Internet Protocol Traffic Analysis with Applications for ATM Switch Design *

Andrew Schmidt and Roy Campbell
University of Illinois at Urbana-Champaign
Department of Computer Science
1304 W. Springfield Avenue,
Urbana, IL 61801
(217) 244-0582

Abstract

Knowledge of data traffic behavior is fundamental to the efficient design and implementation of a virtual circuit packet switch and its supporting policies and mechanisms. In this paper, we examine the behavior of local and wide area *Internet Protocol* (IP) networks, study connection and connectionless protocols, and use our results to estimate the behavior of a virtual circuit network replacement for these networks.

1 Introduction

The traffic analysis presented addresses IP traffic issues for XUNET[5], the network supporting the BLANCA gigabit testbed project at the University of Illinois Urbana-Champaign (UIUC). This research extends the work of other researchers who have made measurements of networks in two areas. First, we measure computer conversations and estimate the number and behavior of virtual circuits that these conversations would establish. Second, we analyze network traffic within a heterogeneous network involving wide area and local area communications, a spectrum of hosts from supercomputers to personal computers, and a variety of users involved in different activities.

Our analysis considers applications of a virtual circuit packet switch to support both a local area network (LAN) and a wide area network (WAN). In a LAN, a virtual circuit packet switch would most likely be located at the hub of a network configured as a star. In a WAN, a switch would interconnect several networks. We examine an existing network of routers interconnecting various regional networks to serve as the WAN model, and a departmental network to serve as the LAN model. Figure 1 shows a diagram of the two applications considered and the conventional network hardware that the switch would replace. In both the LAN and WAN applications the switch replaces an Ethernet.

The first network is a departmental local area network of the University's Computing Services Office (CSOnet), see Figure 1. The network interconnects a heterogeneous collection of over 100 hosts. The great diversity of hosts makes the traffic patterns less biased to certain computer types, network services, or types of network access.

*This work was supported in part by the National Science Foundation grant CDA-8722836, AT&T Bell Laboratories, and the SBC-CNRI.

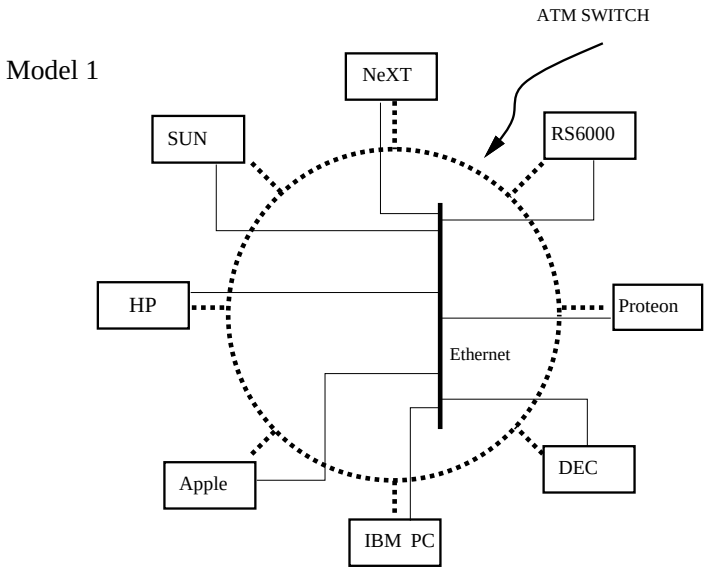

Figure 1: Network Models (LAN and WAN).

The second network modeled is the National Science Foundation Network (NSFnet) backbone site, see Figure 1. The University of Illinois serves as one of a number of Internet backbone sites throughout the United States. It interconnects the NSFnet and several regional networks via a common network located in the Digital Computer Laboratory Building. The segment contains external traffic from the University of Illinois, The National Center for Supercomputing Applications (NCSA), and the regional networks of CICnet and MRnet [9]. Direct access to this network is limited to routers; there are no direct host connections.

NNStat [1] gathered long-term data on packet sizes, types, and applications used. *Tcpdump* [6] and Perl scripts analyzed conversations, and a Sun Sparcstation2 served as the collection host. Samples were taken in periods of twenty-four hours during weekdays in the months of December 1991 and January 1992. Packet loss detected in the monitoring system was under 1%.

2 Characteristics and Results

This section presents measurements of the *Transmission Control Protocol* (TCP) and the *User Datagram Protocol* (UDP) [3]. For the purposes of our study, we concentrate on traffic that would create a virtual circuit in an Asynchronous Transfer Mode (ATM)[7] network. We model conversations as an ISO-OSI transport layer transaction, like a TCP connection. Such conversations consist of a single transaction between two hosts, involve the transfer of data, and have a finite duration. For example, a conversation would include a computer sending electronic mail to another computer or two hosts involved in a Telnet connection.

In more detail, first for TCP and then for UDP, we examine the following properties of conversations: Number of Concurrent Conversations, Conversation Set Up Rates, ATM Segmentation and Reassembly, Connection Lifetime/Bytes Transmitted and Locality of Transfers.

2.1 Transmission Control Protocol

The characteristics of the Transmission Control Protocol as a model for predicting the requirements of a virtual circuit packet switch are presented below. TCP contributed the majority of the packets observed on both the LAN and WAN. On the CSOnet, 60% of the packets were TCP and 25% were UDP. The NSFnet data contained 85% TCP and 15% UDP packets.

2.1.1 Number of Concurrent Conversations

Information about the number of concurrent conversations is important to switch design as it indicates the total number of virtual circuits and virtual circuit identifiers that must be maintained within the switch. The data presented in this section is calculated by incrementing a counter for a TCP SYN packet and decrementing the counter for the connection's FIN. This analysis does not include connections that start and end before and after the 24 hour period of observation. Our findings indicate most connection lifetimes are much shorter than a day, therefore, the results are reasonably accurate.

The activity of concurrent TCP connections is shown in Figure 2 for both the CSOnet and NSFnet. For the CSOnet there are times when no connections are present. Typically, there is a steady state of 20 to 30 active connections. The peak number of concurrent connections is present

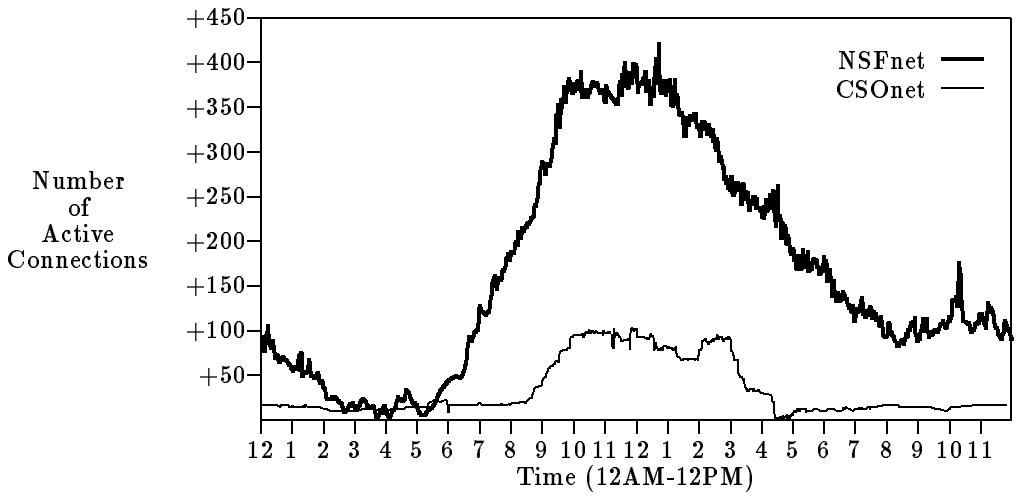

Figure 2: Simultaneous Connections

between the hours of 10 AM and 12 AM with the maximum of 104 occurring just before noon. The mean number of concurrent connections is 55.68, with a standard deviation of 36.54.

