Analysis and Modeling of High Speed Networks: 1993 Annual Status Report*

Hans-Werner Braun
Bilal Chinoy
Kimberly C. Claffy
George C. Polyzos

with contributions by:
Kevin Fall
Nancy K. Groschwitz
Gary A. Hanyzewski
Max Y. Okumoto
Carl D. Scarbnick
Jarom Smith

edited by Joshua Polterock

Applied Network Research
San Diego Supercomputer Center
and
University of California, San Diego

March 7, 1994

Abstract

We present the 1993 annual report on research efforts of the San Diego Supercomputer Center's Applied Network Research group (ANR). Our projects were undertaken in collaboration with the Computer Systems Laboratory of the University of California in San Diego (CSL), as well as a number of other researchers across the country. This year, we continued our research on network analysis and modeling, as well as advanced network technologies including the CASA gigabit testbed. Our research topics include issues of network access, instrumentation for the accurate gathering of performance statistics, accounting, and quality of service in the Internet.

We hope interested parties within the Internet research community will find this report helpful in understanding our research agenda. We also wish to use this report to stimulate further discussion and interaction among researchers pursuing similar projects and goals.

*Any opinions, conclusions, or recommendations in this report are those of the authors and do not necessarily reflect the views of the National Science Foundation, other supporting organizations, General Atomics, SDSC, UCSD or the SDSC Consortium members.

Contents

1 Introduction

This 1993 annual status report of the Applied Network Research (ANR) group at the San Diego Supercomputer Center (SDSC) and the University of California, San Diego (UCSD) represents an update on our research activities and outlines our future directions. Our projects include those funded by the National Science Foundation (NSF), the Department of Defense's Advanced Research Projects Agency (DoD ARPA) via the Corporation for National Research Initiatives (CNRI), and International Business Machines Corporation (IBM).

The National Science Foundation (NSF) specifically funds research into *Analysis and Modeling Tools for High Speed Networks*, NSF Grant NCR-9119473, as well as studies of and engineering support for NSF's National Research and Education Network (NREN) activities, and collaboration in multi-agency activities for Asynchronous Transfer Mode (ATM) networks. NSF and ARPA, via CNRI, fund the CASA testbed project to investigate networking and distributed applications issues across wide-area gigabit-speed networks. Under our IBM collaboration, we study performance issues across the Internet from a systems perspective.

During 1993, the second year since the establishment of the Applied Network Research (ANR) Group, we have continued our work in the area of traffic characterization of national backbones, in particular the NSFNET, and extended our investigations into other strategically selected locations in the Internet fabric. Section 2 presents our work in backbone traffic characterization. Sections 2.1 and 2.2 present modeling studies of operationally collected delay and traffic volume data on the T1 NSFNET backbone. Section 2.2.1 presents a summer research project of an undergraduate student in NSFNET statistics by U.S. state. Section 2.3 presents some example statistics on the T3 NSFNET backbone, in particular those concerning traffic locality, or concentration of traffic among a small subset of possible network addresses. Section 2.4 evaluates the strengths and limitations of some of the operationally collected backbone statistics for forecasting traffic quality and volume. One of our objectives is to develop statistics collection mechanisms with increased relevance to network research and analysis tasks. Sections 2.5 and 2.6 discuss two other issues which constrain the usefulness of operationally collected statistics on wide-area network infrastructures.

An important component of our current research is the modeling of Internet traffic *flows*. In Section 2.7, we discuss the motivation for and the potential applications of such an effort and describe our initial findings. This is an area where we expect to concentrate analysis and modeling work for the third year of the project. In particular, we are developing statistical models of traffic flows based on attributes such as protocol and application.

We have and continue to apply the results of our investigations to operational and planned components of the Internet. Sections 3 discusses SDSC's involvement in the CASA gigabit testbed research project, the MetaCenter network, and our planned connection to a multiagency national ATM network project.

Section 4 presents ANR efforts in NSF's NREN Engineering Group. Section 5 describes the contributions made to the ANR research and analysis effort by three students from the Computer Science Department at the University of California, San Diego. Section 7 lists ANR publications released during the past year.

2 Backbone characterization

During the past year we have undertaken various network analysis and modeling efforts. Although we have for the most part targeted our effort to the NSFNET backbone,[1] we believe our methods apply to wide-area networks in general. We have investigated the suitability of existing and currently collected data (by Merit and ANS) for performance modeling. In particular, we attempted to show associations between network traffic and performance on the NSFNET backbone. However, for various reasons which we discuss in this section, but mainly because of the coarse granularity of the collected statistics, the data do not show correlations between traffic volume and performance. Section 2.1 presents a case study with evidence for why operational data collected at current granularity is insufficient to develop accurate performance models.

We find that collected data does provide a good framework for models of long-term traffic volume on the backbone. Such traffic models are important for forecasting traffic levels on the backbone, and thus for making decisions about future expansion and upgrades. In Section 2.2 we discuss our investigation of the use of time-series (ARIMA) models [2] to model long-term NFSNET traffic volume, and show their strength for forecasting long-term traffic volume. We will present the resulting work in published form at the 1994 IEEE International Conference on Communications (ICC'94) [10].

An area of particular interest in wide-area infrastructures is traffic *locality*, or heavy concentration of traffic among a small subset of possible network addresses. Section 2.3

[1] NSFNET, the National Science Foundation Network, is a general purpose packet-switching network supporting access to scientific computing resources, data, and interpersonal electronic communications.

The NSFNET backbone itself evolved from a 56kbps six-node network in the mid-1980s to today's 45Mbps (T3) network, the current NSFNET extends to a large fraction of the research and educational community, and even into the global arena via international connections.

presents some example statistics from December 1992 highlighting the degree of locality on the T3 NSFNET backbone.

Section 2.4 evaluates the strengths and limitations of some of the operationally collected backbone statistics for forecasting traffic quality and volume. One of our objectives is to develop statistics collection mechanisms with increased relevance to network research and analysis tasks. Sections 2.5 and 2.6 discuss two other issues which constrain the usefulness of operationally collected statistics on wide-area network infrastructures.

Limitations of the operationally collected statistics for certain workload characterization and modeling tasks have led us to undertake more specific measurement studies of various granularities on more limited geographic scales. An important component of this research is the profiling of Internet traffic flows. We describe motivation, applications, and our initial findings in this area in Section 2.7. In particular, we are developing statistical models of traffic flows based on attributes such as protocol or application. We use measured network traffic traces obtained from different types of networking environments. We focus on traffic models related to caching and (soft) state preservation on wide area networks. Such traffic models are central to the development of economical routers, efficient schemes for the transport of IP over ATM, and accounting schemes for IP traffic on wide area networks and internetworks.

The research efforts we describe in this section are all applicable not only to current network infrastructures, but also to planned and emerging advanced Internet technologies, around which we focus Section 3.

2.1 Modeling NSFNET delay using operational data

In evaluating the relevance of the collected node-to-node packet delay statistics on the NSFNET backbone, we investigated the relationship between the amount of traffic present and the measured round-trip delay, and compared the trends in measured values to those predicted by queueing theory. We will only discuss some basic results here; for details refer to [9].

The first analysis examined the correlation between traffic and delay for each of the 26 backbone node pairs for which data was available. Correlations were calculated using the Pearson product-moment coefficient.

We also investigated whether node pairs with higher traffic levels than average had correspondingly higher average delays. To compare delays between different node pairs required that we consider the following different components of the total delay:

1. Processing delay. The delay added by routers and switching equipment.
2. Propagation delay. The transport delay between the nodes.
3. Queueing delay (including transmission time). Additional delay imposed by congestion.

The first two delay sources are fixed for a given node pair and packet size. To compare delays between different node pairs, we needed a way to estimate what portion of the total observed delay is due to processing and propagation delays. We considered the following two alternative methods:

1. Examine the minimum delays observed for each node pair, when the network utilization is very low. These delay times will probably include little or no queueing delay, and should give a reasonable estimate of the total fixed delay for each node pair.
2. Model the observed data and attempt to find weighting factors for fixed delay sources.

For each node pair, the minimum observed delay for that node pair was subtracted from the mean delay. We use the difference as an estimate of the portion of the mean delay resulting from queueing delay. In general, node pairs with higher traffic had higher mean delays.

