speeds up to 45 Mbps) in 1991–92 through a contract with a firm called ANS (Advanced Network and Services, Inc), which is a non-profit joint venture formed by IBM, Merit and MCI in 1990 (see Figure 5). Merit is still responsible for the management of the NSFNET, but it now contracts with ANS for some network management services and ANS, in turn, contracts with Merit to obtain access to Merit's accumulated expertise. In 1991 Merit, IBM and MCI formed another joint venture called ANS CO+RE (pronounced 'core',

⁴²Interview materials; John Markoff, 'A network of networks that keeps scientists plugged in', *New York Times*, 1 January 1992, p 21.

⁴³See Sharon Fisher, 'Whither NREN?', *Byte*, July 1991, pp 181–190.

short for 'commercial' and 'research') as a for-profit subsidiary of ANS. ANS CO+RE sells network services to commercial and research-oriented clients.⁴⁴

There has been some controversy about making the publicly funded Internet available for purely commercial uses by for-profit service providers. The reason given by the government for this particular move was to establish competition for access to the publicly funded networks. But, according to William L. Schrader, president of the firm that runs PSInet (which now must compete with ANS CO+RE): 'It's like taking a Federal park and giving it to K Mart. It's not right, and it isn't going to stand. As a taxpayer, I think it's disgusting.'⁴⁵ In short, one of the larger issues that privatization of the NSFNET raises is: Who should benefit from the subsidies that the federal, state and local governments have given to the building of the NSFNET and the mid-level networks, and how should these beneficiaries be regulated?

Business users have now been granted access to the Internet, albeit for a fee. The Internet inherited most of the users of the ARPANET after the NSFNET was built and enhanced. The building of the NSFNET expanded the circle of users to non-defence government bureaucrats and academics outside the defence and computer science communities. The building of mid-level and regional networks together with the privatization of some network services created the possibility of extending access to the Internet to new business users outside the limited circle of defence contractors. As a result of the rapid growth in the use of the Internet by businesses, commercial users are likely to play a major role in its future development. We turn now, accordingly, to a discussion of the efforts by the US Congress to upgrade the NSFNET backbone of the Internet in the USA through the creation of a National Research and Education Network.

⁴⁴Information supplied to the authors by Jordan Becker of ANS CO+RE and Eric Aupperle of Merit, Inc.

⁴⁵John Markoff, 'Data network raises monopoly fear', *New York Times*, 19 December 1991, p C7.

⁴⁶Here is a typical quote: 'I want to see a day when a school child in Tennessee can come home after class and set down, and instead of playing Nintendo, use something that looks like a Nintendo apparatus and plug into the Library of Congress; and read just not words, but look at pictures and moving graphics presented artfully and imaginatively in a way that captures and holds that child's attention; responds to the child's curiosity so the child can navigate through an ocean of information according to what he or she wishes to explore at the moment.' *High Performance Computing and Communications Act of 1991*, Hearing Before the Subcommittee on Science, Technology, and Space of the Committee on Commerce, Science, and Transportation of the US Senate, 102nd Congress, 5 March 1991, USGPO, Washington, DC, 1991, p 42. See also Albert Gore, 'A national vision', *Byte*, April 1991, pp 188-189; 'Infrastructure for the global village', *Scientific American*, Vol 265, September 1991, pp 150-153.

The National Research and Education Network

In the building of the ARPANET the Defense Advanced Research Projects Agency was the primary actor. The National Science Foundation was the key player in building the NSFNET backbone of the US portion of the Internet. In 1991 the US Congress decided to fund research on a National Research and Education Network (NREN) which will eventually upgrade the NSFNET backbone to gigabit speeds. The leadership of Senator Albert Gore (D-TN) in shepherding the NREN legislation through Congress is widely acknowledged. Gore's leadership has been premised on a somewhat different vision of the future uses of the network than those which are implicit in the growth of the ARPANET, NSFNET and Internet. Gore has stated on numerous occasions that he wanted the NREN to help the USA recover some of its lost international competitiveness, not just by improving the telecommunications infrastructure for academic and business research, but also by helping to develop new information resources for K-12 education and the public at large.⁴⁶ But the final legislation reflects a more elitist vision which coincides more closely with the wishes of the US scientific community than with the more democratic vision of Senator Gore. The story of the politics behind the NREN legislation is a story of how the scientific community's vision prevails.

After much congressional deliberation, the joint version of the Senate bill (S 272) and the House bill (HR 656), the High-Performance

Computing and National Research and Education Network Act of 1991, was passed into law on 11 September. The history of this particular legislative programme begins on 24 June 1986. On that day Senator Gore introduced S 2594, the Supercomputer Network Study Act of 1986. S 2594 required the White House Office of Science and Technology Policy (OSTP) to report to Congress on the federal government's role in promoting supercomputing and high-speed networking.

The OSTP delivered the mandated report to Congress on 20 November 1987. The report itself was prepared by the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), an organization created by Congress in 1976 with the legislation that established the OSTP. But we need to go back a little further to trace the roots of the recommendations in the OSTP's 1987 report.

