

# Deploying ATM in a Data Network: An Analysis of SVC Requirements

*Russell J. Clark*  
*Ronald R. Hutchins*  
*Scott Walker Register*

**GIT-CC-95-14**

March 15, 1995

## **Abstract**

Past and current campus data networks generally utilize connection-less network layer protocols like IP. Future networks are expected to use ATM which provides a connection-oriented service. New ATM switches are coming to market with baseline parameters such as the number of concurrent SVC's supported and the number of call setups possible per second. These parameters will determine the usefulness of this equipment in large campus networks for current applications and topologies. We present an analysis of the effect of these parameters on the ability of a circuit-based network to support the networking requirements of a college campus. We performed this analysis using IP traffic logs from three distinct networks within our campus. We consider four different ATM deployment scenarios and two circuit replacement algorithms. From our analysis we derive expected requirements for SVC setup rates and hold times necessary to support the measured traffic.

College of Computing  
Georgia Institute of Technology  
Atlanta, GA 30332-0280  
rjc@cc.gatech.edu

# 1 Introduction

The communications industry is moving toward the use of Asynchronous Transfer Mode (ATM) as the future workhorse of networking. Unfortunately, very little is known about how to migrate current large scale configurations to this new technology. With ATM, the difference between Tel-Co networking and campus networking is seemingly fading as this migration takes place. An important aspect of this change is that ATM utilizes a Tel-Co style connection-oriented paradigm rather than the connection-less approach of typical local network protocols. This traditional distinction leads to incompatibilities between the new technology and the currently deployed data network systems. Special steps must be taken to ensure that ATM switches are designed to meet the requirements of the traditional data networking community.

Two popular approaches to deploying ATM in the data network environment are Local Area Network (LAN) emulation and Internet Protocol (IP) over ATM [4]. In LAN emulation, the ATM switch provides a Media Access Control (MAC) level emulation of the services in an IEEE 802 LAN [6]. This involves the use of a MAC sub-layer that builds a best-effort, datagram service on top of the the connection oriented ATM service. The advantage of this approach is that it provides support for most popular LAN protocols to operate over ATM. One of the drawbacks of this approach is that most LAN based protocols rely on the broadcast characteristics of 802 LANs for operations like address resolution (i.e. ARP). While it is possible to implement broadcast in the ATM switch, it may be more efficient for the switch to directly support address resolution rather than carry the broadcast traffic.

The IP over ATM approach is a mechanism for transmitting IP traffic directly over ATM switches without the emulation of a MAC layer [5]. In this approach, an ATM virtual circuit is used to directly connect two IP members of a logical subnetwork. Communication to hosts outside of the logical subnetwork requires the use of an intermediate IP router. Both of these approaches involve providing a connection-less datagram service on top of the connection oriented ATM service. This can be done using  $O(N^2)$  Permanent Virtual Circuits (PVC) that interconnect all of the ATM ports and provide communication when it is needed. However, in a large network this will require the maintenance of thousands if not millions of PVCs. The alternative is to use Switched Virtual Circuits (SVC) that are created dynamically between two ports that need to communicate and then removed when that communication is finished. The challenge in a datagram environment is to determine when a communication is finished and it is therefore appropriate to remove a connection.

This research focuses on the evaluation of alternatives for PVC and SVC control policies that provide network connectivity for datagram networks. We gathered IP traffic on three distinct networks within our campus environment: a FDDI backbone and two distinct

Ethernet subnets. We developed simulations and analyzed the traffic to determine the switch VC requirements needed to support the current network traffic patterns. We consider four different ATM configuration scenarios; scenarios that represent different stages of ATM deployment in the current network. We analyze both timeout based and least recently used SVC replacement strategies. We compare the requirements for each scenario and suggest where each would be most appropriate. From our analysis we derive expected requirements for SVC setup rates, circuit hold times, and circuit cache sizes necessary to support the measured traffic.

In the next section we describe the network environment where our measurements took place. In Section 3 we present an overview of our analysis and describe the four different configuration scenarios we studied. The detailed results of our timeout based simulation are presented in Section 4. In Section 5 we present the results of our least recently used simulation analysis. Section 6 includes related work and a summary of our observations.

## 2 The Evaluation Network

In this research we analyzed the traffic patterns of the current campus-wide data network at the Georgia Institute of Technology. An overview of the current network, known as GTNet, is presented in Figure 1. The network is composed of two main FDDI rings and multiple (>150) Ethernet segments. Other technologies (e.g. local FDDI, token-ring) are also found in limited use throughout the campus. The backbone consists of a FDDI ring connecting six routers. The majority of Ethernet subnets are connected directly to one of these six routers. The backbone is connected through one of the routers to a gateway ring that is in turn connected to an external router providing access to the Internet. Several systems that provide high-traffic services (e.g. News, Domain Name Service, e-mail) are connected directly to this gateway ring.

The entire GTNet network consists of approximately 150 subnets and over 14,000 hosts. The network connects a wide variety of systems and the IP, AppleTalk, and Novell Netware protocols are all supported in at least portions of the network.

For our study we analyzed the IP<sup>1</sup> traffic on three specific portions of GTNet: the FDDI backbone, the dial-in subnet, and a departmental subnet. The FDDI backbone is a dedicated subnet connecting the five main campus routers. Other than some dedicated management systems, no end-systems are connected directly to the FDDI backbone. All traffic between subnets connected to different campus routers is carried on this backbone. The dial-in subnet serves as the connection point for the campus modem resources. Virtually all of the traffic on this subnet is associated with one of six terminal servers connecting

---

<sup>1</sup>We did not study SVC requirements for AppleTalk and Novell Netware protocol traffic.

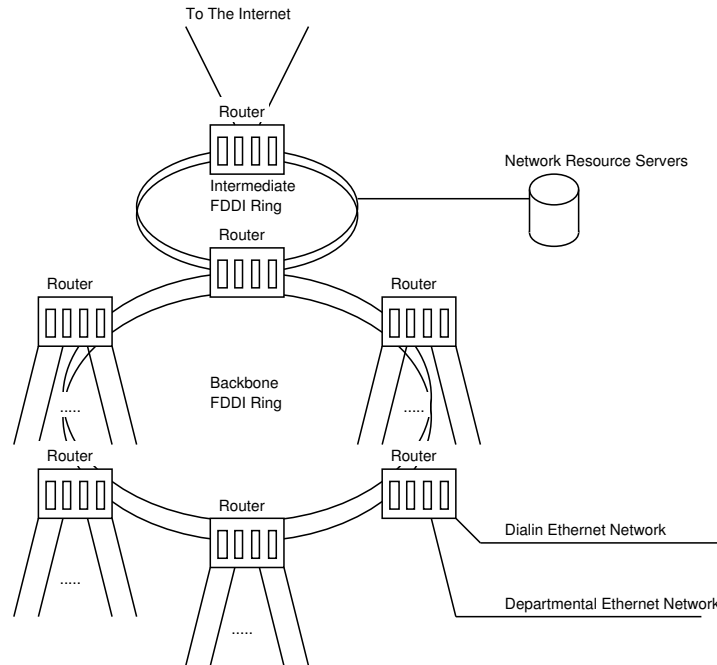


Figure 1: Overview of Georgia Tech Network

a modem pool of 256 modems. Several of the more popular central campus systems are also connected directly to this dial-in subnet so that traffic from dial-in users does not need to cross the router. The third subnet we analyzed is a departmental subnet with approximately 250 hosts. These include office workstations and file-servers.

We collected network traffic information from these three subnets using the NNStat [2] program running on a Sparcstation 1+ from Sun Microsystems. The traffic was monitored during the months of November and December, 1994. Using this configuration we estimate packet losses to have been less than 1% during our monitoring.

