

ABSTRACT

Practical experiments in a wide-area ATM network environment are essential to gain a better understanding of the factors affecting ATM performance. This article summarizes the main lessons learned through the authors' practical experiences with ATM and how these can be applied to performance tuning. In order to achieve maximum performance from an ATM WAN, factors ranging from the transport protocol to the equipment used and the network conditions must be carefully taken into consideration. The authors also discuss some of the software-based throughput measurement tools available for the evaluation of high speed network performance, with particular emphasis on NetSpec.

ATM WAN Performance Tools, Experiments, and Results

Luiz A. DaSilva, Joseph B. Evans, Douglas Niehaus, Victor S. Frost,
Roelof Jonkman, Beng Ong Lee, and Georgios Y. Lazarou
The University of Kansas

Interest in asynchronous transfer mode (ATM) networks has increased steadily over the past few years; however, only recently have the first wide area ATM networks (ATM WANs) been deployed. Practical experiments in this environment are essential for a better understanding of the factors that determine ATM performance. This article summarizes the main lessons learned through our practical experiences with ATM and how these lessons can be applied to improving performance.

When first utilizing an ATM network, it is not uncommon for the user to experience a gap between expectation and reality: throughput that is one or two orders of magnitude below link speed is often reported. The problem is not that ATM networks perform below the advertised level, but that many systemwide parameters must be appropriately set in order to achieve the best possible performance. Some of the most important and well-known [1] factors in ATM performance, which will be discussed in the following sections, are the following.

Protocol Considerations — Most ATM connections currently utilize Transmission Control Protocol/Internet Protocol (TCP/IP) as the transport protocol. When using TCP/IP over high-speed long distance networks, throughput may be limited by factors such as excessive retransmissions due to cell losses at the edge of the network, inefficiency due to maximum transmission unit (MTU) size, and TCP window management [2]. In most cases, these limiting factors can be avoided with careful tuning and occasional modifications to application code.

This research is partially supported by DARPA under prime contract DABT63-94-C-0068 to Sprint Corporation. Extensive information about this research project can be found in <http://www.ittc.ukans.edu/Projects/AAI/>.

Portions of this article were presented at the 1997 IEEE International Conference on Communications (ICC '97).

Host Considerations — When transmitting or receiving over high-speed networks, the host is often the limiting factor. Characteristics to be considered include CPU and memory speed, operating system, system parameters and load, among others.

Network Considerations — Finally, such factors as propagation delay, switch buffer size, link speeds, bandwidth mismatches, traffic policing, and congestion avoidance methods will determine the maximum throughput that can be expected.

This article is organized as follows. In the next section we present the general characteristics of the network used in our experiments. In the third through fifth sections we describe and discuss our results in each of the three main areas listed above. Experiments described in this article utilized two throughput measurement tools: *ttcp* and *NetSpec* [3]. A discussion of the characteristics and merits of these tools is presented in the sixth section. Finally, the last section summarizes our findings.

MAGIC AND AAI NETWORKS

All experiments described in this article were conducted over the MAGIC and AAI networks. The Multidimensional Applications and Gigabit Internetwork Consortium (MAGIC) is a group of industrial, academic, and government organizations participating in gigabit networking research. The MAGIC backbone network (Fig. 1) operates at 2.4 Gb/s, and each site on the network includes LANs or hosts communicating at gigabit-per-second rates [4].

The Advanced ACTS ATM Internetwork (AAI) provides wide-area ATM connectivity. It connects several Department of Defense (DoD) high-performance computing centers and the MAGIC and ATDnet gigabit testbeds. The ATM service is provided by Sprint, and Fig. 2 presents a simplified view of the topology.

PROTOCOL CONSIDERATIONS

The Internet Protocol family, and TCP/IP in particular, have enjoyed great popularity in the last decade. Most studies published about ATM performance utilize TCP/IP [4-8], as do most applications currently run over ATM.

Certain concerns arise when TCP is used over high-speed WANs. The most important of these are now discussed.


RETRANSMISSIONS DUE TO CELL LOSSES

TCP/IP performance is strongly influenced by cell loss rate. The MTU for IP is approximately 192 cells; therefore, losing a single cell has the effect of requiring up to 192 cells to be retransmitted, in some instances aggravating congestion conditions in the network [7].


