Traffic Characteristics of the T1 NSFNET Backbone

Kimberly C. Claffy and George C. Polyzos
`kc@cs.ucsd.edu, polyzos@cs.ucsd.edu`
Department of Computer Science and Engineering
University of California, San Diego
La Jolla, CA 92093-0114

Hans-Werner Braun
`hwb@sdsc.edu`
Applied Network Research Group
San Diego Supercomputer Center
San Diego, CA 92186-9784

Abstract

This paper presents the results of a measurement study of the T1 NSFNET backbone. We first discuss the measurement environment and approach to data collection. We then present measurements results for: long-term growth in traffic volume, including attribution to domains and protocols; trend in average packet size on the network, both over long and medium term intervals; most popular sources, destinations, and site pairs; traffic locality; international distribution of traffic; mean utilization statistics, both of the overall backbone as well as of specific links of interest; and delay statistics.

1 Introduction

The last general overview of measured behavior of the national research and education backbone infrastructure was the landmark study by Kleinrock and Naylor [13] (also in [12]), which presented measurements on the 1973 ARPANET. Since that study, the infrastructure has undergone significant change. The U.S. Department of Defense initially established the ARPANET to enable the network research community to investigate packet-switching technologies. Recognizing the potential of these technologies for the non-defense community as well, NSF later built the NSFNET follow-up, with an explicit mission to foster productivity within the research and education environment. The expanded visibility of operational computer networking has resulted in attracting other agencies and civilian organizations, including on an international scale, to further augment the network with their own network resources.

What is now a pervasive global infrastructure has received little attention in terms of empirical analysis and modeling. As a result, we lack a thorough understanding of the deployed network, and thus the capability to predict, much less secure, its behavior.

Recent studies on isolated aspects of the NSFNET investigated the existence of packet trains on the NSFNET backbone [11], and evaluated specific routing approaches for use on the backbone [9]. Feldmeier [10] studied the estimated performance of a gateway routing table cache. Caceres *et al.* [3] and Danzig *et al.* [8], profiled the characteristics of individual application conversations. Wakeman *et al.* [15] and Asaba *et al.* [1], present analyses of trans-atlantic and trans-pacific traffic, respectively.

Our characterization is based on available data collected since the establishment of the T1 NSFNET backbone in July 1988. We have also selected a specific month, May 1992, to investigate some traffic patterns in more detail. At that time the NSFNET was transitioning from the T1 to the new T3 backbone, and the lower traffic volumes of recent months reflect the gradual migration to the T3 backbone. In both cases we are using raw data sets collected by Merit Network, Inc. according to the procedures described in Section 3. Because the T3 backbone initially did not fully support data collection for all the statistics objects we present in this paper, we restrict ourselves here to the T1 backbone.

It is interesting to note that twenty years after [13], we can say less about certain performance metrics of the current networks than was possible in 1973. Partially responsible is the fact that the ARPANET was an experimental network with facilities that were specifically designed for extensive data collection to support network analysis. Today's environment is markedly different. While in many areas there is far more flexibility in assessing performance of the current infrastructure, the NSFNET project has targeted its instrumentation efforts towards operational rather than research requirements.

In the following section we present a brief description of the environment and the instrumentation for the data collection process. The measured data include: long-term growth in traffic volume, including attribution to domains and protocols; trend in average packet size on the network, both over long and medium term intervals; delay statistics; most popular sources, destinations, and site pairs; traffic locality; international distribution of traffic; mean utilization statistics, both of the overall backbone as well as specific links of interest; and, assessment of downtime for the last few years. The data indicates not only a change in the volume, but also the composition, or cross-section of traffic, over both long and short time horizons.

2 Current network infrastructure

NSFNET, the National Science Foundation Network, is a general purpose packet-switching network supporting access to scientific computing resources and data. Evolved from a 56kbps six-node network in the mid-1980s to today's 45Mbps network, the current NSFNET includes three different levels: the transcontinental backbone connecting the NSF-funded supercomputer centers and mid-level networks, the mid-level networks themselves, and the campus networks. The hierarchical structure includes a large fraction of the research and educational community, and even

---

This research is supported by a grant of the National Science Foundation (NCR-9119473), and a joint study agreement with the International Business Machines, Inc.

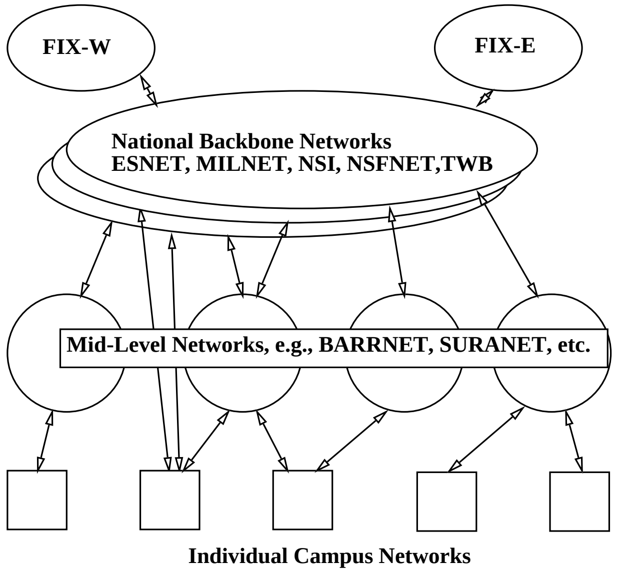

Figure 1: Hierarchical Model of NSFNET Architecture

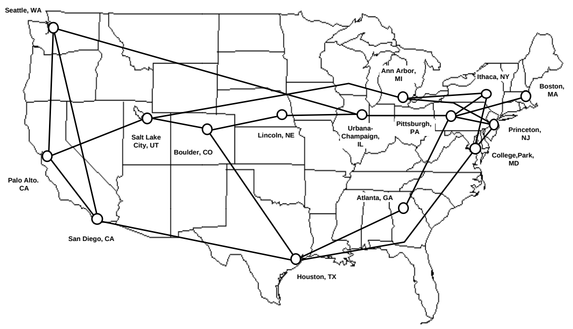

Figure 2: 1992 T1 Backbone Logical Topology

extends into a global arena via international connections. Figure 1 depicts a rough model of this hierarchy. ESNET, Milnet, NSI, NSFNET, and TWB correspond to national backbones of DOE, DoD, NASA, NSF, and DARPA, respectively. Chinoy and Braun [5] describe in detail the underlying topology of the NSFNET backbone in its various stages.

