NASA-CR-199753

NAS2- 13721
IN-62-CR
63888
P-13

RIACS

Research Institute for Advanced Computer Science
NASA Ames Research Center

Experiences with the Bay Area Gigabit Network Testbed

Marjory J. Johnson

(NASA-CR-199753) EXPERIENCES WITH THE BAY AREA GIGABIT NETWORK TESTBED (Research Inst. for Advanced Computer Science) 13 p

N96-15351

Unclas

G3/62 0081591

RIACS Technical Report 95.21

October 1995

Appears in Proceedings of Fifth IEEE Workshop on Future Trends of Distributed Computing Systems, August, 1995

Experiences with the Bay Area Gigabit Network Testbed

Marjory J. Johnson

The Research Institute for Advanced Computer Science is operated by Universities Space Research Association, The American City Building, Suite 212, Columbia, MD 21044, (410)730-2656

Work reported herein was supported by the National Aeronautics and Space Administration under Contract NAS 2-13721 to the Universities Space Research Association (USRA).

Experiences with the Bay Area Gigabit Network Testbed

Marjory J. Johnson
RIACS
NASA Ames Research Center
Moffett Field, CA 94035-1000

Abstract. The Bay Area Gigabit Network Testbed (BAGNet) is a high-performance ATM (155 Mbps) testbed located within the San Francisco Bay Area in northern California. BAGNet is a metropolitan-area network, spanning an area of approximately 50 square miles. There are fifteen sites participating in the testbed, with up to four hosts per site. Although BAGNet is an applications-oriented testbed, much of our effort has been directed towards getting the testbed running and understanding the factors that impact performance of an ATM network. We present some of our experiences in this paper.

1 Introduction

Fifteen organizations within the San Francisco Bay Area in northern California are participating in the Bay Area Gigabit Network Testbed (BAGNet). These organizations include four government laboratories, two major research universities, two research institutes, and almost all the major computer-technology companies in the Bay Area. The goal of the testbed, which was initiated about a year and a half ago, is to develop the infrastructure to support multimedia applications in a large-scale metropolitan-area ATM network environment.

BAGNet is sponsored by Pacific Bell's California Research and Education Network (CalREN) program. Under this program Pacific Bell is providing the basic infrastructure for several ATM testbeds within the state of California, to promote the development of high-performance communication applications. One set of testbeds is located in northern California in the San Francisco area and another set is located in southern California in the Los Angeles area.

For testbeds located in northern California, including BAGNet, the infrastructure provided by Pacific Bell consists of two pairs of ATM switches (one pair at the Palo Alto Central Office in the South Bay and one pair at the Oakland Central Office in the East Bay) connected by a double mesh of SONET OC-3c (155 Mbps) trunks. Since BAGNet actually uses only one of the Oakland switches, there are 4 OC-3c trunks connecting the South Bay and East Bay sections of the testbed. Each individual testbed site is directly connected to one of the Pacific Bell switches via a SONET OC-3c (155 Mbps) link.

BAGNet is a metropolitan-area network; the distance between individual sites ranges from approximately two miles to approximately fifty miles. Each site has at most four hosts directly connected to the testbed. We have established a mesh of permanent virtual channels (PVCs) to provide connections between all possible pairs of hosts. BAGNet is an

IP over ATM network, with AAL5 as the adaptation layer and LLC/SNAP encapsulation. A wide variety of ATM switches, workstations, ATM host adapters, operating systems, and applications are represented in the testbed. Because of this diversity, BAGNet provides a realistic environment that is representative of future operational environments.

In this paper we present some of our experiences with BAGNet. In Section 2 we discuss problems that we have encountered as we set up the testbed and ways that we are resolving these problems. In Section 3 we describe selected testbed applications, and in Section 4 we present experimental performance results. In Section 5 we discuss goals for future testbed experimentation.

