



ATM Network: **Goals and Challenges**

B.G. Kim and P. Wang

Asynchronous transfer mode (ATM) can provide both circuit and packet-switching services with the same protocol, and this integration of circuit and packet-switching services can be beneficial in many ways. Four major benefits of the ATM technique are considered here: scalability, statistical multiplexing, traffic integration, and network simplicity. In the course of achieving these benefits, ATM makes compromises. In this article we assess the benefits and ensuing penalties of these compromises and put them into perspective.

ATM has been accepted in a wide spectrum of telecommunications and data communications communities. Since ATM networks are expected to emerge first in data communications environments, the current status of ATM local-area networks (LAN) and wide-area networks (WAN) is surveyed in the context of connectionless services. ATM and the Internet are described to show the common directions toward integrated services with some resource reservations. Advantages of using the ATM technique have been extensively described [10], but a few counter-arguments cannot be completely ignored [16]—this article considers both attitudes toward ATM networks.

Although ATM was standardized as a transport vehicle of wide-area B-ISDN (Broadband-Integrated Digital Services Network) services, the data communication community quickly recognized it as a potential data communication transport technology. Several factors contributed to a rapid acceptance of the ATM technique in the last several years. For example, the ATM concept already made an inroad into data communications when the cell-based DQDB (Distributed Queue, Dual Bus) protocol was standardized for IEEE 802.6 Metropolitan Area Network (MAN) protocol and its corresponding service, Switched Multimegabit Data Service (SMDS). A MAN is seen as a gateway to wide-area networks (WANs), and adopting a common cell format in both MANs and B-ISDN WANs was a significant development. The argument is then to use the identical cell format in a local-area network (LAN) so that an end-to-end LAN connection across a MAN or WAN can be made in the same transport unit of a cell. The possibility of using the same transmission format across a wide range of speed hierarchy is appealing to both LAN and WAN providers.

Another factor is that LANs are increasingly managed in a centralized fashion by hubs. Increasing use of twisted-pair and optical-fiber media foster centralized hub connections as well, and a switch-based LAN interconnection is seen as an extension (or a replacement) to existing hubs. As

computing power increases, direct connections of high-end workstations to a centralized switch is seen as attractive, especially for video and other high-end applications. In particular, the expected growth of video-related applications makes the connection-oriented ATM technique a suitable choice for this usage.

ATM networks are expected to be deployed in roughly two phases. In the first phase, existing LANs are interconnected by way of small-scale ATM switches resulting in ATM LAN islands. In this phase, an ATM network is used as a private LAN or a customer premise network (CPN) and serves as a backbone of LANs. In the second phase, ATM LAN islands are interconnected as more powerful ATM switches are offered by telecommunication industries for long-distance connections (in the late-1990s). By this time, a true B-ISDN and ATM WAN can be in place. For an early penetration of ATM technology into LANs, it is necessary for ATM to support existing data communication applications including TCP/IP packet transport.

Since ATM is connection-oriented, transport of connection-oriented services such as video can be readily accommodated in ATM networks (with the support of connection-admission and usage-parameter control functions). Here, transport of connectionless packets in ATM is addressed, particularly in terms of transport of TCP/IP packets.

In order to facilitate transport of IP packets over ATM networks, many issues are being addressed in the Internet community. The Internet Engineering Task Force (IETF) has formed a special "IP over ATM" working group to accelerate the development of routing and forwarding IP packets over ATM (sub)networks. Encapsulation of IP packets over ATM AAL5 is discussed in [13]. One approach is to multiplex multiple protocols over a single ATM virtual circuit by carrying a Logical Link Control (LLC) header. In the other approach, each protocol is carried over a separate ATM virtual connection.

Given the flexibility of ATM as a multiplexing and transporting tech-

nology, the working group focused its attention on "the Classical IP over ATM" model, which can be classified into local and end-to-end models. The local model applies ATM as a direct replacement of local LAN segments such as Ethernet. The end-to-end model, on the other hand, concatenates a number of the local models and ATM "wires" acting as interconnection between LAN routers [9]. Transport of classical IP, ARP (Address Resolution Protocol) and InARP (Inverse Address Resolution Protocol) packets over ATM AAL 5 is specified in an environment configured as the Logic IP Subnetwork (LIS) [15]. Each LIS contains at least one ARP server and many clients corresponding to IP systems in LIS.