### 3 Analysis Overview

As with most new technologies, when ATM is deployed in current data networks it will be introduced gradually. In the case of GTNet, ATM is being introduced first as a handful of new ATM networks and second as a replacement for certain highly utilized operational networks. Given this need for partial deployment and integration with the current network, we have designed our analysis of VC requirements around four different switch deployment scenarios. Each of these scenarios will be appropriate for some portions of a campus data network.

The scenarios we analyze differ in the definition of the boundary of the ATM network. Our first scenario is a complete ATM network where every network host is directly connected to an ATM switch through its own dedicated switch port. Communication between hosts involves establishing a virtual circuit through the switch network from one host to the other. This scenario is identified as the *host-nogate* approach, indicating that no gateway device is used. This is a "worst-case" scenario from the point of view of switch VC requirements. While we expect to have several isolated subnets where this approach is valid, it is difficult to foresee any time in the near future where this scenario would apply on a campus-wide basis.

The next scenario we consider is like the first in that hosts are directly connected to ATM ports. The difference is that only the hosts on the subnet of interest are directly connected to the local switch. All other hosts are accessed through a gateway device (e.g. router) connected to one of the switch ports. This scenario would apply when a high use subnet, such as the departmental subnet in our study, is replaced with an ATM switching network. We will refer to this scenario as the *host-gate* approach, indicating the use of a gateway.

The third scenario we consider is the case where ATM is used as a backbone network connecting several routers. An example of this approach, which we refer to as *net-gate*, would be to replace the FDDI backbone of Figure 1 with an ATM switching network. In this approach there are no end-hosts directly connected to the ATM switch.

The fourth scenario, *net-nogate*, is a slight modification of this approach. It represents replacing all backbone and intermediate networks with an ATM switching network. In this case, each subnet in the network is connected via its own ATM port and, as in the *net-gate* approach, all end-hosts are connected indirectly through another network device.

In our analysis we assume that a single virtual circuit between two nodes, either routers or end-hosts, will accommodate all traffic carried between those two nodes. For instance, if there are several concurrent TCP connections between two hosts then the traffic for those connections will be multiplexed over the same circuit. Furthermore, any UDP traffic between the hosts would also use this circuit. This approach eliminates the cost of establishing a new virtual circuit when one is already in place for carrying the datagram traffic. While this assumption is valid for both the LAN emulation and IP over ATM models, it may not always be appropriate. In some cases, it may be necessary to assign a connection to a separate virtual circuit in order to guarantee resources or isolate the traffic for security reasons.

To analyze the switch requirements for our sample data we developed a trace-based simulation of a timeout based circuit replacement algorithm with a variable circuit hold time. In this simulation, if traffic arrives between two addresses where a circuit is not present then a new circuit is created between those addresses. If a circuit is idle for the current circuit hold time then that circuit is removed. We ran this simulation using the

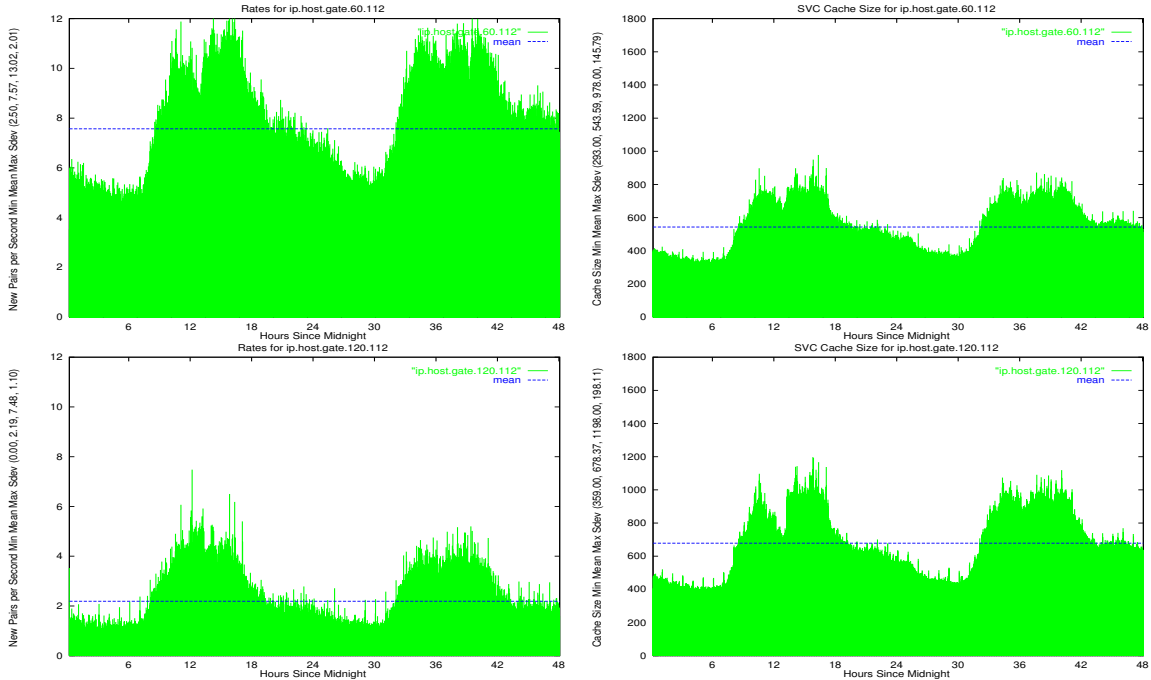


Figure 2: FDDI Backbone Pair Analysis

traffic logs collected from our three study networks. For each traffic log we determined the SVC creation and circuit cache size for each of the four scenarios described above.

An example of the analysis we performed on the GTNet networks is presented in Figures 2 and 3. These figures show the detailed analysis of the *host-gate* scenario for the FDDI backbone. This analysis is from a 48 hour time period covering all of November 22-23, 1994. The graphs on the left side show the observed circuit creation rates for each one minute interval of the day. The graphs on the right side show the corresponding number of active SVCs during each one minute interval. The horizontal line in each graph indicates the observed mean value. The top two graphs in Figure 2 show the results from a simulation using a 60 second circuit hold time. This means that if a pair experience no traffic for 60 seconds, the circuit between that pair is dropped. The next two graphs show the results from a simulation using a 120 second circuit hold time. Subsequent graphs showing 180, 240 and 300 second hold times are presented in Figure 3. The two peaks in each of these graphs are caused by the increase in switch demands during the busy midday periods of the two days we studied. These graphs show that by increasing the circuit hold time we can decrease the number of SVCs that must be created by the switch each second. Subsequently, this longer hold time increases the number of concurrent SVCs that must be maintained by the switch.