2 Setting up the testbed

The fifteen BAGNet sites are NASA Ames Research Center, Apple Computer, Digital Equipment Corporation - Systems Research Center (DEC SRC), Hewlett-Packard Laboratories (HP), International Computer Science Institute (ICSI), Lawrence Berkeley Laboratory (LBL), Lawrence Livermore National Laboratory (LLNL), Pacific Bell, Sandia National Laboratories - Livermore, Silicon Graphics, Inc. (SGI), SRI International, Stanford University, Sun Microsystems Computer Corporation, University of California - Berkeley, and Xerox Palo Alto Research Center (Xerox PARC). Figure 1 illustrates the overall testbed configuration. Figure 2 illustrates how workstations at NASA Ames are connected to BAGNet. The four hosts at NASA are located in three different buildings, one at RIACS, two at the Numerical Aerodynamic Simulation (NAS) supercomputer facility (along with the Fore Systems switch that connects to the Pacific Bell BAGNet facilities), and one at the Central Computing Facility. The testbed workstation at RIACS is named *psyche-atm.riacs.edu.*

The testbed has been functional for about one year, with sites being connected a few at a time. Establishing such a large testbed, while ATM technology is changing so rapidly, has been a challenge.

The heterogeneity of hardware and software that makes up the testbed has contributed significantly to the challenge. Different sites are using different methods of interfacing to the Pacific Bell switches. Most sites connect to BAGNet via local ATM switches, though one site has a workstation directly connected to BAGNet. Several types of ATM switches are represented in the testbed, including Fore Systems, DEC AN2, Newbridge, Synoptics, and Badlan (a non-commercial switch developed by Xerox PARC). Types of host workstations on BAGNet include DEC Alpha, Macintosh, SGI Challenge, SGI Indigo, SGI Onyx, and Sun SparcStations (2, 10, and 20). ATM adapters used by the various workstations include Fore Systems, Efficient, Synoptics, DEC OTTO, and Sun SAHI. Different operating systems include SunOS, Solaris, OSF/1, IRIX, and MacOS.

2.1 Setting up PVCs

Establishing and maintaining testbed connectivity has been a major problem. We are using Permanent Virtual Channels (PVCs) for transmission, since the standard for Switched Virtual Channels is not yet complete. Since it is important that every host be able to talk to every other host, we created a mesh of PVCs. When we first started the testbed, Virtual Paths were not available either, so all routing had to be specified in terms of Virtual Channels. Hence, a large number of PVCs had to be specified in tables in the ATM switches and in host routing tables.