The National Science Foundation originally granted Merit Network, Inc. an award to engineer, build, and manage the T1 backbone network between thirteen NSF sponsored mid-level network sites. The NSFNET backbone is also connected to other networks of US federal agencies via two Federal IntereXchange points (FIX), at the East and West coasts. Merit completed the implementation of the original backbone in June 1988, and about one year later redesigned it to provide multiple paths to each of the 13 nodes. In 1990 the NSF augmented the award to Merit to add one additional T1 node in Atlanta, GA, and subsequently to engineer and implement a T3 network to the 14 T1 sites in addition to two new sites in Boston and Chicago. Some initial implementation of the T3 network was in place by the end of 1990, and the upgrade, including rerouting of traffic, was completed in 1992. Figure 2 shows the logical topology of the final T1 backbone.

United States NSFNET Mid-level Networks
BARRNET (Bay Area Regional Research Network, California)
CERFnet (California Education and Research Federation Network, California)
CICNet (Committee on Institutional Cooperation Network, Mid-West)
CO Supernet (Colorado Supernetwork, Colorado)
CONCERT (Communications for North Carolina Education, Research, and Technology Network)
INet (Indiana Network, Indiana)
JvNCnet (John von Neumann Center Network, Northeast)
Los Nettos (Southern California)
MichNet (Michigan Network, Merit-statewide, Michigan)
MIDnet (Midwestern States Network, Mid-West)
MRnet (Minnesota Regional Network, Minnesota)
NEARnet (New England Academic and Research Network, North-East)
netILLINOIS (Illinois)
NevadaNet (Nevada)
NorthWestNet (Northwestern States Network)
NYSERnet (New York State State Education and Research Network)
OARnet (Ohio Academic Research Network)
PREPnet (Pennsylvania Research and Economic Partnership Network)
PSCNET (Pittsburgh Supercomputing Center Network, Mid-West)
SDSCnet (San Diego Supercomputer Center Network)
SESQUINET (Texas Sesquicentennial Network)
SURAnet (Southeastern Universities Research Association Network, South-East)
THEnet (Texas Higher Education Network)
VERnet (Virginia Education and Research Network)
Westnet (Southwestern States Network)
WiscNet (Wisconsin)
WVNET (West Virginia Network for Educational Telecomputing)

The next layer of branching below the NSFNET backbone includes the mid-level networks.[1] (see Table 2). Mid-level networks connect, sometimes indirectly, to the NSFNET and/or other federal agency backbones, and provide connectivity among sites in the research and academic environment. Federal mission agencies also employ the services of mid-level networks for connectivity to their university-based researchers. The NSFNET backbone allows for peer network connectivity to networks such as national backbones, e.g., Advanced Network Services, AlterNet, Performance Systems International, and SprintLink. The backbone also supports international connections to national backbones of foreign countries.

Local sites attach as clients to mid-level networks. They include universities, research institutions, federal installations, and private commercial organizations.

2.1 Parameters of the T1 backbone

We review a few of the network parameters that affect traffic flow in the NSFNET backbone. All in-

---

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

terfaces to each node on the T1 backbone are of “T1” speed, 1.544 Mbits/second. To access external client networks, each T1 NSFNET backbone node uses Ethernet interfaces, limiting the packet size into a Nodal Switching Subsystem (NSS) to 1500 bytes. Packets typically contain a 20-byte IP header and, in case of TCP, another 20 bytes for the transport protocol.

Packets travel through the network individually and are passed from node to node according to an adaptive routing procedure based on the standard IS-IS protocol. Reassembly of fragmented IP packets occurs at the destination host before packets are forwarded to the application. Each processor on the T1 backbone can buffer up to fifty packets on the output queue of an interface. This buffering contributes to the latency of the delivery of packets to the destination. A T1 NSFNET backbone processor can switch in excess of 1000 packets per second, making the per-packet processing overhead less than 1 millisecond.

3 Data currently collected

In this section we describe the means by which measurements are performed on the T1 NSFNET backbone. Data collection, for the purposes of monitoring, measuring, and analyzing the performance characteristics of the network, is a fundamental requirement for the operation and management of any large-scale network such as the NSFNET backbone.[2] The T1 hardware is PC/RT-based: the switching node architecture consists of multiple processors dedicated to separate functions connected by a common token ring. One processor is dedicated to statistics collection. This processor had to eventually revert to sampling when the traffic volume per node surpassed its processing capability. We describe this situation further in Section 5.

The principal sources of information about the T1 network are the routine collection of three classes of network statistics: internodal delays; interface statistics, which rely on the Simple Network Management Protocol (SNMP)[4]; and packet categorization, performed with an NSFNET Network Statistics set of software, NNStat [2].

3.1 Internodal latency

NSFNET uses the ICMP Echo functionality to record the round-trip times (RTT) between all pairs of backbone nodes. This measurement is performed once every 15 minutes between the exterior interface addresses of the backbone nodes.[3] A backbone node temporarily stores the delay data, transferring it daily to a NOC data collector. From these fifteen minute samples Merit publishes quartile statistics for the monthly internodal delay.

Packet categorization objects collected per node
relative to exterior nodal interface
source-destination matrix by network number (packets/bytes)
TCP/UDP port distribution, well-known subset (packets/bytes)
protocol distribution (e.g., TCP, UDP, ICMP) (packets/bytes)
Packet-length histogram at a 50-byte granularity
packet volume going out of backbone node
relative to entire node
per second histogram of packet arrival rates
NSS (intra-NSFNET) transit traffic volume

3.2 Interface statistics

To maintain data regarding packets and bytes transmitted and received, errors, delay times, and down times, all NSFNET backbone nodes record statistics about the packets which traverse each of their interfaces. Each backbone node, also called a Nodal Switching Subsystem (NSS), runs SNMP servers which respond to queries regarding standard SNMP Management Information Base (MIB) variables. Centralized collection of the data occurs for each backbone interface on each NSS once every 15 minutes. The counters are cleared via only two mechanisms: explicity, when the machine is restarted; and implicitly, when the 32 bit counters overrun. Cumulative counters, retrieved using the SNMP, include those for packets, bytes, and errors transmitted in and out of each interface.[4]

Several in-house utilities allow one to derive multiple statistics from this data, including peak and average link utilizations and the peak fifteen-minute interval of each day.[5]

3.3 Packet categorization

To categorize IP packets entering the NSFNET backbone based on information contained in packet headers, each NSS has a dedicated processor that examines the header of every packet traversing the intra-NSS token ring. NNStat [2] builds statistical objects based on the collected information.

