The National Science Foundation Network

Bilal Chinoy and Hans-Werner Braun*
San Diego Supercomputer Center
San Diego, CA 92186-9784

bac@sdsc.edu, hwb@sdsc.edu

Abstract

This paper describes the National Science Foundation Network, which evolved from a 56 kbps six-node network in the mid-eighties to today's 45 Mbps network, covering a large fraction of the research and education community via attached mid-level networks, and extending into a global arena via international connections.

Introduction

The Internet is a global computer network infrastructure spanning many countries. The U.S. Internet is loosely organized as a three-level hierarchy of networks consisting of Federal agency-sponsored backbones, attached mid-level or regional networks, and campus and institution-wide local area networks connected to these regional networks. This model places the NSFNET backbone at the top level of the connectivity hierarchy, along with the other Federal agency-sponsored networks such as the Department of Energy's Energy Science Network (*ESNET*) and NASA's Nasa Science Internet (*NSI*). The mid-level networks form the second tier of this structure while local campus and institution-wide networks form the third level.

The NSFNET backbone network

The original NSFNET backbone [NSF 87] was designed and implemented to interconnect the six NSF-sponsored supercomputing sites. These were

- *SDSC*: The San Diego Supercomputer Center
- *NCAR*: National Center for Atmospheric Research
- *NCSA*: National Center for Supercomputing Applications
- *CNSF*: Cornell National Supercomputing Facility
- *PSC*: Pittsburgh Supercomputing Center and,
- *JVNC*: John Von Neumann Center

The backbone network used 56 kbps leased lines for the long haul connections, with each of the six sites connecting to an Ethernet local area network.

The original objective of the network was to provide researchers with access to remote supercomputing facilities, but the scope extended rapidly to a broad support role for research and science in general. Other intended benefits were to speed dissemination of research results, enhance education through computer networking, provide an experimental platform for network research, and establish a United States national leadership in networking technology. The existence of the ARPAnet [ARPA 77] as a mature computer network infrastructure used by DARPA program researchers, and the interconnectivity of the NSFNET and the ARPAnet served to further promote this use of the network.

Quickly proving its effectiveness, the rate of growth of traffic and networks using the backbone via mid-level client networks eventually led to network congestion and service degradation. In response, the National Science Foundation solicited proposals to redesign and re- engineer the NSFNET backbone to enable it to sustain the predicted growth of the networking infrastructure. This five year project was awarded in late 1987 to *Merit Inc.*, a consortium of eight universities in the state of Michigan. Merit led an academic-industrial-government partnership with *IBM Inc.* and *MCI Inc.* [1] in the NSFNET backbone reengineering effort.

The redesigned NSFNET backbone network consisted of 13 packet switching nodes interconnected via T1 (1.5 Mbps) dedicated circuits. IBM provided the packet switch hardware and software and MCI provided the leased circuits.

These 13 backbone switch sites included the 6 original supercomputing sites and added 7 regional network interconnection sites.

*both authors were at Merit, Inc. prior to joining the SDSC

[1] The State of Michigan is also a project partner, contributing funds toward the NSFNET backbone project

- *BARRNET*: Bay Area Research Network
- *MERIT*: Michigan regional network
- *SESQUINET*: Texas area regional network
- *NORTHWESTNET*: NorthWest area regional network
- *SURANET*: South East area regional network
- *MIDNET*: Mid-American states regional network, and
- *WESTNET*: Western and mid-western states regional network

Additionally, the backbone connected to other US federal agencies networks via two Federal Inter-eXchange points (*FIX*), one at each coast.

The T1 NSFNET backbone became operational in July 1988, only seven months after the project award was made by NSF.

Mid-level Networks

Mid-level networks occupy the next lower level in the three level connectivity hierarchy [2]. These mid-level networks connect to the NSFNET and/or other federal agency backbones, and provide connectivity for their client networks, which are the campus or local networks. Most mid-level nets have traditionally received assistance from the NSF for managing their organizations and providing network services.

Local Sites

Local and campus networks are the clients of mid-level networks. These include universities, research institutions, federal installations, and other organizations involved in research and education.

Network Architecture

Topology

The NSFNET backbone developers faced the task of providing a robust, high bandwidth network interconnecting 13 sites spread over the continental U.S. The initial topology was to be robust enough to be fault-tolerant to circuit outages and to provide multiple packet-switched paths between any two nodes, yet not to exceed the financial budgetary constraints. Additionally, the initial versions of the T1 packet switch interface adapters were limited to switching 448 kbps, or about one-third of clear channel T1 bandwidth. Therefore the concept of a *logical* (or packet-switched) topology built over a *physical* (or circuit-switched) topology evolved. The initial physical topology is shown in Figure 1 and is termed the *T1 NSFNET phase 1 physical topology*.