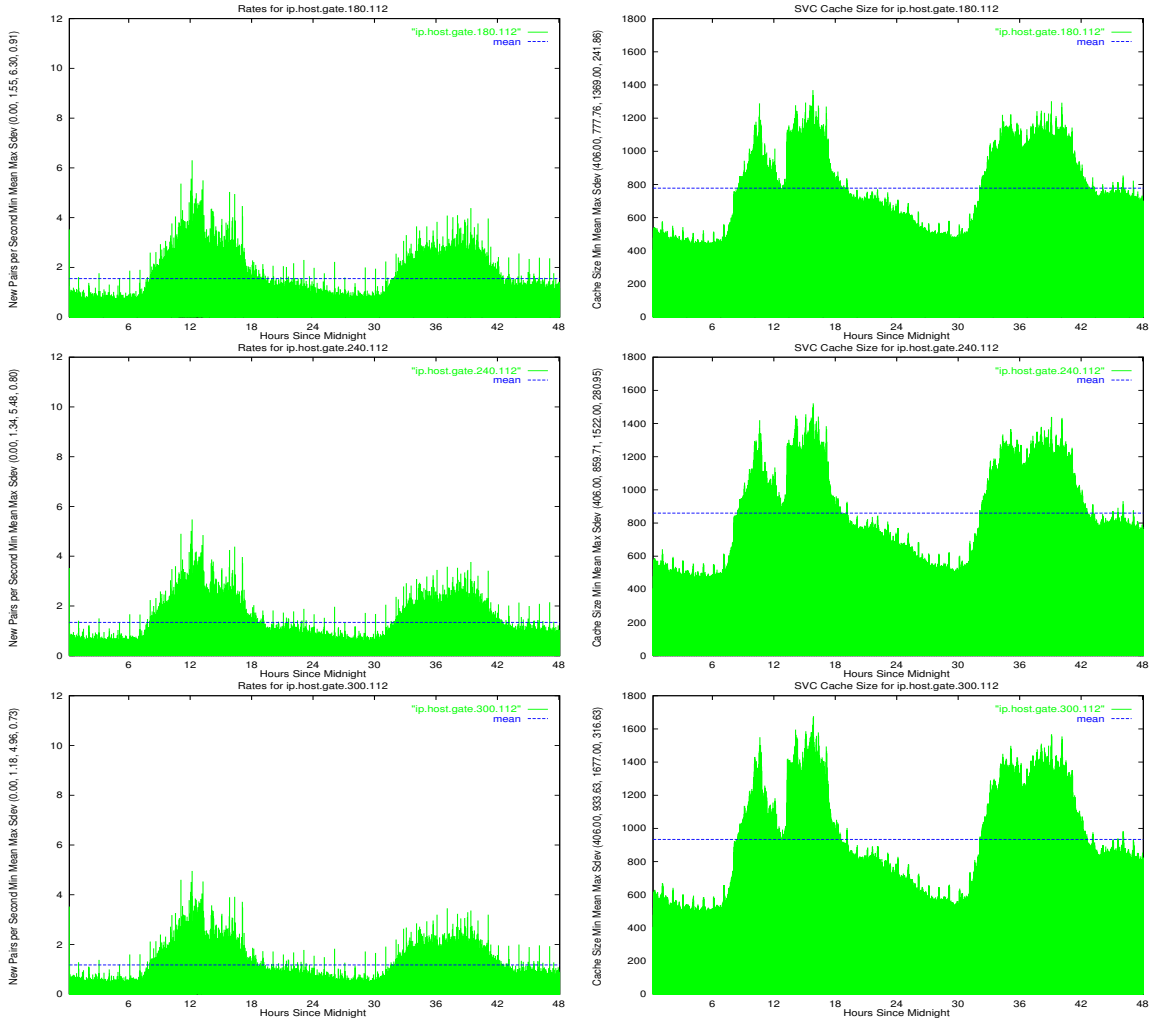


Figure 3: FDDI Backbone Pair Analysis (cont'd)

## 4 Timeout Based Evaluation

In this section we further explore the requirements for SVC hold time and cache size in data networks using various protocol and network configurations. The analysis presented is done using the timeout based circuit replacement strategy. Such a strategy is most appropriate when there is a hold-time cost associated with ATM circuits and it is cost-effective to remove circuits that are not actively used.

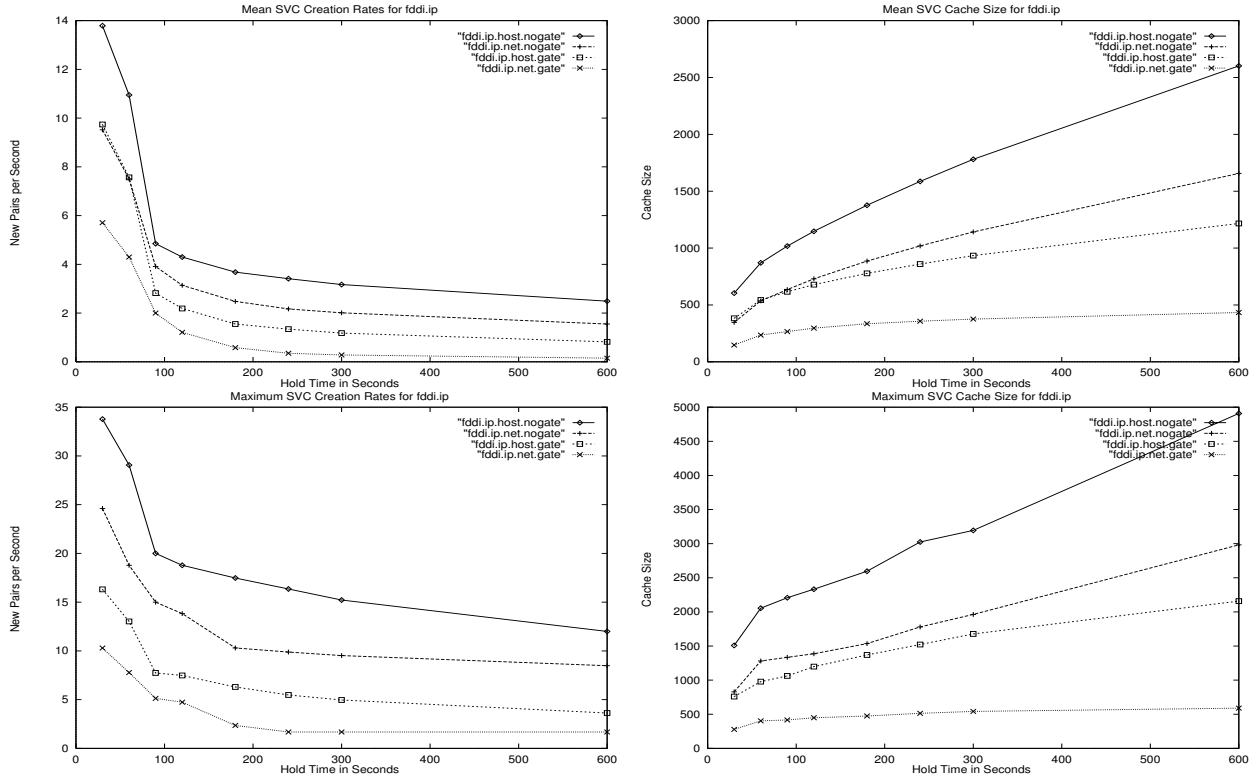


Figure 4: FDDI Backbone Pair IP Summary

## 4.1 FDDI Backbone Analysis

In Figure 4 we summarize the results from our analysis of IP traffic on the FDDI backbone. In this network, the two gateway scenarios involve placing a gateway at the off-campus interface to the Internet. The four plots in each graph represent the four distinct ATM deployment scenarios described above. The first graph presents the mean SVC creation rates as the SVC hold time is increased from 30 seconds to 10 minutes. The bottom, left graph presents the maximum SVC creation rates observed for the analyzed time period. These are derived from the busiest 1-minute interval. The upper, right graph presents the mean cache size or average number of active SVCs observed for various hold times. The maximum cache size is presented in the bottom, right graph<sup>2</sup>.

The first thing we observe in these graphs is that the worst case scenario of host-nogate requires a maximum creation rate of 33.78 SVCs per second. The more likely deployment scenario for the backbone network is net-gate, which has no directly connected end-systems. This scenario sees a maximum creation rate of 10.29 SVCs per second and a mean

<sup>2</sup>Summary statistics are calculated after discarding the first two minutes of the simulation runs. This allows the simulation to reach steady-state behavior.



of 5.71. These rates are significantly lower than the worst case and indicate the marked difference in switch requirements when various deployment scenarios are considered.