A second approach to the problem of estimating the various sources of delay is to attempt to extract the information from the data themselves. We fit the collected delay data to the following model:[2]

$$Delay = \alpha Bytes + \beta Miles + constant + error$$

where (for each node pair):

Bytes	=	Mean number of bytes per 15 minutes
Delay	=	Mean delay over all 15 minute periods
Miles	=	Distance in miles between cities

We had hoped this model would give reasonable estimates for the queueing and propagation delays. The model did a good job of predicting the fixed (processing and propagation) delays. The model was somewhat less successful in predicting the effects of added traffic. Furthermore, the predictions the model made for the effects of changes in traffic were too extreme — the delays predicted for high traffic node-pairs were larger than any actual delays, while the delays it predicted for low traffic pairs

[2] We also fit the data to several non-linear variants of the model, e.g., with delay linearly related to miles but exponentially related to traffic. In all cases, t-values went down, and the constant became a large negative number.

were smaller than the actual delays, indeed, often smaller than the minimum delay for that pair. Thus it seems that the model is giving too much weight to the effect of additional traffic.

Although we did find evidence of a relationship between traffic volume and delay, the relationship appeared to be linear; at the levels of traffic volume observed here, and with the metrics for which data was available, there was no evidence for the curve predicted by queueing theory. Quantifying the relationship between traffic levels and delay proved to be difficult. Two different approaches gave very different results. This study revealed two main problems with the data:

1. Delay data are not representative of the whole 15 min interval. Although we have a byte count for each 15-minute period, we have no information about the fraction of a second when delay was sampled. At that moment, the link might have been extremely busy, or completely quiet. A data point with a large delay, therefore, might be an indication of some other effect of which we have no knowledge, or it might be the result of heavy traffic at the moment when the delay was sampled. Getting more accurate delay estimates, however, would require more complete data.
2. The granularity of the data is too coarse.

2.2 Modeling long-term NSFNET traffic volume using operational data

In planning for the future needs of any complex system, accurate forecasting of the workload is important to assess future capacity requirements and to plan for changes. We have investigated the potential of time-series methods for creating accurate, detailed forecasts of future NSFNET backbone traffic.

A dominant feature of the backbone traffic over the last few years is the overall increase in volume. While we might obtain a rough estimate of future traffic by simply fitting a smooth curve to the data, this method ignores a great deal of information. For example, the traffic may have seasonal components that greatly affect the traffic levels at any given time, which the curve-fitting method would obscure. A more detailed model would take these seasonal trends into account.

One of the advantages of such a model would be that, due to its better match to the data and more precise predictions, one could spot deviations from previous trends more quickly, and generate new and revised predictions. A detailed forecast of yearly traffic patterns also allows for more accurate planning and better decisions. If hardware changes are required at some point in a coming year, the model could estimate when traffic levels are likely to be lowest. If a particular backbone link is approaching maximum capacity, a detailed model of the traffic patterns on that link could predict when traffic will likely exceed link capacity. If the traffic on a less busy link is growing more rapidly than on another, busier, link, models of each could predict when their usage levels would cross (which might have an impact on routing decisions).

Finally, a more accurate model may allow for reasonable predictions several years into the future. Given the enormous growth rate of traffic on the backbone, the ability to make forecasts two or more years in advance has great advantages for planning for future requirements. ARIMA time-series models [2] possess these characteristics.

We used ARIMA modeling techniques to characterize several years of NSFNET packet traffic growth. The data used in this study consist of daily packet totals between all NSFNET backbone nodes, between August 1, 1988 and June 30, 1993. From 1988 through 1990, all data are from the T1 network; from January 1991 to November 1992, the data are the sum of traffic on both the T1 and T3 networks; from December 1992 through June 1993 data are from the T3 network only. All data was collected by Merit Network, Inc. as part of its operation and management of the NSFNET backbone.

To test the ARIMA model's ability to predict across a full year, we fed the model data from 1988 through June 1992, and used them to forecast the subsequent year (i.e., July 1992 through June 1993). Figure 1 and 2 compare the predicted to the actual data for those 12 months. Figure 1 shows all data, from 1988 through 1993; Figure 2 shows only the 12 months of the prediction. The prediction of the model was quite accurate over the entire year, although it was too conservative in predicting the size of the oscillations towards the end of 1992.

We then used the model to forecast NSFNET backbone traffic levels for the coming year. We used all available data, from August 1988 through June 1993, for the prediction. Figure 3 shows the forecast for July of 1993 through June of 1994.

A model capable of forecasting two or more years into the future would have enormous usefulness as a long-range planning tool. To evaluate the ARIMA model's ability to make long-range predictions, the data from 1988 through June of 1992 were used to make a forecast two years into the future (i.e., from July of 1992 through July of 1994). Figure 4 shows this forecast overlaid with the data from Figure 3, which shows a one-year forecast based on all available data (1988-1993).

Figures 1 and 2 illustrate the ability of the model to accurately predict one year into the future. Thus we have some

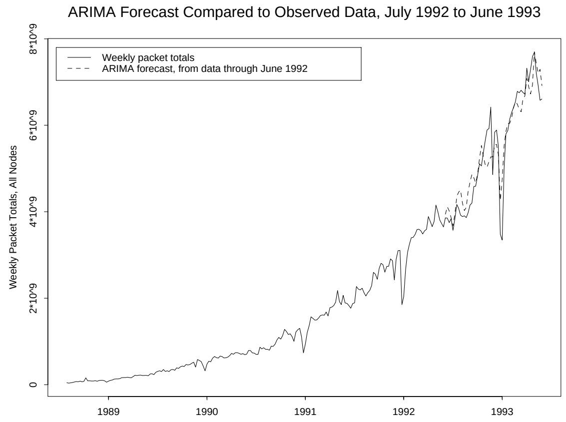

Figure 1: All data compiled between August 1988 and June 1993

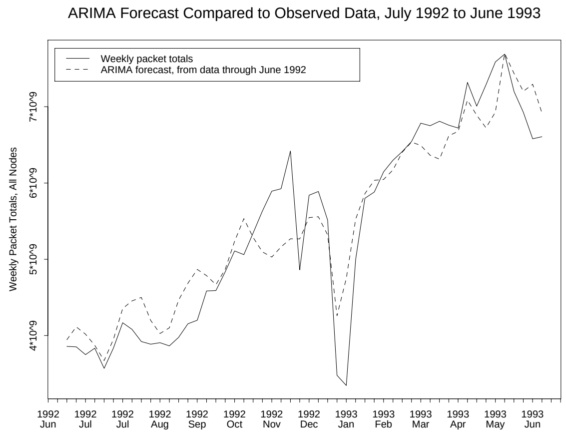

Figure 2: ARIMA 12-month prediction compared to actual data

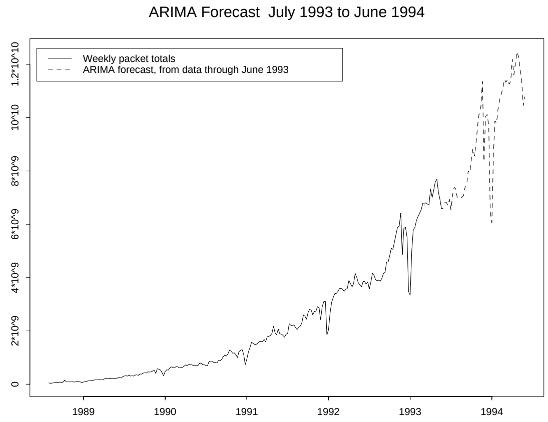

Figure 3: ARIMA 12-month forecast data

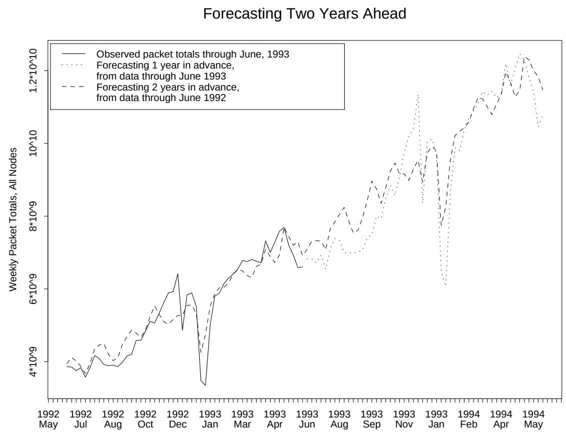

Figure 4: ARIMA two-year forecast compared with one year forecast based on all available data

confidence in the one-year forecast through July 1994.