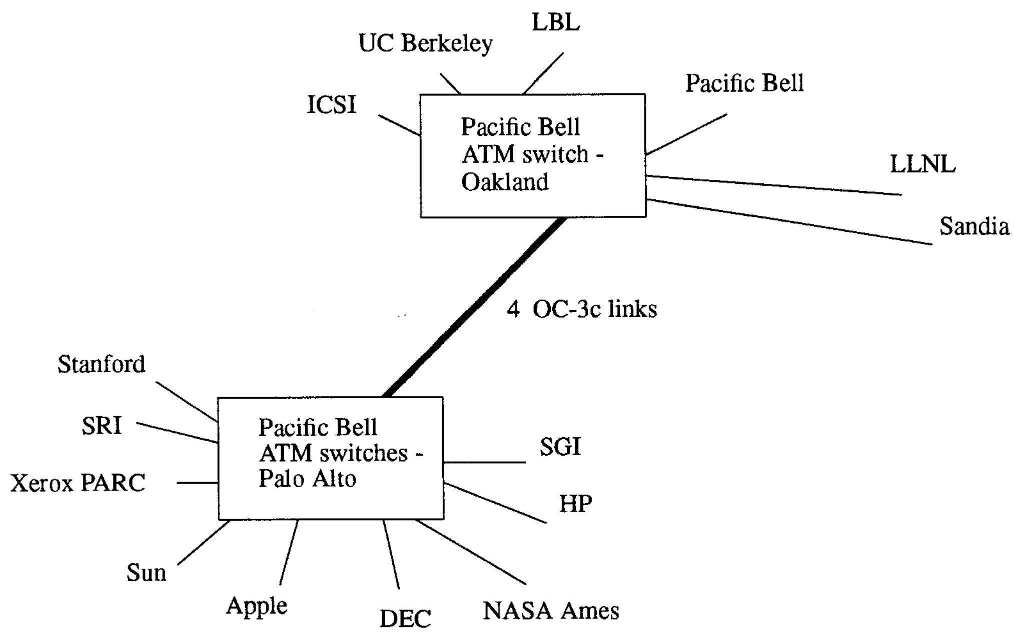

Figure 1. BAGNet Testbed Configuration

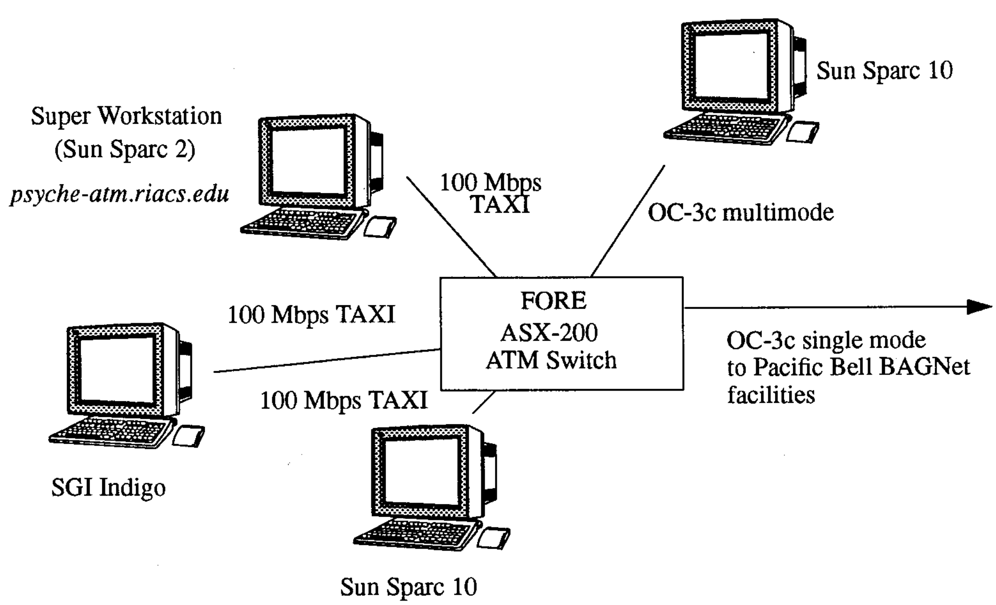

Figure 2. NASA Ames Internal Configuration

Due to limitations in the size of the tables in the Pacific Bell switches, the number of BAGNet hosts per site is limited to 4. If a site desires to connect additional hosts to the testbed, they must be connected via a router. With 15 sites and 4 hosts per site, this gives a total of 60 hosts on BAGNet and 60 X 60 permanent connections to establish. Scripts were developed to set up routing tables in the ATM switches and to map IP addresses to VPI/VCI pairs in the host interfaces. Although these scripts, which were made available via anonymous ftp, greatly facilitated the task of setting up the PVCs, some of the work still had to be done manually.

The debugging effort associated with establishing and maintaining such a large number of permanent connections is tedious. When a connection is not working, it is very difficult to determine the location of the problem. To try to make the debugging chore manageable, we are using a connectivity matrix that is updated once an hour. This matrix shows when two-way communication has been properly established between a pair of hosts. Some erroneous situations that can also be quickly identified from the connectivity matrix are:

- host A sees host B, but not vice versa
- host A doesn't see host B, but several other hosts do

In one obviously ridiculous case, host A claimed to see several hosts that had never even been connected to BAGNet. A quick glance at the connectivity matrix made it clear that host A's tables were improperly configured.

Our initial task was to set up unicast PVCs. However, many of our testbed applications require multicast. We are implementing IP multicast over ATM multicast. Since there is not yet a standard for ATM multicast, we are simulating it by establishing point-to-multipoint (i.e., multicast) PVCs. Each host has one outgoing multicast PVC, connecting this host to all other testbed hosts. All multicast transmission from this host must be mapped to this PVC. Each host must also be configured to receive a multicast PVC from every other BAGNet host. For unicast, a single IP address is mapped to a PVC; for multicast, the multicast IP network (224.0.0.0) must be mapped to a single PVC. Some of the host adapters have had problems supporting this mapping, so progress in setting up the multicast PVCs has been slow. About two thirds of the sites now have at least one host that can receive multicast.