The central agent running the collection software periodically calls out to each of the backbone nodes, logs their statistical objects, and resets the objects. The collection host is an IBM RS/6000 and it collected as much as 50 megabytes of raw statistics daily. Table 3.3 lists the NNStat objects collected operationally on the T1 backbone.

After collecting the data, a specialized software package compiles a monthly matrix of network-number-to-network-number traffic counts, which forms the basis for the publicly available files characterizing traffic across the NSFNET backbone in terms of both individual network numbers and countries.

In all cases, Merit timestamps the collected statistics for the backbone according to Universal Time. The graphs presented in this paper also follow this convention. The tick marks on the x-axis for weekly time series graphs correspond to 0:00 Universal Time for the indicated day of the week.[6]

---

[2] The data for the statistics presented in this report were gathered by separate processors in the NSFNET NSS equipment which aggregate information using the NNStat [2] software package. This compilation, greatly assisted by Merit Network, Inc. and other installations, reflects an effort to capture as much and as accurate data as possible. However, no guarantee is given for the completeness of the data or its accuracy.

[3] Halving this value yields an approximate one-way delay for the 14 by 14 delay matrix. On a relatively uncongested backbone with stable and symmetric routing, such a method of achieving one-way delays is justified. See [7] for more details on the failure of round trip delays to adequately characterize unidirectional latencies across a wide-area network.

[4] Error conditions on the interface include HDLC checksum errors, invalid packet length, and queue overflows resulting in discards.

[5] Daily and monthly traffic summaries are available in reports via anonymous FTP from host *nis.nsf.net*.

[6] As a reference, 0:00 UT is 19:00 EST.

4 Characterization Targets

We now present observed traffic characteristics of the operating T1 network. Kleinrock and Naylor's [13] paper serves as our basic inspiration, yet we do not offer the exactly same set of data since the environment and parameters of interest have changed somewhat. In some cases the graphs presented in [13] to depict network behavior are not possible, or, more specifically, have no meaning, in today's NSFNET. Indeed, often today's environment requires translation of vocabulary applied in Kleinrock's study. For example, today's analogy to an ARPANET "user-HOST" system mentioned above is an external interface to a Nodal Switching Subsystem. Therefore, our statistics may diverge in their precise name, but we intend that the spirit of our effort is the same: to describe characteristics of network behavior, and their changes over time.

Next we briefly discuss the characteristics which Kleinrock and Naylor investigated, and describe our followup parameters:

1. Message and packet size distributions. In today's NSFNET backbone, the term "message" is not applicable.
2. Delay statistics. We present median statistics calculated by Merit, but no dedicated experimental measurements for the isolated purpose of this study.
3. Mean traffic-weighted path length. The T1 network was designed with a maximum diameter of three backbone links, which continued until the installation of the T3 upgraded network. In these architectures, the issue of path length does not hold as much interest.
4. Incest (the flow of traffic to and from hosts at the same local site). We discuss in Section 6.2 why incest does not apply in the current environment.
5. Most popular sites and links.
6. Favoritism (the property which a site demonstrates by sending much of its traffic to a small number of sites).
7. Link utilization.
8. Error rates.

We also investigate a few things which were not as applicable in Kleinrock and Naylor's environment:

1. Attributing longterm growth in traffic volume to domains and protocols.
2. Trend in average packet size on the network, both over long and medium term intervals.
3. International distribution of traffic.

As Kleinrock and Naylor pointed out in their study, such statistics are of more than merely historical interest. In many cases they may lead to change in parameters used for implementation: packet and buffer sizes, number of buffers, channel capacities, fragmentation policies, etc.

Before presenting the statistics in the next section, we briefly discuss the concept of *granularity* in information collection and presentation.

4.1 Aggregation: Time and Space Granularities

In aggregating statistics, one must select granularities along multiple dimensions of time and space. Aggregation involves two parameters: the granularity at which one collects information, and the granularity with which one presents information. For example, a packet count for a fifteen minute interval implies a selected collection granularity. In contrast, the "bucket size" of a histogram defines a granularity of presentation. In both cases, selection of the appropriate granularity for aggregation requires careful consideration.

The T1 NSFNET backbone currently collects statistics through SNMP-based tools at 15-minute intervals, and we restricted ourselves to these data sets with this this time granularity for this study. Such data sets serve only as a starting point. This interval may appropriate to answer questions about high-level distribution of network usage on a daily basis. Other questions, such as analyzing packet interarrival time distributions, or predicting the bandwidth requirements of continuous media data flows, will require a much finer time granularity, perhaps even in the subsecond range.

"Space" granularity in a network is harder to define than time granularity, but we offer a brief description which offers a framework to later sections. Along the space dimension, one might want to focus on specific nodes or links in order to examine behavior such as favoritism or hotspots. Alternatively, when presenting internetwork traffic flows, one might want to develop a matrix of country-by-country traffic flows over time. Other granularities include: backbone node, external interface (of a backbone node), autonomous system, agency, network number, mid-level service provider, host, application, and user. These granularities do not have an inherent order, as a single user or application might straddle several hosts or even several network numbers. As with the time dimension, selecting the optimal granularity depends on the question of interest, and often requires experimentation.

5 Results: Long-Term Trends

Figures 3 and 4 show traffic volume and count of networks configured in the NSFNET backbone, respectively, for the last four years. Figure 3 shows the decreasing trend in traffic volume on the T1 backbone beginning in late 1991, reflecting the migration of traffic to the T3 backbone. The discontinuities in this figure correspond to unavailable statistics. Figure 4 depicts growth, domestic as well as international, in the number of networks configured for the NSFNET backbone.

Figure 5 shows traffic volume on the T1 network attributed to network type (e.g., commercial, research, defense), as documented by the DDN (Defense Data Network) Network Information Center.