The number of concurrent connections on the NSFnet differs from the CSOnet. The first obvious difference is the larger maximum number of concurrent connections. The maximum occurs at 1 PM with a spike of 423 concurrent connections. The mean number of concurrent connections is 209.30, with a standard deviation of 120.22. Another difference is in the utilization of the network. There is a longer period of heavy utilization on the NSFnet lasting from about 8 AM to 5 PM.

Figure 2 indicates that a virtual circuit packet switch might be able to support TCP initiated conversations with a virtual circuit identifier (VCI) table that would need to contain only a few hundred entries. A VCI could also be considered as a route to a distant network. Therefore, an analogy can be draw between the number of routes maintained in a routing table and the number of VCIs that a switch may cache. At the time of our monitoring the routing table on the external router from UIUC contained 3,600 entries. With this information, it would seem that a wide area virtual circuit packet switch would need to support a table of a few thousand entries.

2.1.2 Conversation Set Up Rates

Virtual circuit set up and tear down rates are important to the design of a virtual circuit packet switch and its control software in that the are a potential bottleneck. Much work is currently being done in the area of fast call set up. Our intention here is to try to determine what demands would be placed on a switch and to what degree fast call set up will be needed. An approximation for the required set up rate can be determined by monitoring the arrival rate of TCP SYN packets. This arrival rate gives a lower bound to the maximum number of connections that may need to be established at any given time.

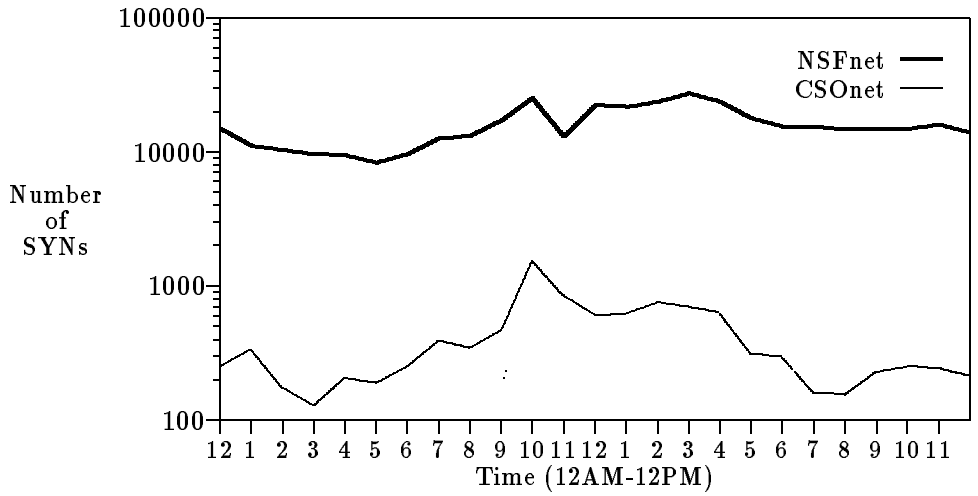

Figure 3: TCP SYNs per Hour

For the CSOnet there is a fairly consistent stream of between 200 to 400 connections starts per hour, see Figure 3. The maximum number of connections starts of 1,549 occurs at 10 AM. The mean number of starts per hour is 457, with a standard deviation of 396. This translates into an average of less than 1 new connection every second. However, a more detailed analysis showed peaks to 14 starts per second.

The NSFnet displays the same prime time usage profile but there are more connections starts per hour. The steady state is between 8,000 to 14,000 connections starts per hour and at peak usage 25,171 connections were started per hour. The mean is 16,340, with a standard deviation of 5,760. This would translate into an average of under 10 new conversations per second, however, the peak number of new conversations starts was recorded 181 per second.

An additional aspect of conversation set up is the latency or delay that a call experiences before it can transmit user data. This is the duration of time between a set up request and the time when the source host receives a conversation established notice (SYN ACK). The data for this analysis was gathered by measuring the difference between the time a TCP SYN from the source host arrived to the time the SYN ACK from the target host arrived and doubling this number, see Figure 4.

For the CSOnet the median time was 16.58 msec, the mean time was 6.59 msec and the standard deviation was 524.76 msec. For the NSFnet the median time was 132.09 msec, the mean time was 10.73 msec and the standard deviation was 1,031.74 msec.

The value of this information to a switch designer is debatable. One could reason that the switch should establish connections at whatever rate is achievable by the switching hardware and that rate would suffice. However, if current networks were to be replaced with ATM virtual circuit packet switches, then the performance shown in Figure 4 would need to be achieved or the user would notice a decrease in performance.

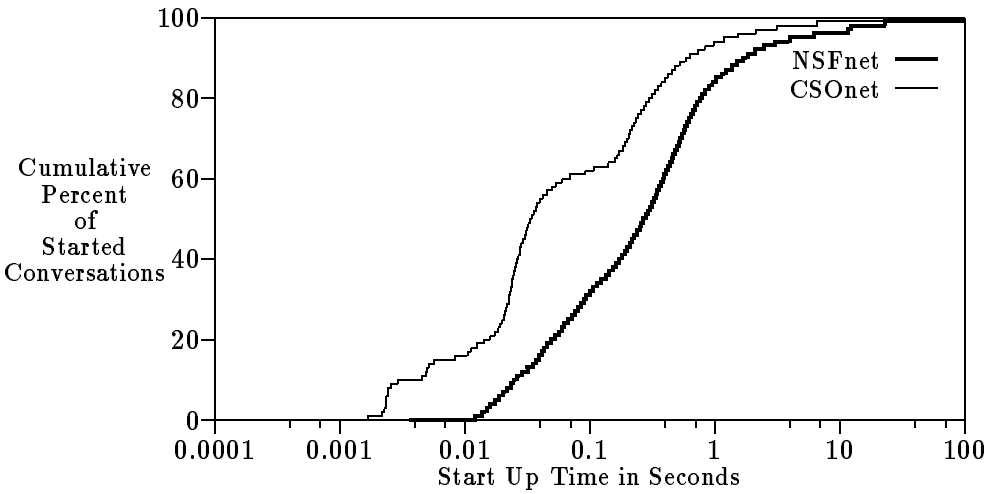

Figure 4: Connection Start Up Latency

2.1.3 ATM Segmentation and Reassembly

If ATM cells are segmented and reassembled (IP-ATM-IP) within a virtual circuit packet switch, the rate of this conversion is an important concern to the switch designer. Analysis of the arrival rate of TCP packets indicates the rate at which this processing would need to be accomplished. The maximum arrival rate is used as the worst case for the IP-ATM conversion. We will use a proposed standard devised by Cooper [4] which specifies that each cell will contain a 44 byte payload of actual IP data. The exact scheme for converting IP packets into ATM cells has not yet been standardized.

In this analysis the arrival rate of TCP packets on the CSOnet was monitored over several 24 hour periods. Typical use of approximately 250K packets was monitored between the hours of 10 AM and 7 PM. However, the maximum occurs nightly from 11-12 PM with 330K packets generated by backups performed over the network. The graph in Figure 5 illustrates the distribution of packet sizes for the 330K packets received and the top curve corresponds to the number of ATM cells required for each packet size group.

For the worst case of 330K IP packets, a total of 1.8 million ATM cells would be generated. This translates into approximately 500 cells per second. Although not graphically illustrated, the maximum arrival rate for TCP packets on the NSFnet is 14 times higher than that of the CSOnet. Therefore, the cell segmentation and reassembly rate would be approximately 7,000 cells per second.