The number of packets which need to be retransmitted can be minimized using selective retransmission. Some second-generation ATM switches implement early packet discard (EPD), which attenuates the impact of cell loss on congestion: whenever a switch determines the need to drop a cell, other cells belonging to the same packet are also discarded.

MTU SIZE

Although the use of small packets minimizes the number of cells to be retransmitted in case of packet discard, it limits throughput due to excessive overhead and increased packet processing. The impact of the resulting additional processing



■ Figure 1. The MAGIC network.



■ Figure 2. AAI network topology.

is greater at the receiving host, responsible for reassembly and checksum verification.

In order to quantify the effects of IP MTU size on throughput, we have set up a local area network (LAN) experiment using a DEC Alpha 3000/700 as the transmitting host generating a steady stream of data at OC-3 rates (155.52 Mb/s). Several different host architectures were used at the receiver, with results shown in Fig. 3. These results only serve to illustrate the potential impact of host architectures on ATM performance and should not be used as a comparison of specific vendor products.

WINDOW MANAGEMENT

A significant factor that tends to limit performance when using TCP over ATM WANs is window management. TCP implements window-based flow control to establish the maximum number of unacknowledged bytes that may be transmitted. In order to ensure continuous transmission, the window size must be at least equal to the bandwidth-delay product for a connection. The typical maximum window size, 64 kbytes or less, severely limits throughput for long distance ATM connections: at OC-3 rates, the bandwidth-delay product for a coast-to-coast connection is on the order of 10^7 bits, and a window size of 64 kbytes limits throughput to about 9 Mb/s.

It is possible to extend the maximum allowed TCP window size (up to several megabytes in some cases) with varying degrees of difficulty on different operating systems. Digital UNIX provides it as a system option; kernel patches are currently available for some versions of Solaris; a source license is needed for changes in the SunOS kernel.

Although most experiments in ATM performance are conducted using tools such as *ttcp*, which generate and transmit data as quickly as possible (in what is commonly known as *full blast mode*), actual user applications will present different traffic behavior and additional processing requirements from the hosts.

Experiments have been conducted in high speed file transfers using the File Transfer Protocol (FTP) [9], with results summarized in Table 1. FTP sits on top of TCP on the protocol stack, making use of services such as reliable transmission; it is therefore affected by window management. Files of various sizes were transmitted between two sites across the United States: the Naval Research Laboratory (NRL) in Washington, DC, and the NASA Jet Propulsion Laboratory (JPL) in California. The sites were connected to an ATM WAN through DS-3 links (44.736 Mb/s). While the maximum throughput obtained with *ttcp* was 23.8 Mb/s, file transfers using FTP achieved only 2.88 Mb/s. The reason for such disappointing throughput is that the unmodified code for FTP makes no allowance for extended TCP windows.

When an FTP call is initiated,

Application	Throughput
<i>ttcp</i> (window size 400 kB)	23.8 Mb/s
FTP (unmodified)	2.88 Mb/s
FTP (modified)	22.4 Mb/s

■ **Table 1.** Variation in throughput with window size for FTP using a transcontinental ATM over a DS-3 link.

limited by disk input/output (I/O). Factors influencing file transfer performance include disk access speed and application and operating system overhead, in addition to protocol factors.

HOST CONSIDERATIONS


Many factors may limit the ability of the host to transmit or receive ATM data. These include:

- CPU — The speed of the CPU will influence, for instance, the processing of incoming or outgoing packets (such as segmentation and reassembly of packets and checksum calculation).
- Backplane (I/O bus) — Support for direct memory access (DMA) and the backplane capacity have significant impact on communications performance.
- Memory bandwidth — This limits the speed of kernel/user memory transfers.
- Disk access speed — This may limit the performance of real applications, such as file transfers and remote queries.
- ATM interface and drivers — Buffering capabilities will determine whether losses occur at the ATM interface, unnecessary copies are made, and so on.
- Operating system — It not only influences the overall performance of the host, but the structure and tuning parameters of different operating systems vary in their support of TCP/IP window sizes, protocol stacks, MTU sizes, and so on.
- Load — Competition for host resources will clearly limit performance.

