QoS and Multipoint Support for Multimedia Applications over the ATM ABR Service

Bobby Vandalore, Sonia Fahmy, Raj Jain, Rohit Goyal, and Mukul Goyal

The Ohio State University

Raj Jain is now at Washington University in Saint Louis, jain@cse.wustl.edu http://www.cse.wustl.edu/~jain/

ABSTRACT Multimedia applications with stringent delay requirements are usually supported by the CBR or rt-VBR service categories of ATM. This article shows that ABR service, which was originally designed for data, can also support multimedia applications under certain circumstances. Issues related to the design of proper ABR traffic management algorithms for such support are presented.

Asynchronous transfer mode (ATM) is a cell-switching connection-oriented high-speed technology. The ATM Forum has defined five service categories: constant bit rate (CBR), real-time variable bit rate (rt-VBR), non real-time VBR (nrt-VBR), available bit rate (ABR), and unspecified bit rate (UBR). The International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) defines similar service categories for ATM. CBR and rt-VBR provide delay and loss guarantees, and can be used to transfer delay- and loss-sensitive unadaptive multimedia applications. Nrt-VBR provides loss guarantees. ABR and UBR are usually used for transferring data applications. UBR is a simple service that gives no guarantees.

ABR provides minimum cell rate (MCR) guarantees and is designed to provide low cell loss for well-behaving sources. It uses closed-loop feedback control to indicate network congestion information to the sources. The sources adjust their allowed cell rate (ACR) based on the network feedback. Feedback is indicated in resource management (RM) cells, which are periodically sent by the source and turned around by the destination. The switches along the path indicate the maximum rate they can currently support in the RM cell. The RM cells in the forward direction are called forward RM (FRM) cells, and those in the backward direction backward RM (BRM) cells (Fig. 1).

Multimedia applications are the key applications to use the large bandwidth provided by high-speed (e.g., ATM) networks. Sample multimedia applications include *videoconferencing*, *video on demand*, *distance learning*, *distributed games*, and *movies on demand*. Some of these applications, specifically those which have video and voice traffic, need delay and loss guarantees. Others, such as the World Wide Web and white boards, may be delay-insensitive.

Compressed video and voice traffic is usually bursty in nature [1], although video streaming products can generate constant-rate video with varying quality. If there are insufficient network resources, CBR and VBR connections may be rejected. In some cases, such as defense applications, it may be more desirable to use whatever bandwidth is available than to get the connection rejected. In such cases, ABR service is preferable. Also, with a CBR or VBR connection, the user may be forced to accept a low-bandwidth connection and cannot use bandwidth even if it becomes available later on. Similarly, if the available bandwidth is reduced for some reason, the network will simply disconnect CBR/VBR connections since currently it is not possible to renegotiate the contract.

ABR service provides MCR guarantees, which can be used to provide an acceptable quality of service (QoS) for video applications. Due to the closed-loop feedback control, usually the queues in the network switches are small and cell loss is low. Most of the queuing occurs only at the end systems. Multimedia applications, specifically video applications, can use feedback information and information about source queues to adjust their rates to efficiently use the available bandwidth [2]. Furthermore, queuing delay can be bounded if part of the bandwidth is reserved for ABR connections.

The remainder of this article is organized as follows. The next section reviews related work. We then define the set of problems to be solved for supporting multimedia applications over ABR service. The article goes on to show how switch algorithms can be modified to provide minimum rate guarantees. We discuss how to reduce the cell loss and delay for multimedia traffic, and then providing support for multipoint connections is discussed. The open issues and problems are examined, and finally, the last section summarizes the article.

RELATED WORK

There are several studies which address the problem of supporting multimedia traffic over CBR and VBR services. The ATM Forum Service Aspects and Applications (SAA) working group has approved specifications on transporting multimedia service over ATM networks. Luo and Zarki use a layered source encoder to generate high-priority and low-priority video traffic [3]. They study two scalable priority schemes which match the distribution of video information contained in the high- and low-priority bitstreams. They show that an interleaved transmission scheme (high-priority cells are interleaved with low-priority ones) performs better than a coupled concatenation scheme (low-priority cells are concatenated after high-priority ones). Aravind *et al.* study packet loss resilience using layered video traffic [4]. They compare the performance of data partitioning, signal-to-noise ratio scalability, and spatial scalability approaches — specified in the Motion Picture Experts Group (MPEG) standards — for generating layered video traffic.