2.1.4 Conversation Lifetimes and Bytes Transmitted

This section analyzes the lifetime of conversations and the bytes transmitted during a connections lifetime. Lifetime information is necessary for several reasons. First, it can be used to give a better understanding of how long a virtual circuit packet switch will typically need to maintain information about a virtual circuit. In addition, it can be used as an indicator as to whether

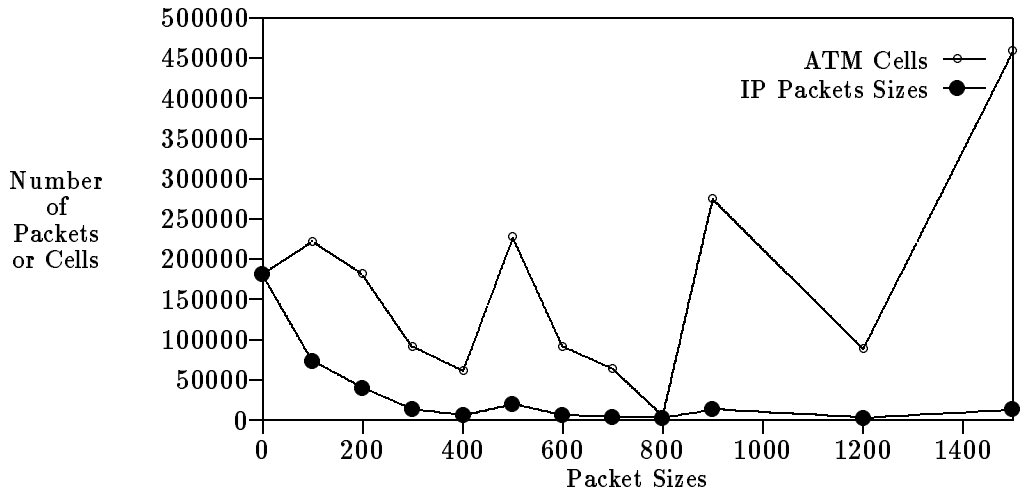

Figure 5: Packet Sizes and ATM Cells Required

a connection for a certain type of application is dormant or has been abruptly disconnected. A connection's lifetime is defined as the time between the reception of a TCP SYN packet to the reception of that connection's FIN.

Figure 6 depicts the lifetime for all TCP connections on the CSOnet and NSFnet. The CSOnet has a high concentration of connections that last less than one second, with 60% ending in under 4 seconds. The median lifetime is 2.48 seconds, mean lifetime is 519.85 seconds with a standard deviation of 838.58 seconds.

Behavior on the NSFnet is similar except that the lifetimes are longer. Although the results are not presented graphically, the delay in connection lifetime is typical of all applications run over the NSFnet. On the NSFnet 60% of the connections are terminated within 12 seconds. The NSFnet TCP connection median lifetime is 12.14 seconds, the mean lifetime is 375.01 seconds and the standard deviation is 1,096.58 seconds.

It has been shown that some applications open TCP connections for relatively long periods of time. A virtual circuit packet switch might therefore need to keep open a large number of permanent virtual circuits. Resource allocation within the switch can be improved if the variation of data traffic over a virtual circuit versus lifetime of the virtual circuit were known.

Figure 7 demonstrates that as Telnet connection lifetimes increase, the total number of bytes transmitted does not likewise increase. The maximum number of bytes transmitted per connection is not associated with the longest connection. This traffic pattern suggests that users transmitting large amounts of data over a link tend to transmit that data within 15 minutes. The longer interactive sessions do not show a great deal of traffic.

The statistical analysis of the NSFnet reveals the mean connection lifetime is 529.13 seconds, with a standard deviation of 1,945.54 seconds, and a maximum of 34.9K seconds. The mean number of bytes transmitted is 13.1K, with a standard deviation of 43.7K, and a maximum of 571K bytes.

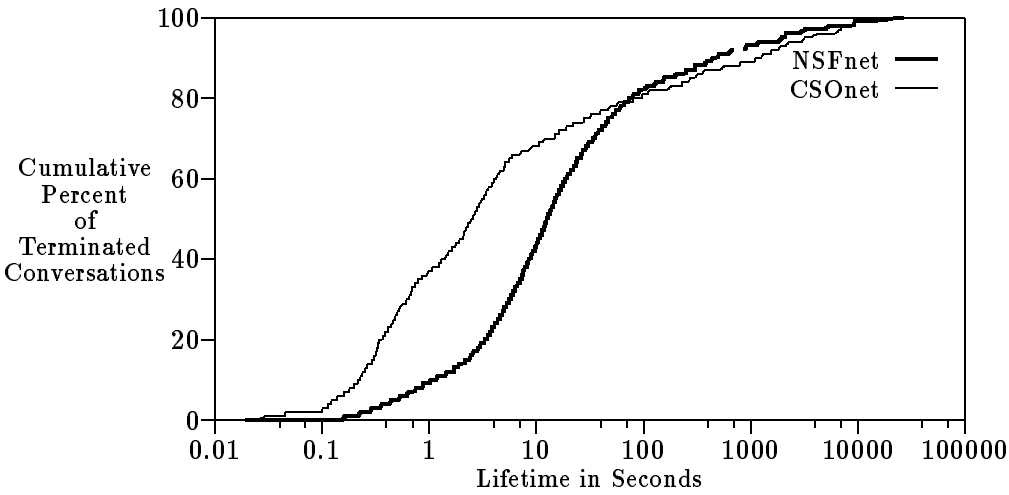

Figure 6: Lifetimes of All TCP Connections

Although not presented graphically, the statistical analysis for the CSOnet is as follows. The mean connection lifetime is 1,611.75 seconds, with a standard deviation of 2,851.88 seconds, and a maximum of 16.6K seconds. The mean number of bytes transmitted is 30.7K bytes, with a standard deviation of 59.5K bytes, and a maximum of 378K bytes.

Unlike interactive applications, it could be assumed that as the lifetime of a bulk transfer application (SMTP) increases the amount of data transmitted would also increase. The data in Figure 8 suggests that this is not always the case. The time of a conversation varies greatly; this is most likely due to differences in host and network speeds.

A statistical analysis of SMTP traffic reveals that for the NSFnet, the mean connection lifetime is 19.92 seconds, with a standard deviation of 62.39 seconds. The mean number of bytes transmitted is 3.4K, with a standard deviation of 7.4K, and a maximum of 83K bytes. Although not presented graphically, the statistical analysis for the CSOnet is as follows. The mean connection lifetime is 3.84 seconds, with a standard deviation of 2.85 seconds, and a maximum of 17.77 seconds. The mean number of bytes transmitted is 1.48K bytes, with a standard deviation of 1.08K bytes, and a maximum of 5.02K bytes.

The findings for the WAN traffic are in agreement with those of previous researchers[2]. However, our findings for the LAN traffic differ substantially from the WAN traffic in that users exhibit very different usage when interacting with local hosts. As the distinction between a LANs and a WANs diminishes, this difference may become less noticeable.

2.1.5 Locality of Transfers

For our purposes, locality is defined as a common destination for data. This is sometimes referred to as locality at the scale of hosts. If TCP connections exhibit locality, hosts and virtual circuit