The combination of the factors above will determine the maximum achievable throughput when using a host as a receiver or transmitter. A performance benchmark ranking

various hosts and operating systems as to their networking capabilities would be very useful; none is available at the moment, partly due to the large number of tunable parameters, whose effects are often nonlinear.

Figure 3 illustrates how different hosts behave differently as the MTU size is varied. In Table 2, we list the maximum throughput we were able to obtain from several architectures over OC-3 links. These results are affected by host configuration and should not be used for comparison among specific vendor products. It is clear that transmit and receive characteristics are not symmetrical, since the two activities require different hardware and software capabilities from the host.



■ **Figure 3.** Variation in throughput (measured at the receiver) with MTU size for various receiving host architectures. Results should not be used as a comparison between specific vendor products.

NETWORK CONSIDERATIONS

Finally, the network conditions must be considered. We will address the factors that relate directly to the network by describing a suite of experiments conducted on an ATM WAN.

The networks used for these experiments are discussed in the second section. It should be emphasized that in these experiments the Sprint commercial ATM network was very reliable and congestion-free. All the effects discussed below were observed at the edge of the network.

THROUGHPUT VERSUS OFFERED LOAD

In defining metrics to help consumers compare various ATM equipment, the ATM Forum is considering throughput in terms of lossless, peak, and full-load behavior [10]. It is useful, then, to obtain the relationship between throughput and offered load, as shown in Fig. 4.

The results were generated from an experiment where traffic between WPAFB and ITTC, NRL and ITTC, and ARL and GSD¹ was simultaneously generated. NetSpec allowed us to vary the offered load by employing application-level pacing to generate constant rate traffic [3]. The three source-destination pairs competed for a common DS-3 link. Therefore, as we increase the offered load, a point is reached where switch buffers overflow and cell losses occur; at this point, a sharp drop in throughput can be seen due to the increased number of retransmissions.

The results refer to the aggregate throughput: the sum of the throughput obtained for each of the three connections. From Fig. 4, we can establish the maximum lossless throughput (the maximum throughput obtained without cell losses shows up in the figure as the point where the plot can no longer be approximated by a single straight line, around 28 Mb/s) and the peak throughput (around 29.5 Mb/s). The full-load throughput is obtained by transmitting full-rate traffic;² in this case, it was measured as 18.4 Mb/s.

A similar plot can be obtained for a single source-destination pair in the presence of rate mismatches. Again, traffic shaping is needed to prevent losses at the point of mismatch. By limiting the transmitting rate, we are able to avoid losses; once the rate is increased above a certain level, throughput drops sharply.

¹ The acronyms refer to sites in Fig. 2.

² Full rate refers to the transfer of packets from local memory to memory on a remote host as fast as the operating system, interfaces, and network allow.

Architecture	Operating system	Transmit throughput (Mb/s)	Receive throughput (Mb/s)
DEC 5000/240	ULTRIX 4.3	54.5	35.8
Pentium-100	Linux 1.3.91	125.0	70.0
SPARC-10/51	SunOS 4.1.3	62.0	99.0
SPARC-20/40	Solaris 2.4	83.0	95.0
SPARC-20/2x125	Solaris 2.4	130.5	115.0
DEC Alpha 3000/400	OSF/1 v3.0	130.6	134.1
DEC Alpha 3000/400	Digital UNIX T4.0	133.2	133.5
DEC Alpha 3000/600	OSF/1 v3.2	133.9	133.5

■ **Table 2.** Maximum throughput observed for hosts used as receiver or transmitter; measurements were conducted over OC-3 connections. These results are affected by host configuration and should not be used as a benchmark for specific vendor products.


Congestion control	TCP versus TCP		UDP versus UDP		TCP versus UDP	
	No	Yes	No	Yes	No	Yes
C to A	57.0	58.5	4.1	44.5	9.5	55.7
B to A	40.5	56.3	5.7	44.2	78.5	41.2
Aggregate throughput	97.5	114.8	9.8	88.7	88.0	90.6

■ **Table 3.** Individual and aggregate throughput for contending TCP traffic streams, UDP streams, and TCP versus UDP streams, with and without congestion control. All values measured in megabits per second (see Fig. 6 for topology).