As we've worked to establish both unicast and multicast connectivity for the testbed, we've encountered various problems with different types of equipment. For example:

- the Efficient adapter supports only 16 PVCs, so the workstations using this adapter cannot be fully connected to all hosts on the testbed
- the Synoptics switch cannot use VCI 32, which corresponds to the IP address of one of the BAGNet hosts
- earlier versions of software for the Fore adapter didn't encapsulate multicast PVCs in LLC/SNAP, causing receiving hosts to crash

2.2 Testbed coordination

With a testbed of this size, coordination is vital. We are using the World Wide Web (WWW) to share information among the various testbed sites, to keep track of testbed status, and to post information to the general public. The URL for the BAGNet home page is http://george.lbl.gov/BAGNet.html [1]. This home page contains an overview of the testbed and a pointer to a more lengthy description of the applications being developed and implemented over BAGNet. In addition most BAGNet sites maintain individual BAGNet pages on the WWW, containing information relevant to that individual site. This information typically includes ping pages for the hosts at that site. The URL for the ping page for psyche-atm.riacs.edu, the RIACS workstation connected to BAGNet, is ftp://riacs.edu/pub/Gigabit_Testbed/psyche-ping.html.

A program called *bagping* [7] is used by many testbed sites to produce ping pages automatically. Pages produced by this program include the following information:

- list of hosts that respond to ping, along with an average ping time
- for each reachable host and for each of several different message sizes, the percentage of ping attempts that are successful
- list of all hosts that have ever been seen from this host, along with the date last seen
- list of hosts responding to multicast pings

Output from *bagping* is automatically sent to the testbed site that maintains the connectivity matrix [5]. A script updates the matrix automatically, once an hour. Other information available on the WWW includes BAGNet status [7], hardware configuration hints [7], and performance data [2, 6].

Several sites have connected fewer than 4 hosts to BAGNet; the total number of hosts is currently around 40. As problems with the PVC tables are resolved, the number of hosts visible to psyche-atm.riacs.edu at one time or another has reached 32. Up to 31 hosts have responded to a single round of ping attempts. The nature of a testbed is to bring hosts up and down, as new software is installed or as a workstation is temporarily disconnected so that it can be used for a different function. Hence, the number of hosts visible at any given time can't be expected to remain constant.

3 Testbed applications

Multiple applications are being implemented on the testbed. The featured application is a teleseminar application, wherein a seminar presented at one testbed site is multicast to all hosts on BAGNet. In addition, subsets of the testbed participants are experimenting with several other multimedia applications. While some of the applications can be run over the Internet in a limited manner, utilization of the higher performance capabilities of the ATM testbed will significantly enhance them.

3.1 Teleseminar application

An essential element of our teleseminar application, of course, is multicast of seminars over BAGNet. However, our concept of a teleseminar is much broader than just electronic distribution of seminars. Our ultimate goal is to provide a rich interactive environment that enables members of a remote audience to participate in a seminar in real time. The intent is to enable teleseminar participants to interact with the speaker (e.g., during a question and answer period), as well as to provide support for side conversations between members of the audience. Our plans also include provision of multimedia networked storage of these teleseminars so that they can be viewed at a later time. Progress toward these goals has been slow, due to the difficulty in setting up reliable connections.

We are currently concentrating our efforts on getting multicast working at all the BAGNet sites. Experimentation with multicast began early this year. We are using M-bone tools, either *nv* or *vic* for video and *vat* for audio. Multicast test sessions over BAGNet have included video of traffic from a camera pointed towards a busy highway, transmission of NASA Select during a shuttle mission, and multicast of selected courses and seminars. About half the testbed sites have participated in these sessions.

Two relatively permanent multicast sessions are maintained on BAGNet for testing purposes. As site switches are configured to support multicast, these sessions can be used to test the setup. One test session uses *nv* video format, while the other session uses higher-rate JPEG video format. The common video format that we have selected for our teleseminar application is full motion JPEG (30 frames per second), with 320 x 240 pixel resolution and 4:2:2 compression ratio. We will use *vat*-style audio encoding at 64 kbps. Both audio and video will be transmitted in accordance with version 2 of the RTP Internet transport standard.

Once all sites are able to receive multicast transmissions, we plan to multicast 100 to 200 hours of seminars weekly.

3.2 Other applications

At NASA Ames Research Center, in addition to the teleseminar application, we are using BAGNet to experiment with the development of collaborative working environments.