A number of studies have addressed the problem of supporting real-time video applications with feedback control.

This research was sponsored in part by Rome Laboratory/C3BC Contract #F30602-96-C-0156.

All our articles and ATM Forum contributions are available through http://www.cis.ohio-state.edu/~jain

0163-6804/99/$10.00 © 1999 IEEE

Kanakia *et al.* show how binary feedback can be used to adjust the traffic generated by a video application [2]. Lakshman *et al.* discuss how to use the explicit rate feedback of ABR service to transport video [1]. The MPEG encoder parameters (such as quantization level) are adjusted to generate traffic which matches the feedback rate. Vickers *et al.* study the problem of transporting unicast and multicast multilayered video traffic using both credit-based and rate-based feedback congestion control [5]. The video sources adapt to changing network conditions by adjusting encoding parameters, and by adding or dropping a video layer. Recently, an algorithm to transport smoothed compressed video over explicit rate networks has been proposed by Duffield *et al.*[6]. Buffers at the source are used to smooth the video traffic, and adaptive encoding is performed based on the explicit rate indicated by feedback. Zheng and Atiquzzaman discuss the traffic management issues for multimedia applications over ATM networks [7].

MULTIMEDIA APPLICATION SUPPORT OVER ABR

Although the studies mentioned above have addressed multimedia transport, the problem of supporting multimedia applications is not yet a fully solved problem. The components of ATM technology are shown in Fig. 2. The main components used to support multimedia applications are the following:

Routing and Signaling — The private network-to-network interface (PNNI) protocol provides routing and signaling capabilities between ATM switches. The current PNNI protocol (version 1.0) supports any cast (connection to one among a group) and QoS-based routing. User-to-network interface signaling currently supports point-to-point connections, point-to-multipoint connections, and leaf-initiated joins.

Traffic Management — Traffic management includes connection admission control (CAC), usage parameter control (UPC), and congestion control. CAC is used to decide whether a connection request should be accepted or rejected, and to decide the parameters associated with that connection. UPC ensures that the sources respect the traffic contract by monitoring and controlling traffic at the entrance to the network. It detects violations and takes appropriate actions (e.g., cell tagging and cell discard). It protects the network and the QoS of other connections from misbehaving connections. For ABR service, feedback and flow control are used to regulate traffic and adaptively share the available bandwidth in a fair manner.

Other components, such as the physical layer interface, testing, network security, and network management are not examined in this article. Routing and signaling are also not addressed here. We only focus on the *traffic management issues for supporting delay-sensitive, loss-sensitive, unicast, and multicast adaptive multimedia applications*. Since delay-insensitive applications have less stringent requirements, they can easily be supported. We assume that the sources are adaptive and can change their rates according to feedback. The problems to be solved for supporting such applications over the ABR service are:

- ABR service provides MCR guarantees which can be used by the multimedia applications to achieve a minimum QoS. Some of the current ABR switch schemes assume MCRs to be zero. These switch schemes have to be modified to support nonzero MCRs.
- ABR service was designed to provide support for delay-insensitive data applications. When supporting multimedia applications over ABR, the delay and the variation in QoS should be minimized.
- Many multimedia applications, such as videocast and distance learning, can be multicast. Therefore, the problem of providing a multicast capability by supporting ABR point-to-multipoint, multipoint-to-point, and multipoint-to-multipoint connections should be addressed.

The next three sections address these three issues.