IMPACT OF COMPETING TRAFFIC

Traffic from different sources competing for the same link will result in a decrease in throughput due to buffer limitations at switches and eventually cause congestion to occur. In order to illustrate this effect, we have set up an experiment involving five source-destination pairs, as depicted in Fig. 5, several of which competing for the same link(s). The values next to each of the five transmitting hosts represent the throughput obtained when each connection was set up individually (top) versus the throughput when traffic was produced simultaneously (bottom), measured in megabits per second. The round-trip time for each connection is provided next to the clocks, and the arrows indicate the direction of the traffic. Every connection was affected, and aggregate throughput decreased 25 percent, from 73.84 to 55.19 Mb/s. If we increase the offered load for each connection in this experiment (i.e., increase the source rate by using Netspec), an even sharper drop in throughput occurs due to cell losses, emphasizing the need for congestion control.

It should be pointed out that TCP's window-based flow control is not sufficient for congestion avoidance, a fact confirmed by the simple experiment described in [11]. The experiment was set up as shown in Fig. 6, with two hosts (one local, one remote) transmitting simultaneously. The DEC AN2



■ **Figure 4.** Throughput versus offered load for three simultaneous connections.



■ **Figure 5.** Effect of competing traffic on throughput in a ten-host experiment.

local-area ATM switch has credit-based flow control capabilities, and each connection was set up with the appropriate number of credits. We have collected results for various combinations of competing traffic streams, using both TCP and User Datagram Protocol (UDP); these are summarized in Table 3.

In all three cases, aggregate throughput is higher when closed-loop congestion control is present at the ATM layer. Even in contention between two TCP traffic streams, a drop of 15 percent in throughput occurs with-


out congestion control; although TCP's backoff mechanism prevents severe throughput degradation, it cannot prevent cell losses from occurring at the switch. Furthermore, congestion control is needed in order to guarantee fairness, as shown in the last two columns of Table 3; without it, UDP traffic (B to A) dominates, taking over most of the available capacity.

SWITCH BUFFERING CAPABILITIES


In order to illustrate the impact of switch buffer size, we conducted an experiment between ITTC and NRL, as shown in Fig. 7. Rate mismatches between OC-3 and DS-3 links exist in both directions.

Using NetSpec, we have generated constant-rate traffic at 35 Mb/s (this takes care of the rate mismatch), varying the size of the burst of data transmitted. The relationship between the burst size and throughput is shown in Fig. 8.

The asymmetry of the results indicates that the switches used at either end have different buffer sizes and allows us to roughly estimate the buffering capabilities as 3.5 Mbytes (GSD switch) and 35 kbytes (NRL switch); these correspond to the knees of the curves in Fig. 8. Early packet discard can be used to mitigate this type of problem.



■ **Figure 6.** Configuration for congestion control experiment.



■ **Figure 7.** Setup for experiment to estimate switch buffer capabilities.

PERFORMANCE EVALUATION TOOLS

In order to characterize performance in an ATM WAN environment, one needs performance evaluation tools capable of conducting large-scale tests that are representative of the traffic generated by current and emerging applications. Two software-based throughput measurement tools have been used in our experiments: *ttcp* and *NetSpec*.

The *ttcp* program can be used as a benchmarking tool to measure throughput of a connection between two hosts using either UDP or TCP. The program was created at the U.S. Army Ballistics Research Laboratory and is in the public domain. It is very useful for quick full-rate experiments involving only two hosts, and for verifying the maximum window size supported by a given host.


However, most of the experiments described in this article require support for carefully controlled multiple simultaneous

```

experiment {
  parallel {
    connection {
      xmtcnt1 = 204.235.72.224;
      rcvcnt1 = 204.235.73.24;
      xmt = 204.235.72.224 (xmtbuf = 52428);
      rcv = 204.235.73.24 (rcvbuf = 262144);
      rep = 204.235.73.24;
      protocol = tcp;
      test = cbr(messagesize = 18600:1, period = 30000);
      delay = 0 s;
      duration = 10 s;
    }
    connection {
      xmtcnt1 = 204.235.68.3;
      rcvcnt1 = 204.235.73.22;
      xmt = 204.235.68.3 (xmtbuf = 262144);
      rcv = 204.235.73.22 (rcvbuf = 262144);
      rep = 204.235.73.22;
      protocol = tcp;
      test = cbr(messagesize = 18600:1, period = 30000);
      delay = 0 s;
      duration = 10 s;
    }
  }
}

```

■ **Figure 9.** Experiment specification example.