Our study was successful in evaluating the feasibility of using time series analysis to make detailed long-range predictions about NSFNET backbone traffic. The close match between predicted and observed traffic levels suggests that this approach is adequate for long-range forecasts and planning with some confidence. One could also use ARIMA models to make predictions about the traffic levels on individual NSFNET backbone nodes or links, rather than on the aggregate. Preliminary model building, however, suggests that it would be necessary to use a separate model order for each node. This should not be surprising, because traffic patterns and growth on individual backbone nodes may be strongly affected by local factors not shared by other nodes, for example, the state of the economy in the region where the node is located, or the addition of a new backbone node that handles traffic formerly handled by this node.

These examples of outside factors do point to one ultimate limitation of the ARIMA model approach, which is that all predictions are necessarily made on the sole basis of previous data. Such models cannot, by their very nature, take into account any outside forces that may fundamentally change the pattern of the data. As access to the Internet becomes more widespread, a greater proportion of users may come from commercial rather than academic institutions. The pattern of use over a calendar year is likely to be quite different for a business than for a university. As new multimedia applications come into wider use, the extremely high volume generated by such applications may affect traffic patterns. Such factors as new technologies, new government regulation, or changes in the national economy may have significant effects that cannot be predicted by the ARIMA approach.

2.2.1 Network vs. traffic growth

We also explored how to best represent NSFNET growth of both the total amount of traffic traversing the backbone and the number of connected and configured networks. We evaluated how much of the traffic increase was due to the addition of new networks vs. an increase in usage by the older, more established networks. Significant observations include:

1. The percentage of newly assigned IP addresses used during their first month is decreasing steadily.

2. As time passes, a smaller percentage of configured networks are responsible for an ever-increasing portion of the new network traffic.

3. The monthly changes in overall volume are mostly due to increases or decreases in activity from established networks, not to increases or decreases in new network activity.

2.3 T3 NSFNET traffic locality

We included in last year's report a characterization of the T1 NSFNET backbone during May 1992. Although no longer in operation, the T1 backbone exhibited many workload characteristics and trends which the T3 NSFNET backbone also exhibits. We followed up the May 1992 study with a study of T3 backbone [5] data for the month of December1992. We provide highlights of the study here.

The study investigated traffic locality at a variety of granularities: network pairs, networks, Administrative Domains (ADs), and nodal switching subsystems (NSSs); Table 1 summarizes traffic flow by these entities. For the purposes of our analysis, we collated traffic data only to or from configured network numbers, and normalized our favoritism calculations to the amount of traffic sent among the 6,131 configured networks of the T3 backbone. Configured network numbers are the only ones for which NSFNET promises a service; traffic to other networks may be viewed as noise that may impose a service impact. There are also networks that are configured but apparently never announced and to which many networks still try to send traffic. [3]

Figure 5 displays the distribution of traffic for December 1992 by Administrative Domain.[4] This figure indicates that the flow characteristics of ADs do not differ in terms of packets and bytes, or between the load into and out of the backbone.

Figure 6 displays the distribution of traffic volume, in terms of bytes offered into the backbone, at various levels of aggregation: network pairs; networks; autonomous systems (ASs); and nodal switching subsystems (NSSs). This figure clearly illustrates the dramatic difference in favoritism with different aggregation granularities. One percent of the 1.3 million network pairs were responsible for approximately 70% of the traffic on the backbone.

[3] Out of date information in the Domain Name System about the addresses of root domain servers is one source of such behavior. Merit has collected statistics on the actual announcement of network numbers to the NSFNET backbone, and found a large fraction of configured networks which are consistently unannounced to the backbone. The recent increase in the use of Class C network numbers has particularly exacerbated this trend. Merit is now trying to collect lists of such network numbers and make the information available to regionals. Regionals could then remove them from their routing database, reducing the impact of the dramatic increase in network number requests for routing configuration. In [5] we discuss further issues arising from non-configured networks on the NSFNET backbone.

[4] An Administrative Domain is a collection of end systems, intermediate systems, and subnetworks operated by a single organization or administrative authority [11]. We use the Autonomous System number to identify a particular Administrative Domain.

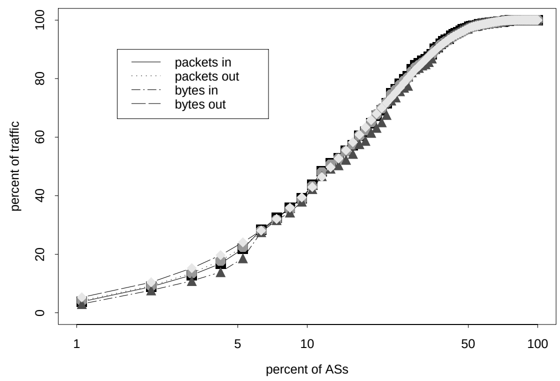

Figure 5: Cumulative distribution of traffic volume into backbone by Autonomous System number

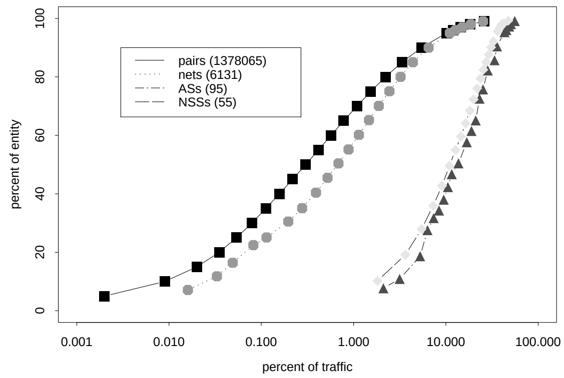

Figure 6: Distribution of traffic aggregated at various levels

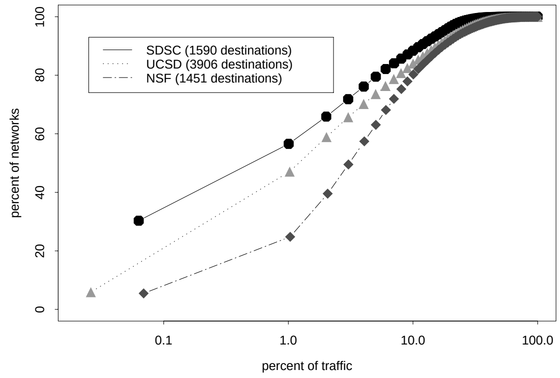

Figure 7: Cumulative proportion of December traffic sent from source to destination

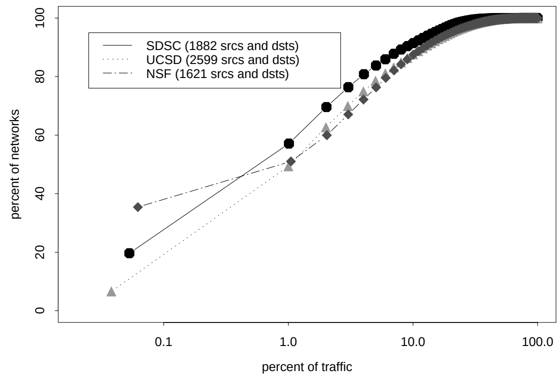

Figure 8: Cumulative proportion of December traffic exchanged between 3 networks and rest of network

The traffic was more evenly distributed among the more coarsely aggregated groups, such as Administrative Domains.

Figure 7 illustrates the impact of the "favorite site" effect reported by Kleinrock and Naylor for the 1973 ARPANET [13]. The graph plots the favoritism for three networks: the proportion of traffic to each source's n favorite destinations for the month of December 1992. The connected dots in the graph mark the unit percentiles of cumulative traffic volume. We selected these sites to demonstrate the significant difference in the degrees of favoritism exhibited by the networks at UCSD, SDSC, and the NSFNET backbone. 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.