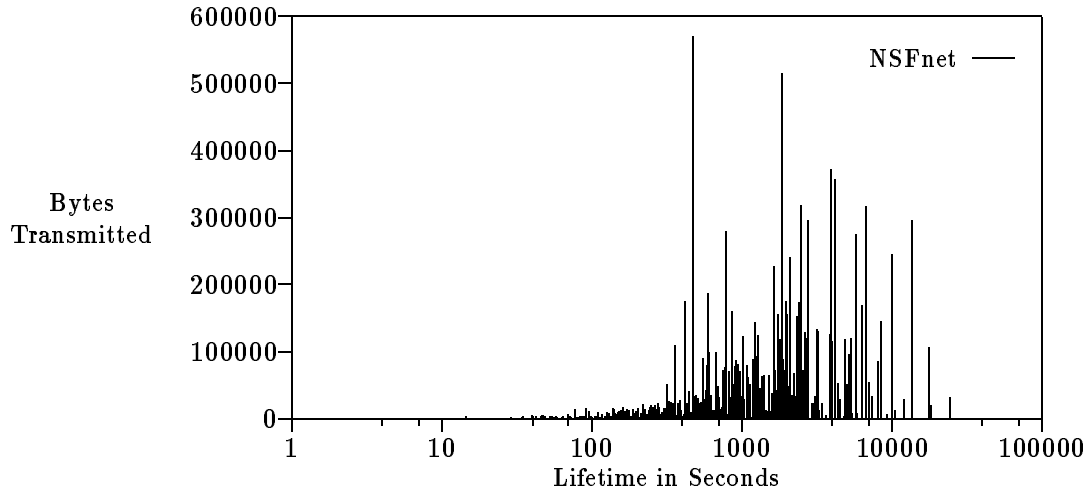

Figure 7: Telnet Lifetime vs Bytes Transmitted

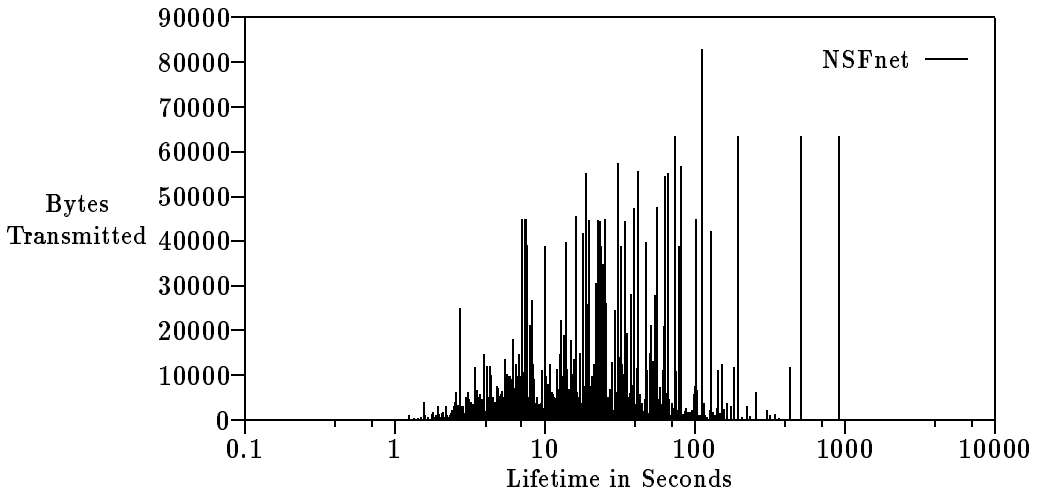

Figure 8: SMTP Lifetime vs Bytes Transmitted

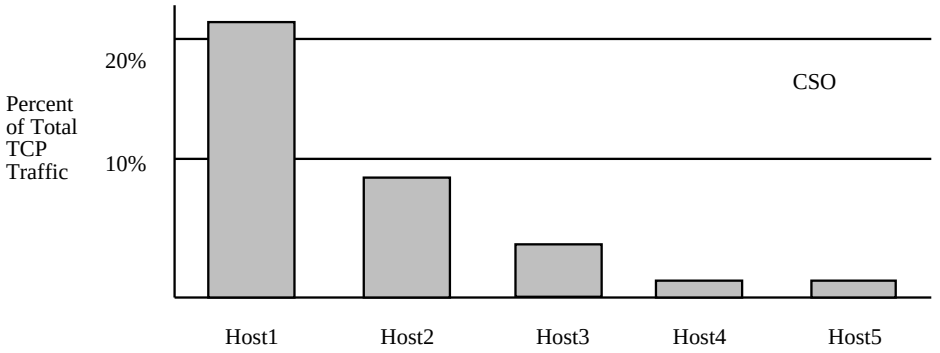

Figure 9: Locality of TCP Traffic for CSO.

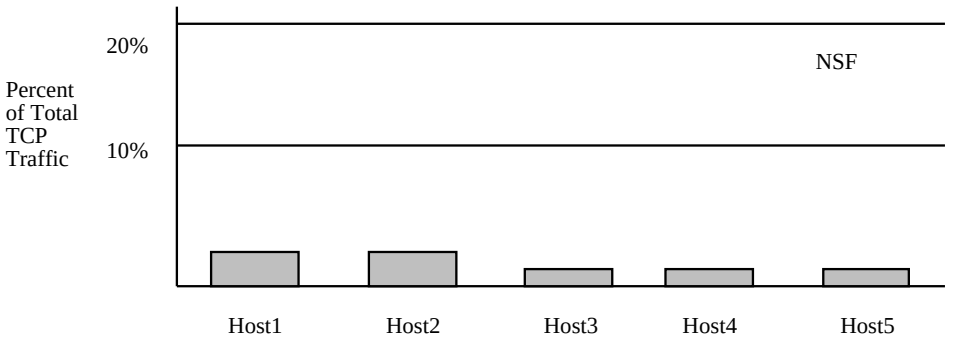

Figure 10: Locality of TCP Traffic for NSF.

packet switches could establish one connection to a distant network and multiplex several TCP conversations across that connection.

As shown in Figure 9, almost 40% of the TCP packets are addressed to only four hosts on the CSOnet. The high degree of locality is also present on the NSFnet as seen in Figure 10. Almost 15% of the TCP packets are sent to only five hosts.

A more detailed locality analysis of the CSOnet is shown in Figure 11. This graph represents all of the connections initiated within a 24 hour period. The hosts are numbered from 1 to 106 and are sorted first by the number of other hosts to which they communicate and then by the number of connections to those hosts. Every host initiating a connection is graphed.

The results of this analysis are as follows. A source host connects, on average, to 2.36 destination hosts, with a standard deviation of 3.68 connections. The maximum number of TCP connections that a source host made to N destination hosts was 29. The mean number of conversations between connected hosts is 9.75, with a standard deviation of 16.00 conversations. The maximum number of conversations between a host and its targets is 98. It can be seen in the graph that 70% of the hosts initiated connections to only 1 other host.

A more detailed locality analysis of the NSFnet is shown in Figure 12. This graph represents all of the connections initiated within a 24 hour period. The hosts are number from 1 to 147 and are sorted first by the number of the other hosts with which they communicate and then by the number of connections to those hosts. Due to the large number of host connections on the NSFnet, only 0.85% are simultaneously graphed. The graph is a sample of every 180th host that appeared

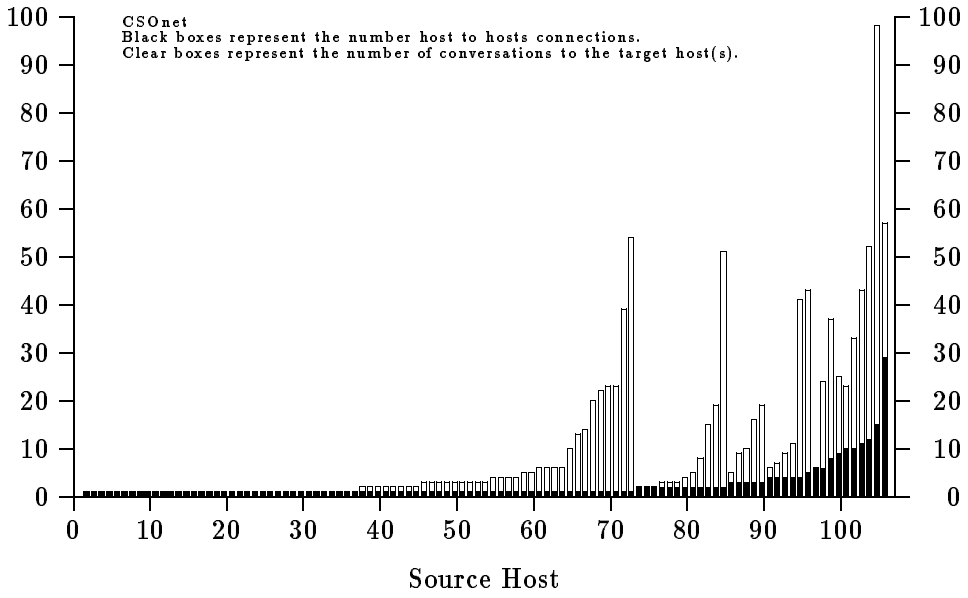

Figure 11: Connections on a Host to Host Basis (CSOnet)

in the data file.