In 1982 the Panel on Large Scale Computing in Science and Engineering issued a report, later called the 'Lax Report' (because the Panel was chaired by Peter Lax), which noted that the US research community was seriously lacking in access to high-performance computing. Jointly funded by the NSF and the Department of Defense, the Panel's report recommended that there be a new national supercomputer programme. In 1983 the FCCSET made the Lax Report its point of departure and formed a Panel on Supercomputers to examine what federal policies could be adopted to 'advance the development and use of large-scale computers'.⁴⁷

The OSTP's 1987 report echoed the Lax Report and the FCCSET's 1983 report in asserting that the USA needed to be concerned about growing international competition in supercomputers and in highly capable computer networks, citing efforts of the Europeans and the Japanese that threatened to leave the USA behind. It recommended that the government establish a long-range strategy for basic research on high-performance computing (HPC), to encourage joint business-university-government research in advanced software technology and to coordinate the building of a research network 'to provide distributed computing capability that links the government, industry, and higher education communities'.⁴⁸

The National Research Council of the National Academy of Sciences (NAS) issued a report in 1988 entitled *The National Challenge in Computer Science and Technology* which strongly reinforced the message delivered by the 1987 OSTP report. The NAS report argued that government funding of advanced computing research was necessary for preserving US competitiveness in an industry which accounted for as much as 10% of the GNP and almost 10% of all capital investment.⁴⁹ On 18 October 1988 Senator Gore introduced S 2918, the National High-Performance Computer Technology Act. No action was taken on this proposed legislation in 1988, but Gore reintroduced the bill on 18 May 1989 as S 1067, the High-Performance Computing Act.

Hearings were held in the summer of 1989 by House and Senate subcommittees on S 1067 and a similar bill aimed at funding HPC research by the Department of Energy. The Senate hearings included a discussion of the NREN, highlighted in an opening statement by Senator Gore, and representatives of a number of computer companies, universities and government agencies strongly supported the idea itself, if not always the specifics of the legislation.⁵⁰ Most of the testimony supported the views articulated by the OSTP and NAS reports that there was a danger of falling behind the international competition in

⁴⁷Office of Technology Assessment, *Supercomputers: Government Plans and Policies*, USGPO, Washington, DC, March 1986, p 22.

⁴⁸Executive Office of the President, Office of Science and Technology Policy, *A Research and Development Strategy for High Performance Computing*, Washington, DC, 20 November 1987, p 2.

⁴⁹National Research Council of the National Academy of Sciences, *The National Challenge in Computer Science and Technology*, National Academy Press, Washington, DC, 1988, p 7.

⁵⁰*National High-Performance Computer Technology Act of 1989*, Hearings before the Subcommittee on Science, Technology, and Space, of the Committee on Commerce, Science, and Transportation of the US Senate, 21 June, 26 July and 15 September 1989, USGPO, Washington, DC, 1989.

supercomputing and in computer-related scientific research in the absence of a federal HPC programme.

Senator Gore and his staff called upon a number of experts who wanted the goals of the HPC programme to be more ambitious, to extend access to advanced computer networks to a broader public. For example, Robert Kahn, by then the President of the Corporation for National Research Initiatives (CNRI), articulated the democratic educational vision fairly specifically: 'I believe there will be a real utility in the network for the educational system at virtually every level. Furthermore, there is a clear utility to the rest of society as well.'⁵¹ Kahn also referred to the NREN as part of a national 'information infrastructure' that would include systems like the CNRI's Digital Library System that would allow users to access library information anywhere in the country without knowing where the information was.⁵²

On 8 September 1989 Dr D. Allan Bromley, the President's Science Advisor and Director of the White House OSTP, released a report endorsing the creation of a Federal High-Performance Computing Program.⁵³ This report elaborated on points made in the 1987 OSTP report, but put a greater stress than the earlier report on the need to use existing supercomputing capabilities to 'expedite solutions to US scientific and technical challenges'.⁵⁴ This was still consistent with the scientific community's vision of the future of the network but demonstrated how the entire scientific community (and not just the computer scientists and military contractors) now saw an increased stake in shaping the direction of future national HPC and networking programmes.

Despite strong support from Bromley for HPC and the NREN, the Bush administration did not call for new funding at this time. Presumably Bromley was still fighting against the forces opposed to 'targeting' high-technology industries led by White House Chief of Staff John Sununu, CEA Chairman Michael Boskin and Budget Director Richard Darman. On 8 June 1990 the NSF announced that it would fund five gigabit testbed projects with \$15.8 million over three years. The passage of the Defense Department Appropriations Bill in October 1990 was also a notable event, because in that bill the Congress authorized \$20 million for supercomputing and high-speed network research at DARPA as a sort of protest against White House resistance to funding of HPC and NREN programmes.