By classifying traffic according to the TCP/UDP port it uses, one can attribute traffic to specific end-user applications. Figure 6 shows traffic volume on the network distributed by application type for the

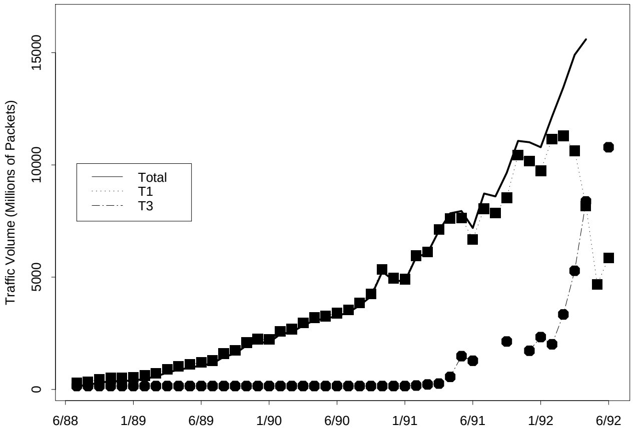

Figure 3: Traffic Volume on the NSFNET Backbone

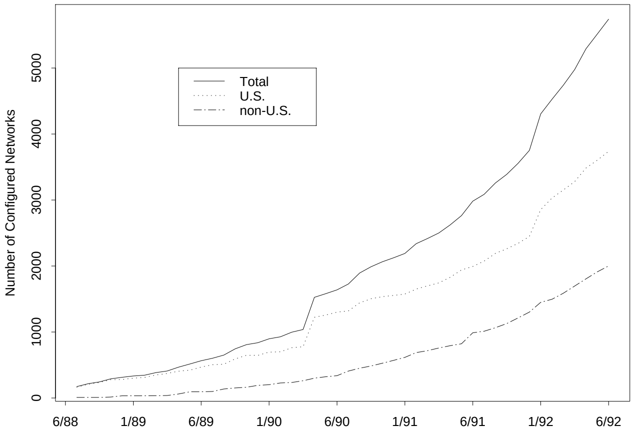

Figure 4: Number of Networks Attached to the T1 Backbone

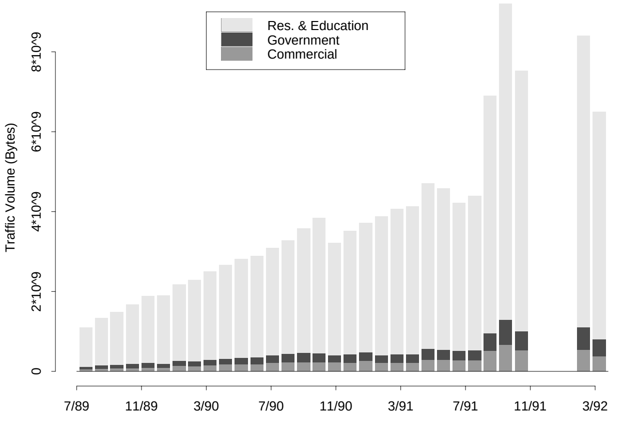

Figure 5: Traffic Offered by Network Type

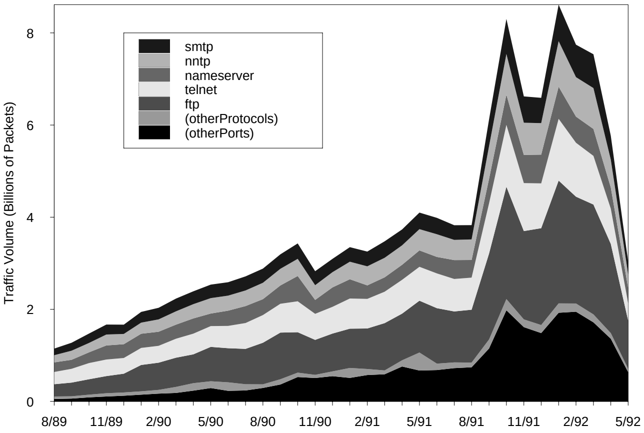

Figure 6: Traffic Offered by Protocol

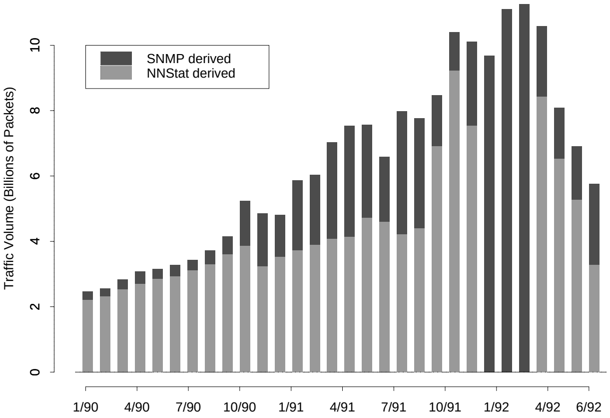

Figure 7: T1 NSFNET Backbone Packet Totals

top seven application categories in May 1992. This figure illustrates how the composition of applications has also changed over the longer term. The “other protocol” category corresponds to applications using a transport protocol other than TCP or UDP. The “other port” category corresponds to non-standard or not well-defined ports. Notice that this last category has grown much larger over the years, reflecting an increasingly diverse evironment, and the diminishing ability of Merit to track individual new applications which often use non-standard or not well-defined ports. Some methodology for classification based on the degree of interactivity, or quality of service requirements, of network traffic, will become increasingly important as the environment evolves.

A further notable event occurred in September 1991. During the 1990-91 time frame, significant discrepancies between the SNMP-based traffic counts and those derived by means of NNStat emerged, as shown in Figure 7. It became clear that the processor collecting the NNStat data was unable to keep up with the total nodal traffic flow.

Responding to these concerns, Merit deployed a sampling technique which captures only one out of fifty packets. This reduced the discrepancies signifi-

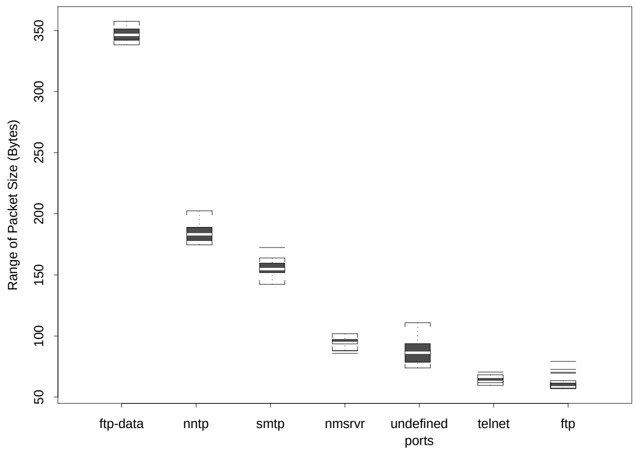

Figure 8: Mean Monthly Packet Sizes

cantly, and results in a significant jump in packet volume on the graph in Figure 6. We explore the effects of sampling on NSFNET backbone statistics in [6].