The next-generation IP needs to be flexible enough for various link technologies and subnetworks on an end-to-end basis [5]. The classical IP over ATM approach invests in routers and bridges to extend them as intelligent switches. These intelligent switches then phase in and out the different link media and subnets. For instance, a bridge can learn about outgoing virtual channels for each incoming cell and intelligently keep track of ATM connection status. An opposing approach is proposed in a "lightweight subnet model" [9], which attempts to eliminate all or most of the IP header (and its corresponding functions) from the fact that a binding is established during call setup time between two end systems. Since a direct mapping between TCP/UDP and ATM virtual circuits (VCs) already exists, most of the overhead can be saved. Continuing proposals and studies are expected in the Internet community for transport of TCP/IP over ATM networks.

In the second phase of ATM deployments, as ATM LAN islands are connected across WANs, routing of IP packets over ATM WAN has to be provided [4]. The CCITT (now called the International Telecommunications Union or ITU) has recommended two general approaches: indirect and direct. In the indirect approach, connections among routers (switches) are pre-established in (semi-) permanent VCs so that the destination address of an IP packet is

immediately mapped to a VCI/VPI at a LAN router. In the direct approach, routing of IP packets is performed (directly) by an ATM network by provisioning several connectionless servers (CL servers). As an IP packet is segmented, the first ATM cell of the packet contains the destination IP address. When the first ATM cell arrives at a CL server, the routing table is consulted to obtain a VCI/VPI routing label for the next router or the destination. Subsequent cells belonging to the same IP packet can be forwarded with the VCI/VPI selected for the first cell. Thus, forwarding of IP packets can be achieved without costly reassemblies at CL servers.

The Internet

Whereas ATM is connection-oriented (CO), the Internet experiences substantial successes in providing various user services on the basis of CL routing of datagram packets. Some future services, however, may require CO services. Video and audio teleconferencing are good examples.

Resource reservation protocols suitable for CO user services have recently been proposed in the Internet community. Of particular interest are two reservation protocols with the capability of multicast connections to a selected group of receivers. In a multicast connection, information is sent in a single piece but is delivered to many receivers as the network copies and forwards the information to multiple outgoing links. In ST-II (Stream Protocol), a multicast routing tree is constructed with the sender at the root and a series of exchanges of control messages [19].

Once established, ST-II agent nodes (executing ST-II protocol functions on top of IP) become nodes in the multicast routing tree and send multiple copies to links leading to subtrees. In RSVP (ReSeRVation Protocol), on the other hand, the data source sends the *path* packet to the multicast receivers [22]. The path traversed by the *path* packet becomes an implicit route for data. However, resources are not reserved until a reservation request initiated by receiver makes its way toward the source in a reverse path of the *path* packet. RSVP is expected to be useful for such ap-

plications involving diverse types of receiver devices and frequent changes of receiver connections. For example, in video conferencing, there can be a range of receiver devices and associated capabilities. Also, members may dynamically join and leave the conference. Instead of having the sender keep track of connections to all receivers, some burden is shared by the receiver in requesting resource reservations. Still, further studies may be needed in providing CO services over the Internet.

ATM Protocols in ISO Reference Model

From the preceding discussions, one can recognize that major objectives of ATM and the Internet are converging. Both attempt to provide integrated CO and CL user services but differ in how such services are to be delivered. ATM relies on CO network service whereas the Internet evolves around a CL routing principle. As ATM LANs are prepared to make inroads into CL data services, it may be helpful to clarify the correspondence between ISO Reference Model ATM layers [10]. This relates transport of IP packets over an ATM network and can put different communication functions into a proper perspective. The architectural framework is presented here by separating user and control planes in the B-ISDN (ATM) protocol architecture. In fact, most confusion is caused by mixing the two planes.

Protocol hierarchies of CO user services are illustrated in Figure 1 for control and user planes. For CO service, connection is established in the control plane by a signaling network before data is exchanged in the user plane. As the connection is established, VCI/VPI pairs are allocated along the connection path for the purpose of routing ATM cells along the path, which corresponds to the routing function in the ISO network layer in Figure 1(a).

As a new connection path is allocated, the signaling network informs the affected switches of appropriate VCI/VPI mappings for individual switches. In the control plane many network services may be offered, such as various information and intelligent

user services (credit card calls and 800 services, for example). They are offered as a part of the network applications on top of the signaling network, and are depicted in the figure as a shaded box. Once the connection is established, user data can be transported over the established path as shown in Figure 1(b).