NASA scientists in many disciplines routinely analyze large amounts of image data. At RIACS we are creating a workstation environment that will enable scientists at different geographical locations to analyze data jointly, while conducting a video-conferencing session at their workstations. Elements of this environment include a video-conferencing tool developed at Sandia [4], a tool for sharing X-window applications, and a software package for data analysis. This application is not expected to be bandwidth intensive, except for the transfer of large image files over the network during interactive work sessions. Low and consistent delay are necessary to support the video-conferencing and to enable the interactive work. We plan to collect traffic statistics during a typical work session, develop a traffic model of collaborative applications, and evaluate the scientific benefits of this type of environment. Sandia National Laboratories - Livermore is our testbed partner in this project. We expect ATM to enable this type of application by providing low and predictable delay and by facilitating interactive transfer of large files.

A novel application being explored by NASA is to use networks to provide remote access to NASA wind tunnels during testing of new aircraft designs. Currently, a team of engineers and scientists must spend several days at the wind tunnel site to conduct a test. Cameras and sensors are set up inside the wind tunnel, and there are monitors in a control room which allow on-site personnel to monitor the test. Network technology can now enable engineers and scientists at their home site to monitor wind-tunnel tests by viewing results in real time on their own desktop workstations. In this case, the video of the model inside the wind tunnel must be very high quality, to show the details of air flow during the test. Functionality tests have been conducted using AEROnet, NASA's proprietary network connecting researchers to the NAS supercomputer facility. Participants are enthusiastic about the new working paradigm, but higher bandwidth is needed to support the application. We plan to test the application over BAGNet between NASA and SGI, whose InPerson tool is being used to provide video conferencing for the application.

Another project at NASA Ames is the development of a video-on-demand client/server application with LLNL and Apple Computer. Other applications involving other testbed participants include the development of distributed digital multimedia libraries, distributed medical imaging, and remote instrumentation.

4 Performance

Tools that have been used to analyze performance of BAGNet include the simple ping (using messages of various sizes) and the *ttcp* and *netperf* [3] benchmark tools. We have also made some interesting observations while doing file transfer with ftp. Ping results have been useful in determining basic connectivity and host-to-host response time, while *ttcp* and *netperf* tests have been directed towards identifying and resolving throughput degradation problems.

We decided during initial planning for the testbed that we would use "best effort" transmission, i.e., we would not limit the amount of bandwidth used by individual PVCs. Otherwise, with so many permanent connections to provide, each PVC would have been allocated less than 1 Mbps bandwidth and BAGNet would not have been a high-performance testbed at all. We knew that we would eventually have to address congestion problems, but we decided to deal with the issue when it arose.

4.1 Throughput

The *bagping* program described in Section 2 uses a variety of ping message sizes, ranging from the default 56 bytes to 5K bytes. The percentage of replies is tallied for each of these message sizes, using a total of 20 trials and 100 ms as a timeout value. When the message size was greater than 3K, there was a dramatic drop in the percentage of successful pings. With the largest message sizes, this percentage dropped to almost 0. Throughput statistics collected by testing with *ttcp* and *netperf* revealed the same phenomena, i.e., throughput dropped dramatically when the message size was greater than 3K.

Some testing was conducted between LLNL and DEC to figure out what was happening. With the aid of personnel from Pacific Bell, it was determined that bandwidth policing on the Pacific Bell switches was causing the performance degradation. Bandwidth policing, a policy instituted several months after the testbed was established, limits

Packet Size (bytes)	Packets		Throughput (Mbps)	
	sent	rcvd	srce	dest
4096	20520	14715	67.24	48.22
6144	17202	13	84.54	0.06
9140	7698	0	56.27	0.00

Table 1: Throughput with bandwidth policing enabled (50 cells per PVC)

Packet Size (bytes)	Packets		Throughput (Mbps)	
	sent	rcvd	srce	dest
4096	19948	19583	65.37	64.17
6144	17461	17072	85.82	83.90
9140	7814	7720	57.13	56.44

Table 2: Throughput without bandwidth policing

the bandwidth used by a single PVC by restricting buffer space. In this case the buffers were limited to 50 cells per PVC. Clearly, large messages were simply overflowing the buffer, causing the throughput to drop to zero. The problem essentially disappeared when bandwidth policing was disabled, and the PVCs simply shared the buffer pool for the associated port. Tables 1 and 2 (taken from [6]) show the dramatic improvement in throughput between two DEC Alpha workstations, one at LLNL and one at DEC, when bandwidth policing was disabled. Because of our experiences with BAGNet, Pacific Bell has agreed to disable bandwidth policing on all the PVCs.