A circuit-switched multiplexer and transmission resource manager was used to superimpose the logical topology over the physical topology. This circuit switch demultiplexed each physical T1 link into three 448 kbps logical links. A logical link could now be set up between any pair of nodes by circuit switching this link through the network. By attaching packet switches at the end points of these logical links, a richer logical or packet-switched topology was created. This is called the *T1 NSFNET phase 1 logical topology*, seen in Figure 2 [3].

The drawback of this scheme was that even though aggregate bandwidth between any pair of nodes on the network was T1, any single user could not get more that 448 kbps through the backbone. This was not a serious concern in the initial stages because most access (regional and state-wide) networks did not have T1 bandwidth themselves. Also, transport layer protocol performance was being improved to fully utilize the larger bandwidths the re-engineered NSFNET afforded over the previous 56 kbps network. While there were undoubtedly some applications that were constrained because of this bandwidth limitation, most end users were satisfied with network performance.

As time went on it became increasingly apparent that as traffic continued to grow, the 448 kbps per logical link would be a limitation. Additionally, the phase 1 network did have spur nodes that had no physical circuit redundancy. This meant that a circuit outage at the spur point would isolate the attached leaf node. Operational experience showed that such isolations were a significant fraction of total backbone outages. Additional circuits from those leaf nodes to other points in the connectivity graph were needed to address the redundancy issue. A third important factor in the network evolution was that the T1 interface adapters on the packet switches were now capable of driving a clear-channel T1 link, thereby obviating the need for a circuit multiplexer.

Therefore, in 1989, after approximately one year of network operations, the partnership redesigned the backbone topology. T1 links were added so as to maintain the backbone diameter at 3 hops, the circuit-switched multiplexers were removed and the newer T1 interface cards were used in the packet switches. The resulting topology is illustrated in Figure 3, the *T1 NSFNET*

[2] Mid-level networks have also been called "regionals", reflecting their geographical span, but we will use the term *mid-level* to reflect their hierarchical position in the architecture.

[3] In fact these logical links could be changed dynamically in response to packet switch failures or other outages. This allowed for a dynamically configurable logical topology