The mean number of connections from a source host to a destination host is 2.66, the standard deviation is 14.48, and the maximum is 1,024. The mean number of conversations between these hosts is 38.67, with a standard deviation of 452.73. A total of 18,722 hosts initiated connections. On the NSFnet 80% of the hosts initiated connections to only 1 other host.

The information presented above gives the switch designer an indication of the load that a user might place on the switch. This information can also be used for multiplexing several TCP connections over one virtual circuit. Locality of UDP data transfer could be important in reducing set up overhead when UDP packets are sent over a virtual circuit network. In the next section we examine the UDP data we have collected.

2.2 User Datagram Protocol

For our analysis of UDP traffic measurements we studied the rate of arrival of UDP packets and the locality of UDP transmission. Protocols like UDP are more difficult to model as conversations than protocols that establish connections like TCP. Connectionless protocols may be difficult to support on a virtual circuit network if the set up and tear down rate of virtual circuits is slow. Nevertheless, many common network services like Sun's Network File System (NFS) and the Simple Network Management Protocol (SNMP) use UDP. Thus, connectionless traffic can be expected as a portion of the traffic received by a virtual circuit packet switch replacing current networks.

In terms of conversations UDP packets are *stateless*. It is impossible to determine if a given packet is at the beginning, middle or end of some conversation between two hosts. A simple technique to support connectionless protocols on ATM networks is to set up and tear down virtual circuits for each packet.

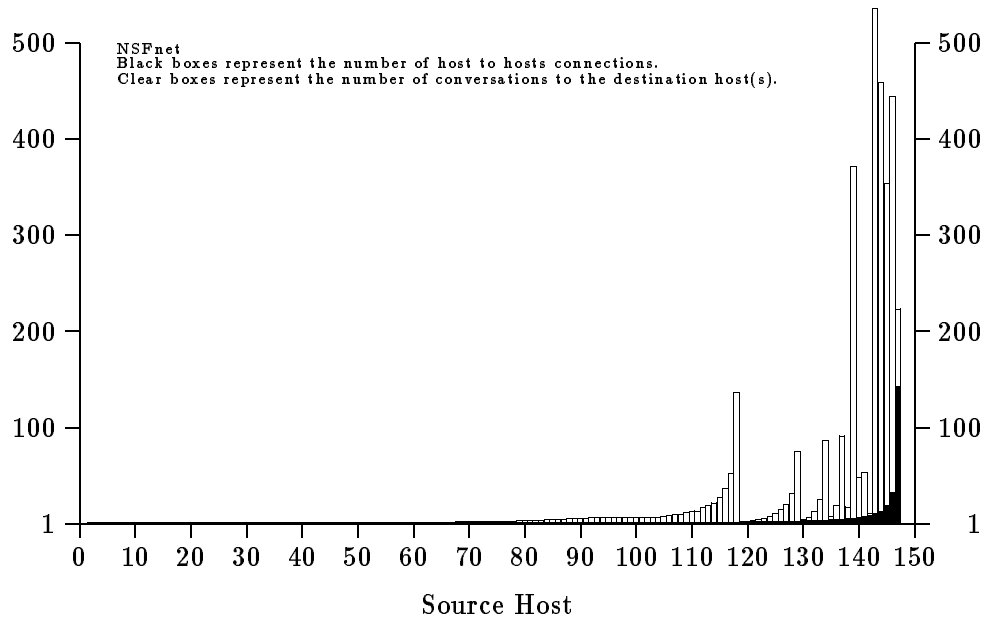

Figure 12: Connections on a Host to Host Basis (NSFnet)

Figure 13 presents the UDP packet arrivals per hour on the CSOnet and NSFnet. If a virtual circuit is set up and torn down for each UDP packet in the networks measured, the number of virtual circuits per hour needed would correspond to the number of packets per hour shown in Figure 13. On the CSOnet, there is a constant flow of approximately 100K UDP packets per hour and a peak usage of 203K UDP packets per hour. The mean number of UDP packets is 108K packets, with a standard deviation of 33.8K. Analysis of the individual UDP packets types showed a large number of NFS packets.

The NSFnet has a constant flow of approximately 300K UDP packets per hour and a peak usage of 881K UDP packets per hour. The mean number of UDP packets is 400.9K packets, with a standard deviation of 122.1K. Unlike TCP traffic, the number of UDP packets per hour are not several orders of magnitude higher on NSFnet than on the CSOnet. This could be due to the fact that the services offered by UDP are typically restricted to departmental LANs or perhaps local domains.

Like TCP, UDP traffic increases during prime user working hours. However, UDP is characteristically more irregular and bursty. Consequently the simple approach of establishing a unique virtual circuit for each UDP packet sent over an ATM network would require an average 250 new virtual circuits set up per second on an ATM replacement for NSFnet. Even though UDP traffic made up only 15% of the total volume of packets, the total number of UDP virtual circuits set up per second could be much higher than TCP's 10.

One potential technique for reducing the number of virtual circuits created for UDP traffic would be to multiplex multiple UDP packets with similar destination hosts over the same virtual circuit. If UDP traffic exhibits locality, this technique could significantly reduce the number of virtual circuit call set ups and tear downs.

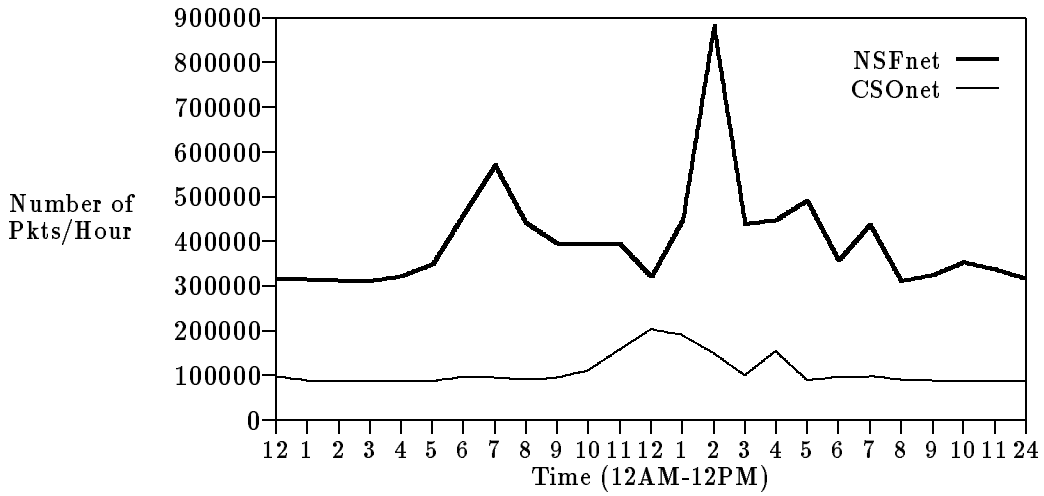

Figure 13: UDP Packets on Received/Hour

Analysis of UDP packet locality on the CSOnet and NSFnet indicates not only that it exists but that it is a prominent characteristic of a heavily loaded network. Figure 14 demonstrates that approximately 50% of the UDP packets are destined to only five hosts on the CSOnet. Figure 15 indicates that almost 30% of the UDP packets are destined to only five hosts.

Thus, our measurements indicate that multiplexing UDP packets over a single virtual circuit would be a significant optimization. Multiplexing would reduce both the total number of virtual circuits and the overhead of virtual circuit set up and tear down. This analysis is in agreement and related to previous research [8].

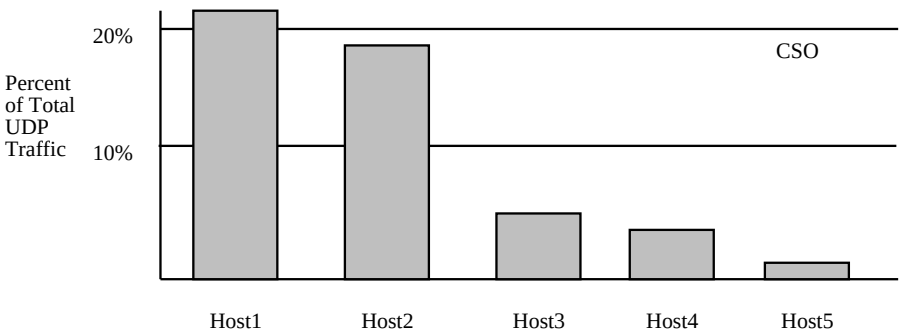

Figure 14: Locality of UDP Traffic on CSOnet.