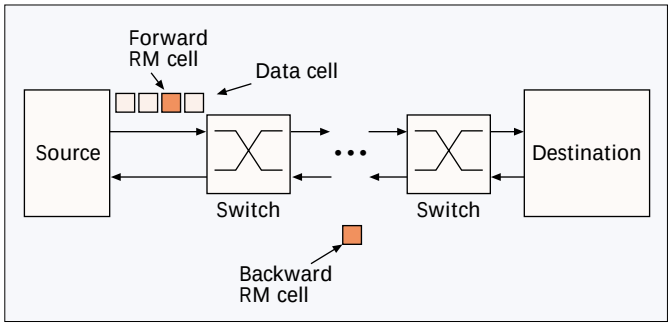


Figure 1. *An RM cell path.*

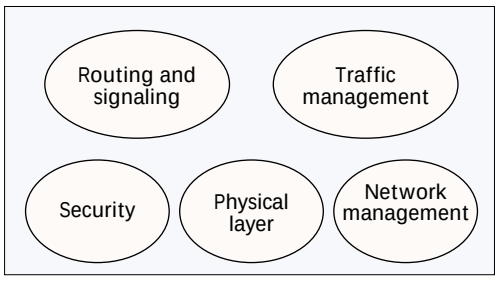


Figure 2. *Components for supporting multimedia applications.*

MCR SUPPORT FOR ABR SWITCH ALGORITHMS

The ATM Forum Traffic Management Specifications version 4.0 addresses ABR traffic management issues related to ATM end systems in detail. The behavior of the ABR switch is not standardized by this specification. ABR switch algorithms use either binary or explicit-rate feedback to provide congestion information to the end systems. The switches monitor the load on the network and divide the available bandwidth fairly among the competing sources. Several distributed switch algorithms which provide explicit-rate feedback control have been proposed and proven to give fair and efficient allocations.

When the requested MCRs are zero, the switch schemes attempt to satisfy the max-min fairness criteria. For nonzero MCRs, several similar fairness criteria can be used as defined in the ATM traffic management specifications. A generalized definition of fairness is given in [8], in which the excess bandwidth is divided proportional to a predetermined weight associated with each ABR connection. Other fairness criteria, such as giving MCR plus equal shares, or MCR plus MCR-proportional shares, are special cases of this general definition of fairness. The MCR guarantee of ABR service can be used to achieve a minimal QoS of multimedia applications.

Switch algorithms can be modified to support nonzero MCRs. As an example, the ERICA+ switch algorithm was modified to support non-zero MCRs. The ERICA+ switch algorithm operates at the output port of an ABR switch. It periodically monitors the load, capacity, and activity of connections, and calculates an overload factor. The monitoring interval is called the *averaging interval*. The algorithm also keeps track of the maximum allocation given as feedback during the previ-

ous averaging interval. The available capacity divided by the number of connections gives an estimate for the fair share of a connection. If the link is not overloaded, the algorithm gives the explicit feedback rate as the maximum of the current cell rate divided by the overload factor and the maximum previous allocation. If the link is overloaded, the maximum of the current cell rate divided by the overload factor and the fair share is given as the explicit rate. ERICA+ also uses dynamic capacity estimate scaling to control queuing delay.

Before discussing the modifications to switch algorithms, we define the term *excess fair share* (EFS)and the notion of *activity level*. The EFS is calculated by dividing the excess available bandwidth in proportion to the predetermined weights for each connection. The general weighted fair share is the sum of the MCR and EFS. The activity level is defined as the ratio of the actual source rate divided by the fair share of the source. The value of the activity level is less than one if the source does not use its fair share. The effective weight of a connection is its activity level multiplied by the predetermined weight. The EFS is calculated by dividing the excess available bandwidth in proportion to the effective weights.

Let $SrcRate(i)$ be the source rate of virtual connection (VC) i. Let MCR(i) be the MCR of VC i. Let ERz(i) be the explicit rate calculated by an algorithm assuming zero MCRs. Let $ER_{nz}(i)$ be the explicit rate to be calculated for nonzero MCRs. The modifications to a switch scheme to support MCR guarantees are as follows:

- The problem with nonzero MCRs is reduced to one with zero MCRs. This is done by subtracting the MCR from each connection's current source rate ($SrcRate(i)$ – MCR(i)) to obtain the excess rate over the MCR of each source.
- The switch algorithm for zero MCRs is applied to these excess rates to obtain the explicit rate ($ER_z(i)$ is obtained).
- The MCR of the connection is added to the above explicit rate and indicated in the BRM cells.