■ **Figure 8.** Throughput versus UDP burst size.

connections and traffic shaping. A more sophisticated tool, *NetSpec* [3], was used in these experiments. *NetSpec* provides a block-structured language for specifying experimental parameters and support for controlling performance experiments containing a large number of connections across a LAN or WAN, and is specifically designed to maximize the reproducibility of the experiments [12].

Some of the features of *NetSpec* that provide enhanced support for large scale WAN experiments, as compared to other software-based performance measurement tools, are:

BLOCK-STRUCTURED LANGUAGE

Block-structured language enables the user to fully describe an experiment, including one or more connections. Parameters for each connection include the addresses of transmitting and receiving hosts, test traffic type, transport protocol, test duration, data rate, and packet and burst size.

The language makes each experiment easily reproducible. An example of a script for *NetSpec* 2.0, used for an experiment described previously, is shown in Fig. 9.

PARALLEL AND SERIAL CONNECTIONS

Parallel and serial connections are segregated by block statements into groups whose execution is performed in parallel or in series; these connection groups can be arbitrarily nested.

A possible use for this capability is to emulate different models of network use. Consider a serial group of connections that is part of a parallel group. The connections in the parallel group might represent long-lived connections, while the members of the serial group would represent short-lived connections established and discarded during the experiment [12].

Experiments with more than 100 simultaneous connections have been conducted. The parallel connection feature is illustrated in Fig. 9.

TRAFFIC TYPES

In addition to the full-rate traffic supported by *ttcp*, *NetSpec* supports constant-rate and "random-rate" traffic.

The constant-rate mode is implemented by periodically transmitting blocks of fixed size. The random-rate mode allows the user to specify the boundaries for a uniformly distributed random block size and interblock delay.

Checklist	Pointers to tools
1 What is the round-trip time for the connections being used?	ping
2 If using TCP/IP, is the window (socket buffer size) at least as large as the bandwidth-delay product for the link?	ttcp, netconfig, application code
3 Can the MTU size be increased?	ifconfig
4 Is there significant load on the hosts used for transmission and reception?	top, ps
5 What is the maximum throughput obtained with host/operating system?	ttcp, Netspec
6 If your application accesses the hard disk, what is the maximum access speed?	disktest
7 Are there rate mismatches in the path used? If so, are you pacing the traffic adequately?	atmarp, atmconfig, NetSpec
8 Is there any competition for the link(s) or switch(es) used?	SNMP
9 Which route is being followed?	netstat-rn
10 Are there any cell or packet losses? If so, what is the cause?	netstat, atmstat, SNMP
Notes: ping, ifconfig, and netstat: standard UNIX tools ttcp and Netspec: performance measurement tools atmarp, atmconfig, and atmstat: ATM driver-specific tools available for FORE ATM drivers netconfig, top, and disktest: custom tools for UNIX available in the public domain	

■ **Table 4.** Checklist for debugging/performance tuning in ATM networks.

SINGLE CONTROL POINT

An arbitrary number of connections can be set up, used, and closed from a single controller process running on any host in the network.

NetSpec currently supports TCP and UDP as the transport protocol, and has been successfully deployed for the following architectures: SunOS 4.x, Solaris 2.x, Irix 5.x and Digital UNIX 3.x. It has proven invaluable in conducting tests such as those described previously.

In future versions of NetSpec, the specification language will be made more object-oriented. Support for new traffic types, multicast connections, and data collection at intermediate points in each connection are also being added.

CONCLUSIONS

This article illustrates how proper evaluation of protocol/host/network factors is essential in order to obtain optimal performance in ATM WANs. In Table 4, we provide a checklist for factors that affect performance in such an environment, as well as some useful UNIX commands and utilities that can be used to check and/or modify the appropriate parameters. While the list is certainly not comprehensive, we believe that by following these steps one is able to avoid many of the pitfalls which commonly plague the novice (and sometimes even the experienced) user.

For most of the experiments described in this article, traffic was generated using either ttcp or NetSpec, and therefore is considerably less bursty than might be expected in most applications. Therefore, in practical applications one may expect less favorable results than those reported here.