Table 1: Summary statistics on traffic locality on the T3 backbone in December 1992

Percent of (1378065) Network pairs responsible for traffic					
percent of traffic	25%	50%	75%	90%	97.5%
number of site pairs	740	4173	21132	74795	224015
percent of site pairs	0.054	0.303	1.53	5.43	16.26
Percent of (6131) Networks responsible for traffic					
percent of traffic	25%	50%	75%	90%	97.5%
number of total nets	7	42	150	399	1022
percent of total nets	0.114	0.685	2.446	6.508	16.669
Percent of (95) ADs responsible for traffic					
percent of traffic	25%	50%	75%	90%	97.5%
number of total AD's	6	13	24	34	45
percent of total AD's	6.32%	13.68%	25.26	35.79	47.37
Percent of (55) NSSs responsible for traffic					
percent of traffic	25%	50%	75%	90%	97.5%
number of total NSSs	3	7	12	17	22
percent of NSSs	5.45	12.73	21.82	30.91	40.00

Distribution of destination networks for three source networks					
Percent of (3906) Networks to which UCSD Sent Traffic					
percent of traffic	25%	50%	75%	90%	97.5%
number of networks	10	39	149	353	666
percent of networks	0.256	0.998	3.815	9.037	17.051
Percent of (1590) Networks to which SDSC Sent Traffic					
percent of traffic	25%	50%	75%	90%	97.5%
number of networks	1	10	60	182	350
percent of networks	0.063	0.629	3.774	11.447	22.013
Percent of (1451) Networks to which NSF Sent Traffic					
percent of traffic	25%	50%	75%	90%	97.5%
number of networks	16	45	116	254	530
percent of networks	1.103	3.101	7.994	17.505	36.527

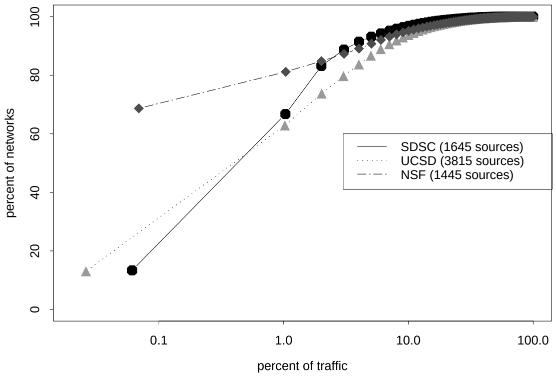

Figure 9: Cumulative Proportion of December traffic sent from source to destination

Figure 8 illustrates a similar statistic for traffic both sourced and destined from and to these three networks: the "favorite source" effect. Figure 9 illustrates the analogous measure for traffic both sourced and destined to these three networks. One can see a definite difference in workload profile among these three networks; the supercomputer center tends to exhibit a higher degree of favoritism than the other networks, and the campus communicates with a significantly larger number of sites.

Given 6,131 configured networks, which can all communicate with each other, over 37.5 million net pairs could conceivably communicate. Of these, we recorded traffic from only 3.67% (1.38 million) of them. Furthermore, 1% of those networks was responsible for 70% of the traffic for the month of December 1992. In other words, the top 15,000 net pairs collectively consumed approximately 70 times the bandwidth they would have consumed had the traffic matrix among communicating net pairs been uniform, and approximately 400,000 times the bandwidth they would have consumed had the traffic matrix among all configured net pairs been uniform.

Table 1 includes summary statistics on traffic locality for these three networks. Because we derive these data from statistics that the NSFNET backbone provides, they reflect only traffic relative to the NSFNET backbone and do not consider traffic that does not traverse any NSFNET NSS (e.g., regional network bypasses).

2.4 Address Growth and Application Diversity

In Section 2.2 we presented a model of long-term T1 NSFNET traffic volume growth. We have also evaluated current T3 NSFNET backbone statistics for their ability to allow one to forecast other aspects of network traffic. We focused on two aspects: (1) the IP network address structure, and how the status of an IP address relates to the evolution of available Internet network numbering space; and (2) the growth in application/service diversity on the Internet as measured by TCP/UDP port numbers. We focus on the limitations of these two methodologies, both of which were initially designed to support short term engineering and planning needs, such as routing and tracking the rough cross-section of traffic. Suboptimalities in architecture and implementation prevent the use of these methodologies for some long-term forecasting and planning objectives. For example, the Internet architecture makes it inherently difficult to track many applications by TCP/IP port number. The situation will pose a serious obstacle to long-term planning with the growth of real-time continuous media applications, which are able to consume significant fractions of the available bandwidth for long periods of time. A complete version of our study is available in [5]. We highlight only the main points here.

The current operational Internet infrastructure is generally strongly focused on real-time operational and near term engineering requirements. As a result, the community may acknowledge the importance of gathering valid and useful statistics for analysis, but generally this collection must take second priority to the immediate operational requirements of the network service providers. Many times the statistics are used more for public relation purposes, such as to demonstrate the ever-increasing trend for higher bandwidth demands, or for simple periodic reports, rather than for data sets which are key to engineering, planning, and network research.

Although typical operational statistics allow some tracking of Internet growth, they limit our ability to forecast capacity requirements in a network with ever richer functionality. Today's Internet aggregates traffic from among many clients, with various applications, with various associated service qualities. This situation threatens to deteriorate further as network service providers try to respond to market demand in the face of technological advances in speed and ubiquity with little resources left to devote to rational investigation of the infrastructure itself. Two changes in today's environment demand a more rational approach to performance assessments and capacity planning. First, the underlying link fabric, at ever higher performance, comes at an ever higher expense. Second, the capacity requirements of new, high bandwidth applications, in particular, the real-time demands of digital continuous media, implies that workload on the network will increase super-linearly with added clients.

To investigate beyond such traditional metrics of network usage we quantify, for the current NSFNET environment, aspects of network ubiquity, as measured by IP network numbers, and the multiplicity of services, as measured by port usage statistics. The collected statistics indicate superlinear growth of IP network numbers, and therefore Internet clients, over the last several years. The trend is clearly continuing on a global scale; international clientele as of the end of 1993 account for over 40% of the network numbers known to the U.S. infrastructure. As the need to attribute network usage intensifies, e.g., for accounting and billing purposes, the currently available data sets will seem even more inadequate. Deployment of network number aggregation techniques (e.g., CIDR) [14], which hide the interior structure of a network cluster, will further aggravate the situation.

Tremendous growth also exists in application/service diversity on the Internet as measured by TCP/UDP port numbers. The ever-increasing diversity in Internet application profiles, whose complexity will increase further with the newer continuous-flow multimedia applications, will require reassessment of network flow mechanisms such as queuing management in routers. Even within the non-continuous flow paradigm, subcategories of traffic such as interactive, transaction, or bulk traffic, may exhibit performance requirements which are different enough to justify adaptive queue management.

During 1993, ANS deployed software for the NSFNET service that allowed more flexibility with the port distribution assessments, though the inherent difficulty with the Internet model of application attribution remains. Furthermore, the recently established InterNIC[5] activity may allow greater flexibility in maintaining accurate databases of network number and traffic type statistics. Concerted attention to such activities will help foster an Internet environment where network planning and traffic forecasting can rely on more than the traditional traffic counters used in the past.

2.5 Traffic sampling

We do not underestimate the difficulty of changing the statistics collection methodologies of operational wide-area networks. Parameterizing tool requirements for collection is challenging enough prior to the implementation of a large scale infrastructural network, and a serious undertaking if it must be retrofitted after the network has been operating for years. Statistics collection in modern networking environments inevitably involves cost-benefit tradeoffs.

Developing tools is not the only obstacle; another problem in current environments is the huge amount of data which statistics collection generates. Recent dramatic increases in the speed of wide area backbones pose obstacles to complete statistics collection; managers of high-speed networks are under tremendous pressure to optimize resource usage to fulfill the data collection objective. Sampling offers a strategy to alleviate these pressures.

In 1993 we undertook an empirical investigation of sampling methodologies and their impact on network traffic characterization. We simulated several sampling algorithms on a packet trace collected on an NSFNET backbone node, exploring the effect of different parameters of sampling, such as: (1) time-driven vs. event-driven methods; (2) random vs. deterministic selection patterns; (3) the granularity, or sampling fraction; (4) the interval, or length of time over which we sample. We used as assessment targets in this study the distribution of packet sizes and packet interarrival times.

The critical element in our study [6] of sampling methodologies is a measure of how accurately various methods of sampling can answer questions related to wide area network traffic characteristics. We used several statistical metrics which indicate the similarity between two distributions, and compared the sampled traces to the parent population.

Our results revealed that the time-triggered techniques did not perform as well as the packet-triggered ones. Furthermore, the performance differences within each class (packet-based or time-based techniques) are small. We used a statistical metric related to the χ^2 statistic, the ϕ^2 metric, to characterize the degree of association between the sample distributions and the population. This metric does not provide absolute characterization of sampling performance, and is not conducive to rigorous hypothesis testing, but it does demonstrate that one technique is generally superior to another across sampling fractions and sampling intervals.

The sampling methodologies we describe in [6] can be extended and applied to characterizations of network traffic that are based on proportions, e.g., TCP/UDP port distribution. More difficult would be to characterize the goodness of fit of the sampled source-destination traffic matrix, mainly because of its large size and because many traffic pairs generate small amounts of traffic during typical sampling intervals.