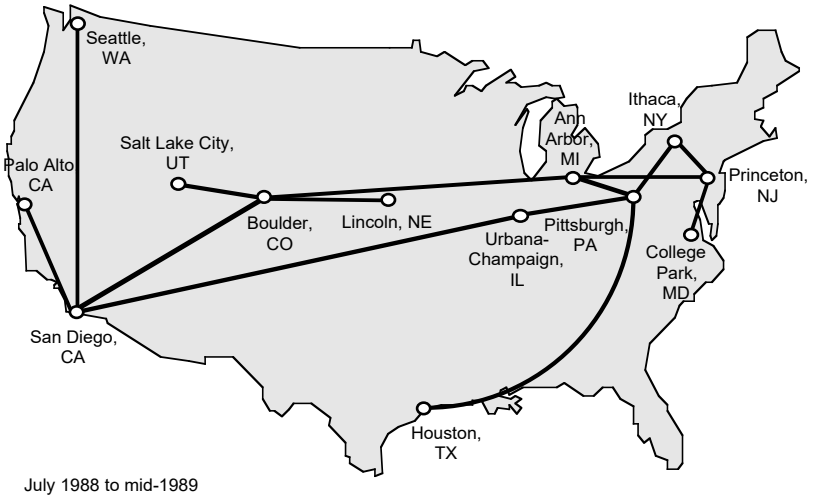

Figure 1: NSFNET Phase I physical topology

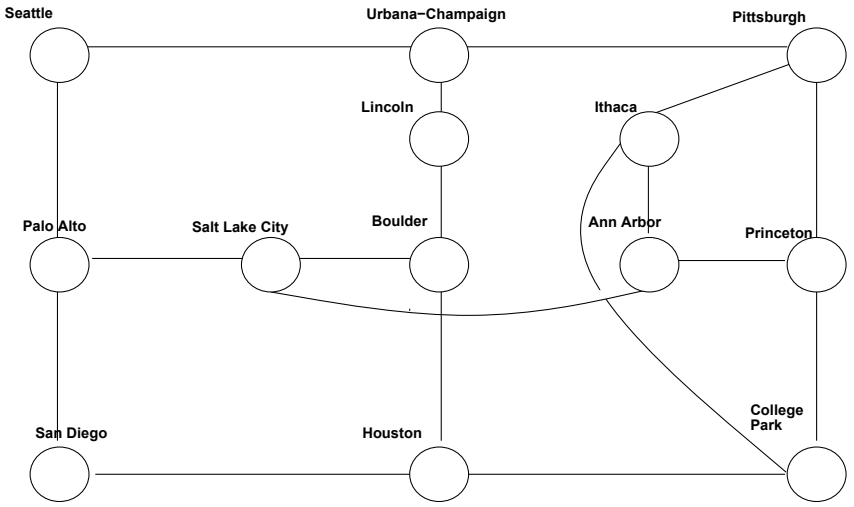

Figure 2: NSFNET Phase I logical topology

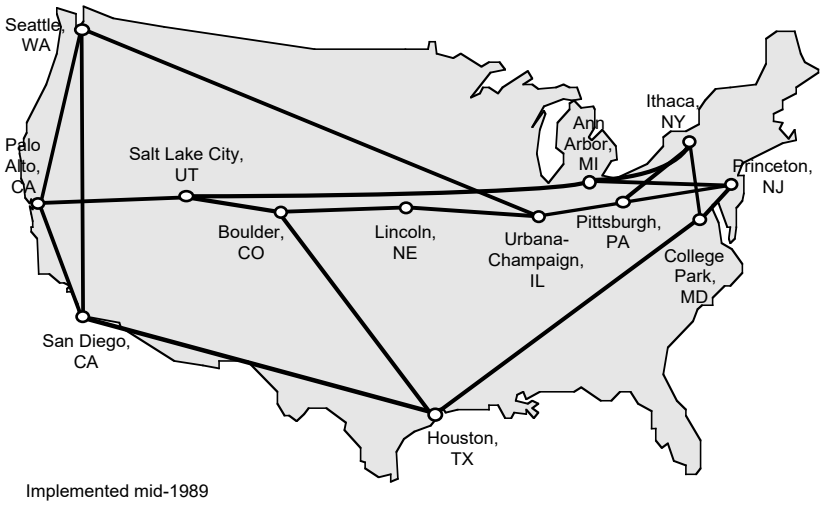

Figure 3: NSFNET Phase II logical topology

phase 2 topology. Note that the distinction between logical and physical topologies was not needed any more, as no circuit multiplexing was required.

Packet Switch

Each NSFNET backbone site houses a switching node (termed *NSS*, for Nodal Switching Subsystem). The NSS is composed of 9 IBM RT/PCs connected by dual 4 Mbps token rings. The NSS is thus modeled as a loosely coupled multiprocessor system, with the token rings forming the main data and control bus for the system. The RT/PCs run a flavor of 4.3 BSD Unix as their operating system, with enhancements to provide high-speed packet forwarding, routing and network management features. The RT/PCs are functionally of three types.

- *RCP* (Routing Control Processor)

  One machine, designated the RCP, runs software to perform the distributed routing algorithm to calculate routing tables and disseminate them to the other machines in the NSS. The RCP also performs network management functions by responding to management queries on behalf of the other machines in the NSS.

- *PSP* (Packet Switching Processor)

  These machines are designated as switching processors, with their function being to forward IP datagrams based on their destination addresses. Forwarding tables are computed on the RCP and loaded to each PSP. There are two flavors of PSPs.

  - PSPF: This machine attaches to the synchronous T1 lines via a CSU/DSU (channel service unit/data service unit). These machines forward packets to and from the T1 lines. There is one PSPF per T1 link coming into a NSS. Since there are typically three links into a NSS, three machines were designated as PSPFs.
  - Exterior-PSP: This is the packet switch that interfaces to the attached regional network via an ethernet. This is the exterior interface the backbone provides to the attached regional nets. Typically there is one E-PSP per NSS, but more machines have been used at high traffic nodes to split offered load across multiple interfaces.

- Packet Monitor

  One machine per NSS is dedicated to monitoring the dual token rings and characterizing packets switched on the rings by the contents of their headers [Braden 88] [4]. This information was invaluable in performing network capacity planning and performance studies, as well as debugging Internet system wide problems [5].

The *NTP* (Network Time Protocol) [Mills 88] allowed the hierarchical synchronization of backbone packet switch clocks. High accuracy external reference clock sources were used to synchronize the RCPs at each NSS. The PSPs at each node synchronized their clocks by using the local RCP as a reference, thus forming a hierarchical structure of clock synchronization. Clock synchronization is an important feature for network management schemes that rely on clock offsets to signal anomalous conditions.