$$ER_{nz}(i) = \text{MCR}(i) + ER_z(i)$$

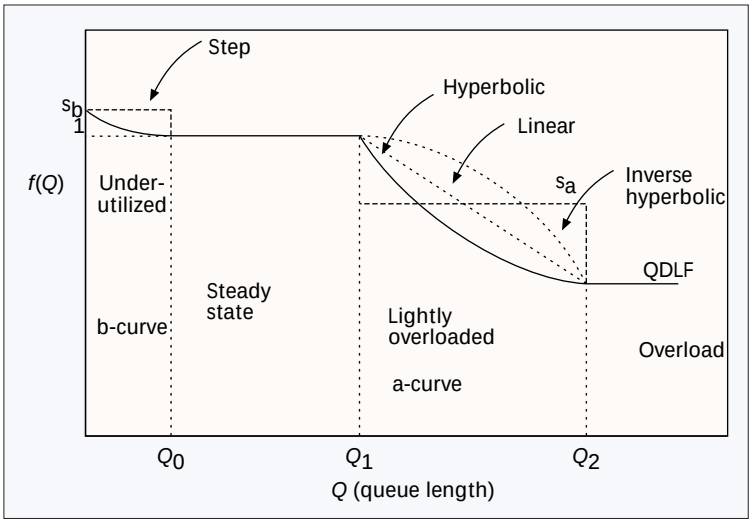
Simulation results of the above algorithm show that the MCR guarantee is provided, and allocations converge to the desired general fair allocation [8]. A proof for convergence of the modified algorithm is also given in the same reference.

CONTROLLING DELAY AND LOSS

ABR service is designed for data traffic with no delay guarantees. This section discusses approaches to minimize delay, loss, and variation in QoS.

During startup the queue length is dependent on the value of initial cell rate (ICR) times round-trip time (RTT). Other competing sources also influence the initial queue buildup. During the steady state period, due to the closed-loop feedback control of the ABR service, network queues are usually small, and queues build up only at the end systems. The cell loss is low for the ABR service. Multilayer encoding and forward error correction techniques can be used to make the application more loss-resilient.

The queuing delay of an ABR connection can be controlled by using a dynamic queue control function. For example, the ERICA+ switch scheme uses a queue control function $f(Q)$ to calculate bandwidth to be allocated as $f(Q)$ times the available bandwidth. The value of $f(Q)$ is dependent on the switch queue length (Q). A controlled queue length, and hence controlled delay in steady state, can be achieved using an appropriate queue control function.



■ **Figure 3.** *Queue control functions.*

A good queue control function should exhibit the following properties:

- Sources should be encouraged to increase their rates if the queue is very small (below a desired length). This implies that $f(Q) > 1$ in the range $0 \le Q < Q_0$.
- Maintain constant queue length at steady state (i.e., $f(Q) = 1$ in the range $Q_0 \le Q < Q_1$).
- Decrease the source rate under lightly loaded conditions (i.e., $f(Q) < 1$ in the range $Q_1 \le Q < Q_2$).
- Under heavy overload, only part of the available bandwidth should be used in order to leave enough capacity to drain large queues. However, a minimum portion of the available bandwidth should always be used for forwarding normal traffic. This can be achieved by limiting the $f(Q)$ value to a threshold. This threshold is known as the *queue drain limiting factor* (QDLF). This implies that $f(Q)$ = QDLF for the range $Q_2 \le Q < \infty$.

Figure 3 shows step, linear, hyperbolic, and inverse hyperbolic queue control functions. A thorough simulation and analysis of these different queue control functions and their trade-offs is given in [9]. Based on the analysis, the inverse hyperbolic queue control function was found to have the best performance. The linear queue control function, which is simpler to implement, performs satisfactorily in most cases. Due to the bursty nature of compressed video and voice traffic, the network might not achieve a steady state. By carefully choosing the parameters of the queue control function, the queuing delay can be minimized. Moreover, buffering is used at the source and/or destination to remove jitter in delay.

Most of the end users of multimedia applications are humans. Human users expect constant QoS. Hence, it is desirable to minimize the variation in the QoS perceived by the human user. Currently, an RM cell is generated for each *Nrm* – 1 data cells sent. According to current traffic management specifications, the value of the *Nrm* parameter is usually set to 32. At high link speeds (155.52 Mb/s or more), this would give rise to frequent feedback from the network. We assume that the adaptable video source can change rates only at frame boundaries by changing parameters such as quantization levels. Therefore, the frequent feedback due to low *Nrm* is not useful for the adaptable source. By setting a larger value of *Nrm* (say 256), the feedback rate is made comparable to the source frame rate. Disadvantages of using a higher value of *Nrm* are that it decreases the responsiveness of the network to

new sources, and may increase the effective RTT, resulting in larger end-to-end delay.