The ATM layer performs the mapping of incoming VCI/VPI pair to the outgoing pair, but is not capable of determining autonomous routes for individual cells or packets. In other words, the ATM layer transports data in fixed-size cells according to the predetermined addressing information. It can now be reasoned that AAL and ATM layers constitute ISO data link layer. AAL layer is regarded as a part of the data link layer because of the segmentation and reassembly (SAR) function (including the error-checking after the reassembly). The ATM layer is also regarded as a part of the data link layer (not the physical layer) because of the variable bandwidth that can be dynamically allocated to the VCI/VPI ATM path.

For CL service, corresponding layer structure is depicted in Figure 2. The control plane similar to Figure 1(a) is not shown and only the user plane is shown. Recall that CL service can be implemented in a wide-area ATM network in an indirect (permanent VC) or a direct (CL server) manner. The shaded box in the network layer function is present only for the direct approach, which provides Internet-like CL routing. With the indirect approach, VC paths are allocated among routers (switches) ahead of time and network routing function is not necessary.

ATM Goals and Realities

ATM is based on connection-oriented network paths and requires compromises to support both CO and CL user services. The purpose of this section is to present a balanced view of merits and capabilities of ATM—discussions here are more relevant to the ATM WAN environment. Four issues are presented: scalability, statistical multiplexing, traffic integration and network simplicity. It should be pointed out that ATM indeed possesses all these capabilities and more,

such as easy management; the focus here is to present the extent of ATM capabilities.

Scalability

Assertion. Scalability is indeed one of the most valuable properties of ATM. The key factors contributing to the scalability are a switch-based architecture and the common cell structure

detection/resolution times. Users can access ATM networks via a variety of physical connections irrespective of media types and applications. Within the limit of the physical link bandwidth, an arbitrary bit rate can be allocated to a user. Furthermore, the bandwidth remains allocated for the entire connection. At the same time, from the viewpoint of systems, the network bandwidth can be provisioned in a scaled manner with the switch-based configuration. As the network load increases and as more

VC connection is established on an end-to-end basis and cells cross sub-network boundaries transparently. This allows for data to be transported in the same format over the entire network span regardless of data rates at intervening subnetworks.

Counterpoint. Although increased bandwidth demands can be handled in a flexible manner in ATM networks, the maximum bandwidth available to a user is still limited to the link bandwidth. This appears similar to present data communication networks where various communication interfaces exist such as RS-232, Ethernet, FDDI, etc. Similarly, the network bandwidth will be limited by the size of the switch. The aggregate switch bandwidth grows in discrete step, say, from a 4×4 to an 8×8 switch, etc. This may be comparable to the present networking hierarchy offered at discrete quantities of 1, 10, 16, and 100 Mbps. The observation can be further extended to the speed hierarchy in WAN. Synchronous Digital Hierarchy (SDH) and Synchronous Optical Network (SONET) specify discrete speed hierarchies in units of 155 Mbps and 45 Mbps. Namely, individual bandwidth can be allocated in a scaled fashion, but the aggregate network bandwidth is offered in discrete steps.

The uniform cell format across a range of subnetworks is indeed appealing and can facilitate the hardware processing to gain switching speeds. However, a communication network consists of a wide range of transport speeds, from low-speed terminals at a few Kbps to backbone networks with data rates ranging from hundreds of Mbps to several Gbps. This hierarchical organization stems from the fact that not all traffic is routed across the speed boundaries. In fact, a rule of thumb is that 80% of traffic is routed within a subnetwork, whereas the remainder is transferred to other subnetworks and the higher-speed backbone networks. ATM networks are also organized in hierarchical fashion, by which many 155Mbps ATM networks feed into higher-capacity networks. Although ATM cells are processed in hardware, costs of processing ATM cells at Gbps speed may become prohibitively high