From these graphs we also observe that increasing the SVC hold time between 0.5 and 10 minutes always results in a reduction in the rate of SVC creations and an increase in the number of concurrent SVCs in use. The interesting point appears with the 90 second hold time samples. This point forms the “knee” in the creation rate graphs for all four scenarios. For the first three samples, 30, 60 and 90 seconds, increasing the hold time results in a significant reduction in creation rates. Increasing the hold time beyond 90 seconds results in a much less pronounced reduction in the creation rates.

The conclusion we can draw from this analysis is that the majority of active pairs will experience an idle-time of less than 90 seconds before traffic is transmitted again. In other words, if a pair is not heard from within 90 seconds there is a high probability that it will not be heard from again within 10 minutes. An SVC management algorithm that holds circuits for longer than 90 seconds is likely to see minimal benefit from this longer hold time. Using the 90 second hold time, our simulation results indicate that for the net-gate scenario, a 2.00 mean SVC creation rate is required with a maximum rate of 5.13 circuits per second. This requires a mean cache of 265.74 circuits and a maximum of 417. In the worst case scenario of host-nogate, the 90 second hold time requires a mean SVC creation rate of 4.85 with a maximum of 20.00 per second. This requires a mean cache of 1017.33 circuits and a maximum of 2207.

While the ATM LAN implementations we are interested in will not distinguish between various transport protocols, it is interesting to analyze the behavior of traffic generated by these protocols in order to better understand the causes of the IP level behavior we observed. In Figure 5 we summarize the analysis for various IP payloads on the FDDI backbone. We graph the results for the TCP, UDP, and ICMP protocols. Additional protocols are grouped into the *Other* category. Here we only present the protocol comparisons for the net-gate scenario. The other scenarios show similar trends relative to their respective total IP characteristics.

From these graphs we observe that a knee appears at the 90 second hold point similar to the IP graphs. Also, we see that the total number of circuits required to carry the individual protocols is higher than that required when traffic is multiplexed over the IP channel. For instance, at the 90 second hold time, the sum of mean SVC cache sizes for all protocols is 331.59. This is 25% higher than the mean of 265.74 required for the combined IP traffic. This is because several traffic pairs have simultaneous conversations using different protocols. We also observe that the network load created by a protocol is not necessarily a good indication of the number of VC's required to support that traffic. Table 1 lists the total packets observed per protocol on each network over the sample time periods. On the backbone there is a large percentage of *Other* traffic, even more than for UDP. On the backbone, this traffic is mostly due to the use of tunneled IP multicast traffic

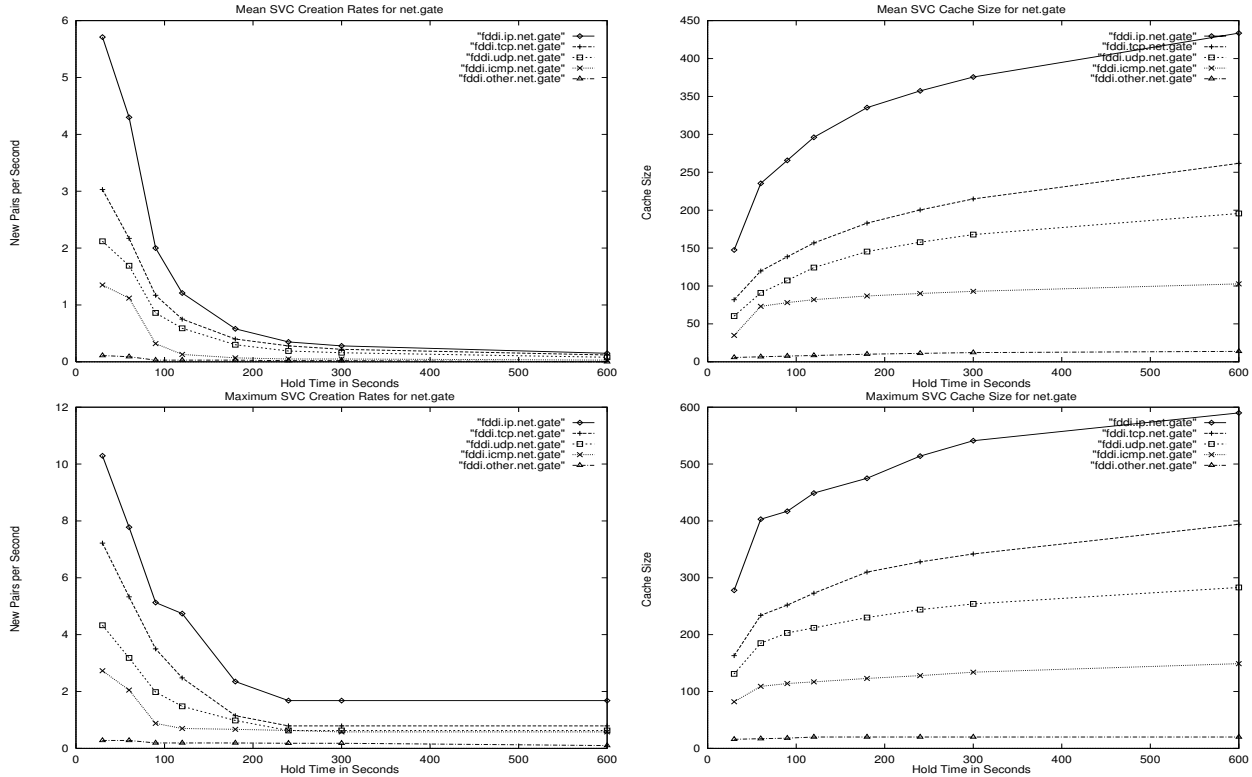


Figure 5: FDDI Backbone Pair Protocol Summary

between MBONE routers [3]. While this is significant traffic, it is primarily between only three hosts and requires only two VC's. On the other hand, the rather small ICMP packet counts account for a fairly large number of VC's. This is because of the existence of several management systems in the network that periodically *ping* or *traceroute* hundreds of hosts throughout the network. These short-lived conversations don't produce much traffic but require a great deal of circuit management overhead.

Protocol	Backbone	Dial-in	Department
TCP	119,886,579	69,273,673	14,904,436
UDP	20,616,747	652,607	7,021,376
ICMP	4,378,216	650,803	4,244,454
Other IP	54,363,150	17,241	16,625
Total IP	199,244,692	70,594,324	26,186,891

Table 1: Protocol Packet Counts.