Cell loss and delay guarantees for multimedia ABR connections can also be provided by using a separate queue for them, and reserving a part of the bandwidth for these connections. Appropriate scheduling algorithms need to be used. The ABR parameters and switch algorithm parameters can be different for these connections. The averaging interval length can be large (order of feedback delay) if a large value of *Nrm* is used. Since these connections do not have a steady state, the parameter Q_1 can be chosen as $Q_1 = Q_0$. To provide a good transient response, a hyperbolic function can be used in the range $Q_1 \leq Q < Q_2$.

MULTIPOINT CONNECTION SUPPORT

Many multimedia applications are multicast, so it is important to provide scalable and efficient multipoint support for ATM. Traffic management for multipoint connections is a complex problem, due to the presence of multiple senders with different traffic characteristics, multiple receivers with different QoS requirements, and different bottlenecks along the multipoint connection path. Another difficulty in ATM multipoint support is that higher protocol layers must handle the multipoint communication functions, since lower layers are connection-oriented and not inherently broadcast. The signaling and routing working groups at the ATM Forum have defined point-to-multipoint connection setup and are finalizing UBR multipoint-to-point connection setup.

This section discusses the problems and current solutions for supporting multipoint ABR connections. The main feature of ABR is its feedback control mechanism. For point-to-multipoint ABR connections, the main issue is how to consolidate the feedback information at branch points. In multipoint-to-point connections, fairness definition within the sources in the same connection and among connections is difficult. In addition, feedback has to be regulated at the merge points since RM cells of different sources may be indistinguishable. Multipoint-to-multipoint connections combine the issues of point-to-multipoint and multipoint-to-point connections.

POINT-TO-MULTIPOINT CONNECTIONS

The current ATM signaling and routing standards only support point- to-multipoint VCs. Multicast trees are constructed through root- initiated joins. Leaf-initiated joins are also allowed by the latest standards.

Flow control for point-to-multipoint connections is an extension of that for point-to-point connections. A number of studies have addressed the problem of choosing appropriate values for the source end system parameters for point-to-multipoint connections to avoid transient queues. In addition to this, several algorithms tackle feedback implosion. The problem of "feedback implosion" arises if the feedback information provided to the sender increases in proportion to the number of leaves (receivers). To solve this problem, consolidation of feedback information from the different leaves of the tree is necessary (Fig. 4).

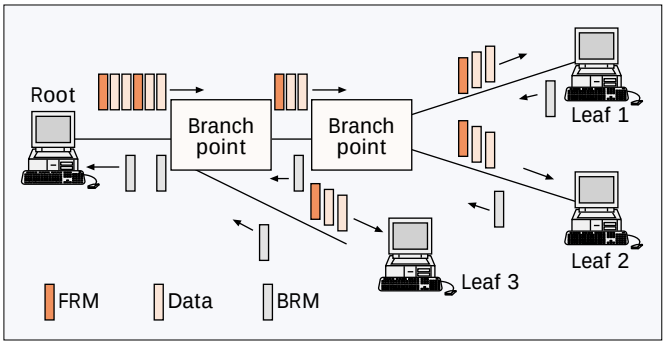


Figure 4. *Point-to-multipoint connections.*

The earliest proposal to solve the feedback consolidation problem works as follows. The forward RM cells are multicast to all branches, and also turned around and sent back to the sender, carrying the minimum explicit rate value denoted by all branches. This approach was used to extend various point-to-point rate allocation algorithms. The extension preserves the fairness and efficiency of the corresponding point-to-point algorithm. This method, however, suffers from the "consolidation noise" problem caused by the asynchrony of feedback arrival and rate changes. Consolidation noise occurs when feedback is not received from some of the branches when an RM cell is to be returned to the sender. Several variations of the early algorithm have been proposed to solve the feedback consolidation problem.