On 24 October 1990 a revised version of S 1067 was passed unanimously by the Senate, but House passage was delayed because the Department of Energy wanted to coordinate the entire HPC programme, against the wishes of both the White House and Senator Gore.⁵⁵ The Department of Energy was trying to carve out a future role for its national laboratories in Oak Ridge and Los Alamos, facilities that were facing the prospect of major cuts thanks to decreased emphasis on nuclear weapons research. On 24 January 1991 Senator Gore and seventeen other co-sponsors introduced S 272, the High Performance Computing Act of 1991. This bill was very similar to S 1067. On 28 January 1991 Congressman George Brown (D-CA) introduced HR 656, which was nearly identical to S 272.

On 4 February 1991 the presidential budget request was released. This time it included funding for a High Performance Computing Initiative which would increase federal spending on HPC R&D by \$149 million from \$489 million in FY91 to \$638 million in FY92. Apparently

⁵¹*Ibid*, p 279.

⁵²*Ibid*, pp 283 and 300.

⁵³Executive Office of the President, Office of Science and Technology Policy, *The Federal High Performance Computing Programme*, USGPO, Washington, DC, 1989.

⁵⁴*Ibid*, p 1.

⁵⁵The revised version of S 1067 combined elements of S 1067 with S 1976 – the Department of Energy High-Performance Computing Act – but apparently it did not satisfy the DoE's concerns about management of the programme. See *High-Performance Computing and Communications Act of 1991*, Hearing, 5 March 1991, p 42.

Table 2. Estimated cost to the federal government of HPC/NREN programmes (by fiscal year, in millions of dollars).

Agency	1992	1993	1994	1995	1996
National Science Foundation	46	88	145	172	199
NASA	22	45	67	89	115
NIST	3	4	6	8	10
Total	71	137	218	269	324

Source: High Performance Computing Act of 1991, report submitted by Ernest Hollings of the Senate Committee on Commerce, Science, and Transportation, 16 May 1991, USGPO, Washington, DC, 1991, p 13.

the logjam in the White House on spending for HPC had been broken. On the following day Senators Johnston, Walup, Domenici, Ford, Briggman and Craig introduced S 343, the Department of Energy High Performance Computing Act. This act authorized federal funding for the DoE's part of the High Performance Computing Initiative.

The Science, Technology, and Space Subcommittee held a hearing on S 272 on 5 March 1991. On 11 July 1991 the House approved an amended version of HR 656 by a voice vote. A Brown-Gephardt amendment which contained Buy American provisions was approved also by voice vote. On 18 July 1991 the Senate approved the Veterans Administration, Housing and Urban Development, and Independent Agencies Appropriations bill (HR 2519), which increased funding for NSF research programmes by almost 14% and provided a large increase for computer research and NSFNET.

On 11 September 1991 the Senate amended and passed both S 272 and HR 656 and sent both back to the House. Major provisions of the bill followed a compromise which gave two-thirds of the funding to the National Science Foundation, and overall responsibility for managing the NREN to that agency. The FCCSET would continue to play a role in planning the NREN itself.

The HPC/NREN bill was an authorization bill and not a budgetary appropriation bill. According to a report by the Office of Management and Budgets (OMB) the new legislation would authorize a total \$650 million of new spending by the NSF, \$388 million by DARPA and \$31 million by the National Institute of Standards and Technology (NIST), an arm of the Department of Commerce, from FY1992 to FY1996 (see Table 2).

The gigabit testbeds

Further evidence for the central role of the scientific community in the political coalition behind the HPC/NREN Act of 1991 can be found in the way the gigabit testbeds were designed. The testbeds predate the passage of the 1991 Act because in April 1990 the National Science Foundation and DARPA decided to give the Corporation for National Research Initiatives (CNRI) grants totalling approximately \$15.8 million to begin to plan them. When these grants were issued, CNRI took the lead of a national research project with participants from the universities, national laboratories, supercomputing centres and major private companies.

The main goal of the testbeds is to provide information for planning the upgrading of the NREN (the successor to the NSFNET under the HPC/NREN Act) to gigabit speeds, presumably by the target data of 1996. The three gigabit wide area networks of the testbeds will be among the first broadband transport systems operating in the USA.⁵⁶ The testbeds have been designed to accelerate the development of

⁵⁶MCI is building a broadband network in Texas using SONET links mainly for carrying multiplexed voice messages. Other telephone carriers are deploying ATM switches and SONET links for this purpose.

commercial gigabit WAN equipment and software. The main rationale behind the testbeds, from the perspective of the network scientists, is to form a proving ground for technologies that will permit Internet traffic to grow at current rates without degrading performance – that is, to help build pipes wide enough to contain the projected overloads.

The involvement of private computer and telecommunications firms, national research laboratories, specific scientific research efforts (mostly university-based) and regional supercomputing centres ensures that a broad set of network applications will be built into the various testbed experiments and that these experiments will influence future commercial offerings. The gigabit applications developed with government funding have to be somehow connected with the 'Grand Challenges', as defined in the HPC/NREN Act, but the business-funding applications do not have this same restriction.