For a detailed explanation of the testing that was conducted, including several graphs that illustrate the detrimental impact of bandwidth policing, see [2].

At NASA Ames we have experienced similar throughput problems while using ftp over BAGNet. Last summer, before Pacific Bell had implemented the bandwidth policing and while traffic on BAGNet was still very light, we conducted some ftp tests between psyche-atm.riacs.edu at RIACS and a BAGNet host at Sandia (a Sun Sparc 10). Files selected for the tests were typical binary image-data files that an Earth scientist might want to transfer across the network as part of his daily scientific activities. One file was 350 megabits; the other was 800 megabits. The average *get* was 10 Mbps and the average *put* was 13 Mbps. While these numbers don't seem very impressive, loopback testing on psyche-atm.riacs.edu (a Sun Sparc 2 clone) shows that its maximum achievable throughput is about 30 Mbps. Based on this information, these ftp results seem reasonable.

However, about six months later, we had some entirely different experiences with ftp. While trying to ftp a file from Sandia to RIACS over BAGNet, the transfer rates to *get* a 2.5 Megabyte file ranged from 7 to 39 kbytes/sec. The average, taken over several attempts, was 24 kbytes/sec. This was much too slow to support our videoconferencing application. At the same time, ftp over the Internet was considerably faster. Upon examination of the traffic at the NASA Ames switch that interfaces to BAGNet, we saw that there was considerable multicast traffic being sent to psyche-atm.riacs.edu. When we shut off the multicast traffic at the switch, the file transfer rates increased by an order of magnitude, but they were still significantly slower than the earlier rates. In the near future we will investigate the cause of the degraded file transfer rates. Bandwidth policing is likely to have contributed to the problem.

4.2 Delay

Ping pages posted on the WWW give a general indication of typical round-trip delay over BAGNet. As expected, ping times are much shorter over BAGNet than over the Internet. For example, the time to ping a host at Xerox PARC from psyche-atm.riacs.edu over BAGNet is consistently around 2.5 ms, while the time to ping the same host at Xerox PARC over the Internet varies widely from approximately 8 ms to 45 ms. The average of 20 pings over the Internet was approximately 20 ms.

The *bagping* program has been running on psyche-atm.riacs.edu for several months. Round-trip ping times to the other BAGNet hosts have been fairly consistent. When congestion is present, it is immediately obvious from the longer-than-normal recorded times. Ping pages from hosts at other sites, generated within the same time frame, generally confirm that when congestion occurs, it affects all the sites equally.

Many factors contribute to delay, including type of site interface to BAGNet, type of host, type of host ATM adapter, type of operating system, and the number of ATM switches in the transmission path. The influence of several of these factors can be readily observed from the ping pages. As expected, it takes longer for psyche-atm.riacs.edu to ping hosts in the East Bay than the South Bay, because of the extra ATM switch in the transmission path. The time required to ping other hosts in the South Bay ranges from 1 to 2 ms, while the time to ping hosts in the East Bay ranges from 3 to 4 ms. When there is congestion, delay to East Bay hosts is affected more than the delay to South Bay hosts. This is because of the back-to-back Pacific Bell switches separating the South and East Bay hosts. In addition, the trunks connecting the South Bay and the East Bay hosts are shared by the other CalREN testbeds; we have no way of determining how much congestion BAGNet traffic might be encountering from these other testbeds.

One unexpected observation from the ping pages has been that it often takes longer to ping a switch itself than to ping hosts connected to that switch. For example, SRI has four IP addresses on BAGNet; one of these is their local Fore switch. Traffic addressed to the switch sometimes apparently backs up, while, at the same time, traffic addressed to other workstation hosts connected to the switch is unaffected.

Two excerpts taken from psyche-atm.riacs.edu ping pages illustrate the situation. The first shows typical ping results from the four BAGNet hosts at SRI; all response times are approximately 2 to 3 ms.

May 3, 1995, 9:14 a.m.

bagnet0.bagnet.sri.com	2.93 ms
bagnet1.bagnet.sri.com	2.36 ms
bagnet2.bagnet.sri.com	2.29 ms
bagnet3.bagnet.sri.com	1.90 ms

The second excerpt, taken from the next series of pings an hour later, shows that considerable congestion had developed at the first host.