2.6 Aggregation granularity

In the previous section we discussed the effect of sampling mechanisms on modeling efforts. Another aspect of statistics collection which constrains subsequent modeling efforts is their granularity. We present a brief discussion on the meaning and implications of the granularity of statistics.

[5] Established by the National Science Foundation, the InterNIC is a collaborative project of three organizations working to provide network information services. General Atomics provides Information Services; AT&T provides Directory and Database Services; and Network Solutions, Inc. provides Registration Services.

Aggregation, either in time or space, 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 an interarrival time histogram defines (when the arrival times of all packets have been captured) the granularity of presentation. In both cases, selection of the appropriate granularity for aggregation depends on the question one is addressing, and sometimes defines the range of questions one can address.

An example of time granularity is the current 15-minute aggregation interval for most statistics collection on the T3 backbone. Such a coarse granularity may be appropriate to answer questions about high-level distribution of network usage on a daily basis. Other questions, such as analyzing the dynamics of packet interarrival time distributions, greater insight into bursty behavior, or prediction of bandwidth requirements for continuous media data flows, will require a much finer time granularity.

An example of granularity along the space dimension is the geographic focus of a particular measurement: one might want to explore a specific node or link to examine behavior such as favoritism or hotspots. Alternatively, when presenting internetwork traffic flows, one might want to develop a model of traffic flows according to policy requirements, such as flows among countries over time. Other granularities include traffic by: multibackbone environment (e.g., of different agencies), single backbone, backbone node, external interface of a backbone node, backbone client service provider, Administrative Domain, IP network number, host, end user, and application. 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, the appropriate granularity depends on the question of interest. Figure 10 presents some of the possible layers of interest to the Internet community. We provide examples of data aggregated at different levels of aggregation in this model.

2.7 Internet traffic flow profiling

Limitations of the operationally collected statistics for certain workload characterization and modeling tasks have led us to undertake more specific measurement studies of various granularities on more limited geographic scales. An important component of this research is the profiling of Internet traffic flows; in this section we provide motivation and structure our flow profiling methodology. A more complete description will be available in a forthcoming report [7].

During the early years of the Internet, one could characterize the traffic as clearly defined interconnections of hosts to a common backbone network, the ARPAnet. Deployment of the NSFNET in the mid-1980's brought a more hierarchical architecture, intended to support an operational national backbone network interconnecting mid-level networks of regional scope, which themselves support connectivity to campuses. Far more successful than anyone imagined, the NSFNET spurred connectivity well beyond the auspices of its original charter, acting as a catalyst to interoperability of networks in the research, education, commercial, and government arenas across the globe. The resulting international fabric transports many thousands of traffic flows simultaneously.

In the process, however, the strict hierarchical architecture has given way somewhat to the flexibility of what has become a complex web of global interconnection. Although certain loose hierarchical components still exist within the global flow weave, such as the U.S. NSFNET, the general case is quite arbitrary. Within the weave, individual flows follow a variety of structure, and their fleeting and diverse nature frustrates attempts to model, or even define, them.

We have developed a methodology for profiling Internet traffic flows which draws on previous flow models. Our methodology for modeling flows differs from previous studies that have concentrated on end-point definitions of flows, by mainly focusing on TCP flows delimited by SYN and FIN packets. Several other studies have examined Instead, we focus on the IP layer and define flows based on traffic satisfying various temporal and spatial locality conditions, as observed at internal points of the network. This approach to the definition and characterization of network flows has significant applications to various central problems for networking based on the Internet model. Among them, optimization of caches for routing, feasibility studies and optimization of routing based on quality-of-service considerations, usage based accounting, and optimization of the transport of IP traffic over an ATM fabric. In this paper we concentrate on metrics characterizing individual flows including, volume in packets and bytes per flow, and flow duration, at various granularities of the definition of a flow, such as by destination network, host-pair, or host and port quadruple. Our measurements demonstrate

- the brevity of a significant segment of IP flows
- that the number of host-pair IP flows is not significantly higher than destination network flows, and
- that schemes that rely on traffic caching would significantly benefit if they make caching decisions taking into account higher layer information.

Our measurements have implications for: performance requirements of routers at Internet "hotspots"; general and

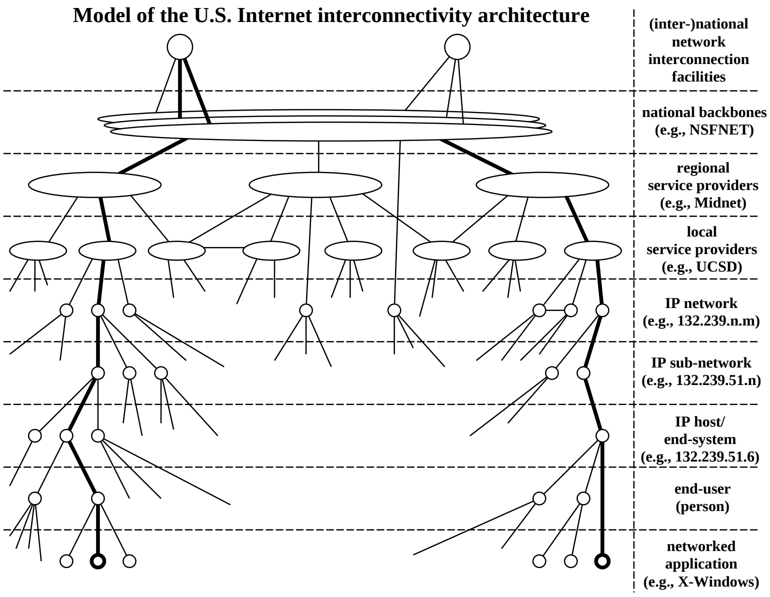

Figure 10: Model of U.S. Internet interconnectivity architecture

specialized flow-based routing algorithms; future usage-based accounting requirements; and traffic prioritization. In the coming year, we plan to continue research in the area of Internet traffic flow profiling.

3 Emerging Network Technologies

3.1 CASA Testbed Facilities

In 1993 the CASA network facilities started to take shape. Two CASA sites, JPL and CalTech, were already connected by a HIPPI fiber extender. The link between SDSC and CalTech was tested and turned over to us in September 1993, and the first pair of HIPPI/SONET (H/S) gateways were installed. Los Alamos National Labs (LANL) designed and manufactured these gateways, which allow the interconnection of HIPPI local area networks over SONET links. This section provides a brief functional description of these gateways.

3.1.1 Motivation

The CASA gigabit testbed focus has been to profile and test the performance of large-scale, distributed applications running across supercomputers interconnected by a high-speed net. The design calls for the transport of HIPPI data across wide area networks, and we achieve this by packing HIPPI data directly into SONET payloads. The SONET lines used in CASA are clocked at OC-3 (155 Mbps), and we stripe 8 OC-3 pairs to achieve an aggregate bandwidth of approximately 1.2 Gbps, which is sufficient to interconnect at 800 Mbps, the HIPPI speed.

The HIPPI standard and architecture has been designed for local area use with distances up to 25 meters. A key challenge in the CASA design was to preserve the HIPPI flow control necessary for correct local operation while providing wide area, serial interconnection. This task requires terminating HIPPI connections locally at the H/S gateway and not transporting the HIPPI flow control information across the wide area network. The H/S gateways themselves provided elastic buffering to allow for speed matching and data handling across the wide area. Figure 11 illustrates this scheme, showing that two machines communicating across CASA would involve three distinct connections: two locally at the sender and receiver and one "connection" between the H/S gateways across the wide area. The result is to preserve local HIPPI flow control semantics while allowing wide area connectivity at the HIPPI rate of 800 Mbps.

We present some initial findings in the next section.