Figure 1. ATM layers for connection-oriented services

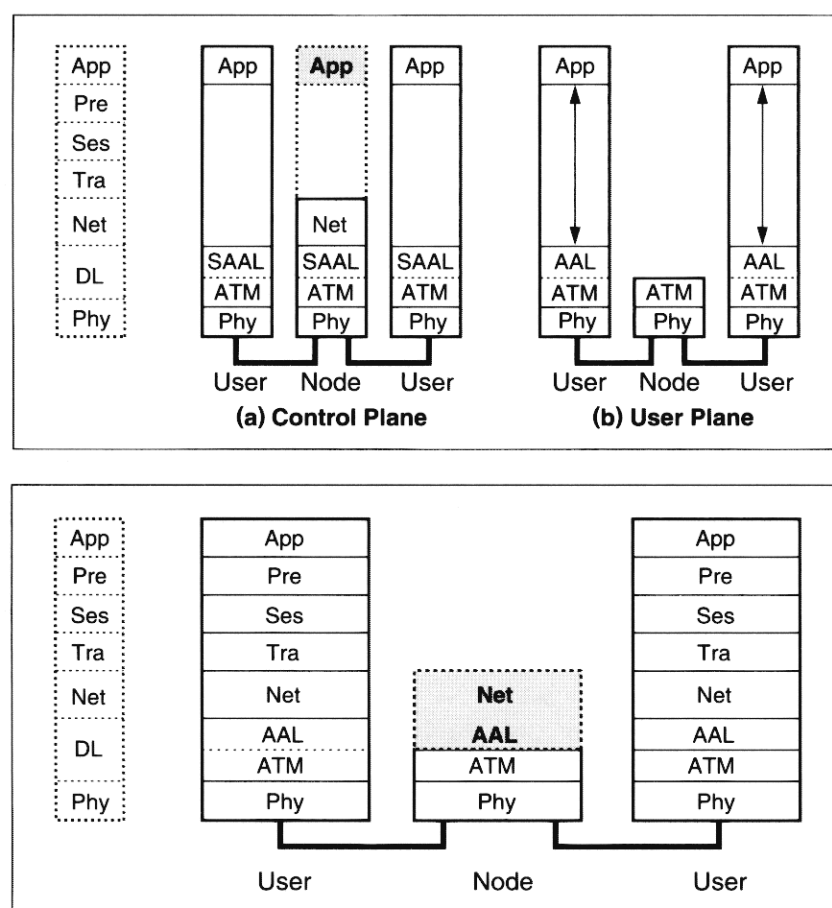


Figure 2. ATM layers for connectionless services

across all ATM system components. Conventional LAN technologies (Ethernet, FDDI, etc.) are limited by the propagation delays involved in coordinating the sharing of the link bandwidth. For example, increasing Ethernet speed reduces efficiency due to correspondingly long collision

subnetworks need to be connected, more switch ports can be added in an incremental fashion.

The contribution of the common cell structure is significant as well. Whereas present data communication networks re-encapsulate packets as network boundaries are crossed (say, from an Ethernet to FDDI), costly packet re-encapsulation is avoided in ATM networks. An ATM

(although all optical processing may provide a breakthrough). The cell format that is considered reasonable at one level of the speed hierarchy may not appear attractive at others, especially at high speeds [14].

Statistical Multiplexing

Assertion. A variable bit-rate (VBR) source is characterized by having different degrees of activities during a connection, and many applications envisioned for future ATM networks exhibit such behavior. A VBR video source, for example, has a peak rate at scene changes but a significantly lower rate as temporal and spatial compressions are performed after a scene change. Characteristics of VBR traffic may be represented by long-term average and peak cell emission rates, among others. Since not all VBR sources are expected to generate cells at their peak rates, the bandwidth less than the peak rate can be allocated to a VBR source. This allows for more sources to be admitted than the number of sources admitted by peak rates. At the same time, statistical variations in traffic load from individual sources can be smoothed out as many sources are multiplexed, resulting in better utilization of the shared resources.

Although shared-medium networks, such as IEEE 802 LANs and FDDI, are based on the same principle of statistical multiplexing, they allocate the entire link bandwidth to one user on a temporal basis. Since contention among users has to be coordinated in a fair manner, the protocols tend to be complicated (and less scalable). In contrast, the allocated bandwidth is available to the user all the time in ATM and at the same time statistical variations in VBR traffic are exploited to improve the resource utilization.

Counterpoint. Statistical multiplexing is shown to be paradoxical in that the dedicated bandwidth may not be fully utilized for fear of sudden surge in bursty traffic [12]. Depending on the degree of burstiness, significant overengineering may be necessary. A rule of thumb is to allocate the maximum of 80% or 85% of the link bandwidth, limiting the utilization of link capacity. Also, the statistical multi-

plexing works well when many random sources are multiplexed [18]. As high-bandwidth applications involving video traffic are developed, there may be simply too few traffic sources to fully exploit the benefits of statistical multiplexing.