[4] The software used to do this was a modified form of the **NNStat** package developed at the Information Sciences Institute, and publicly available

[5] For example, this information was used to quantify Domain Name System (DNS) traffic source, destination and volume which was long suspected to be misconfigured. In fact, this process decreased the DNS traffic substantially

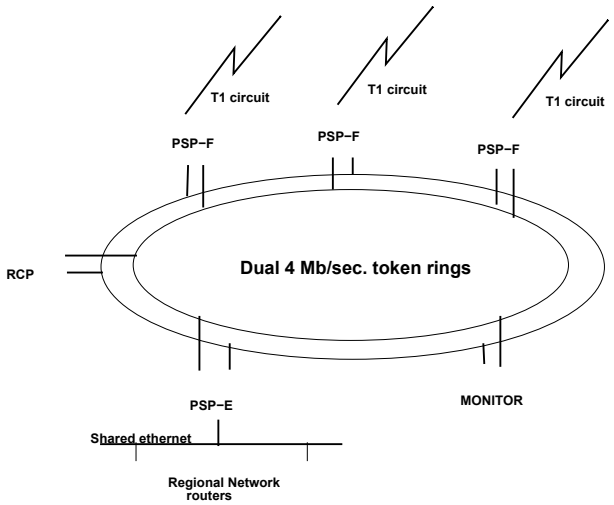

Figure 4: Nodal Switching Subsystem

A functional diagram of a typical NSS is shown in Figure 4.

Routing considerations

The TCP/IP Internet is organized as an interconnection of *Autonomous Systems* (AS). An AS is a collection of internetwork routers managed and administered by a single authority or organization. As an example, all the IP network numbers and routers used by the Pittsburgh Supercomputing Center are grouped as AS 204.

Using this scheme, the backbone is modeled as a single AS into which regional network AS(s) attach. At each connection point, one or more regional network ASs interact with the backbone in order to share routing and reachability information. We will separate the discussion of routing into two sections.

The NSFNET *intra-domain* routing or internal backbone routing handles the computation, maintenance and propagation of optimal routes between all entry/exit points on the backbone.

The *inter-domain* routing handles the exchange of reachability and routing information between ASs. For the NSFNET system, the interdomain protocols are used to share information between the backbone network and the attached regional networks.

Intra-domain routing

The NSFNET backbone uses the Intermediate System to Intermediate System (IS-IS) routing protocol to perform intra-domain routing [Rekhter 88] [6]. The IS-IS standard, while developed with OSI networks in mind, was adapted for use in IP networks. This protocol uses the Djikstra Shortest Path First (SPF) algorithm to compute optimal routes, with each link in the topology graph assigned a cost. This link cost is configured in a manner which ensures minimum packet switch hop routes between any two points on the backbone. Multiple equal cost paths between nodes are supported by the routing protocol.

The IS-IS protocol is one of a class of routing protocols known as the *link-state* protocols. In this approach applied to the NSFNET, each RCP periodically exchanges its state of the connectivity graph with all other RCPs. In this fashion each RCP maintains a database of link state. This ensures rapid convergence toward a steady state in the event of a link or node failure.

Inter-domain routing

The inter-domain routing interaction at the regional network attachment points allows routing and reachability information to be shared between the interacting ASs. The operation of the inter-domain protocols is a point-to-point peer interaction, with one *routing session* established between peer routers. The process of communicating this information is known as *advertising* routes. The attached AS advertises to the backbone the list of networks it can route packets to. Conversely, the backbone (local RCP) advertises the networks the backbone can route to (via the other NSS attachment points) to the regional AS. Network information is thus shared between the backbone and client ASs.

Reachability and Routing protocols

A reachability-based interdomain routing protocol does not carry any information about the path to a particular net, but only if the net is currently reachable. Routing-based protocols, on the other hand, contain qualitative information about the path to a destination network in addition to its current state of reachability. Examples of this information are the number of AS hops or the full AS path.

The protocol initially used for inter-domain routing in the NSFNET system was *EGP* (Exterior Gateway Protocol) [Mills 86]. This is a reachability protocol with no provisions for qualitative metrics. The rapidly growing connectivity and the lack of a globally engineered connectivity hierarchy made path selection critical, and inter-domain path information metrics became important.

[6] This ANSI routing protocol is in the OSI standardization process

In response to this, the *BGP* (Border Gateway Protocol) was designed and implemented in 1990 [Lougheed 90]). BGP has provisions for retaining the AS path of the route as well as explicit routing loop suppression and other features. It is important to note that the EGP protocol was one of the first inter-domain protocols to be designed and worked very well for its target environment.

All backbone client networks, then, exchange routing and reachability information with the NSFNET backbone via BGP or EGP [7]. Mid-level networks typically use their own intra-domain routing protocol within the confines of their AS boundaries.