The combination of private, academic, government and scientific participants provides a kind of insurance that there will be commercial spinoffs from the network technologies developed in the testbeds. Federal government funding is aimed at reducing the risk for private firms and helping them to train personnel in high-performance computing and gigabit network technology. One can interpret the testbeds, in short, as government–industry joint ventures or R&D consortia for the development of gigabit WAN technologies.

The CNRI oversees the work of five different testbeds: (1) Aurora, (2) Blanca, (3) Casa, (4) Nectar, and (5) Vistanet. All except Vistanet receive NSF funding. Unlike the other four, Aurora does not involve a supercomputing centre. It is based in the Northeast with the main participants being Bellcore, IBM, MIT and the University of Pennsylvania. The telecommunications carriers associated with Aurora are Nynex, Bell Atlantic and MCI. Blanca is a national effort to further research in basic network technologies. The primary research participants for the Blanca testbed are: AT&T Bell Labs, Lawrence Berkeley Laboratory, the National Centre for Supercomputing Applications, the University of California at Berkeley, the University of Illinois at Urbana–Champaign and the University of Wisconsin at Madison. The collaborating telecommunications carriers are: Ameritech, Astronautics, Bell Atlantic, Norlight and Pacific Bell. The Casa research team consists of Los Alamos National Laboratory in New Mexico, California Institute of Technology, the Jet Propulsion Laboratory in Pasadena and the San Diego Supercomputing Center; MCI, Pacific Bell and US West are the telecommunications carriers. The principal research participants in Nectar are Carnegie-Mellon University and the Pittsburgh Supercomputing Center; the collaborating telecommunications carriers are Bell Atlantic/Bell of Pennsylvania. Vistanet is based in North Carolina. Its main research participants are: BellSouth, GTE, MCNC (formerly the Microelectronics Center of North Carolina), North Carolina State University and the University of North Carolina at Chapel Hill. The telecommunications carrier for Vistanet is BellSouth.

Aurora, Blanca and Casa are all developing gigabit WANs that cross state boundaries, while the Nectar and Vistanet networks are contained within a single state. The gigabit testbeds exist as proving grounds for the computing and network technologies that will permit the explosive growth of traffic on the Internet to continue. That purpose permeates all the testbed programmes. The testbeds focus, in particular, on the following common goals: creating flexible but robust high-speed net-

works and developing new technologies for distributed and parallel computing over those networks. These goals are being pursued in a way which is consistent with preserving the strengths of the older Internet computing environment: that is, permitting a multiplicity of types of equipment to be interconnected so as to promote as much competition as possible among alternative vendors in order to give low prices and maximal flexibility.⁵⁷ In pursuit of these goals, the testbeds share more specific objectives: eg comparing ATM with other broadband switching technologies, creating interfaces between HiPPI LANs and ATM/SONET systems, addressing issues created by real-time or near real-time network applications, and modifying or replacing the TCP/IP protocols to deal with problems of real-time computing over high-speed networks.

Another common theme underlying all the testbed research is the need to protect the supercomputing centres from federal and state budget cuts by widening their research programmes to include advanced networking and gigabit applications of interest to legislators. In doing this, the supercomputing community, DARPA and the NSF are responding to the need to win political support from academic scientists and engineers who are not directly involved in computing and network research. They are also responding to the perceptions of legislators that the connection between expenditures on networks and supercomputers and US economic competitiveness is too abstract (a problem for other 'big science' programmes as well). The legislators therefore want the supercomputing programmes to produce more tangible results, not just in advancing basic research but also in directly promoting commercializable new technologies.

Many of the gigabit applications involve 'big science' priorities in medicine, chemistry, earth sciences, meteorology, seismology, and 3-D rendering and visualization, thanks to the FCCSET and OSTP focus on 'Grand Challenges' in science and technology in their influential reports on high-performance computing. This is, of course, wholly appropriate for NSF-funded research projects and typical of US R&D policy, but it does not really address the more short-term concerns of legislators about strengthening US competitiveness.

The participation of private firms – mainly computer and telecommunications equipment companies and telecommunications carriers – in the testbeds is evidenced in applications research on teleconferencing, chemical plant management and distance collaboration. The work on visualization and particularly volume-rendering in science is likely to be easily adapted to business applications, as is the research on video broadcasting and multicasting. The participation of private firms is pushing the scientists and engineers to make sure that existing computer and telecommunications equipment works with the new networks and that the 'not-invented-here' syndrome does not dominate their thinking.

Towards universal access to the Internet

The discussion has focused so far on the development of the Internet and how the US portion of the Internet backbone (the NREN) is likely to be upgraded to gigabit speeds. As such, it has been a study of an evolution in telecommunications technology by a 'technical elite' from the upper echelons of academia, government and business. But that elite realizes that the future of the Internet will soon involve much wider

⁵⁷For an interesting discussion of goals for the evolution of the Internet see David Clark, Lyman Chapin, Vint Cerf, Robert Braden and Russ Hobby, 'Towards the future Internet architecture', RFC 1287, December 1991. This is available from any repository of Internet documents, such as nis.nsf.net, by anonymous file transfer protocol (FTP).

access, and they have planned the testbeds to develop technologies that will enable the Internet to adapt successfully to the explosive growth in usage that it has experienced in the last few years. Their strategy is to find rough equivalents to the things that allowed the Internet to thrive in its evolution from 56 kbps to 45 Mbps data transfer rates and from a few thousand to 4–5 million users.