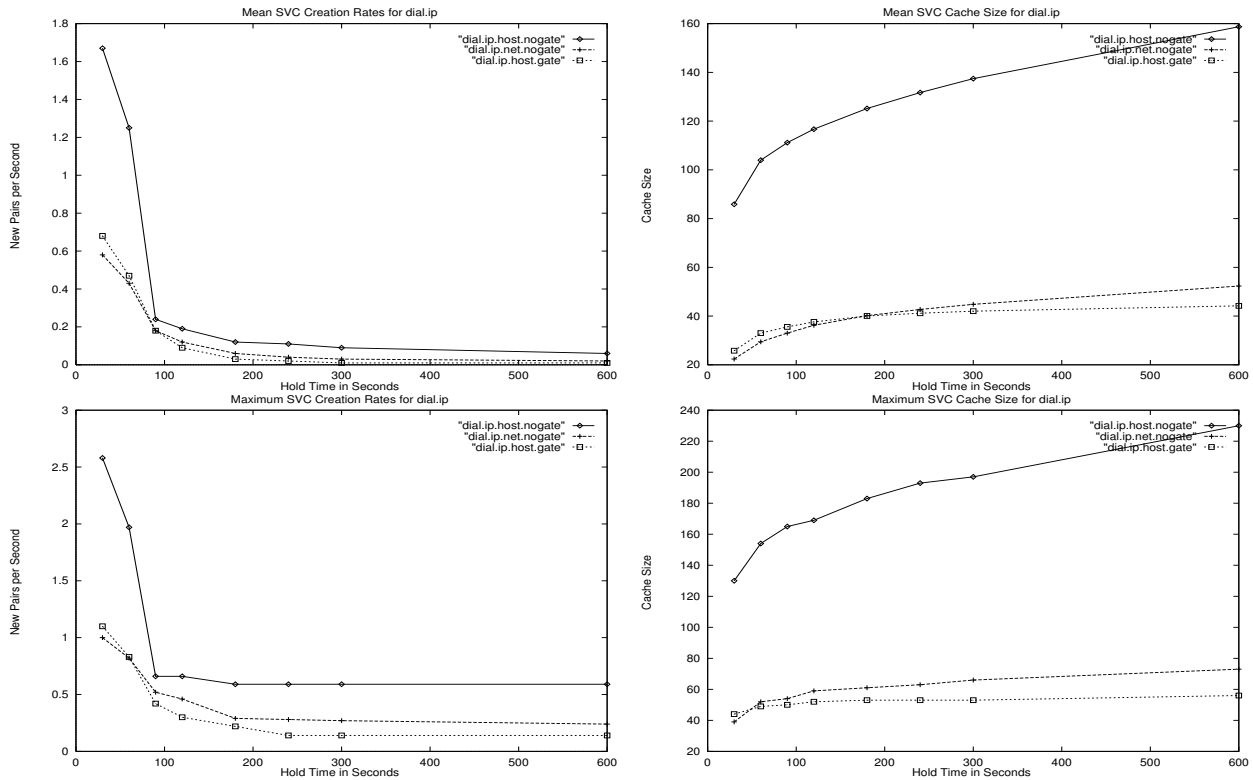


Figure 6: Dial-in Ethernet IP Summary

## 4.2 Dial-in Subnet Analysis

We analyzed the IP traffic on the Ethernet subnet serving the dial-in modems using logs from the 48 hour time period covering all of November 28-29, 1994. We summarize the results of this analysis in Figure 6. This traffic requires fairly low VC creation rates and cache sizes. In this network, the gateway scenarios involve placing a gateway at the current router interface connecting this subnet to the campus backbone. For this subnet, and the subsequent departmental subnet, we do not include the net-gate analysis. Since these subnets are connected to only one router interface they do not carry any cross-traffic between other networks. With all traffic on these subnets being either to or from either a host on this subnet or the router interface, the net-gate scenario would combine all traffic into a single SVC based on the common subnet address.

As in the backbone scenarios, there is a knee in the creation rate graphs for the dial-in subnet at around 90 seconds. If a switch holds connections beyond 90 seconds idle time, there will be little benefit on this subnet. Using the 90 second hold time, our simulation results indicate that for the host-gate scenario, a 0.18 mean SVC creation rate is required with a maximum rate of 0.42 circuits per second. This requires a mean cache of 35.66 circuits and a maximum of 50. In the worst case scenario of host-nogate, the 90 second

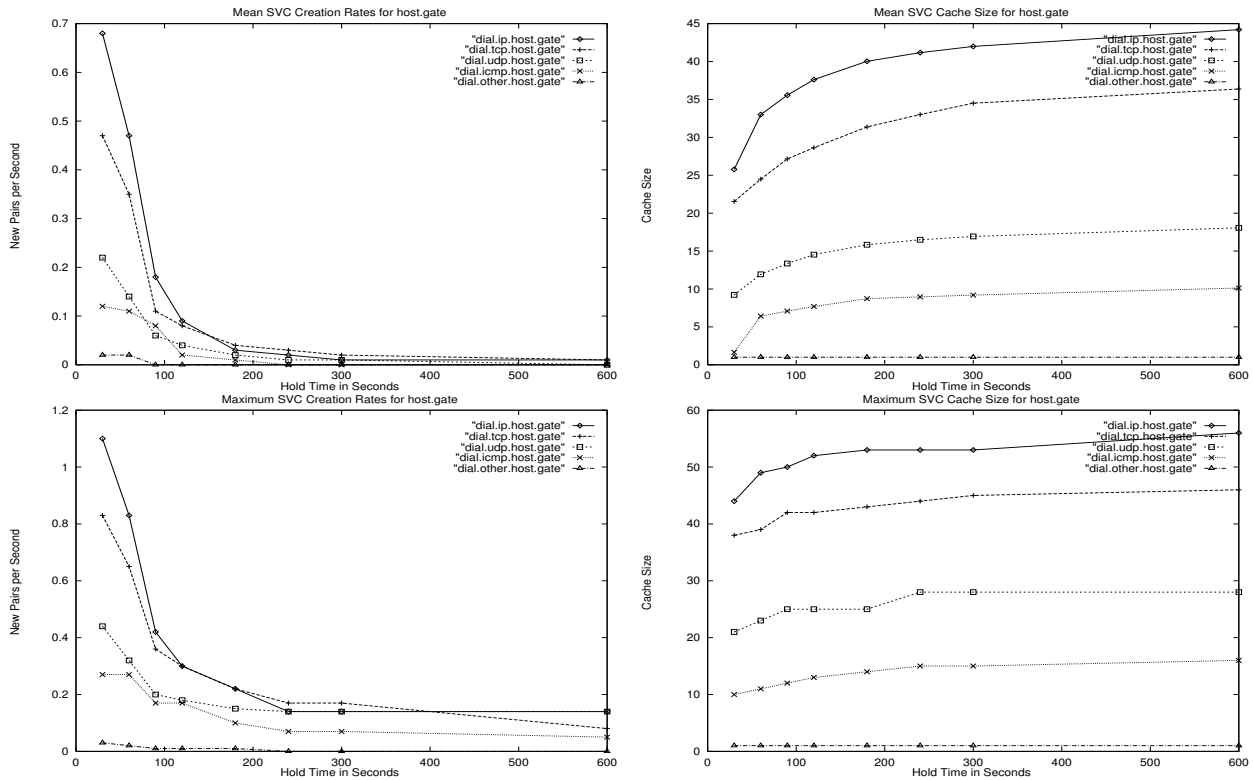


Figure 7: Dial-in Pair Protocol Summary

hold time requires a mean SVC creation rate of 0.24 with a maximum of 0.66 per second. This requires a mean cache of 111.17 circuits and a maximum of 165.

When we look at the protocol distributions for the dial-in subnet, (see Table 1) we observe that 98% of the traffic is TCP. This is mostly because of Telnet sessions from the dial-in terminal servers. It is interesting to note that while the packet count on this subnet is quite high, the mean packet is quite small (85 bytes, compared to 250 on the departmental subnet)<sup>3</sup>. In Figure 7 we present the per-protocol SVC analysis for this subnet. The mean SVC creation rate (upper, left graph) shows an interesting result. For the circuit holding times between 120 and 300 seconds, the SVC creation rates for TCP are actually a little higher than for the combined IP traffic. This occurs because the combined IP traffic includes other protocol packets that keep the circuits open while the TCP connections are idle. Without this additional traffic, a three minute idle timeout allows a few more circuits to drop out just before they are needed again.

<sup>3</sup>This is because of the large number of single character messages commonly generated by interactive Telnet sessions.