ATM requires connection admission control (CAC) and usage parameter control (UPC) protocols to determine the amount of bandwidth to be allocated to a connection and to prevent excess traffic beyond the allocated bandwidth, respectively. Determination of bandwidth required for a connection request is not trivial [1, 7]. The UPC algorithm may produce excessive cell losses for a small violation [21]. Details of CAC and UPC protocols suitable for statistical multiplexing of diverse traffic types may not be completely understood for some time.

Traffic Integration

Assertion. With a uniform cell format, data from different sources can be readily integrated in ATM networks. By dedicating resources for a brief cell time per source, data from different sources appears to be transmitted concurrently. This is similar to time-sharing computing systems where each waiting job is given a fixed quantum of processor time. The unit of traffic integration is an ATM cell. A cell can be easily inserted into a cell stream from a large collection of data. In data networks, on the other hand, a short packet can be delayed until the transmission of a long packet is completed. Furthermore, a true traffic integration is taking place at the ATM layer as cells from different AAL classes, signaling (control), and management data are all mixed in the same cell format.

Counterpoint. Cells are integrated at ATM layer. However, it is not always desirable to mix cells with sufficiently different QOS requirements. For example, cells with low-loss requirement may have to be sent ahead of cells that can tolerate delays. This implies that a scheduling algorithm is necessary to distinguish cells with different QOS requirements, perhaps based on AAL types. ATM does not specify a scheduling algorithm and assumes an FIFO (First-in, First-out)

policy with priorities specified in the CLP (cell loss priority) bit. Since no QOS-based scheduling is performed, the ATM cell stream delivered to a receiver can be substantially different from what entered the network.

For example, periodic cells from a constant bit-rate (CBR) source can appear as clusters of cells when the delay variance tolerance is high [3]. In order to integrate cell streams with QOS requirements, it appears that network nodes have to exercise judicious scheduling algorithm to treat cells differently. This demonstrates that although cells are integrated as they enter the network, they have to be segregated within the network in order to satisfy different QOS requirements. This argument leads to complex network functions as well.

For certain applications, loss-based traffic integration may not be practical. When a reliable transport service is to be offered, a loss of a single cell forces retransmission of the entire data. The efficiency of a reliable transfer service may be measured by the bandwidth-delay (BwD) product. The BwD product is equivalent to the number of bits that can physically fill up the connection pipe from the sender to the receiver. For a wide-area ATM network connection, the BwD product can be sufficiently large that the retransmission may result in discarding a large amount of data already in the connection pipe. For the reliable transport service, a large BwD product can have a serious performance implication.

Network Simplicity

Assertion. In a high-speed network, network nodes need to be simple to catch up with fast communication speed. In ATM, functionalities in network nodes are simplified in three ways. First, by taking advantage of lower bit error rates in optical fiber, transmission errors are not monitored at network nodes (except for ATM header checking). Error handling is performed only at the network boundary nodes or at user-end devices. Secondly, since ATM has a fixed-size cell and guarantees cell ordering, the boundary of a data frame is transparent to the receiver. Frame-delimiting functions are not

necessary at network nodes. Finally, the routing of cells is made simple by pre-allocating routing labels for the entire fixed-path across the network. This eliminates needs for reassembly and re-encapsulation in the network.

In summary, the network simplicity is achieved by eliminating VC-level functions associated with user traffic control. Instead, node functions are limited to those for transport of ATM cells admitted into the network from various VCs. As the result of simplified node functions, ATM functions can be implemented in hardware, further improving on the processing speed. In the Internet, for example, more processing requirements are added to maintain resource reservations [22], and network nodes can become bottlenecks in high-speed networks. By reducing overhead at network nodes, ATM may be able to allow for very high-speed transmission with low levels of delays and delay jitters.

Counterpoint. An ATM network is envisioned as one in which network intelligence is located at boundary nodes while transit nodes provide data transfer along pre-established paths. This open-loop strategy is an example of an out-of-band (OOB) control strategy, whereby control and data information may be carried in different paths in the network. Due to the lack of feedback information, a problem in a data path may not be recognized immediately. For example, when congestion develops in a path it can go unnoticed for some time, resulting in the loss of a large amount of data. Furthermore, the promise of the network simplicity may not be fulfilled as more complicated functions become necessary in ATM. It was pointed out that a certain form of intelligent scheduling may be necessary to satisfy different QOS requirements of connections. Also, for applications requiring reliable transport services, the lack of error checking within the network may force costly retransmissions.