Future areas of research include more extensive studies on available bit rate service, the impact of signaling, and the characterization of real-life application performance (video, voice, multimedia, etc.).

ACKNOWLEDGMENTS

The experiments reported in this article would not have been possible without the active contributions of many individuals and organizations. Among these are Rick Lett from Sprint Government Systems Division, Daniel Bromberg at the NASA Jet Propulsion Laboratory, Suresh Bhogavilli and Javad Boroumand at NASA Goddard Space Flight Center, Kenneth Finnegan at the Naval Research Laboratory—Stennis, and Chuck Thacker and Hal Murray at DEC.

Sprint's Broadband Operations Center (BBOC) also provided valuable assistance monitoring many of these tests.

The following trademarks are used in this article: Digital UNIX, DEC Alpha 3000, and DEC AN2 are trademarks of Digital Equipment Corporation; SunOS, Solaris, Sun SPARCstation, and Sun SPARCcenter are trademarks of Sun Microsystems.

REFERENCES

- [1] A. S. Tanenbaum, *Computer Networks*, 3rd ed., Upper Saddle River, NJ: Prentice Hall, 1996.
- [2] J. D. Cavanaugh and T. J. Salo, "Internetworking with ATM LANs," available at <http://www.msci.magic.net/docs/magic/ip-atm.ps>.
- [3] R. Jonkman, "Netspec: A Network Performance Evaluation and Experimentation Tool," available at <http://www.ittc.ukans.edu/netspec/>.
- [4] G. Y. Lazarou et al., "Using Measurements to Validate Simulation Models of TCP/IP over High Speed ATM Wide Area Networks," *Proc. ICC '96*, Dallas, TX, June 1996.
- [5] O. Bonaventure et al., "TCP/IP and the European ATM Pilot," *Proc. 1995 Conf. Network Protocols*, Sept. 1995.
- [6] J. B. Evans et al., "A 622 Mb/s LAN/WAN Gateway and Experiences with Wide Area ATM Networking," *IEEE Network*, vol. 10, no. 3, May/June 1996, pp. 40-48.
- [7] B. J. Ewy et al., "TCP/ATM Experiences in the MAGIC Testbed," *Proc. 4th IEEE Symp. High Performance Distributed Comp.*, available at <http://www.magic.net/tcp/overview.html>, Aug. 1995, pp. 87-93.
- [8] A. Romanow, "TCP over ATM: Some Performance Results," *IEEE JSAC*, vol. 13, no. 4, May 1995.
- [9] L. A. DaSilva, R. Lett, and V. S. Frost, "Performance Considerations in File Transfers using FTP over Wide-Area Networks," *Proc. 22nd Int'l. Conf. Resource Mngmt. and Performance Eval. of Enterprise Comp. Sys. (CMG96)*, San Diego, CA, Dec. 1996.
- [10] R. Jain, "Frame-Level Performance Benchmarking Work at the ATM Forum," *Proc. 1996 DARPA Wksp. Wide-Area ATM Performance*, Lawrence, KS, June 1996.
- [11] H. Zhu et al., "Performance Evaluation of Congestion Control Mechanisms in ATM Networks," *Proc. 21st Int'l. Conf. Resource Mngmt. and Performance Eval. of Enterprise Comp. Sys. (CMG95)*, Nashville, TN, Dec. 1995.
- [12] R. Jonkman et al., "Netspec: A Network Performance Evaluation Tool."

ADDITIONAL READING

- [1] A. Battou, "Measurements of UNI 3.0 Signalling Setup Times," *1996 DARPA Wksp. Wide-Area ATM Performance*, Lawrence, KS, June 1996.

BIOGRAPHIES

LUIZ A. DASILVA [StM] (ldasilva@itcc.ukans.edu) received his B.S. and M.Sc. degrees from the University of Kansas in 1986 and 1988, respectively. His thesis addressed the transmission of speech over packet-switched networks. He has been working as a researcher and doctoral candidate at the Information and Telecommunication Technology Center at the University of Kansas since 1995. Prior to this, he worked at IBM from 1989 to 1994. His current areas of research include performance evaluation of gigabit networks (with the ACTS ATM Internetwork project) and the pricing of multi-service commercial networks.