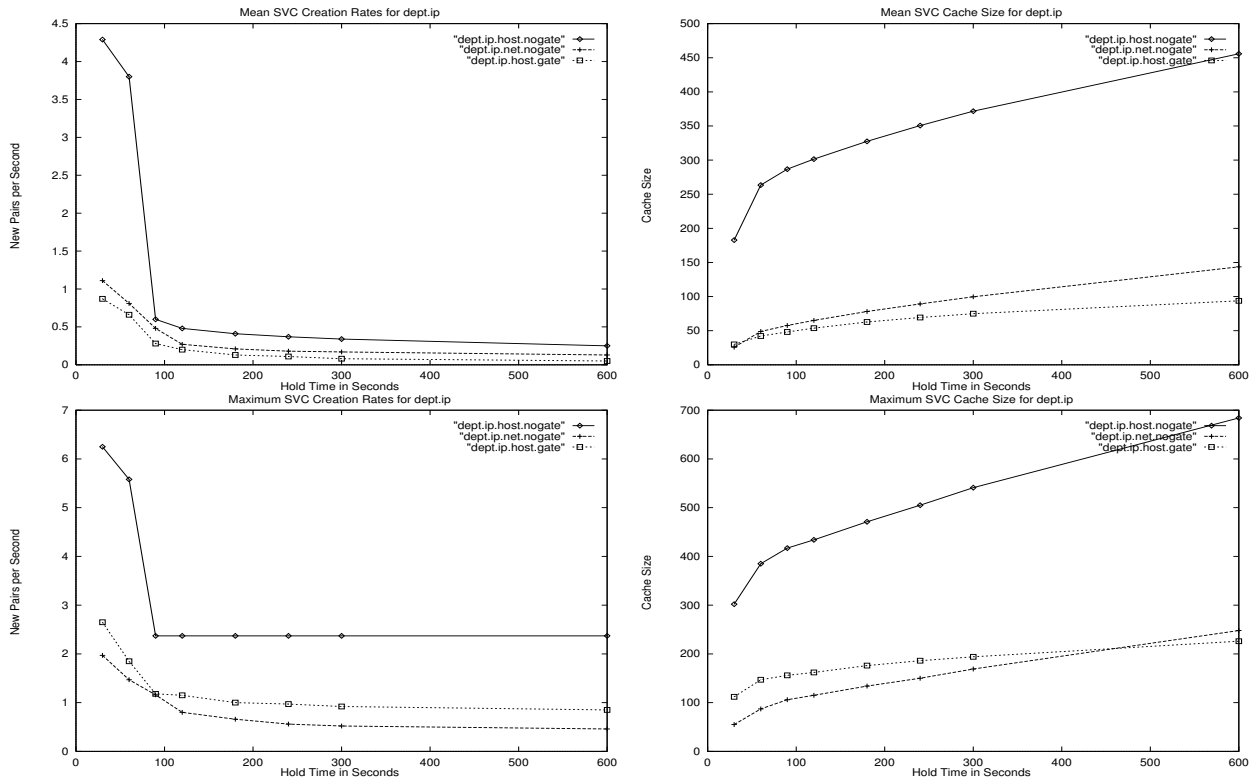


Figure 8: Departmental Ethernet IP Summary

### 4.3 Departmental Subnet Analysis

We analyzed the IP traffic on the departmental Ethernet for the 48 hour time period of December 7-8, 1994. As in the dial-in subnet, the gateway scenarios involve placing a gateway at the current router interface connecting this subnet to the campus backbone. The results from this analysis are presented in Figure 8. This traffic requires somewhat higher SVC creation rates and cache sizes than the dial-in subnet. Again, we do not include the net-gate analysis for this subnet.

As in the previous two networks, there is a knee in the creation rate graphs for the departmental subnet at around 90 seconds. Using the 90 second hold time, our simulation results indicate that for the host-gate scenario, a 0.28 mean SVC creation rate is required with a maximum rate of 1.18 circuits per second. This requires a mean cache of 48.13 circuits and a maximum of 156. In the worst case scenario of host-nogate, the 90 second hold time requires a mean SVC creation rate of 0.60 with a maximum of 2.37 per second. This requires a mean cache of 286.69 circuits and a maximum of 417.

The protocol distributions for the departmental subnet, (see Table 1) show a much higher percentage (27%) of traffic from UDP than the other two networks (10% backbone and

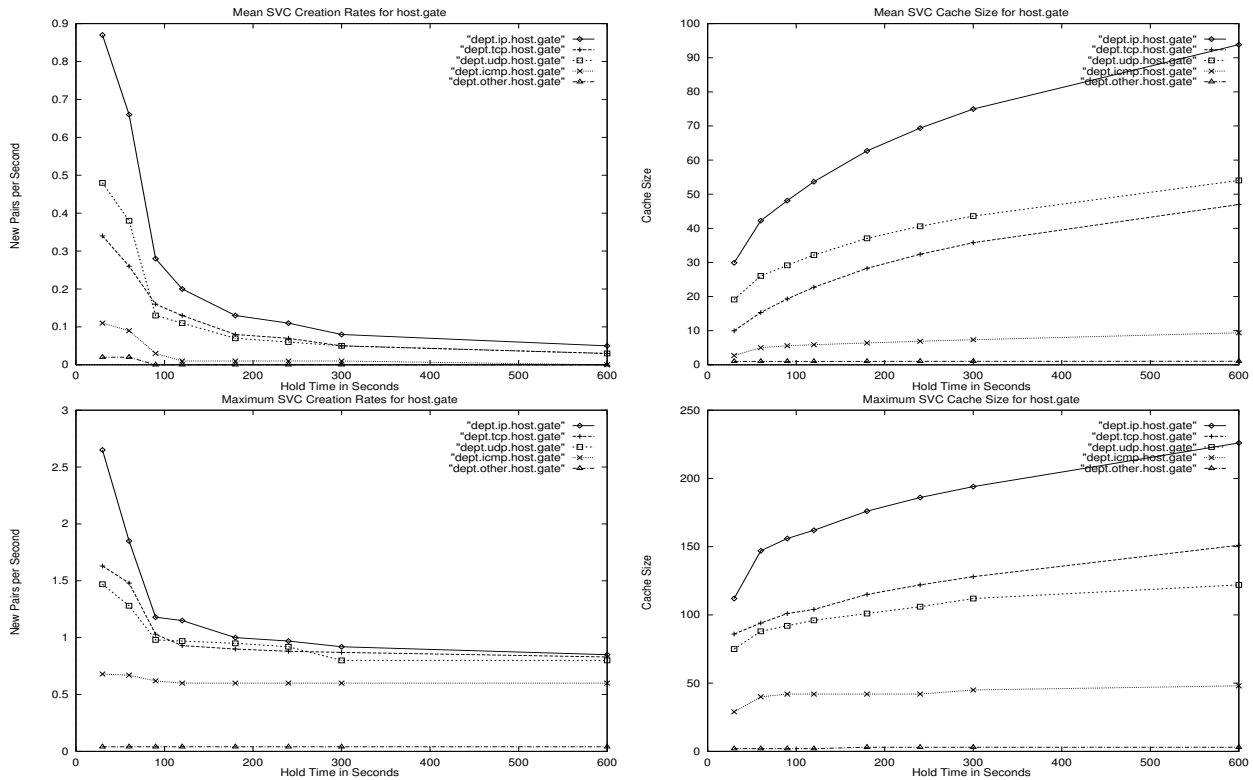


Figure 9: Departmental Pair Protocol Summary

0.9% dial-in). This is largely due to the heavy use of NFS file servers on this subnet. In Figure 7 we present the per-protocol SVC analysis for this subnet.

It is interesting to note that while the total number of packets carried on the departmental subnet in a 48 hour period is much less than that carried by the dial-in subnet, the demand placed on the switch for SVC creations is much higher on the departmental subnet.