Finally, provision of multicast connections may be more difficult in the ATM environment. The connectivities from the source to receivers can be seen as a routing tree with the source at the root and receivers at the

leaves. Connections of individual data paths may be simple in ATM, but signaling nodes now have to maintain multicast routing trees for such applications. Namely, the routing intelligence (signaling nodes) and the routing execution (transit nodes on data paths) are performed in different parts of the network. This can lead to significant delays when frequent changes are to be made in multicast connectivities, as in a video distribution service, for example.

Conclusion

ATM and the Internet show convergence to a network supporting diverse traffic types and QOS requirements. Migration of connectionless service into ATM LANs and WANs has been examined and arguments on some of the merits of ATM technology have been presented. ATM technology has many benefits including scalability, traffic integration, statistical multiplexing and network simplicity, and we foresee numerous and rapid developments of ATM products and services. By providing both connection-oriented (CO) and connectionless (CL) services over the CO routing technique, ATM has to make some compromises.

In summary, ATM possesses many attractive features as a networking technology. For successful ATM, not only development of new applications but also migration of existing applications into ATM networks is important—evaluation of ATM technology in the context of applications may be a valuable exercise. ■

References

1. Anick, D., Mitra, D., and Sondhi, M.M. Stochastic theory of a data-handling system with multiple sources. *B.S.T.J.*, 61, 8 (Oct. 1982).
2. Atkinson, R.J. Default IP MTU for Use over ATM AAL5. Internet Draft, draft-ietf-atm-mtu-07.txt, Feb. 1994.
3. ATM Forum. ATM user-network interface specification. Version 3.0, Prentice-Hall, New York, 1993.
4. Box, D.F., Hong, D.P., and Suda, T. Architecture and design of connectionless data service for a public ATM Network. In *Proceedings of INFOCOM '93*, (San Francisco, March 1993).
5. Brazdziunas, C. IPng support for ATM services. Internet RFC 1680, 1994.
6. Video Codec for audiovisual services at $p \times 64$ Kbit/s. CCITT Recommendation H.261, Geneva, 1990.
7. Choudhury, G.L., Lucantoni, D.M., and Whitt, W. On the effectiveness of effective bandwidths for admission control in ATM networks. *ITC 14* (June 1994), 411–420.
8. Cidon, I., Gopal, I., and Segall, A. Connection establishment in high-speed networks. *IEEE/ATM Trans. Networking 1*, 4, 469–481 (Aug. 1993), 469–481.
9. Cole, R.G. IP over ATM: A framework document. draft-ietf-atm-framework-doc-00.txt, 1994.
10. De Prycker, M. *Asynchronous Transfer Mode: Solution for Broadband ISDN*. Second edition, Ellis Horwood Limited, 1993.
11. Fraser, A.G. Early experiments with Asynchronous Time Division Networks. *IEEE Network 7*, (Jan. 1993), 12–26.
12. Gechter, J. and O'Reilly, P. Conceptual issues for ATM. *IEEE Network 3*, 1, (Jan. 1989), 14–16.
13. Heinanen, J. Multiprotocol encapsulation over ATM Adaptation Layer 5. Internet RFC 1483, July 1993.
14. Jain, N., Schwartz, M., and Bashkow, T.R. Transport protocol processing at Gbps rates. In *Proceedings of SIGCOMM '90*. (Philadelphia, Penn., Sept. 1990).
15. Laubach, M. Classical IP and ARP over ATM. Internet RFC 1577, Jan. 1994.
16. Lea, C.T. What should be the goal for ATM? *IEEE Network 6*, (Sept. 1992), 60–66.
17. Ohta, M. Conventional IP over ATM. draft-ohta-ip-over-atm-00.txt, 1994.
18. Partridge, C. Gigabit networking. Addison-Wesley, Reading, Mass., 1993.
19. Topolcic, C. Ed., Experimental Internet stream protocol, version 2 (ST-II). Internet RFC 1190, Oct. 1990.
20. Williams, M.I. ATM—What does it mean? In *Proceedings of Networkshop '93*. (Melbourne, Australia, Dec. 1993).
21. Wilts, R., Witters, J., and Petit, G.H. Throughput analysis of a usage parameter control function monitoring misbehaving constant bit rate sources. In *Proceedings of the Second Workshop on Performance Modeling and Evaluation of ATM Networks*. Bradford, U.K., (July 1994).
22. Zhang, L., Deering, S., Estrin, D., Shenker, S., and Zappala, D. RSVP: A new resource ReSerVation Protocol. *IEEE Network*, (Sept. 1993), 8–17.

