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.

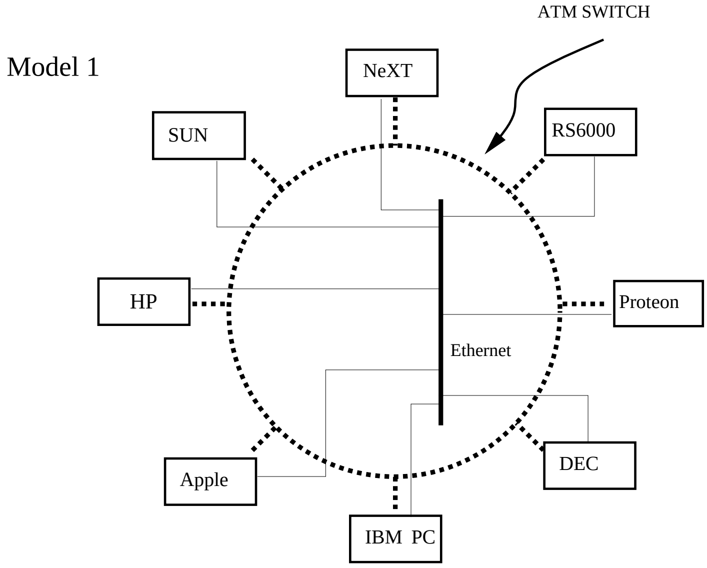

Computing Services Office Network

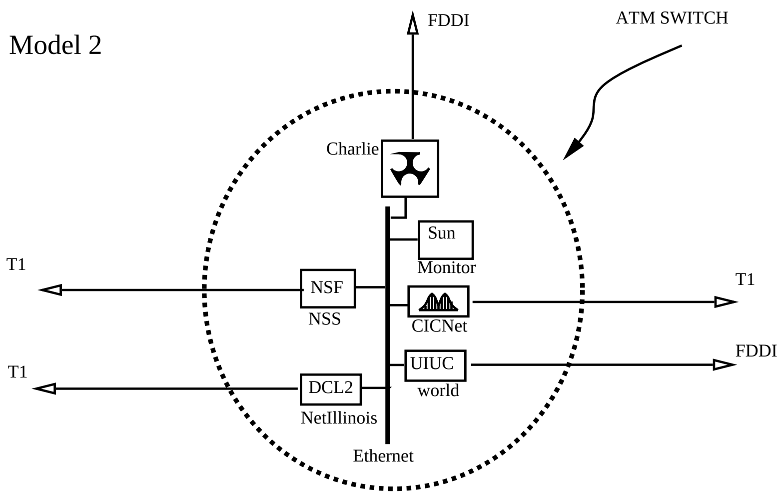

NSFnet Connection at UIUC

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

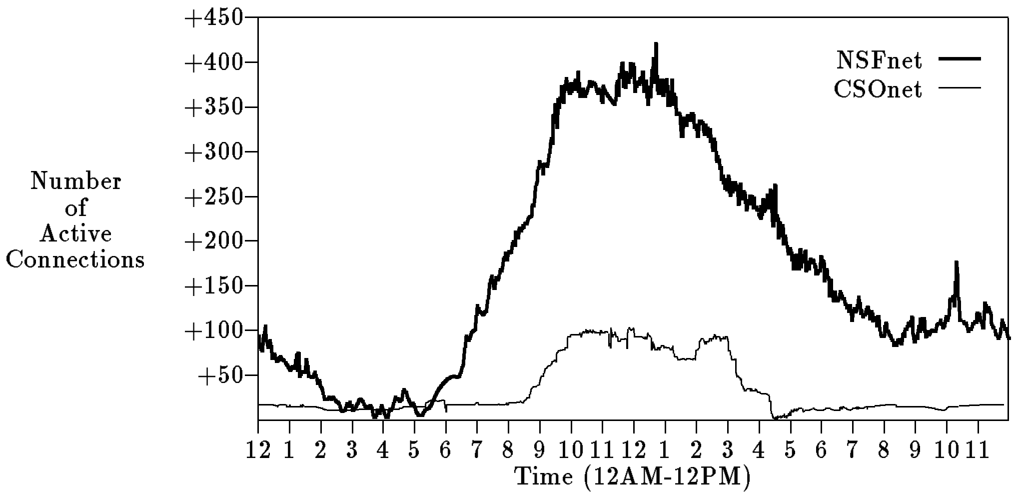

Figure 2: Simultaneous Connections

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.