#### 4.4 Traffic Distribution

The IP traffic on the FDDI backbone exhibits a highly skewed distribution where a small number of pairs account for a large portion of the traffic. From the November 22 traffic logs there were 12,170 distinct pairs in the host-gateway analysis generating a total of 81,483,088 packets. Of these pairs, the top 10 busiest pairs generated 35,037,727 (43%) of the total packets carried.

We evaluated the possibility of using PVCs to connect the most popular network traffic pairs. We created a simulation and ran it using PVCs for the top 10, 20 and 30 address pairs. We discovered that the effect of these PVCs on the required SVC creation rates was

negligible. This effect is so limited because the most popular traffic pairs are never idle long enough that their circuits are removed. While the use of PVCs in this manner may facilitate traffic management, there is little benefit to be gained in reduced demand on the SVC creation mechanisms.

## 4.5 Maximum Circuits

While analyzing the data for this research, we considered what the total circuit requirements would be if there was a permanent circuit established between every pair of hosts that communicate using GTNet. With approximately 14,000 hosts we can estimate that  $\frac{14,000(13,999)}{2} = 97,993,000$  circuits would be required to connect them all. However, circuits are really only needed between the hosts that will eventually communicate. To try and determine roughly how many circuits would be required, we ran our simulation using an infinite hold time. In this simulation, no circuit is ever removed once it is created. These simulations were run using eight consecutive days of IP traces from the FDDI backbone. The results of this simulation are presented in Figure 10. Since these simulations are only looking at backbone traffic, the results do not include counts for source-destination pairs within a subnet, that is, they do not include intra-network traffic. They also do not include traffic between subnets attached to the same router. Such traffic stays within the router and never reaches the FDDI ring. We estimate that the actual number of potential host-host traffic pairs on the backbone is no more than 81,000,000 based on the distribution of hosts around the network. Despite these caveats in the data we gathered, these results do give us an indication of the long term trends in traffic pairs on the total campus network.

The maximum value in each graph represents the total number of distinct pairs that used the backbone over the evaluation period. For the host-gateway analysis we see that 28,383 distinct pairs communicated over the backbone during the eight day period. This is less than .04% of the total possible pairs on the backbone. The net-gate analysis indicates that a total of 2,408 GTNet subnet pairs were heard from. This is out of an estimated 9,375 total possible subnet pairs under the current configuration, a utilization of 26%.

The host-nogate simulation reaches a maximum of 192,419 pairs. This represents the total number of distinct pairs in the entire Internet that utilized the GTNet FDDI backbone over this 8 day period. Similar to the host-nogate, the net-nogate analysis indicates the total number of Internet network pairs that utilized GTNet backbone. This number reached 91,975 after 8 days.

For all four graphs in this study we can identify eight regions of rapid increase in the cache size. These regions correspond to the busy periods of each day. We also observe that even after eight days there is some continued growth in the number of new pairs that communicate. The two gateway scenarios, host-gate and net-gate, appear to be leveling

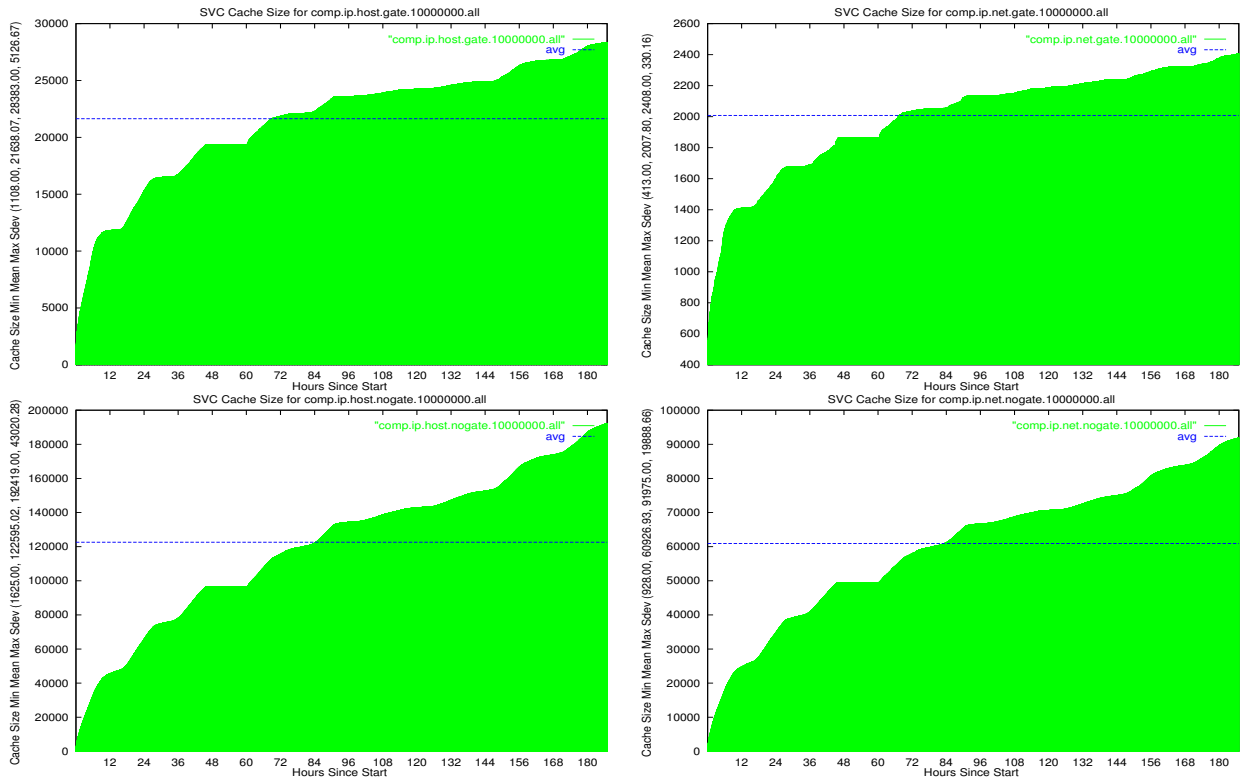


Figure 10: Infinite Hold Time Analysis

out after about four days while the nogate scenarios are still increasing dramatically after eight days. This would seem to indicate that most of the on-campus pairs that communicate, do so at least once every four days. The communication with systems outside the campus is, however, not nearly so stable.

## 5 Least Recently Used Caching

In the previous section we presented the results from a timeout based replacement strategy. In a LAN environment, where circuit hold-time costs may not apply, it may be more appropriate to consider a Least Recently Used (LRU) strategy that only closes circuits when there is no space to create a needed circuit. In order to compare the SVC creation requirements for an LRU algorithm to those of a timeout algorithm we created a simulation using LRU and analyzed the SVC requirements for several LRU cache sizes. For this analysis, we focused on the FDDI backbone traffic. The departmental subnets each require less than 1000 SVCs for all of the simulations presented in Figures 6 and 8. This means that an LRU algorithm with a cache size of at least 1000 SVCs would sufficiently support all active connections with minimal SVC creation requirements.