JOSEPH B. EVANS (evans@itcc.ukans.edu) received the B.S.E.E. degree from Lafayette College in 1983, and the M.S.E., M.A., and Ph.D. degrees from Princeton University in 1984, 1986, and 1989, respectively. Since 1989 he

has been at the University of Kansas, first as an assistant professor and since 1994 as an associate professor of electrical engineering and computer science. Prior to this, he held a postdoctoral position in the Network Systems Research Department of AT&T Bell Laboratories in Holmdel, New Jersey, where he was involved in the design of a high-performance integrated network. While at Princeton, he was awarded an AT&T Bell Laboratories Graduate Fellowship for 1984-1988; during this time, he was also a part-time employee of Bell Laboratories, working in the field of speech processing algorithms for packet networks. He was on sabbatical at the Cambridge University Computer Laboratory and Olivetti & Oracle Research Laboratory in Cambridge, England for the 1996-1997 academic year, working in the area of mobile computing and communications systems. His current research interests include high-speed (gigabit) networks, mobile networking, digital signal/speech processing, special-purpose computer architecture, and VLSI implementations.

DOUGLAS NIEHAUS (niehaus@ittc.ukans.edu) has been an assistant professor in the Electrical Engineering and Computer Science Department at the University of Kansas since 1993. His interests include real-time and distributed systems, operating systems, ATM networks, performance measurement, and programming environments. He received his Ph.D. in computer science from the University of Massachusetts at Amherst, where his thesis addressed the design and implementation of real-time systems. He was a system programmer at Bell Laboratories and AT&T Information Systems from 1981 to 1986, and at Convergent Technologies from 1986 to 1987.

VICTOR S. FROST (frost@ittc.ukans.edu) [SM] is the executive director for Research for the Information and Telecommunications Technology Center at the University of Kansas. From 1987 to 1996 he was director of the Telecommunications and Information Sciences Laboratory. He was a Guest Editor for *IEEE Communications Magazine* (March 1994) and the *IEEE Journal on Selected Areas in Communications* (May 1995) and is an associate

editor for *IEEE Communications Letters* and the *ACM Transactions on Simulation and Modeling of Computer Systems*. His current research interest is in the area of communication networks and he is conducting research using the MAGIC and AAI wide area testbeds. Dr. Frost received a Presidential Young Investigator Award from the NSF in 1984. He is a member of Eta Kappa Nu and Tau Beta Pi and a senior member of the IEEE. He served as Chair of the Kansas City section of the IEEE Communications Society from June 1991 to Dec. 1992. He received B.S. (1977), M.S. (1978), and Ph.D. (1982) degrees from the University of Kansas. In 1982 he joined the faculty of the University of Kansas, where he is currently the Dan F. Servey Distinguished Professor of Electrical Engineering and Computer Science.

ROELOF JONKMAN (rjonkman@ittc.ukans.edu) has received his B.S.CoE from the Hogeschool, Enschede, the Netherlands, in 1994. He has been working as a researcher since 1994 at the Telecommunications and Informations Sciences Laboratory at the University of Kansas. His areas of interest include network benchmarking, network performance analysis, wide area ATM networks, security, and systems performance.

BENG ONG LEE (blee@ittc.ukans.edu) received his B.S.E.E. in 1995 from the University of Kansas. Currently he is a graduate research assistant at the Information and Telecommunication Technology Center. His research interests include high-speed (gigabit) networks, and traffic source modeling.

GEORGIOS Y. LAZAROU [StM] (glaz@ittc.ukans.edu) is a doctoral student in electrical engineering and a researcher at the Information and Telecommunication Technology Center at the University of Kansas. He received his B.E. and M.E. degrees in electrical engineering from City College, City University of New York in 1992 and 1993, respectively. His areas of research are open loop control and modeling and simulation of high-speed ATM networks, and performance analysis of TCP/IP over ATM with the AAI project. He is a member of Eta Kappa Nu, Tau Beta Pi, and Phi Theta Kappa.

IEEE/ACM Transactions on Networking

*A bimonthly archival journal,
reflects the multi-disciplinary nature
of communications networks*

*Presents the highest quality papers
on both applied contributions and
theoretical research contributions*

Yearly Subscription Rate \$36

For more information, call +1 800 678 IEEE (US only); +1 732 981 0060
(Product Code: 019-159)