Policy-based routing

A critical observation made during the operation of the 56 kbps NSFNET backbone was that some control on the believability of routing information was required and that there was a strong requirement for the *controlled* exchange of routing information between multiple independently administered domains. The process of preventing information not explicitly required or requested by an Autonomous System from reaching it is known as *firewalling*.

To this end, the T1 NSFNET backbone RCPs consult a mapping table that specifies allowable network numbers and their parent ASs. Only these networks are actually incorporated into backbone routing tables and disseminated to the other regional networks. The mapping table also includes the notion of a connectivity priority per network/AS pair, using IP network numbers and AS identifiers. This priority mechanism facilitated a pre-designed multilevel backup path capability in cases where a network number is part of multiple domains. The outbound dissemination process to the attached mid-level networks is also similarly controlled, both in terms of AS identifiers as well as IP network numbers.

The primary benefit of this process was that only networks that were mutually agreed upon by all parties concerned were exchanged via the inter-domain protocol interaction. This meant that the backbone always carried a well-defined set of networks in its routing tables. This greatly simplified network trouble-shooting and planning. The NSFNET Network Operations Center utilized a process by which client network managers could submit an electronic request to add or change information in this policy database. This information was verified for authenticity and was then incorporated into the backbone configurations, at which time packets sent to the NSFNET backbone destined for those networks could be handled appropriately.

[7] Mid-level networks still using EGP are in the process of transitioning toward BGP

The datagram forwarding process

IP datagram forwarding is the process by which an IP router accepts a datagram arriving on an input interface and, based on information in the datagram IP header, decides to which of its outgoing interfaces the packet should be *forwarded*. Most IP networks, including the NSFNET backbone, forward packets based on the destination IP network address. This is done by extracting the IP destination address from the datagram and consulting the router's forwarding tables for a match. The forwarding tables generally map the destination address to the outgoing interface or the *next hop* router.

The NSFNET backbone routers perform the forwarding process by consulting two separate tables. The first table is the mapping between destination networks and the autonomous system to which they belong. This is known as the *network* table. Having determined the appropriate exit or destination AS, the routing software must consult a second table known as the *regional* table, to locate the next hop interface for that AS. Therefore, the traditional network → next hop lookup has been modified to be a network → AS → next hop scheme.

The motivation behind this dual-lookup scheme comes from the fact that the NSFNET NSS is a collection of machines of which only one (RCP) performs the routing calculations. The RCP downloads this information to all the other packet switches. When the route to a particular AS changes, only the AS → next hop mapping on the packet switches needs to be changed and not the entire set of networks that are associated with the AS. Route change overhead is thus reduced.

Priority queuing

Packets arriving at a router's input interface are placed in a queue to await servicing. Under light loads the length of this queue is normally zero, and packets are serviced (forwarded) as soon as they arrive. If there is a non-zero packet queue then the backbone routers use the *FCFS* (First Come First Serve) queuing discipline.

Queuing routing and network management packets behind other traffic could potentially affect the operation of the backbone network itself. To avoid this potential disruption all backbone routers manage dual queues, one for normal traffic and the other for all priority traffic. Packets are taken off the priority queue to be serviced before the normal queue is serviced. FCFS is still used within each queue. High priority packets are any packets that are either sourced from or destined to a backbone IP address. Therefore all user traffic is considered lower priority than management traffic.

Traffic and workload characteristics

Offered load to the NSFNET backbone increased dramatically since the first month of operation, as seen in Figure 5. This graph plots the aggregate of all traffic entering the NSFNET backbone network from all entry points.

Additionally, the number of networks connected to the NSFNET system also grew rapidly. Starting from approximately 300 networks, the NSFNET backbone carries more than 5000 total networks in its routing tables. This trend is plotted in Figure 6, which shows the count of all networks included in the NSFNET backbone routing admission database. The rise in the number of foreign networks attests to the international impact of the NSFNET backbone.

The variation in offered load over the course of a day is plotted in Figure 7. This graph shows 3 distinct periods in the daily variation. A low utilization period occurs between the hours of 6 am and 3 pm GMT. Network load increases steadily between 3 pm and about 11 pm EDT, after which time it begins to decrease. This pattern has remained consistent since the inception of the network.

Figure 8 shows the distribution of traffic by application type, both by packets as well as by bytes. Interactive services, file exchange and mail services make up the majority of this distribution. Note that the larger packet size used by file transfer protocols manifests itself as a larger piece of the distribution when considered by bytes as opposed to packets switched on the network. This distribution has also remained remarkably consistent to date, and Internet researchers expect that in the future a larger fraction of traffic will be windowing or remote file system applications as the networking infrastructure becomes faster and more robust.