3.1.2 Work to date and results

Link Layer Performance

We have conducted several basic performance tests to calibrate the behavior of the CASA network and understand its performance characteristics. Using IOSC HIPPI testers, we measured the available HIPPI-to-HIPPI band-

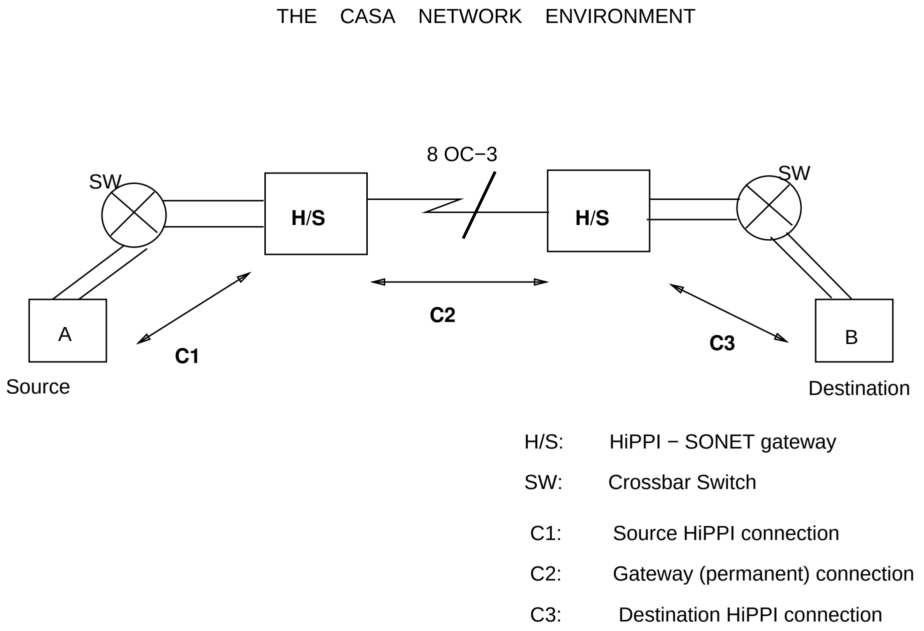

Figure 11: A HIPPI conversation across CASA is composed of three separate connections: C1, C2, and C3.

width at approximately 132 Mbps over a single OC-3 fiber pair. In striped mode, we used a maximum of 6 OC-3 pairs and measured approximately 660 Mbps. Note that the H/S gateways optionally implement Forward Error Correction (FEC) and use one channel to redundantly transmit data, reducing the number of available data channels to five.

We measured the round trip delay across the link at less than 3 ms. Straight-line distance between SDSC and CalTech is approximately 150 miles, and transmission and propagation latencies for fiber-optic gear is about 1 ms for 125 miles. The latency is almost entirely the speed of light and transmission delay across the wide area facilities.

TCP Performance

We have a number of supercomputers connected to the CASA network. These include the CRAY C90 and Intel Paragon at SDSC, the Intel Delta and Paragon at the California Institute of Technology (CalTech), and a CRAY Y-MP at Jet Propulsion Laboratories (JPL). The CRAY Y-MP and C90 are capable of driving a HIPPI channel at or close to full bandwidth, using full sized IP packets (64 KB) because these machines are limited by per-packet interrupt latencies.

The delay-bandwidth product of the SDSC-CalTech link is approximately 0.3 MB, or about four packets of 64 KB each in flight. We therefore use the TCP window scaling option for TCP, allowing the CRAYs to use the entire available TCP bandwidth if possible. Table 2 displays tests conducted using FTP, transferring a 10 MB file using 64 KB packets and large windows.

We make one important observation. Large supercomputers such as the SDSC CRAY are timeshared machines, with a CPU scheduler controlling the sharing of CPU resources. The effect of scheduling on a network-intensive job such as a large bulk transfer is quite significant; interrupting such jobs adds significant delay to packet transmissions and degrades performance. The simultaneous availability of dedicated supercomputers can pose a significant logistics and coordination issue.

We stress that these are only initial figures and much work remains in understanding the dynamics of TCP in the CASA HIPPI environment.

X Windows-based H/S Performance Tool

Researchers at the Los Alamos National Laboratory have designed the HIPPI-SONET gateway to maintain a number of counters that track performance as well as control the operation of the gateway. These counters are accessible via a serial line (RS-232) interface. We have designed and implemented a UNIX front end that allows the display and control of these registers across a TCP/IP network, via an X Windows graphical user interface.

Using this front end, we can easily configure the gateway to use any combination of 1 through 8 OC-3 fiber pairs, thereby providing fine grained control of link band-

Transmit/Receive	C90 (SDSC)	Y-MP (SDSC)	Y-MP (JPL)
C90 (SDSC)	191.4 (Dedicated)	36.4	35
Y-MP (SDSC)	40.3	107.2 (Dedicated)	-NA-
Y-MP (JPL)	33	-NA-	107.2 (Dedicated)

Table 2: CASA throughput results: Values displayed in MBytes/s with TCP using *ttcp*. 15.482 GBytes sent during loopback dedicated tests. 9.289 GBytes sent during inter-machine dedicated tests. TCP window scale factor: 4. TCP socket buffer: 378 KB. Available CASA link Bandwidth: 520 Mb/s.

width. We are using this capability to study congestion and blocking in the CASA environment. We can also switch on or off the FEC feature, allowing us to study the effect of link errors on throughput and to gauge the effectiveness of the FEC scheme. Figure 12 shows the user interface.

3.1.3 Future Work

Currently, we are studying the interaction between TCP congestion control schemes and the HIPPI flow control that operates at a lower layer. Our work should provide insights into the effectiveness of TCP in the CASA environment. We also plan to conduct studies to determine the effectiveness of the FEC capabilities of the HIPPI/SONET gateways. The bit and frame error rates across the SONET links will be a key parameter in our study.

3.2 The NSF MetaCenter

In 1992, the directors of the National Science Foundation supercomputing centers: Cornell Theory Center (CTC), National Center for Supercomputing Applications (NCSA), Pittsburgh Supercomputing Center (PSC), and the San Diego Supercomputer Center (SDSC), formed the nucleus of a collaboration based on the concept of a national MetaCenter for computational science and engineering.

3.2.1 Motivation

Each NSF-funded supercomputer center manages and operates its own computational environment. Environments include the supercomputers themselves, associated disk, storage archive, printing, and visualization systems. A structure allowing a user at any of the centers to use remote facilities in a seamless fashion would enhance productivity potential. We call such an integrated environment the NSF MetaCenter.

One objective of the MetaCenter in place is to offer users an integrated work environment in which computational jobs and files can easily migrate among the consituent supercomputer centers. The MetaCenter file and archive system would provide a consistent view regardless of the center at which the file is actually stored or accessed. This abstraction requires software that maintains a single system image of the various services offered at the centers. A key requirement of such a MetaCenter is a high-speed network, called the National Computational Environment Network (NCEnet).

An approximate analysis of the requirements of the MetaCenter quickly leads to the conclusion that the current network infrastructure between the centers is not adequate for the NCEnet network [8]. Typical NCEnet applications not only require large amounts of bandwidth but also facilities for bandwidth management and allocation. The centers therefore require a high performance network interconnecting the centers, with adequate capacity to address the needs of the MetaCenter applications, as well as bandwidth management to ensure that applications get the necessary network resources.

3.2.2 Work to date

The NSF centers jointly initiated a project to characterize the performance and estimate bandwidth available to supercomputer-based applications that utilize the current intercenter network. Each center installed a "probe" workstation on the FDDI ring that serves as the interconnection between the NSFNET backbone and the attached regional network. We are currently using the probes to

1. Perform a trend analysis of TCP performance between the centers, to characterize the "headroom" available to high bandwidth applications. The general purpose network allows the sharing of network resources among a large number of customers. Such sharing is possible, for the most part, because individual applications consume only a fraction of the total bandwidth. We anticipate that MetaCenter applications will require a base level of resources higher than current Internet applications. It is important, therefore, to characterize the amount of unused resources or "headroom" that the NSFNET

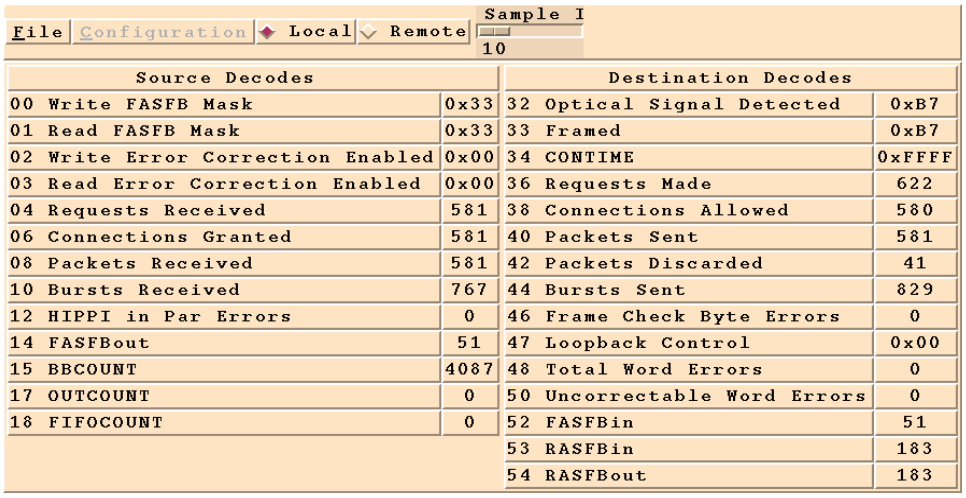

File | Configuration | Local | Remote | Sample I 10

Source Decodes		Destination Decodes	
00 Write FASFB Mask	0x33	32 Optical Signal Detected	0xB7
01 Read FASFB Mask	0x33	33 Framed	0xB7
02 Write Error Correction Enabled	0x00	34 CONTIME	0xFFFF
03 Read Error Correction Enabled	0x00	36 Requests Made	622
04 Requests Received	581	38 Connections Allowed	580
06 Connections Granted	581	40 Packets Sent	581
08 Packets Received	581	42 Packets Discarded	41
10 Bursts Received	767	44 Bursts Sent	829
12 HIPPI in Par Errors	0	46 Frame Check Byte Errors	0
14 FASFBout	51	47 Loopback Control	0x00
15 BBCOUNT	4087	48 Total Word Errors	0
17 OUTCOUNT	0	50 Uncorrectable Word Errors	0
18 FIFOCOUNT	0	52 FASFBin	51
		53 RASFBin	183
		54 RASFBout	183