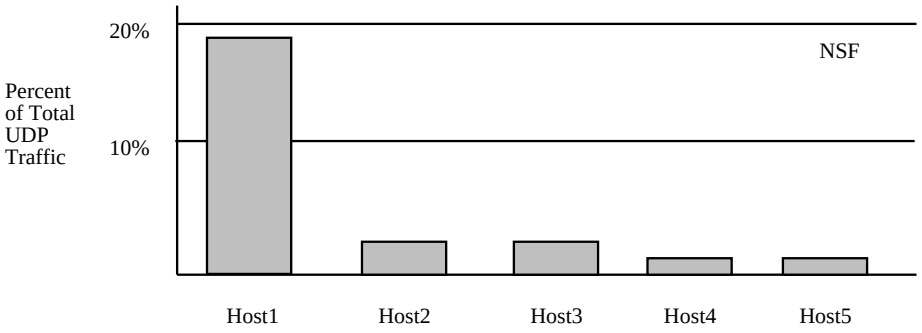

Figure 15: Locality of UDP Traffic on NSFnet.

3 Conclusions

We present measurements of IP traffic on two networks, a LAN and a WAN. We examine properties of the traffic that would impact a replacement ATM network and the design of its virtual circuit packet switches. The analysis estimates the number of virtual circuits, their set up rates, the rate of IP-ATM segmentation and reassembly, conversation lifetimes, bytes transmitted in a lifetime, and locality of transfers in such a network.

The following summarizes our results:

- Number of Concurrent Conversations:
  - CSOnet TCP. 20-30, a peak of 104, mean of 55.68.
  - NSFnet TCP. A broad range with a peak of 423, mean of 209.30.
- Rates of Conversation Set Up:
  - CSOnet TCP. 0.6-0.8 per second, a peak of 14 per second.
  - NSFnet TCP. 2.22-3.89 per second, a peak of 181 per second.
- Latency of Conversation Start Up:
  - CSOnet TCP. Median time 16.58 msec, mean 6.59.
  - NSFnet TCP. Median time 132.09 msec, mean 10.73.
- ATM Processing Overhead:
  - CSOnet IP. A maximum of 500 cells per second.
  - NSFnet IP. A maximum of 7,000 cells per second.
- Conversation Lifetimes:
  - CSOnet TCP. Median 2.48 seconds. Mean 519.85 seconds.
  - NSFnet TCP. Median 12.14 seconds. Mean 375 seconds.
- Bytes Transmitted over Lifetime:

  - CSOnet Interactive (Telnet):
    * Mean connection lifetime 1,611 seconds.
    * Mean number of bytes transmitted 16.6K.
  - CSOnet Bulk Transfer (SMTP):
    * Mean connection lifetime 3.84 seconds.
    * Mean number of bytes transmitted 1.48K.
  - NSFnet Interactive (Telnet):
    * Mean connection lifetime 529 seconds.
    * Mean number of bytes transmitted 13.1K.
  - NSFnet Bulk Transfer (SMTP):
    * Mean connection lifetime 19.92 seconds.
    * Mean number of bytes transmitted 3.4K.

- Locality of Transfer:
  - CSOnet TCP. 40% and UDP. 40% of traffic to 4 hosts.
  - CSOnet TCP. Mean Number of Host to Host Connections = 2.65
  - CSOnet TCP. Mean Number of Host to Host Conversations = 12.26
  - NSFnet TCP. 15% and UDP. 30% of traffic to 5 hosts.
  - NSFnet TCP. Mean Number of Host to Host Connections = 2.66
  - NSFnet TCP. Mean Number of Host to Host Conversations = 36.67

The Internet measurements provide a better understanding of conversations and the performance requirements of an ATM virtual circuit switch in a LAN and WAN environment. In particular, the results indicate that there is a minimum number of VCI's that should be supported by a switch and that the rate of connection set up has large variance. Based on these measurements, we plan to determine the effect of IP traffic on ATM switch design issues such as buffer allocation in call set up, the adaptation layer design and performance, and congestion control. The ultimate goal of our research is to enhance the protocol and control software of ATM switches like XUNET.

4 Acknowledgments

We would like to thank the University of Illinois at Urbana Champaign Campus Area Networking group for their cooperation with this study.

References

[1] Robert T. Braden and Annette L. DeSchon. NNStat 2.5, September 1990. Internet Statistics Collection Package.

[2] R. Caceres, P. Danzig, S. Jamin, and D. Mitzel. Characteristics of Wide-Area TCP/IP Conversations. In *ACM SIGCOM*, September 1991.

[3] Douglas Comer. *Internetworking with TCP/IP.* Prentice Hall, 1990.

[4] E. Cooper, O. Menzilcioglu, R. Sansom, and F. Bitz. Host Interface Design for ATM LANs. *Proceeding of the 16th Conference on Local Computer Networks*, October 1991.

[5] A. G. Frazer, C. R. Kalmanek, A. E. Kaplan, W. T. Marshall, and R. C. Restrick. Xunet 2: A Nationwide Testbed in High-Speed Networking. In *Proc. of IEEE INFOCOM '92*, May 1992.

[6] Van Jacobson, Craig Leres, and Steve McCanne. Tcpdump, June 1989. Available via anonymous ftp to ftp.ee.lbl.gov.

[7] Steven E. Minzer. Broadband ISDN and Asynchronous Transfer Mode (ATM), September 1989. IEEE Communications Magazine.

[8] Jeffrey C. Mogul. Locality at the Scale of Processes. In *ACM SIGCOM*, September 1991.

[9] Dennis Cheng-Hong Sng. Network Monitoring and Fault Detection on the University of Illinois Campus Computer Network. Master's thesis, University of Illinois, 1990.

5 Appendix

This section contains graphs showing detailed plots for several of the topics covered in the body of the paper.

The graphs are organized as follows:

1. Bulk Transfer Lifetimes
   - SMTP Figure 16.
   - FTP Figure 17.
2. Interactive Lifetimes
   - Telnet Figure 18.
   - X Windows Figure 19.
3. Bytes Transmitted for Various Lifetimes
   - CSOnet SMTP Figure 20.
   - CSOnet Telnet Figure 21.
   - CSOnet X Windows Figure 22.
   - NSFnet X Windows Figure 23.
4. Packet Lengths
   - All IP Figure 24.
   - TCP Figure 25.
   - UDP Figure 26.