An interesting question was whether the offered load to the backbone displayed any *concentration* properties. That is, whether a few networks sourced (and sank) a large fraction of the total traffic. Figure 9 shows the time variation of the percentile of all networks that offer a fixed amount of the total traffic. From the graph it is seen that 25% of the traffic on the backbone is sourced by only about 1% of networks, 50% by about 3% of the networks and 75% of traffic by about 9% of networks. Furthermore these numbers have remained fairly steady or have decreased over the period shown. This decrease can be attributed to the fact that newer networks added are, for the most part, smaller volume contributors to the aggregate traffic. This high degree of traffic concentration is an important feature of Internet traffic and can be taken advantage of in a variety of different schemes [Chinoy 92].

Network Management

Merit and its partners have adhered to the network management standards put forth by the Internet community. However, when the T1 NSFNET backbone became operational there were no clear standards for IP network management protocols and techniques. The first signs of the protocol standardization came when the Internet community developed the *SGMP* (Simple Gateway Management Protocol). This was a request-response type protocol which specified a management client and an agent which responded to queries. The backbone routers incorporated SGMP agents and the Merit Network Operations Center (NOC) managed the network using a SGMP graphical client developed by IBM Research. Soon thereafter the *SNMP* (Simple Network Management Protocol) was introduced and gained community acceptance [Case 89]. The backbone packet switches ran SNMP agents which responded to standard Management Information Base (MIB) queries using the SNMP protocol. The client was a graphical SNMP program which later evolved into the IBM XGMON management system. Merit began development of its own graphical network monitoring system called the *Internet Rover*, which became the primary tool for NOC operation.

Additionally, the NOC used the *ping* (Packet Internet Groper) to monitor the connectivity of packet switches. Ping generates an ICMP echo request packet targeted at the monitored entity and expects an ICMP echo reply packet back indicating that the entity is network reachable. Merit also developed tools to combine both ping as well as SNMP to monitor the network.

The *traceroute* tool proved to be an invaluable diagnosis tool to determine routing and network path related problems. Traceroute uses the ICMP Time to Live (TTL) field to determine the path between source and destination. Each router along the path returns an error message to the source host, thereby identifying itself.

The Merit NOC is staffed 24 hours a day, 7 days a week by network operators who respond to network alarms and problems. The NOC is supported by an Internet Engineering team and has the combined resources of the IBM and the MCI NSFNET support personnel at their disposal. The NOC uses a *trouble ticketing* system based on the IBM Information Management Software package to track all problems and issues with the network to ensure the quickest and most informed response.

Network security

To improve *router* security a variety of steps were taken. Since the operating system of the NSFNET routers is