The Internet of the future will have to deal with a much larger group of users, from much more diverse user communities. The Internet Activities Board, for example, is thinking about making the Internet capable of handling 1 billion networks.⁵⁸ In order to service these users adequately without degrading the performance of the network, the current megabit backbone will have to be upgraded to a gigabit backbone – that is, they have to build larger pipes.

Because the kinds of applications on the network are moving in the direction of greater diversity and interactivity, some requiring real-time performance that makes the network much more transparent to its users, many of the older ways of managing the network have to be rethought and re-engineered. Some applications will have to be granted higher priority for network delivery than others. For example, videoconferencing applications cannot be delivered the same way as e-mail, because it does not matter if there are delays in receiving e-mail, but real-time video signals break down quickly if a certain number of video packets do not get to their destination on time. Similarly, there is likely to be more multicasting in the future networks than there is now, and this needs to be dealt with at the level of network architecture.

To get a feel for the growing diversity of user communities on the Internet, we examine two important and relatively new user communities in this section: public libraries and K-12 schools. Not only are employees of libraries and schools using the network, but also the users of libraries and the student bodies of schools. Since there are over 30 000 public libraries in the USA, and many millions of library users, this is likely to be an important source of demand for network capacity in the near future.⁵⁹ The numbers are considerably larger for K-12 schools.

Interconnecting the nation's public libraries

The interconnection of public libraries would vastly increase the use of the Internet by the general public. Over 200 large university libraries are already on the Internet, according to the Coalition for Networked Information.⁶⁰ Many major government and university library catalogues and databases are already accessible on-line over the Internet.

For example, all major university libraries use the Internet to access centralized cataloguing systems like OCLC and RLIN. The MELVYL system in California merges all the on-line public access catalogues (OPACs) of the California system to greatly ease bibliographic searches of works in all the collections and to facilitate interlibrary loans. Because book and journal prices have been rising rapidly in recent years, there have been great incentives to reduce acquisition expenditures through greater use of interlibrary loans. The MELVYL system is accessible to all library users via terminals in the libraries and via dial-up systems. Most users have switched from using the paper version of the card catalogue to the electronic version, because the latter is more accurate, up to date and reasonably easy to use.⁶¹

There is no question that there is strong demand for library intercon-

⁵⁸The network capacity desired by various members of the IAB varies from 10 million to 10 billion, so one should not take the 1 billion goal too literally. This is one of the IAB's highest priority items because the current system is running out of unique names for the addresses of networks being added to the Internet. The current system works on the basis of a 32-bit address space, and the IAB is looking into several alternatives, including moving to a 64-bit address space.

⁵⁹See Jack Kessler, 'Library use in the United States of computer, networks, and broadband: an evolution, a retrogression?', Department of Information Science, University of California, Berkeley, March 1992, Appendix A1.00.

⁶⁰Communication to the authors from Vint Cerf, March 1992. See also Lois M. Kershner, 'A public library perspective on the NREN', paper delivered at a conference on Libraries and the National Research and Education Network, June 1990, pp 19–20.

⁶¹Gary Strong, Kathy Hudson and John Jewell, 'Electronic networking: California state and public libraries', paper delivered at a conference on Libraries and the National Research and Education Network, June 1990, p 20; and Kessler, *op cit*, Ref 59.

nection across universities, but the same factors which have pushed universities to become part of the network are likely to operate on other public libraries, even though their acquisition budgets and their cataloguing costs are more modest. There is a growing trend towards the supplementing of print publication methods with electronic (and particularly digital) ones. Some new 'publications' are circulated only in electronic form. Librarians wishing to help users get information from these sources must have access to the networks. The growing demand for the availability of audio and video tapes, compact discs and CD-ROMs to library users has been reflected in library acquisitions as well. As high-performance network technology becomes more accessible and affordable, even smaller public libraries will be able to send large text files and even video images over the network along with the more limited text and symbolic data they now provide.⁶²

The main obstacles to making library interconnection a useful network service are of both a technical and a legal nature: the legal problems are significantly harder to deal with than the technical.⁶³ For instance, current copyright laws do not allow the unlicensed use of the full text of library materials over networks. Unlicensed use constitutes an infringement of the 'display rights' of copyright owners. New intellectual property protection methods and guarantees are needed to create incentives for owners of copyrights to make their property available over the networks.⁶⁴ The main technical problems concern the interconnection of existing library computer systems (some of which are based on proprietary or OSI interconnection standards) with the TCP/IP-based Internet. The Internet itself is likely to evolve towards transparency to both TCP/IP and OSI systems, especially as it moves towards gigabit transfer rates, so this problem may not be very important beyond the immediate future. Thus we conclude that the real problem is protecting intellectual property rights.⁶⁵

There is already a strong demand for library interconnection on the Internet. This demand is sure to increase as network technology gets cheaper and easier to use. By connecting to the Internet, libraries will be better able to serve their existing users. In fact some new users (the handicapped and the geographically isolated) may begin to use library services *only* after the libraries are brought onto the network. But before this potential can be realized, public funding will be needed to cover the expense of adding new network linkages and some new methods for protecting the intellectual property rights of the creators of library materials need to be devised.