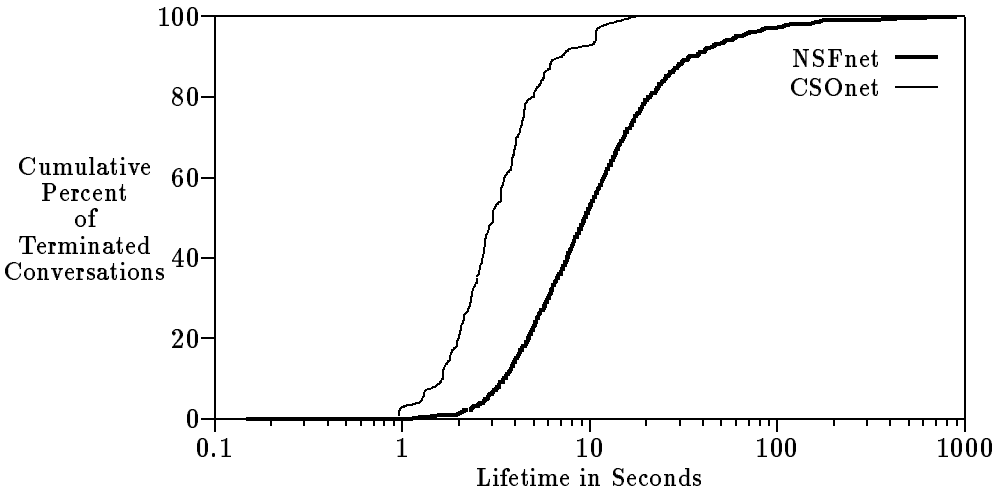

Figure 16: SMTP Lifetimes for CSOnet and NSFnet

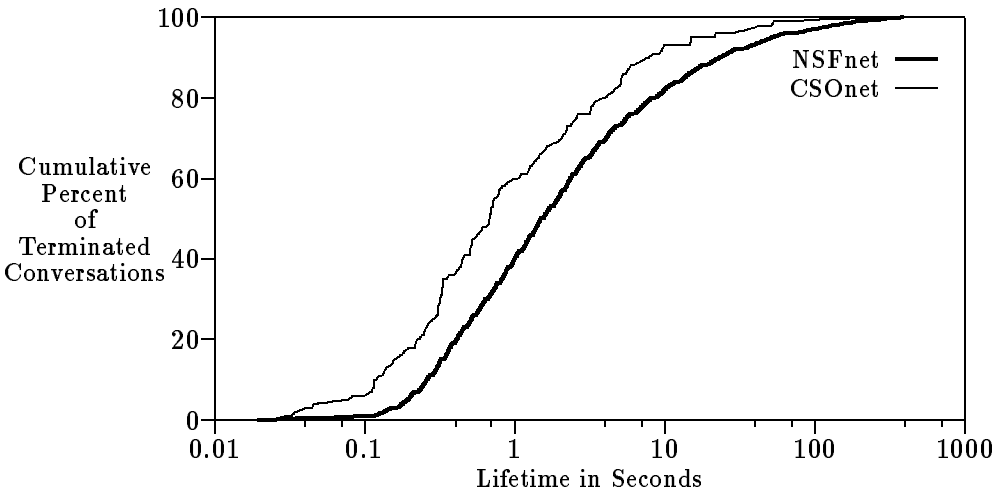

Figure 17: FTP Lifetimes for CSOnet and NSFnet

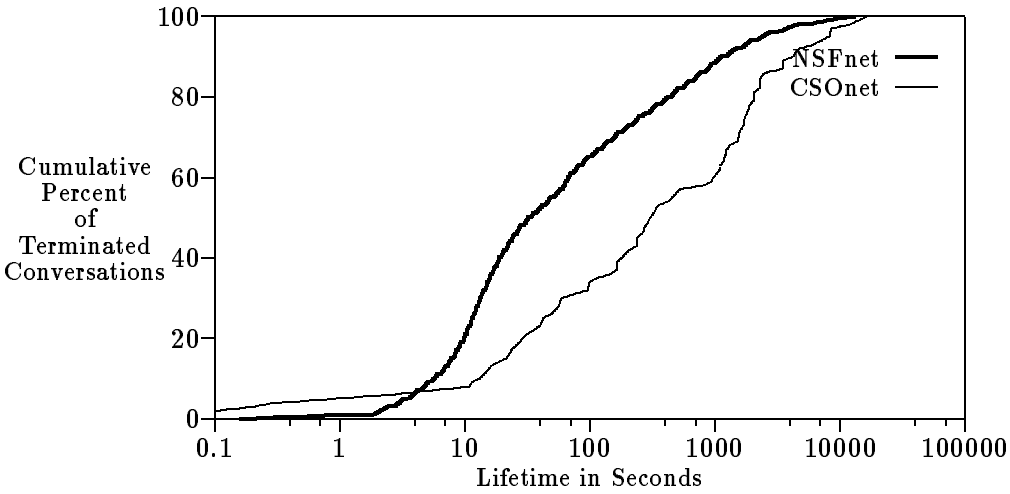

Figure 18: Telnet Lifetimes for CSOnet and NSFnet

Figure 19: X Windows Lifetimes for CSOnet and NSFnet

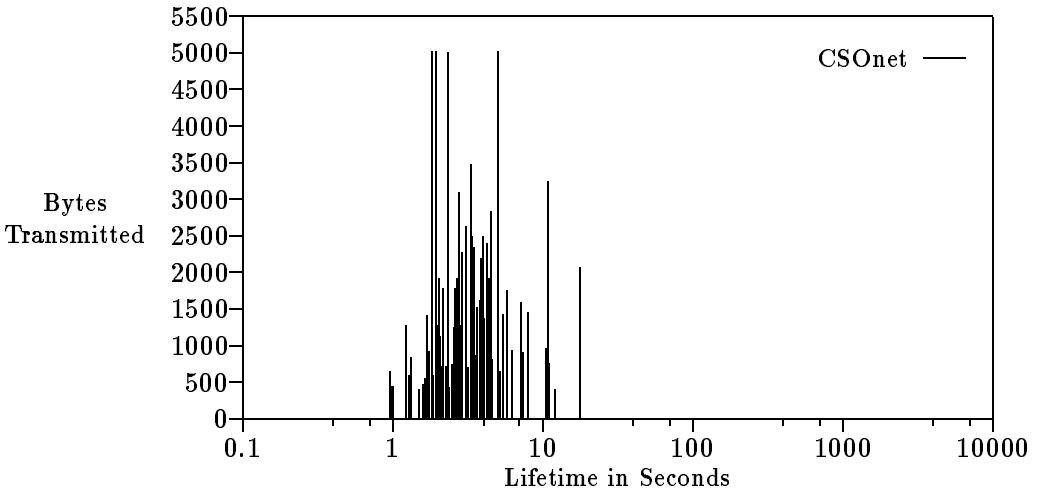

Figure 20: SMTP Bytes *vs* Lifetimes for CSOnet

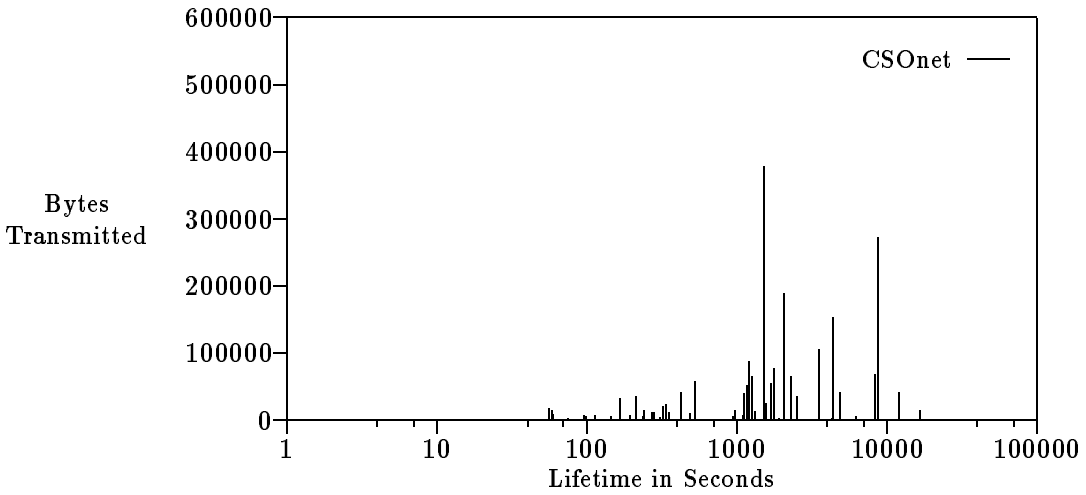

Figure 21: Telnet Bytes *vs* Lifetimes for CSOnet