Figure 12: The X Window based user interface for the HIPPI/SONET gateway.

provides, as this will determine the performance of MetaCenter applications.

2. Study packet flow characteristics of NSFNET traffic across the wide area network to better understand the resource demands that applications impose on the net.

3. Study routing system fluctuations across the NSFNET. The dynamics of routing information determine the level of network stability and reachability. We have an ongoing effort to use the probe workstations to collect routing update data to better understand the routing update process.

3.3 Interagency ATM Connection

In 1992, SDSC received an NSF award "Interagency ATM Connection at the San Diego Supercomputer Center" to collaborate with DOE and NASA to explore the suitability of ATM technologies for building and operating interagency networks. Unfortunately, DOE and NASA have been held up in their implementation because of a successful legal protest against the procurement. Consequently, we been unable to start this project and are awaiting resolution.

4 NSF NREN Engineering

We continued our efforts with the Division of Networking and Communications Research and Infrastructure (DNCRI) office at the National Science Foundation on matters related to engineering and planning of the NREN, in the context of the NSF NREN Engineering Group (NEG). The NEG effort involved collaborative work with the National Science Foundation's DNCRI, as well as Peter Ford of the Los Alamos National Laboratory (LANL).

The primary objective of NSF's NREN Engineering Group is to facilitate a graceful transition to the Interim Interagency National Research and Educational Network (IINREN). This work encompasses several aspects, both technical and non-technical, of network planning, design and architecture.

Braun participated in several activities related to the Engineering and Operations Working Group (EOWG) of the Federal Networking Council (FNC), as well as its Federal Engineering Planning Group (FEPG). He also participated in the Intercontinental Engineering Planning Group (IEPG), with the expectation of fostering international collaboration. Issues of particular concern is the ongoing design of Network Access Points (NAP), ANSnet instrumentation, regulatory investigations, and network accounting issues.

As part of a collaborative effort between the federal networking activities and the commercial sector, Braun, in collaboration with Paul Mockapetris of ARPA and others, investigated whether a need existed for short term measures to ease restrictions of electronic mail exchanges among commercial service providers. This was based on the perceived need that current restrictions on the federal network components would hinder progress for free electronic mail exchanges among commercial service providers. Commercial service providers seem to have enough flexibility to deal with the situation without external help, thus no action was required by the FNC.

We also presented issues surrounding network analysis activities to a technical audience of the NSFNET regional network community in June 1993, with the hope of stimulating more interest in analysis across boundaries of local service providers.

In the previous year, we had met with NTIA officials to discuss regulatory issues, but it appeared that the time was not ripe for extensive efforts in this area. Therefore, we have focused on laying groundwork towards requirements and instrumentation, to better prepare for a time when public policy and regulation become a more important issue. As such, we have written three papers [3] [12] [1] that we hope will foster further discussion in this area. In May 1993, we presented one of these papers [4] at a Harvard JFK school workshop *Public Access to the Internet.* The paper, "Network Analysis in Support of a Public Internet", and presentation emphasized how current and future statistics collection can (or cannot) support the expanding policy concerns of the diverse, homogeneous, and increasingly interconnected infrastructure.

We continue to work with NSF, Merit and ANS to determine how the T3 NSFNET component of the ANSNET backbone network affects requirements for NSFNET statistics collection techniques and requirements. We published and presented two papers [3] [5] at the Inet '93 conference in August 1993 in San Francisco. We still hope these activities lead to multi-agency support for aggregation of network statistics data from multiple service providers. The ANR group continues to discuss these issues with other agencies and service providers.

4.1 Network service policies

We collaborated with UCSD economist Roger E. Bohn of the International Relations and Pacific Studies School and Steven Wolff of the National Science Foundation on a policy paper regarding precedence-based services on the NSFNET [1]. In this section we give a brief description of this study.

The current architecture and implementation of the Internet assumes a vast aggregation of traffic from many sources and stochastic distribution of traffic both in space (traffic source) and time (burstiness of traffic volume). Given this general assumption, Internet components typically have little if any ability to control the volume and distribution of incoming traffic. The network, particularly from the perspective of the router, is vulnerable to significant consumption of networking resources by high-volume applications, with possibly little stochastic behavior, from a few users. This often impacts the overall profile of network traffic as aggregated from many clients. An example is the continuous flows introduced by real time applications such as packet audio, video, or rapidly changing graphics.

Most architecture and instrumentation for accounting and traffic control in the Internet reflect its historical status as a typically government bulk-funded service for the academic community. It has been a research environment with usage-insensitive costs that are often transparent to the end-users. As a result, the current Internet architecture is not conducive to allocating network resources among multiple entities or at multiple qualities of service. Presently, there is no incentive to limit the use of bandwidth, and the appetite for bandwidth is growing far faster than the government can support. As traffic grows, and new applications with fundamentally different traffic characteristics come into widespread use, resource contention will become a problem.

This situation creates a time window where applications exist on a network not designed for them, but before an appropriately architected network can augment the current infrastructure and cope with the new type of workload. We propose a scheme for voluntarily setting Internet traffic priorities by end-users and applications, using the existing 3-bit Precedence field in the Internet Protocol header.

Our proposal has three elements. First, network routers would queue incoming packets by IP Precedence value instead of the customary single-threaded FIFO. Second, users and their applications would voluntarily use different and appropriate precedence values in their outgoing transmissions according to some defined criteria. Third, network service providers may monitor the precedence levels of traffic entering their network, and use some mechanism such as a quota system to discourage users from setting high precedence values on all their traffic. All three elements can be implemented gradually and selectively across the Internet infrastructure, providing a smooth transition path from the present system. The experience we gain from an implementation will furthermore provide a valuable knowledge base from which to develop sound accounting and billing mechanisms and policies in the future.

We recognize that an IP Precedence-based service, even in conjunction with accounting and/or charging, is not going to protect the Internet from applications that do not fit the basic premise of the current Internet: stochastic sharing. That is, traffic priorities, or even accounting, are not a solution to this problem. But accounting can help this problem in the transitory phase of the current environment. In addition to billing applications, maintaining statistics for the variety of multiple services can also yield a better understanding of network demands, a valuable knowledge base to develop interim policies, capacity planning, and fairer and potentially more competitive cost recovery in the future.

5 Student involvement with ANR

Four UCSD students, Kimberly Claffy, Nancy Groschwitz, Jarom Smith, and Kevin Fall have made significant contributions to the ANR groups efforts in analysis and modeling of high speed networks.

1. *Kimberly Claffy* is expected to complete her Ph.D degree in 1994 under Professor George Polyzos in the Computer Systems Laboratory within the Department of Computer Science and Engineering at UCSD. Her thesis work, described throughout this status report, was planned as an integral component to the original NSF Network Analysis proposal which funds a portion of ANR efforts.

2. *Nancy K. Groschwitz* has completed her M.Sc. degree in CS at UCSD performing research under the direction of Prof. Polyzos and writing a thesis titled "Traffic and Delay Patterns on the NSFNET Backbone." The paper "A Time Series Model of Long-Term NSFNET Backbone Traffic," co-authored by Groschwitz and Polyzos, is based on a chapter of her M.S. thesis and has been accepted for presentation at the *IEEE International Conference on Communications* (ICC'94).