Figure 8 displays boxplots for several well-known Internet protocols. These boxplots illustrate the range of the mean monthly packet size for the indicated protocol, over the last two years. The outside ends of the boxplots indicate the range of the mean packet sizes; the inner boxes show the middle half of the data, and the line in the middle of the plot represents the median.

This figure reflects a quite visible distinction among three types of data transfer: interactive, transaction-oriented and bulk transfer. Current implementations of interactive applications frequently send end-to-end packets with single or few character payloads. Transaction-type protocols generally exchange short, multi-character lines, while bulk data transfer mechanisms typically use full-size packets for the payloads.[7] As Figure 8 illustrates, the mean packet size is relatively short for the telnet protocol, somewhat longer for transaction-oriented protocols such as NNTP and SMTP, and much larger for bulk transfer protocols such as FTP-data. Frequent transmission of zero-payload acknowledgements during TCP connections dramatically reduces the average packet sizes for many individual TCP applications.

6 Results: A Closer Look

We now focus on a specific month, May 1992, to investigate traffic patterns in more detail. During this month some 980 billion bytes were carried through the T1 network by some 5 billion packets to a total of 4,254 networks. On the average the entire network was accepting some two thousand packets per second and carrying roughly 366,000 bytes (almost 3 million bits) per second among end sites. The mean packet length was 186 bytes/packet, although such a statistic does not well reflect the bimodal distribution of packet sizes.

We presented earlier a long term perspective of packet size. While we do not have data to present statistics on complete packet size distribution. we do have some indication of packet size for the month of May as well as a selected week in May. Figure 9 shows the distributions of mean packet sizes for the 1500 site pairs who exchanged the most traffic in May 1992. The middle and lower graphs of the same figure show the same distributions for the middle and bottom 1500 site pairs, respectively.

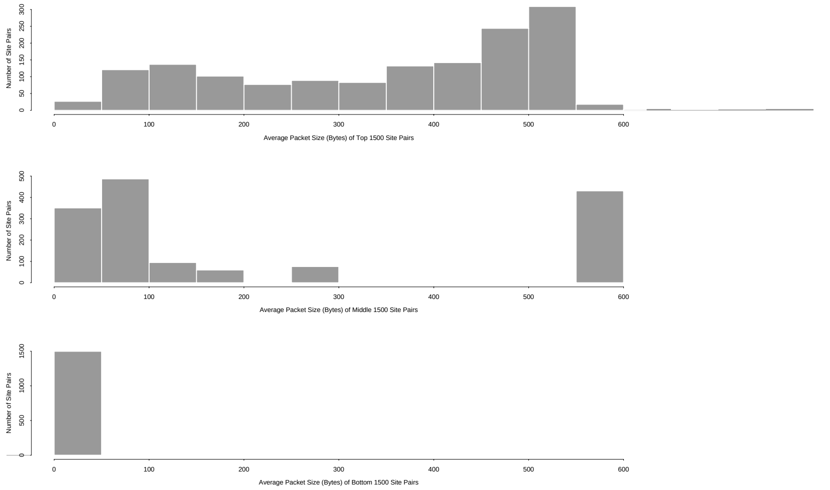

Figure 9: Distribution of Mean Packet Sizes for Network Number Site Pairs (showing the top, middle, and bottom range of traffic volume) of T1 NSFNET Backbone in May 1992.

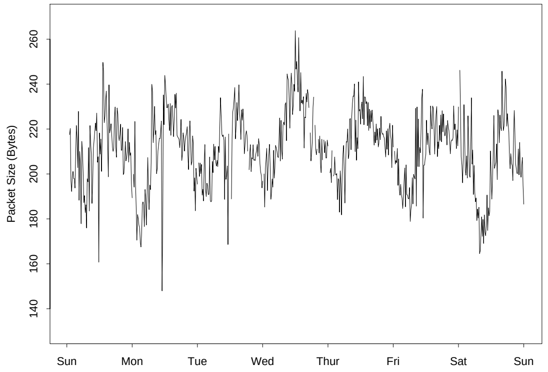

Figure 10: Average Packet Size on T1 NSFNET Backbone (10 to 17 May 1992)

Figure 10 shows the average packet size for each fifteen minute interval over the course of a week. This graph is consistent with the hypothesis that the usage of bulk transfer applications, using larger packet sizes, intensifies during the off-peak hours, or alternatively, that interactive activity, generally characterized by smaller packet sizes, drops off during off-peak hours. Both of these phenomena would result in the same effect on the graph. However, Paxson [14] confirms the latter behavior on several Internet data sets, while finding that the majority of non-interactive connections do occur during peak hours.

---

[7] Full-size packets, of the Maximum Transmission Unit size, are frequently configured around 576 bytes for many Internet hosts.

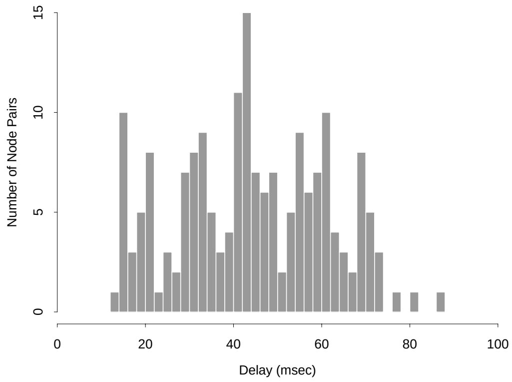

Figure 11: Distribution of one-way Median Delays Between Backbone Nodes (May 1992)

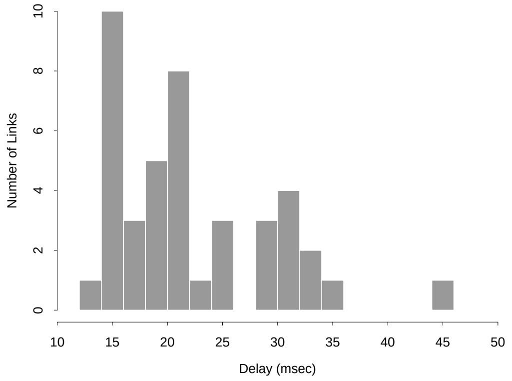

Figure 12: Distribution of one-way Median Delays across Backbone Links (May 1992)

6.1 Latency

The NSFNET backbone collects a node-to-node latency matrix. Figure 11 shows the distribution of median delays between any two nodes, including non-adjacent ones. Figure 12 shows the same statistic for only adjacent nodes, corresponding to the backbone links themselves rather than a full node-to-node matrix.