Interconnecting the nation's schools

A few local and state educational networks are now linked to the Internet and a number of computer bulletin boards are available to educators. Teachers, as well as students, have been very successful with their Internet endeavours. But progress in this area has been relatively modest to date because of the failure of most educational policy makers to realize the potential.

The Internet has been used to provide interconnection between local high schools and major universities. For instance, a Pacific Bell grant funded a link between Davis Senior High School (in the city of Davis, CA) and the University of California at Davis (UCD). Through this link UCD provided an expanded curriculum for the high school. The Davis link increased opportunities for multilingual and disabled students to

⁶²Kessler, *ibid.*

⁶³Strong, Hudson and Jewell, *op cit*, Ref 61, p 20.

⁶⁴Edwin Brownrigg, 'Developing the information superhighway', paper delivered at a conference on Libraries and the National Research and Education Network, June 1990, p 7.

⁶⁵It should be noted that this is one of the major foci of the work at the Center for National Research Initiatives.

receive personalized instruction. Teachers were able to get new ideas and information from university instructors in specialized areas.⁶⁶

Another important experiment in K-12 connectivity is the Texas Education Network (TENET) funded by the Texas Education Agency (TEA). Texas has 1050 school districts, 6400 public school campuses, 200 000 teachers and more than 3.2 million students. After an aborted attempt in 1989–90 to build their own network, the TEA decided to use the Texas Higher Education Network (THEnet), which is a mid-level network for the NSFNET system, as the principal means for interconnecting K-12 schools. TENET uses dial-ups of THEnet university nodes via 800 lines to get access to the Internet. The school terminals, which are mainly Macintosh or MS-DOS microcomputers, can use the Internet for electronic mail, bulletin boards, USENET conferencing systems and access to university databases. As of May 1991 TENET had 10 000 accounts, with 50 new accounts requested per day.⁶⁷

Through various educational networks such as the FrEdMail Network (Free Educational Mail Network) and various discussion groups on BITNET, e-mail availability has been extended nationally and to a lesser degree internationally to teachers and students in K-12 schools. FrEdMail has expanded to over 120 nodes including nodes in Australia, Canada and Ireland.⁶⁸ One of the goals of the FrEdMail network is to 'Promote and foster the development of a low-cost, community-based, distributed electronic data communications network . . .'⁶⁹ FrEdMail has provided an easy and inexpensive way to set up an individual Internet node.⁷⁰ The minimal hardware set-up requires: an Apple II computer (but not Apple IIc), a modem, a telephone line and the \$60 FrEdMail communications software package. The only other major expense involved is the cost of telephone charges. These costs are kept low by calling in off-peak hours.⁷¹

E-mail over FrEdMail and BITNET has allowed teachers to collaborate on various instructional ideas and projects nationwide. Students are now able to adopt electronic penpals in the USA and in other nations through FrEdMail.⁷² Some studies have indicated that the sending of electronic mail to penpals over computer networks has improved students' reading and writing skills.

In addition, various bulletin boards and discussion groups have been established for K-12 schools. For example, on the FrEdMail network educators can choose from more than 10 different bulletin boards specifically dedicated to mathematics, social sciences, foreign languages, ideas and collaboration for projects and activities, and projects for disabled students.

Other educational bulletin boards are accessible through direct dial-up connections. For example, the OERI of the US Department of Education runs its own bulletin board for public access of announcements and files of data and information. Also, the New York City Public School System operates NYCNET, which is no longer accessible by an 800 number because it became too popular.⁷³

As high-speed network technology becomes cheaper and more readily available, distance-learning – which requires video transmission to be fully effective – can be provided. Currently, handicapped children suffer from the lack of distance-learning services, so they are likely to be major beneficiaries of inexpensive broadband educational technology.⁷⁴ Also, high-speed networks may make it possible to enhance the curriculum at high schools across the nation by giving them access to state-of-the-art

⁶⁶Janet Meizel, 'Electronic networking at Davis Senior High School', paper delivered at a conference on Libraries and the National Research and Education Network, June 1990, pp 15, 16.

⁶⁷Tracy LaQuey Parker, 'TENET – the Texas Education Network', University of Texas, May 1991, manuscript sent to the authors via Internet electronic mail, 19 May 1992.

⁶⁸Bob Shayler, 'FrEdMail Network Inter-System map', *FrEdMail News*, Winter 1991, p 19.

⁶⁹Goals of the FrEdMail Foundation', *FrEdMail News*, Winter 1991, p 3.

⁷⁰It should be noted that a FrEdMail node is only for electronic mail and is not a full-function Internet node.