May 3, 1995, 10:37 a.m.

bagnet0.bagnet.sri.com	48.70 ms
bagnet1.bagnet.sri.com	2.24 ms
bagnet2.bagnet.sri.com	2.23 ms
bagnet3.bagnet.sri.com	1.89 ms

This situation persisted throughout the next day and a half; other hosts' ping pages revealed the same situation. The host bagnet0.bagnet.sri.com is actually the local switch at SRI. SRI said that occasionally they would notice traffic to the switch backing up, while traffic directed through it remained unaffected. Rebooting the switch corrected the problem.

Low and predictable round-trip delay is important for support of collaborative work. As we mentioned in Section 3, at RIACS we are experimenting with a collaborative work environment that includes a tool for sharing X-window applications. We have experimented with this tool over BAGNet between RIACS and Sandia and over the Internet between RIACS and the University of Arizona. The volume of network traffic generated by this tool is minimal, because only commands are transmitted over the network. Response was instantaneous over BAGNet, while the tool behaved somewhat sporadically when we used it over the Internet. We were never quite sure how to interpret lengthy delays over the Internet; we couldn't tell whether the delay was due to the network or whether the tool itself wasn't working properly.

5 Future testbed experimentation

BAGNet was originally planned to be a two-year testbed. Over half of that time is now gone, and much of it has been spent setting up the testbed infrastructure. We are hoping to be able to extend the time frame for the testbed, so that we will have adequate time to accomplish our application goals.

Our primary short-term goal for the testbed is to get multicast functioning at all the sites. Then we will be able to proceed with our plans for the teleseminar application.

In addition, we have just initiated a collaboration with Bellcore to capture data for traffic analysis, so that we can see what a data stream for a real ATM application looks like. Bellcore is interested in general traffic analysis to help them understand how to manage data transmission for commercial-service offerings; several of the individual sites plan to use the collected data to model specific applications. We are also interested in observing and analyzing interactions between applications that are sharing the testbed.

Finally, we will continue to install software upgrades for ATM switches and host adapters, as they become available; such upgrades should add functionality and improve performance of the testbed. It is unlikely that Switched Virtual Circuits will be available within the testbed time frame, which is a disappointment for us. We have recently been notified that we can now use Virtual Paths, but we will probably not pursue this option, since setting up and debugging another set of communication paths would be a diversion from our primary application-oriented goals.

Acknowledgments

The success of BAGNet is due to the concerted efforts of many people, too numerous to name individually. We owe special thanks to Bill Johnston, LBL, who led our efforts to establish the testbed and who is still our primary testbed coordinator, and to Robert Kahn, CNRI, who strongly encouraged our group to pursue the teleseminar application when we first started planning the testbed. We also wish to acknowledge Berry Kercheval, Xerox PARC, who assumed primary responsibility for the task of assigning PVCs and developing scripts to install the routing tables in ATM switches and host adapters. Many others who have played major roles in testbed activities are either listed below in the references, or can be quickly identified by spending a few minutes browsing through BAGNet documents on the World Wide Web.

References

[1] "BAGNet Home Page," *World Wide Web document*, http://george.lbl.gov/BAGNet.html.

[2] L. Berc, "Debugging an ATM packet loss mystery," *World Wide Web document*, http://chocolate.pa.dec.com/mystery/mystery.html.

[3] R. Jones, "Netperf Homepage," *World Wide Web document*, http://www.cup.hp.com/netperf/NetperfPage.html.

[4] R. Mines, J. Friesen, C. Yang, "DAVE: A Plug and Play Model for Distributed Multimedia Application Development," *Proceedings of the ACM Multimedia '94 Conference*, Sept., 1994.

[5] A. Periyannan, "BAGNet Ping Matrix," *World Wide Web document*, ftp://atg.apple.com/pub/bag-matrix/matrix.html.

[6] D. Wiltzius, "LLNL ATM Performance Data," *World Wide Web document*, http://www.llnl.gov/bagnet/perf-llnl.html.

[7] D. Wiltzius and L. Berc, "BAGNet Information & Status," *World Wide Web document*, http://chocolate.pa.dec.com/status.html.

RIACS
Mail Stop T041-5
NASA Ames Research Center
Moffett Field, CA 94035