6.2 Traffic Locality

Traffic locality, or favoritism, is a specific type of geographic flow, relative to a particular observation point. Kleinrock and Naylor [13] measured levels of *incest* within the ARPANET, where they define incest as traffic which travels zero hops because it enters one ARPANET IMP and exits the same IMP.

These traffic locality issues have led us to explore several aspects of the non-uniformity of traffic flow. In Figure 13 we plot the cumulative distribution of messages sent from and to the n busiest source and destination networks (note the log scale). Over 50% of the traffic is generated by the busiest 31 of the 4254 site networks (0.7%), and over 50% travels to the 118 most popular (2.8%) destinations. Specific site pairs show even more marked favoritism: 46.9% of the total traffic on the backbone travels between 1500 (0.28%) of the 560,049 site-pairs.

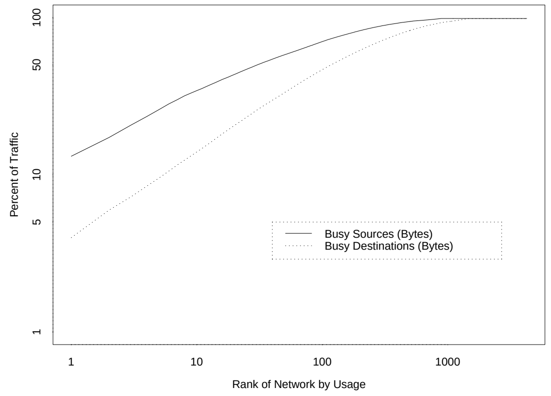

Figure 13: Cumulative Distribution of Traffic for 200 Busiest Sources and Destinations

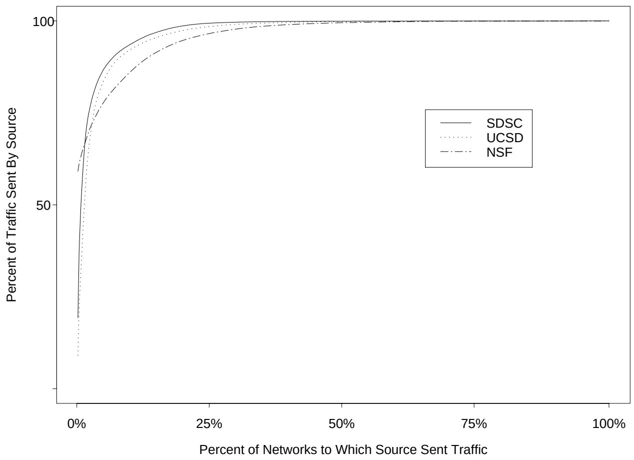

Figure 14: Cumulative Distribution of Traffic to Favorite Destinations

Figure 14 illustrates the "favorite-site" effect. The graph plots the source-favoritism for three networks: the percent of traffic to each source's n favorite destinations. We selected these sites to demonstrate the significant difference in the degrees of favoritism exhibited by UCSD, SDSC and NSF's local networks.

For UCSD, 90 percent of the traffic goes to the 63 most favored (6.7%) of the 933 sites. The 90th percentiles for SDSC and NSF were 62 out of 933 total sites (6.6%), and NSF 61 out of 458 (13%) sites, respectively. Note that favoritism at each source site involves a separate set of most popular destination sites, since each source need not have the same set of favorites.

6.3 International Distribution of Traffic

As described in Section 3, Merit also provides a report attributing monthly traffic volume to individual countries. Figure 15 uses this data to illustrate the global use of the NSFNET infrastructure, in particular for transit traffic among non-US sites. Note that the non-US data line follows the axis on the right side of this figure. Analysis of traffic among non-US networks

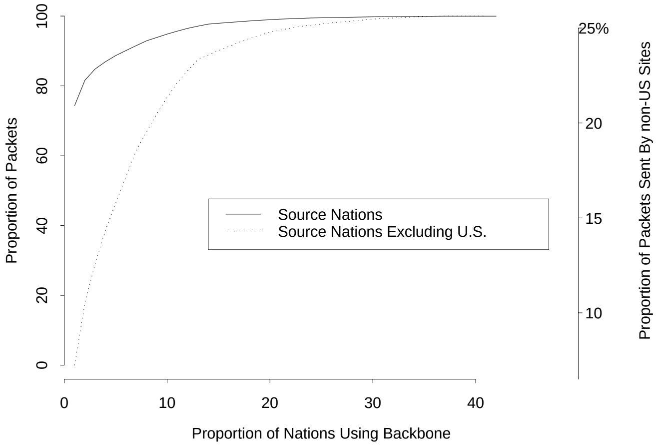

Figure 15: May 1992 Cumulative Distribution of Packet Volume into T1 NSFNET Backbone by Country (May 1992)

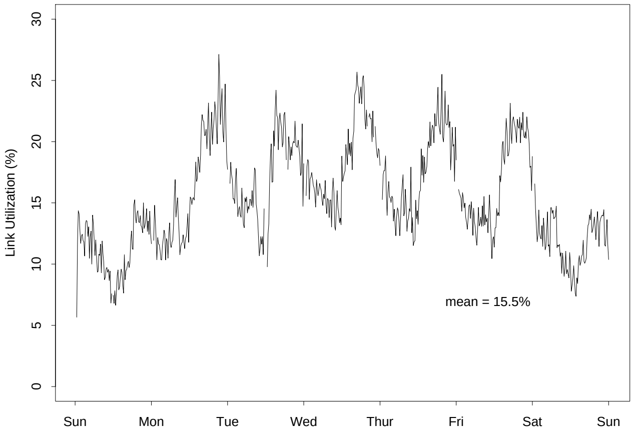

Figure 16: Mean Utilization of T1 Backbone Links (10 to 17 May 1992)

which use the United States infrastructure is not only of research interest, but also of obvious interest to a carrier considering the provision of IP services.

6.4 Traffic Volume and Utilization

We now discuss other global measures of network behavior. We plot most figures in this section over the course of the week 10 to 16 May 1992 to enable a closer focus on daily as well as weekend/weekday cycles.

The internal traffic on backbone links is one measure of the effectiveness of the network design and use. In Figure 16 we show the link utilization averaged over the entire network on a 15 minute basis for the week. The maximum 15 minute average line load was approximately 27.12% percent and corresponded to an internal network flow of roughly 0.41 Mbps. The maximum 15 minute line load for a single link on the network was approximately 89% and corresponded to an internal network flow of roughly 1.34 Mbps. The daily cycles and weekend lulls are unsurprising. As the NSFNET's international clientele grows, we may see some mitigation of these daily cycles.