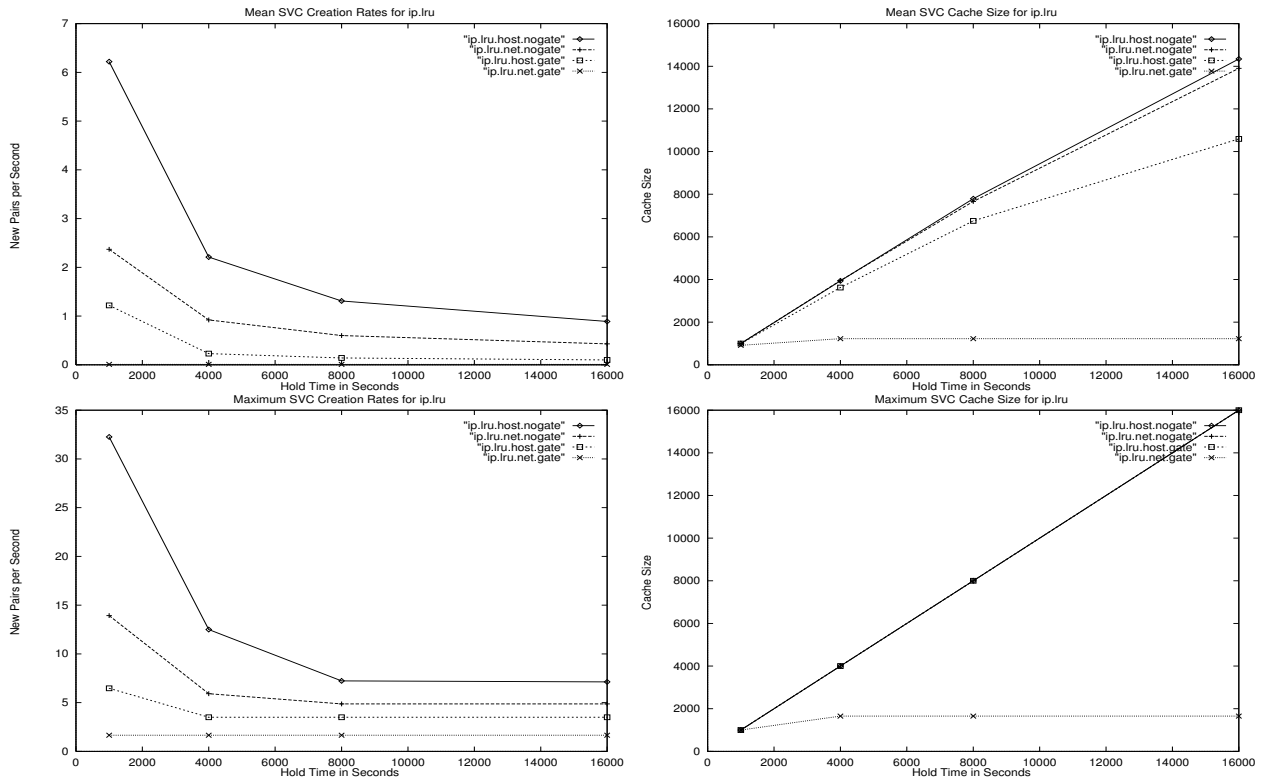


Figure 11: FDDI Backbone LRU Summary

The results of our LRU simulations for the backbone are presented in Figure 11. We simulated the four deployment scenarios using LRU cache sizes of 1000, 4000, 8000, and 16000 circuits. As might be expected, the SVC creation rates using LRU are generally lower than with a timeout based approach. For the net-gate scenario, the mean creation rate is 0.01 for all cache sizes. This is because the total number of circuits used by this scenario in the 48 hour period is only 1655. Even the smallest cache size of 1000 circuits is sufficient to support this scenario with minimal creation rates. The other scenarios require more total circuits and therefore require higher creation rates when the total number of concurrent SVCs is limited.

For all but the net-gate scenario, the maximum cache size used in the LRU simulation is the maximum available. The mean cache size is also close to this maximum for all but the 16000 maximum size runs. We also see that for all but net-gate the 1000 circuit size is too small to effectively accommodate the demand. This is seen in the large creation rates for 1000 and is confirmed by the mean cache size values presented in Figure 4. The “working set” of circuits appears to be larger than 1000 for these scenarios.

With a cache size of 16000, all of the scenarios have a mean circuit creation of less than 1 and a max of less than 10. It appears that a cache size of 4000 may be sufficient to satisfy

the current demands since this gives all scenarios a mean creation rate of less than 3 and a max less than 15.

## 6 Conclusions

As the networking community considers the use of ATM in data networking, it is important that we determine the requirements for ATM switches in the local and campus network environment. In particular, it is important to establish the specifications necessary to support the connection-less, datagram traffic typically found in data networking today. In making these determinations it is important to consider the particular deployment scenarios that will be used for ATM and make sure the equipment selected meets the requirements of the scenario in which it will be used.

Previous research in the determination of virtual circuit requirements for data networking has not considered these different configuration options. Saran and Keshav performed a cost-based analysis of circuit holding time policies for IP traffic in ATM networks and found that Least Recently Used (LRU) performs well [7]. They present evidence that LRU can be effective in the presence of either a paging cost model (no hold time cost) or a holding cost model. Schmidt and Campbell analyzed TCP and UDP traffic separately over LAN and WAN segments to determine connection setup rates [8]. Both of these previous studies have focused on measuring host-to-host connections in the network and do not consider the different deployment scenarios we evaluated.

In this paper we have identified several requirements for ATM switches that will be deployed in current LAN environments. From our analysis the current most demanding network scenario on GTNet is the host-nogate configuration of the FDDI backbone. This scenario requires that a switching network using a timeout based SVC algorithm with a 90 second hold time should be able to sustain about 5 SVC creations per second with bursts up to 20 per second. This scenario requires a mean of around 1000 concurrent circuits with a maximum of just over 2000 circuits. If fewer concurrent circuits are available then the switch will need to support slightly higher SVC creation rates. In the net-gate backbone scenario or host-gate subnet scenarios, SVC creation rates of 2 per second sustained and 5 per second peak should be sufficient. This shows the marked difference in requirements when the deployment scenario is considered.

We discovered that an idle timeout period of 90 seconds is probably the most appropriate for the IP traffic on each of the three distinct networks we studied. This consistency among the networks is somewhat surprising, especially when one considers the varied protocol mix on the three subnets. We also observed that packet counts are by no means a reliable predictor of SVC requirements in a network. It is important to consider the communication patterns when determining SVC needs. The actual observed pairs on

GTNet are only a small fraction of the total possible number of pairs. This observation reinforces the assumption that SVCs should be used in the LAN environment rather than establishing PVCs between all possible network pairs.

From our analysis of a LRU replacement strategy we determined that a cache size of 4000 circuits is sufficient for our current traffic characteristics. Most of the currently available ATM LAN equipment we have studied will cache at least 4000 VCs with circuit establishment rates between 10 and 100 per second. According to our analysis these parameters should be sufficient for our current network requirements. On the Fore ASX 100 switch, SVCs are removed at the same time the ARP entry is removed for a destination. This occurs after 15 minutes of idle time [1]. We have observed that a hold time of just 90 seconds would probably be sufficient for this purpose.

## References

- [1] E. Biagioni, E. Cooper, and R. Sansom. Designing a practical ATM LAN. *IEEE Network*, pages 32–39, March 1993.
- [2] B. Braden and A. DeSchon. NSFnet statistics collection system (NNStat). *Release 3.2*, December 1992.
- [3] S. Casner. Frequently asked questions (FAQ) on the multicast backbone. *file://-venera.isi.edu/mbone/faq.txt*, May 1993.
- [4] D. Ghosal, H.J. Chao, D. Saha, and S.K. Tripathi. IP on ATM local area networks. *IEEE Communications*, pages 52–59, August 1994.
- [5] M. Laubach. Classical IP and ARP over ATM. Request for Comments (Experimental) RFC 1577, Internet Engineering Task Force, jan 1994.
- [6] P. Newman. ATM local area networks. *IEEE Communications*, 32(3), March 1994.
- [7] H. Saran and S. Keshav. An empirical evaluation of virtual circuit holding times in IP-over-ATM networks. In *Proceedings of IEEE INFOCOM '94*, pages 1132–1140, June 1994.
- [8] A. Schmidt and R. Campbell. Internet protocol traffic analysis with applications for ATM switch design. *Computer Communication Review*, 23(2), April 1993.