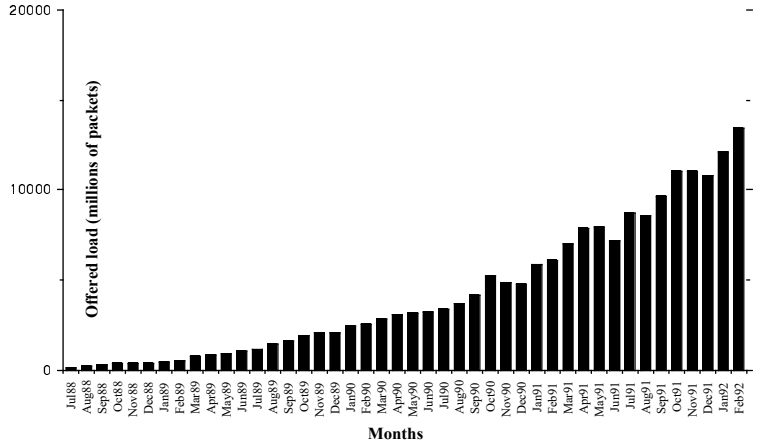

Figure 5: NSFNET growth in traffic

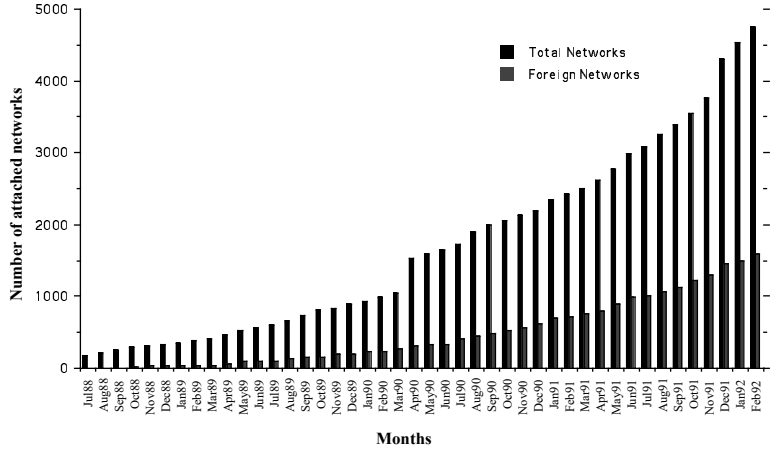

Figure 6: NSFNET growth in attached networks

Unix BSD 4.3 based, the known security flaws in the operating system were corrected. In addition, the *inetd* superserver was modified so that unless certain requests to spawn a daemon came from a known and trusted address (which was table configured), the service could not be started from the superserver. In addition, all unwanted services like the *tftp*, *smtp*, and the *finger* daemon were turned off.

To improve *network* security a series of filters were implemented. These could be turned on or off on a router interface basis to prevent traffic originating from or destined to a particular network from using the backbone.

Source filtering was used on all external interface to prevent traffic coming from the backbone address into the backbone. Since no legitimate traffic would enter the backbone pretending to be a backbone router, this prevented spoofing of backbone addresses, ensuring routing integrity.

Destination filtering could be used in an emergency to prevent all traffic destined to backbone addresses from entering the network. This was seen as an emergency measure only. Steps were taken to ensure that legitimate management traffic from the MERIT NOC would still be able to monitor the network in this state. Note that regional and client networks using SNMP do not need to query any backbone addresses in order to monitor the backbone network, so that traffic would not be affected. In addition to the above measure, the backbone routers have two levels of password based access control. These passwords are changed regularly. All critical information is logged to disk files and these files are also checked regularly for any anomalies.

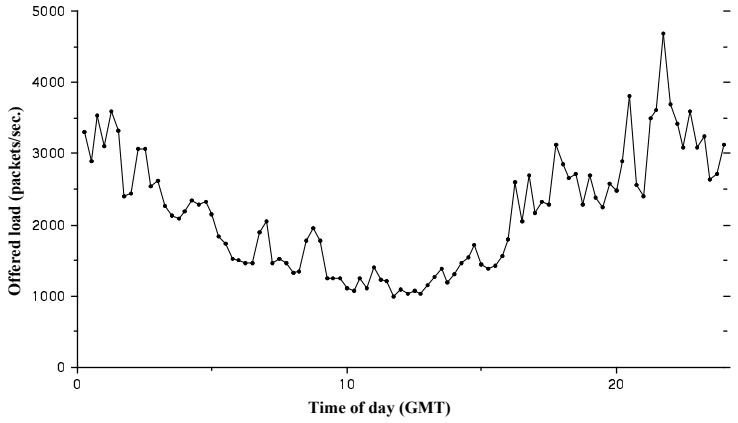

Figure 7: Typical load variation over a day

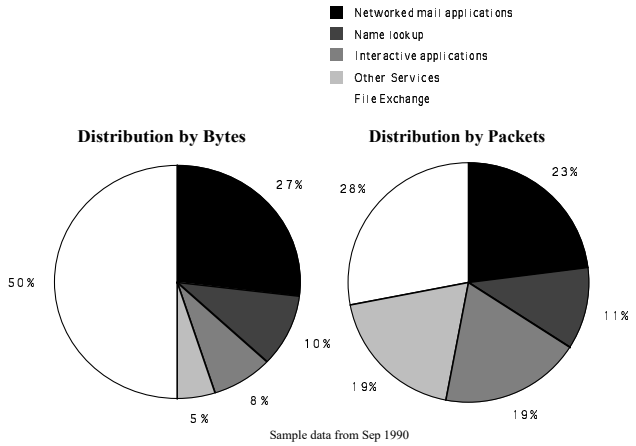

Figure 8: NSFNET Major Application Distribution

The T3 NSFNET backbone

The Merit proposal for the re-design of the NSFNET backbone foresaw the rapid growth in Internet usage and explicitly proposed a migration to a T3 (45 Mbps) based backbone. This step was also important in order to preserve U.S leadership in the wide-area Internet infrastructure. Specific planning for the T3 backbone began in the summer of 1989 with initial network requirements and architectures. It was decided to implement the T3 network as an *overlay* network to the existing T1 based backbone, and migrate traffic from the T1 to the T3 networks. In this architecture the T3 backbone is a logically separate subnetwork attached to the T1 backbone at a few selected locations. An important feature of the T3 network was the *co-location* of IBM provided T3 packet switches at MCI *POPs* or Points of Presence. This allowed for a more robust topology and effectively eliminated transit traffic from traversing an end-node. Details of this architecture can be found in [Chinoy 92]. After some initial difficulties involving the new T3 packet switches were resolved, the T3 network began to offload traffic from the T1 backbone in the fall of 1991. The current plans show the dismantling of the T1 network by the end of 1992.