Continued on page 109

ogy design of ATM virtual private networks. UCLA CSD Technical Report, Nov. 94.

6. ETSI ETR. CBDS over ATM. Chester Version. May 1993.
7. Gerla, M., and Kleinrock, L. Flow control: A comparative survey. *IEEE Trans. on Comm.* (Apr. 1980), 553-574.
8. Monteiro, J., Gerla, M., and Fratta, L. Statistical multiplexing in ATM networks. *Performance Evaluation* 12, 3 (June 1991), 157-167.
9. Report on the Broadband ISDN Protocols for Providing SMDS and Exchange Access SMDS, *Bellcore SR-NWT-002076*, 1 (Sep. 1991).
10. Tranchier, D.P., Boyer, P.E., Rouaud, Y.M., and Mazeas, J.Y. Fast bandwidth allocation in ATM networks. *XIV ISS Proceedings* (Yokohama, Oct. 1992).
11. Walters, S.M., and Ahmed, M. Broadband virtual private networks and their evolution. *XIV ISS Proceedings* (Yokohama, Oct. 1992).

About the Authors:

SHIVI FOTEDAR is a doctoral student in the computer science department at UCLA. Current research interests include ATM architecture and protocols, traffic management and performance analysis.

MARIO GERLA is a professor in the computer science department at UCLA. Current research interests include ATM traffic management and congestion control, wireless communication, and all-optical networks.

PAOLO CROCETTI is a member of the technical staff at Italtel, Milan, Italy. Current research interests include ATM traffic management and video on demand.

LUIGI FRATTA is a professor in the Electrical Engineering department of the Polytechnico di Milano, Italy. Current research interests include ATM traffic management and performance analysis, and wireless networks.

Authors' Present Address: Computer Science Department, University of California, Los Angeles, CA 90024; email: {gerla, fotedar}@cs.ucla.edu

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© ACM 0002-0782/95/0200 \$3.50

Continued from page 38

23. Trajkovic, L., and Golestani, S. Congestion control for multimedia services. *IEEE Network* 6, 5 (Sept. 1992), 20-26.
24. Turner, J. Managing bandwidth in ATM networks with bursty traffic. *IEEE Network* 6, 5 (Sept. 1992), 50-58.
25. Yazid, S., and Mouffah, H.T. Congestion control methods for B-ISDN. *IEEE Commun. Mag.* 30, 7 (July 1992), 42-47.

About the Author:

RONALD J. VETTER is an assistant professor of computer science at North Dakota State University. Current research

interests include high-performance computing and communications, multimedia systems, and high-speed optical networks. **Author's Present Address:** North Dakota State University, IACC Building, Room 258, Fargo, ND 58105; email: rvetter@plains.nodak.edu

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© ACM 0002-0782/95/0200 \$3.50

Continued from page 44

About the Authors:

B.G. KIM is an associate professor at the University of Massachusetts at Lowell. Current research interests include control and traffic management functions in ATM networks. **Author's Present Address:** University of Massachusetts at Lowell, Lowell, Mass., 01854; email: kim@cs.uml.edu

P. WANG is a senior software engineer at the Open Software Foundation. Current research interests include distributed

computing and its applications. **Author's Present Address:** Open Software Foundation, Cambridge, Mass., 02142; email: pwang.osf.org

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© ACM 0002-0782/95/0200 \$3.50

Continued from page 100

About the Authors:

MICHAEL PERLOFF is a senior member of the technical staff at GTE Government Systems Corporation. Current research interests include simulation and modeling, communications network analysis, network protocols, and network management.

KURT REISS is a senior member of the technical staff at GTE Government Systems Corporation. Current research interests include broadband network performance, ATM traffic management, and network design.

Author's Present Address: GTE Communication Systems Division, 77 A Street, Needham Heights, MA 02194; email: {reiss.kurt, perloff.michael}@mail.ndhm.gtegsc.com

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© ACM 0002-0782/95/0200 \$3.50