⁷¹Interview with Jim Edson, FrEdMail Network SysOp at Indiana University, Bloomington, IN, 25 November 1991.

⁷²*Ibid*; see also Jean W. Pierce, 'Computer networking for educational researchers on Bitnet', *Educational Researcher*, January–February 1991, pp 21–23.

⁷³Pierce, *ibid*.

⁷⁴Charles E. Crume and Cleborne D. Maddux, 'Educational computer networks: an overview', *Educational Technology*, July 1990, p 30.

educational materials, thus producing a higher-quality and more equitable secondary school system.

It should be noted, however, that this is not the way we are heading now. The main debate in secondary school applications of broadband technology currently is the expanded use of 'educational' TV with lots of commercial advertising such as that currently provided by the Whittle Corporation's Channel 1 system. Many K-12 educators are concerned about the diversion of funds from teacher salaries and basic educational equipment towards high-tech electronic systems. The latter are likely to be used successfully only in schools where there has been some considerable expenditure on the training of teachers and students in basic computing and networking skills. So one should not see the interconnection of K-12 schools as a solution for the current ills of the public school system, but rather as an opportunity to reorganize the schools and to upgrade the skills of both teachers and students at (hopefully) a relatively low cost.⁷⁵

Free-nets and the national public telecomputing network

Some argue that the USA needs a concept broader than education internetworking at the K-12 level and national library interconnection. One such broader concept is 'free-netting', or community-based networks. According to this democratic vision, each free-net provides: e-mail; universal access to information in the areas of health, education, government, technology, the arts, recreation and the law; and public access to on-line library catalogues. The Cleveland Free-net, a pioneer in the free-netting area, contains bulletin boards for over 300 different special-interest groups for all city residents. Cleveland's 'electronic city', which opened in July 1986 with about 500 logins a day, has been very popular. By June 1990 the system received over 5000 logins each day.

Some visionaries would like to see the spread of the free-net idea to other cities, eventually to form a National Public Telecomputing Network (NPTN) through a federation of free-nets. Clearly this would require a substantial increase in public funding of infrastructure development at the local level.⁷⁶ It should be noted, however, that the federal government pays only around 10% of the total costs of maintaining the Internet in the USA (\$60–100 million per year), with universities and local and state governments picking up the rest of the tab (over \$600 million per year).⁷⁷ Just as important as the increased infrastructure costs would be the modification of user interfaces to make network services easier to use. The FrEdMail experience suggests that this is possible, as does the experience with foreign national data services like those connected with the French Minitel system. But there has to be a decision to harmonize the user interfaces of the diverse data services on a large and heterogeneous network for such a strategy to succeed.⁷⁸

Proposals for a national public network

The most recently publicized democratic vision for the future of the networks is the proposal of the Electronic Frontier Foundation for a national public network (NPN). The Electronic Frontier Foundation (EFF) was started in the summer of 1990 by Mitchell Kapor and John Perry Barlow. Kapor was the founder of Lotus Development Corporation. EFF is also supported by one of the founders of Apple Computer Corporation, Steve Wozniak, and by major figures in the computer industry like Esther Dyson and David Farber. EFF is pushing for an

⁷⁵For an excellent set of recommendations, see John Clement, 'Networking for K-12 education: bringing everyone together', unpublished paper, Educom, Washington, DC, 3 May 1991.

⁷⁶The information on free-netting and the National Public Telecomputing Network is based on two sources: Tom Grundner, 'Free-netting: the development of free, public access community computer systems', paper delivered at a conference on Libraries and the National Research and Education Network, June 1990, pp 17–19; and Kathleen L. Maciuszko, 'Community online systems', in Allen Kent, ed, *Encyclopedia of Library and Information Science*, Marcel Dekker, New York, NY, 1992. Additional information was obtained in an electronic mail discussion with Tom Grundner on 22 November 1991.

⁷⁷These rough estimates were confirmed in the written comments of Vint Cerf to the authors on an earlier draft, March 1992.

⁷⁸See Jeffrey A. Hart, 'The Teletel/Minitel system in France', *Telematics and Informatics*, Vol 5, May 1988, pp 21–28.

NPN that provides universal and inexpensive access to data through a Narrowband Integrated Services Digital Network (N-ISDN) built and maintained by the telephone companies.⁷⁹

An enormous debate has sprung up in the computer and telecommunications communities about this proposal. The main problem, according to some telecommunications engineers, is that the NPN may result in excessive prolongation of the life of the copper cabling in the telecommunications infrastructure. You can get universal N-ISDN quickly only if you do it over existing copper wires. N-ISDN requires less sophisticated switching equipment than broadband ISDN (B-ISDN). Thus, according to the NPN opponents, the NPN may delay the transition from copper to fibre and from current switching technology to broadband switching, thus making it impossible for the USA to keep up with the state of the art in telecommunications technology.