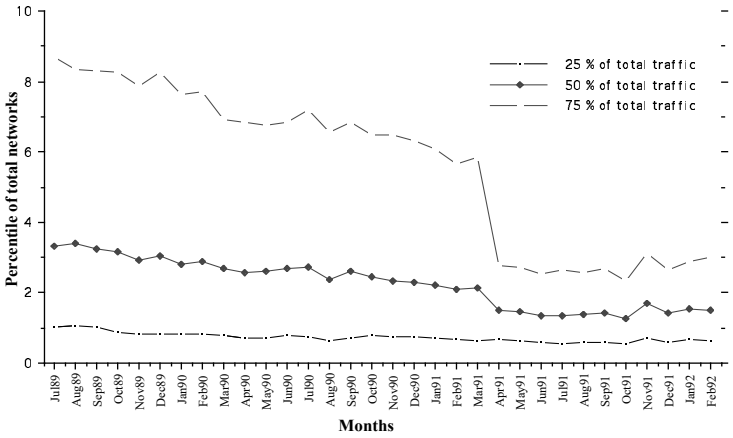

Figure 9: NSFNET Traffic Concentration by Network

Landmarks

The following is a chronological list of notable landmarks in the NSFNET project to date.

- NSF award received (Nov 1987)
- NSFNET T1 network operational (July 1988)
- ARPAnet decommissioned (Summer 1988)
- Packet video over NSFNET demonstrated (October 1988 at InterOp '88)
- Direct interagency connections established, which evolved to the FIXes (Dec 1988)
- Backbone topology redesigned (June 1989)
- *CA*NET* (Canadian National Network) attached (May 1990)
- IDNX multiplexers removed from network (Feb 1990)
- International link to Geneva, Switzerland (March 1990)
- IP connectivity at T3 between Washington D.C and Ann Arbor, MI (March 1990)
- OSI CLNP capability demonstrated (May 1990)
- First cross-country T3 IP connectivity (Dec 1990)

Acknowledgements

We would like to acknowledge Kim Claffy (SDSC) and Elise Gerich (Merit) for the valuable suggestions and feedback they gave us during the course of writing this paper.

References

[ARPA 77] McQuillan, J., Walden, D., "The ARPA Network Design Decisions", Computer Networks, Vol. 1, No. 5, Aug 1977, pp 243-289.

[NSF 87] Mills, D.L., Braun, H-W., "The NSFNET Backbone Network", Proc. SIGCOMM 1987 Workshop, Aug 1987, pp 191-196.

[Postel 81] Postel, J., "Internet Protocol", RFC 791, Sep 1981.

[Postel 81] Postel, J., "Internet Control Message Protocol", RFC 791, Sep 1981.

[Mills 86] Mills, D.L., "Autonomous Confederations", RFC 975, Feb 1986.

[Clark 89] Clark, D.D., "Policy Routing in Internet Protocols", RFC 1102, May 1989.

[Rekhter 88] Rekhter, Y., "The NSFNET Backbone SPF based Interior Routing Protocol", RFC 1074, Oct 1988.

[Braden 88] Braden, R., DeSchon, A., "NNStat - Internet Statistics Collection Package, Introduction and User Guide", ISI Research Report ISI/RR-88-206, USC ISI, June 1988.

[Braden 88] Braden, R., "A Packet Monitor Program", ISI Research Report ISI/RR-88-207, USC ISI, June 1988.

[Mills 88] Mills, D.L., "Network Time Protocol (v1): Specification and Implementation", RFC1059, Jul 1988.

[Braun 89] Braun, H-W., "The NSFNET Routing Architecture" RFC 1093, Feb 1989.

[Case 89] Case, J., et. al. "A Simple Network Management Protocol", RFC1098, April 1989.

[Rekhter 89] Rekhter, Y., "EGP and Policy Based Routing in the new NSFNET Backbone", RFC 1092 Feb 1989.

[Lougheed 90] Lougheed, K., Rekhter, Y., "A Border Gateway Protocol", RFC 1163, June 1990.

[Chinoy 92] Chinoy, B., Rekhter, Y., "Injecting inter-AS routes into intra-AS routing: A Performance Study", Computer Communications review, Vol 21, No. 1, Jan 1992.

[Chinoy 92] Chinoy, B., Smith, P., "Inside the T3 Network", Merit/NSFNET Link Letter, Mar/Apr 1992.