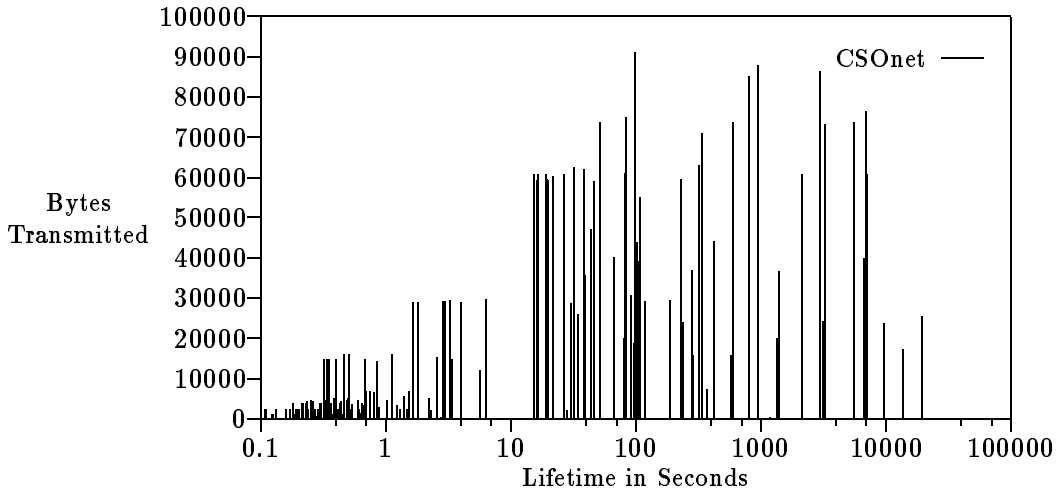

Figure 22: X Windows Bytes *vs* Lifetimes for CSOnet

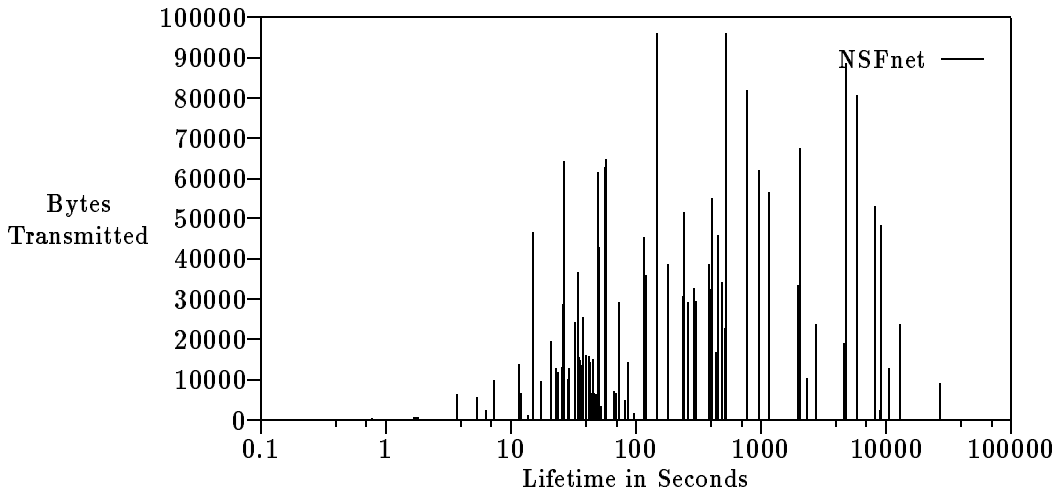

Figure 23: X Windows Bytes *vs* Lifetimes for NSFnet

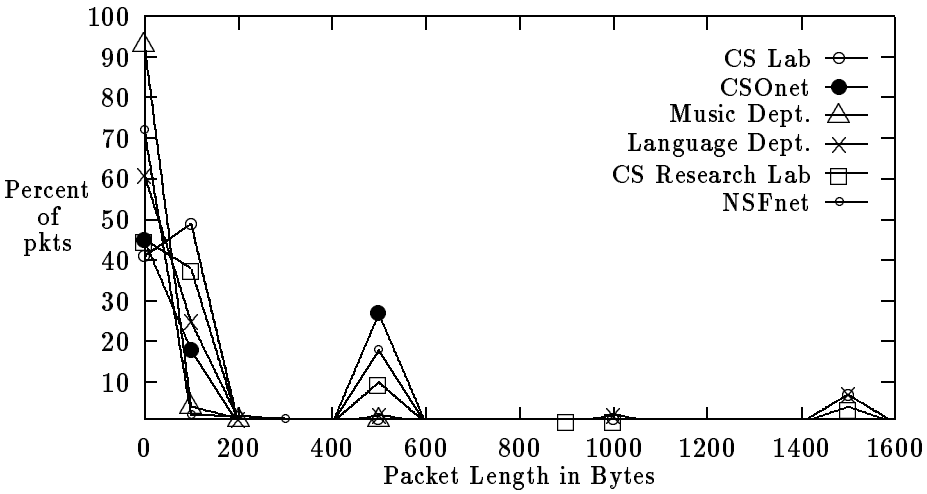

Figure 24: All IP Packet Lengths for Various Campus Networks

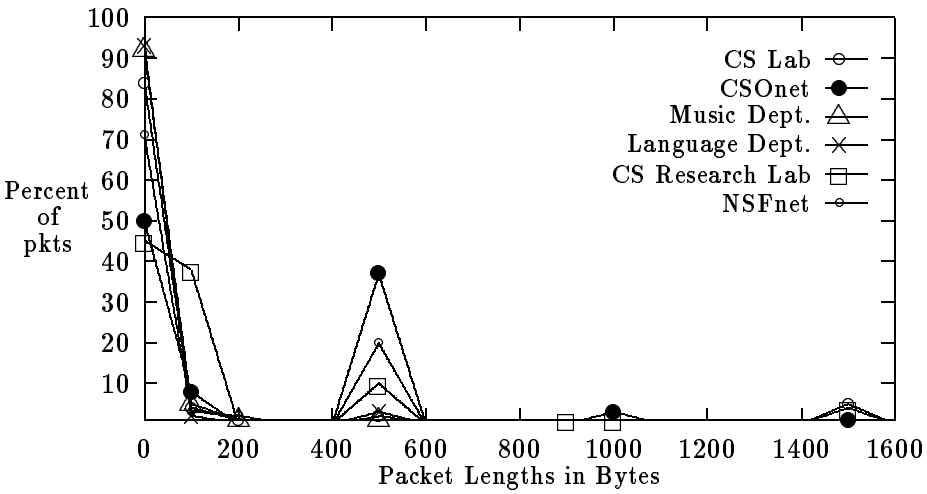

Figure 25: TCP Packet Lengths for Various Campus Networks

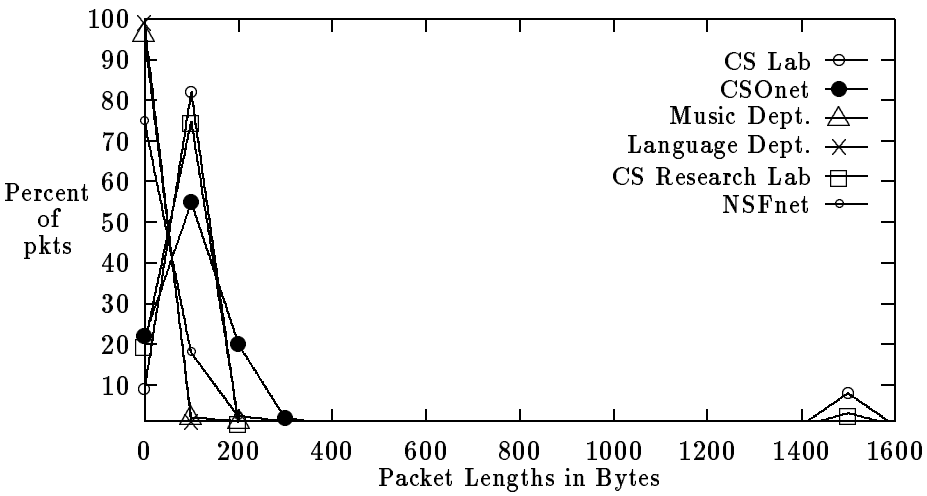

Figure 26: UDP Packet Lengths for Various Campus Networks