The NPN makes sense to the network engineers only as a method for broadening access to the Internet, and then only if the prices of N-ISDN services are low enough to convince subscribers that it is worth paying a premium over their current telephone rates to get access to the Internet. According to these critics of the NPN, broad deployment of N-ISDN services might delay moves towards broadband ISDN (B-ISDN) because N-ISDN requires investments in N-ISDN central office switches and related equipment which will have to be amortized. In addition, N-ISDN will require new investments in copper wires and coaxial cabling, just at a time when optical-fibre cabling is beginning to become economically competitive with copper. Thus one would want to move only gradually and incrementally towards universal N-ISDN connectivity, and with an eye towards easing the transition to broadband to the kerb and neighbourhood as soon as possible.

Summary

In light of the current strong demand for and relatively small costs involved in using the Internet for educational purposes, there appears to be great potential benefit connected with continuing to subsidize public library interconnection and the extension of Internet access to K-12 schools. Experiments with K-12 networking, such as the Texas Educational Network and the FrEdMail Network, make this reasonably clear. Internet connectivity is considerably cheaper and provides a much broader access to information than does the building of dedicated library and school networks. More importantly, the extension of network services to library users and schoolchildren and their teachers is likely to result in a large increase in demand for network services generally. Even if these users stick to simple applications like e-mail and file transfers, their numbers are so large that they will create a major increase in demand for network services.

Conclusions

There have been five main visions of networking embodied in the history of the building of the ARPANET, the NSFNET and the Internet, and the NREN and gigabit testbeds: (1) the military vision, (2) the computer science vision, (3) the elite academic vision, (4) the business vision, and (5) the general educational vision. The military vision was married to the computer science vision during the early years of the ARPANET.

⁷⁹For elaboration see 'The national public network begins now, and you can help to build it', *EFFector Online*, 6 November 1991. This newsletter is available via the Internet by sending an e-mail request to eff@eff.org.

The elite academic vision became dominant with the decommissioning of the residual ARPANET and its replacement with the NSFNET and the Internet. It is the dominant vision in the High Performance Computing and National Research and Education Network Act of 1991. Most of the research in the gigabit testbeds is consistent either with the computer science or the elite academic visions, even though the underlying rationale is to make the Internet capable of adapting to its currently explosive rate of traffic growth.

The business vision was incorporated in the decision to allow commercial enterprises to interconnect with the Internet through private Internet interconnection services firms. It has also entered into the planning for future networks through the public and private funding of business-oriented gigabit applications research.

The general educational vision was implicit in the initial proposals for the NREN, particularly in the speeches of Senator Albert Gore, and explicit in the proposals by the Electronic Frontier Foundation for a national public network. This vision has taken concrete form in experiments like the Texas Education Network and the FrEdMail Network and efforts to interconnect public libraries across the nation. The public educational vision is somewhat inconsistent with the other four visions because it requires more user-friendly interfaces.

The history of the Internet is a history of the incorporation of more and more inclusive visions of participation in the benefits of computer networking. Traffic is growing exponentially, users are becoming more numerous and diverse, and applications are moving in the direction of real-time collaboration over the networks. What is needed now is the commitment to develop the technologies that make the Internet capable of dealing with this explosive rate of growth in users and the greater diversity of applications and user communities on the network.

The history of the Internet suggests also that there is a viable alternative to the vision for the future of broadband networks that has been promulgated by the telephone companies. The telephone companies have been arguing and investing on the basis of their belief that the country needs to move to broadband capability by allowing them to provide video (mostly one-way cable TV) services to their subscribers. This will give them, and their competitors in the cable TV industry, they argue, an incentive to lay more fibre more rapidly and closer to homes, factories and offices than would otherwise occur. It will also give them an incentive to develop faster the broadband switches and other technologies that will make the public switched networks more capable and eventually cheaper to operate.

The main sticking point to all this, however, is that the type of service that the telephone companies wish to offer to justify higher subscriber fees – cable TV – is a one-way service that does not require switching. While it may make sense to make the cable companies and the telephone companies compete with one another for both telephone and cable services, on the premise that competition is generally better than monopoly, even if the monopoly is regulated, nevertheless, how one gets from telephone-company-supplied cable TV service to a national broadband public switched network is by no means clear. Most importantly, the subscribers who watch one-way cable TV are unlikely to learn from the cable programming how to take advantage of the interactive services they will eventually get when the full broadband network is finally available.

The history of the Internet provides a useful and perhaps more realistic alternative vision of the transition to broadband networking. The users of the Internet, who are expanding rapidly in numbers and diversity, are learning how to use the services that will be more widely available when national broadband networks are in place. Some of the current users of the Internet will have a chance to innovate new broadband services via experiments like the gigabit testbeds discussed above. It is easier to subsidize and cross-subsidize in the Internet user community than in the current public switched networks, and to train users how to utilize the new services that will eventually be available with broadband networks. These differences between the Internet and the public switched networks may make the Internet a superior transitional vehicle towards broadband networking. New users of the Internet and its more capable successors – in the schools, businesses and government agencies of the country – are more likely than the telephone companies' new cable TV customers to contribute to increases in national productivity. They may not get you to universal broadband networks as fast, but they are quite likely to get you there more productively.