3. *Kevin Fall* is expected to complete his Ph.D degree in 1994 under Professor Joseph Pasquale in the Computer Systems Laboratory within the Department of Computer Science and Engineering at UCSD. As part of the CASA team at SDSC, he investigates protocol and performance issues with respect to the HIPPI environment on the gigabit network.

4. *Jarom Smith*, a UCSD undergraduate student majoring in cognitive science, has been involved in the project through the NSF sponsored program Research Experiences for Undergraduates (REU) during the Summer of 1993. Jarom worked on the following three projects and prepared corresponding internal reports [15] [17] [16]:

   (i) NSFNET Traffic Visualization by US State, (ii) NSFNET Backbone Growth Trends: Do More Networks Mean More Traffic?, and (iii) The Impact of Multicasting Across the NSFNET.

5. For this coming third and last year of the NSF award we expect to have two more undergraduate students (through the REU program) to work on visualization and preparation of a toolkit integrating tools developed for the project.

6 Summary

Characterization of the traffic composition and performance of wide area data networking environments continues to be an important focus of the activities of the Applied Network Research group. As the emphasis of the ANR group includes aspects of the operational infrastructure as well as academic research activities, the complexities of both areas have influenced the work scope that we describe for 1993.

We reported on results as well as work in progress for multiple studies in the area of analysis and modeling of both research and operational wide area networking. For 1994 we expect to continue in areas of flow analysis and their impact on the infrastructure, on gigabit networking on the CASA testbed, and on supporting NSF with their evolution plans for the NREN and the infrastructure at large.

7 ANR Publications

ANR has published, submitted, or otherwise made available several papers and reports of research activities during 1993. Several papers included collaborative efforts with outside people, who co-author those papers. Among the co-authors are Steve Wolff (NSF/DNCRI), Bob Aiken (then at NSF/DNCRI, now at DOE), Roger Bohn (UCSD economist), Peter Ford (LANL), Yakov Rekhter (IBM) and Paul Bosco (MIT).

ANR's 1992 annual progress report:

"1992 Annual Report of the Research Progress of the Applied Network Research (ANR) Group", Feb 1993, H.-W. Braun, B. Chinoy, K. Claffy, and G. C. Polyzos, SDSC Report GA-A21224

Several papers emphasize our activities in the area of flow analysis:

"A framework for flow-based accounting on the Internet", July 1993, H.-W. Braun (SDSC), K. Claffy, and G. C. Polyzos, to appear in Proceedings of IEEE SICON '93. SDSC Report GA-A21358

"Internet traffic flow profiling", K. Claffy, G. C. Polyzos and H.-W. Braun, Jan 1994, SDSC Report GA-A21526, UCSD Report CS93-328, submitted for publication.

Routing and addressing concerns related to the Internet evolution and its technology:

"Improving the Routing and Addressing of the Internet Protocol", March 1993, P. Ford (LANL), Y. Rekhter (IBM), and H.-W. Braun, SDSC Report GA-A21256, IEEE Network, March 1993, IBM T.J. Watson Research Center, Research Report RC 18744 (82023)

"CIDR and the Evolution of the Internet" H.-W. Braun, Peter S. Ford and Yakov Rekhter, SDSC Report GA-A21364, Proceedings of INET'93, Republished in ConneXions Sep 1993 (InterOp93 version)

Papers related to policy aspects:

"NSF Implementation Plan for Interim NREN", May 1992 B. Aiken (NSF), H.-W. Braun (SDSC) and P. Ford (LANL). Editor: K. Claffy (SDSC). SDSC Report GA-A21174, Journal on High Speed Networking, 2(1), 1993

"Network analysis issues for a public Internet" K. Claffy and H.-W. Braun, May 1993 SDSC Report GA-A21350, Proceedings of Harvard JFK School workshop: "Public Access to the Internet"

"Mitigating the coming Internet crunch: multiple service levels via Precedence", R. Bohn, H.-W. Braun, K. Claffy and S. Wolff, submitted for publication, November 1993.

"Network Analysis in Support of Internet Policy Requirements", May 1993, H.-W. Braun and K. Claffy, SDSC Report GA-A21253, Proceedings of INET'93.

Aspects of network analysis and modeling:

"Long-term traffic aspects of the NSFNET", May 1993, K. Claffy and H.-W. Braun and G. C. Polyzos, SDSC Report GA-A21238, Proceedings of INET'93.

"Traffic Characteristics of the T1 NSFNET Backbone", K. Claffy, G. C. Polyzos, and H.-W. Braun, Jan 1993, SDSC Report GA-A21019 UCSD Report CS92-237, Proceedings of INFOCOM'93

"Application of Sampling Methodologies to Wide-Area Network Traffic Characterization", K. Claffy, G. C. Polyzos, and H.-W. Braun, May 1993 SDSC Report GA-A21239 UCSD Report CS93-275, Proceedings of ACM SIGCOMM '93

"Measurement Considerations for Assessing Unidirectional Latencies", K. Claffy, G. C. Polyzos, and H.-W. Braun, Jan 1993 Journal of Internetworking, v4n3. September 1993, UCSD Report CS92-252, SDSC Report GA-A21018.

"NSFNET traffic visualization by U.S. state", Jarom Smith, NSF REU project at SDSC, September 1993

"NSFNET backbone growth trends: do more networks mean more traffic?" Jarom Smith and Bilal Chinoy, NSF REU project at SDSC, September 1993

"On the impact of multicasting across the NSFNET," Jarom Smith, NSF REU project at SDSC, September 1993

"A Time Series Model of Long-Term NSFNET Backbone Traffic" N. K. Groschwitz and G. C. Polyzos, Proceedings IEEE International Conference on Communications (ICC'94) New Orleans, LA, May 1994.

K. Claffy, H.-W. Braun and G. C. Polyzos, "Tracking Long-term Growth of the NSFNET", March 1994, to appear in Communications of the ACM, 1994.

Internet routing aspects:

"Dynamics of Internet Routing Information", B. Chinoy, Mar 1993, Proceedings of ACM SIGCOMM '93

References

[1] R. Bohn, H.-W. Braun, K. Claffy, and S. Wolff. Mitigating the coming Internet crunch: multiple service levels via precedence. submitted for publication, SDSC TR GA-A21530, November 1993.

[2] G. Box and G. Jenkins. *Time Series Analysis: Forecasting and Control.* Holden-Day, 1970.

[3] H.-W. Braun and K. Claffy. Network analysis in support of Internet policy requirements. In *Proc. of INET '93*, June 1993. SDSC GA-A21253.

[4] K. Claffy and H.-W. Braun. Network analysis issues for a public Internet. In *Proceedings of* Public Access to the Internet. McGraw-Hill, May 1993. SDSC Applied Network Research TR GA-A21350.

[5] K. Claffy, H.-W. Braun, and G. C. Polyzos. Long-term traffic aspects of the NSFNET. In *Proc. INET '93*, August 1993.

[6] K. Claffy, G. C. Polyzos, and H.-W. Braun. Application of sampling methodologies to network traffic characterization. In *Proc. ACM SIGCOMM '93*, pages 194–203, September 13-14 1993.

[7] K. Claffy, G. C. Polyzos, and H.-W. Braun. Internet traffic flow profiling. UCSD TR-CS93-328, SDSC GA-A21526, November 1993.

[8] H. David Lambert et. al. Metacenter networking: A white paper. Submitted to the National Science Foundation Division of Networking and Communications Research and Infrastructure, May 1993.

[9] N. K. Groschwitz. Traffic and delay patterns on the nsfnet backbone. Master's thesis, Department of Computer Science and Engineering, University of California, San Diego, 1993. M.Sc. Thesis.

[10] N. K. Groschwitz and G. C. Polyzos. A time series model of long-term NSFNET backbone traffic. In *Proc. IEEE International Conference on Communications (ICC'94)*, May 1994.

[11] S. Hares and D. Katz. Administrative domains and routing domains a model for routing in the Internet. RFC 1136, 1989.

[12] B. Kahin, editor. *Building Information Infrastructure*, 1989.

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

[14] Y. Rekhter and T. Li. An Architecture for IP address allocation with CIDR. Technical report, IBM TJWRC, February 1993. Internet Draft.

[15] Jarom Smith. Nsfnet traffic visualization by u.s. state. NSF REU project at SDSC, September 1993.

[16] Jarom Smith. On the impact of multicasting across the nsfnet. NSF REU project at SDSC, September 1993.

[17] Jarom Smith and Bilal Chinoy. Nsfnet backbone growth trends: do more networks mean more traffic? NSF REU project at SDSC, September 1993.