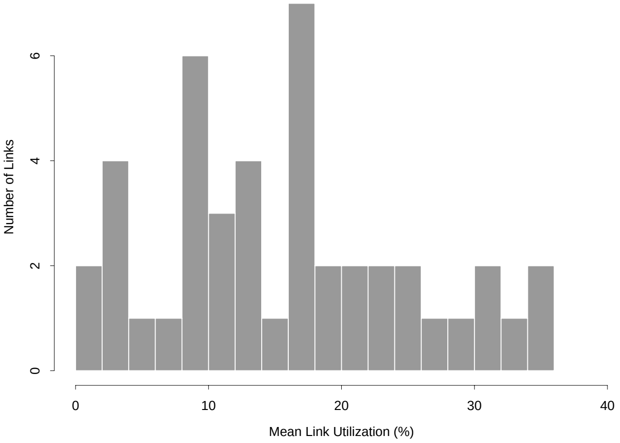

Figure 17: Distribution of Mean Utilizations of T1 Backbone Links (10 to 17 May 1992)

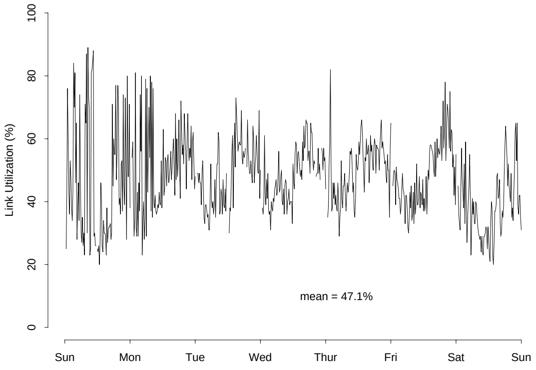

Figure 18: Utilization of Most Heavily Used Link in Each Fifteen Minute Interval (10 to 17 May 1992)

Figure 17 shows the distribution of the mean utilizations of all the links on the backbone. Note that the T3 is now assuming much of the traffic loads, so utilization is much lower now than it was several months ago.

Figure 18 illustrates the utilization of the most congested link during each 15 minute interval. Note that this graph does not depict a single physical link, but rather a consolidated "virtual" link composed of the link which was the busiest during each of the 96 fifteen-minute intervals of the day.

Figure 19 presents a comparison of utilizations for the link with the highest average utilization. The upper graph presents the link itself; the middle graph presents the reverse direction of the same link; and the lower graph presents the difference in utilization between the above two directions.

6.5 Reliability

A reliability metric for an infrastructure as pervasive and complex as the NSFNET backbone must consider many parameters, including nodal interface link

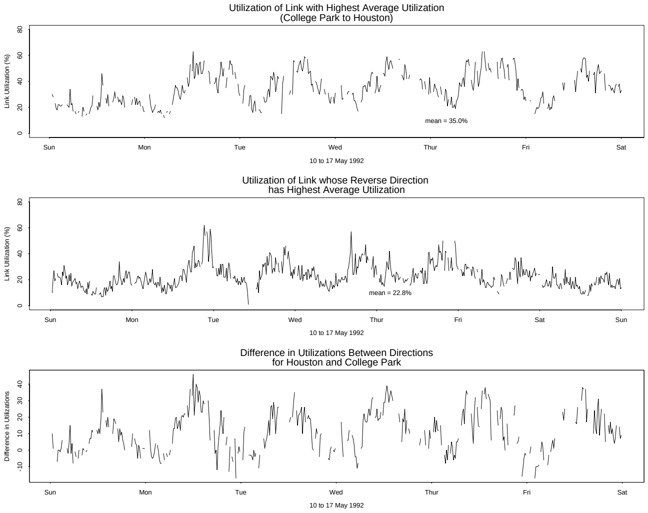

Figure 19: Link with Highest Average Utilization (College Park to Houston)

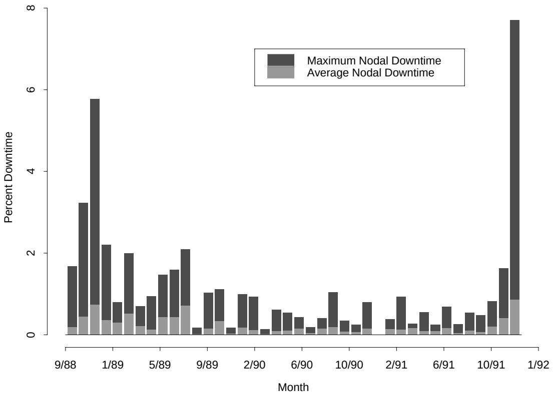

Figure 20: Maximum and mean nodal downtime for T1 NSFNET backbone nodes.

errors, bit error rates as perceived by DSU/CSUs,[8] full and partial node outages, routing configuration errors, or other unanticipated events. We offer two examples of views of individual reliability aspects of the backbone.

Network error events on the T1 backbone are characterized as either Class 1 (full node outage) or Class 2 (service reduction). The Class 1 percentages of node reliability, as viewed by the Network Operations Center (NOC), are included in publicly available monthly reports. Figure 20 plots the maximum and mean nodal downtime for all T1 NSFNET backbone nodes from 1988 to 1991.

Figure 21 shows communication errors as seen by serial line interfaces. The graph counts the number of links during each fifteen minute interval that show link errors during that observation period.

Figure 22 provides a histogram of the same data, showing that it is far more common for a fifteen-minute observation interval to show at least one link with errors than none. Note that the bucket corresponding to **1** errors, could correspond to a single link

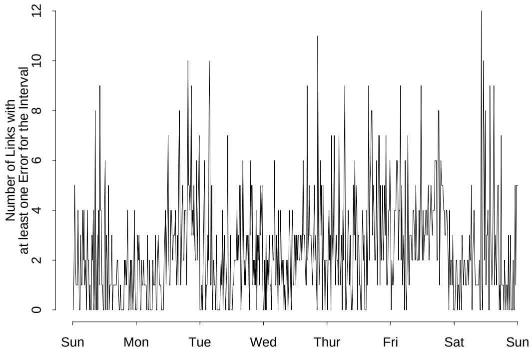

Figure 21: Link Errors throughout the T1 NSFNET backbone, 10 to 17 May 1992

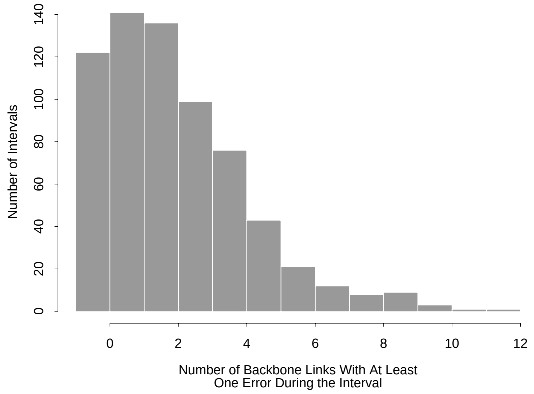

Figure 22: Link Error Histogram throughout the T1 NSFNET backbone, 10 to 17 May 1992

---

[8] Data Service Units and Channel Service Units are machines which interface digital links with nodal switching subsystems.

having a single bit error for a full fifteen-minute interval.

These are very limited snapshots into selected aspects of NSFNET backbone "reliability", and we are not in a position to draw conclusions or make reliability assessments from these limited data sets. A responsible investigation of of backbone reliability require richer data sets.

7 Conclusions

We have presented some traffic characteristics of the T1 NSFNET backbone. We have included both long term characterizations, essentially for the lifetime of the T1 network, as well as more detailed results for the month of May 1992. The long-term data are presented on a monthly basis and were obtained from publicly available summaries of measurements published by Merit Network, Inc. We can make the following observations from the presented data.

Traffic both in packets and bytes and the number of networks (in the sense of assigned IP address families to campus and other networks served by the T1 backbone) is steadily increasing since the network installation. At approximately the end of 1991, the T1 traffic volume dropped off, as the NSFNET project began to divert traffic to the T3 backbone. The increase in traffic volume in bytes seems quadratic, while the increase in networks served appears linear.

Most of the traffic volume is within the research community, and the highest volume applications are file transfer (using the FTP protocol) and network news distribution (using the NNTP protocol). Lately, a considerable proportion of the traffic cannot be directly attributed to protocols and applications because of the proliferation of protocols using non-standard TCP/UDP port numbers or other transport protocols.

Monthly summaries of mean packet size do not reveal any particular trend, but fifteen-minute data shows daily cycles which are compatible with the hypothesis of bulk transfer applications, using larger packet sizes, intensifying during the off-peak hours, or correspondingly, that interactive activity, generally characterized by smaller packet sizes, drops off during off-peak hours. Delay statistics (the monthly median of sample packet delays obtained through `ping` at fifteen-minute intervals) reveal that typical end-to-end delays on the backbone do not exceed 100 milliseconds and typical link delays do not exceed 45 milliseconds.

As was true two decades ago in the ARPANET environment, traffic favoritism is high. For example, 0.28% of the (customer/campus/site) network pairs generate 46.9% of the traffic. Link utilization is high, even following the diversion of a considerable proportion of the traffic to the T3 network. The mean overall utilization for the month of May 1992 was 15.4% while 5 nodes had more that 30% mean utilization for the month. Over fifteen-minute intervals, utilization of highly utilized links typically exceeded 50% and sometimes 80%. The most heavily used link for the month, College Park to Houston, had utilization almost always exceeding 20% (for fifteen-minute intervals) and more than 50% during the peak hours of the day. Interestingly, the reverse direction, Houston to College Park, had almost uniformly lower utilization.

The available data hold further potential for analysis which can lead to a better understanding of traffic on the network. However, there are also limitations of the data which make exploration of some interesting questions problematic, in particular those involving correlations between instantaneous network performance and traffic intensity and characteristics. Complicating the task are the enormous difficulties with methodological data collection at such large scale, and in an operational environment. These problems explain, in our view, the lack of other similar studies for wide area networks. They also render any traffic data particularly valuable and significant.

We expect to continue our research and refine our methodologies in the process of applying them to new realms. In particular, we hope to perform similar investigation of the T3 networking environment, as well as the CASA gigabit network project and its planned infrastructure.

Acknowledgements

We would like to express our appreciation for the cooperation and assistance we received from Merit Network, Inc. and Advanced Network Services (ANS).

References

[1] T. Asaba, K. Claffy, O. Nakamura, and J. Murai. An analysis of international academic research network traffic between Japan and other nations. In *Inet '92*, June 1992.

[2] R.T. Braden and A. DeSchon. Nnstat: Internet statistics collection package, Introduction and User Guide. Technical Report RR-88-206, ISI, USC, 1988.

[3] R. Caceres, P. Danzig, S. Jamin, and D. Mitzel. Characteristics of wide-area TCP/IP conversations. In *SIGCOMM '91*, September 1991.

[4] J.D. Case, M. Fedor, M.L. Schoffstall, and C. Davin. Simple Network Management Protocol (SNMP). Internet Request for Comments Series RFC 1157, 1987.

[5] B. Chinoy and H.-W. Braun. The National Science Foundation Network. Technical Report GA-A21029, SDSC, 1992.

[6] K. Claffy, H.-W. Braun, and G. Polyzos. Application of sampling methodologies to wide-area network traffic characterization. UCSD Technical Report CS93-275, January 1993.

[7] K. Claffy, H.-W. Braun, and G. Polyzos. Measurement considerations for assessing unidirectional latencies. *Internetworking: Research and Experience*, to appear in 1993.

[8] P. B. Danzig, S. Jamin, R. Caceres, D. J. Mitzel, and D. Estrin. An empirical workload model for driving wide-area TCP/IP network simulation. *Internetworking: Research and Experience*, 3(1), 1991.

[9] M. Davis. Analysis and optimization of computer network routing. Master's thesis, U. of Delaware, 1988.

[10] D. Feldmeier. Estimated performance of a gateway routing table cache. Technical report, M.I.T. Laboratory for Computer Science, 1988.

[11] S. Heimlich. Traffic Characterization of the NSFNET National Backbone. In *Proceedings of the 1990 Winter USENIX Conference*, 1988.

[12] L. Kleinrock. *Queueing Systems, Volume II: Computer Applications*. Wiley, 1976.

[13] L. Kleinrock and W. E. Naylor. On measured behavior of the ARPA network. In *AFIPS Proceedings, 1974 National Computer Conference*, volume 43, pages 767–780. John Wiley & Sons, 1974.

[14] V. Paxson. Empirically-Derived Analytic Models of Wide Area TCP Connections. Unpublished Paper, 1992.

[15] I. Wakeman, D. Lewis, and J. Crowcroft. Traffic analysis of trans-atlantic traffic. In *Inet '92*, June 1992.