Fahmy *et al.* discuss a scheme where the switch forwards BRM cells back to the sender only if BRM cells have been received from all branches [10]. However, if the network is overloaded, the switch quickly passes BRM cells returning from the branches without waiting for all branch BRM cells to arrive. This approach eliminates the consolidation noise, while maintaining fast transient response in case of overload.

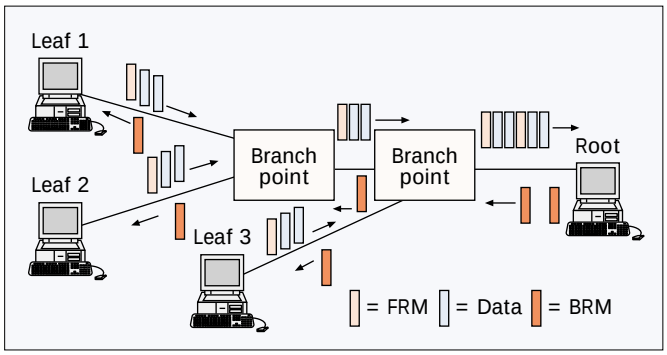


Figure 5. *Multipoint-to-point connections.*

MULTIPOINT-TO-POINT CONNECTIONS

Multipoint-to-point connections are supported only for the UBR service in the baseline text of the current ATM Forum specifications. Traffic management rules for multipoint-to-point connections are still under study. Bandwidth management at the merge point is important since the merged traffic is the sum of all traffic from the leaves (Fig. 5).

Defining the fairness criteria for multiple senders in the same connection can be difficult. Fairness can be defined at the source, VC, or flow level. A flow is defined as a VC coming on an input link. All fairness definitions are equivalent for point-to-point connections, where each VC has one source and comprises one flow. For multipoint-to-point connections, a VC can have multiple sources and several flows (where each flow may contain merged traffic of several sources). Four different types of fairness can be defined for multipoint-to-point connections as follows [11]:

- *Source-based fairness:* Bandwidth is fairly divided among active sources as if they were the source in point-to-point connections, ignoring group memberships.
- *VC/source-based fairness:* Bandwidth is fairly divided among active VCs. The bandwidth of each VC is fairly allocated among active sources of that VC.
- *Flow-based fairness:* Bandwidth is fairly allocated among active flows.

- *VC/flow-based fairness:* Bandwidth is fairly divided among active VCs. The bandwidth of each VC is fairly allocated among active flows of that VC.

More work is needed to extend traffic management rules for multipoint-to-point connections. Rigorous performance analysis to evaluate complexity, fairness, transient response, delay, and scalability trade-offs is required.

MULTIPOINT-TO-MULTIPOINT CONNECTIONS

In the presence of multipoint-to-multipoint connections, switches must combine the branch point and merge point algorithm functionalities. Multipoint-to-multipoint connections are still under study. Careful performance analysis is necessary to evaluate congestion control for many-to-many VCs with different topologies, traffic characteristics, bottlenecks, and RTTs.

OPEN ISSUES

More extensive performance analysis is required for multimedia applications over ABR service. The following issues should be resolved to provide support for scalable and adaptable multimedia applications over ABR service:

- The problem of mapping the traffic characteristics of multimedia to the traffic parameters of ABR service (e.g., MCR) has to solved. Efficient connection admission policy and policing mechanisms for such traffic have to be developed.
- Support for renegotiation and heterogeneity in multipoint connections would provide a more flexible service.
- Efficient schemes for supporting multipoint-to-point and multipoint-to-multipoint connections have to be developed. The schemes should have low overhead and provide fair allocation of the available bandwidth.
- Hierarchically encoded video (multilayer) streams can provide scalable video. In the event of congestion, some of the layers need to be dropped. Currently, ATM connections support only dual priority using the cell loss priority (CLP) bit of cell header. Developing a multipriority discard scheme is necessary to support scalable and adaptable video.
- Guidelines for setting ABR source parameters to provide support for multipoint connections transporting multimedia traffic should be developed.
- The performance of multimedia multipoint connections should be comprehensively analyzed using different topologies. Scalability requirements should be analyzed by studying the performance using a large number of end systems.

SUMMARY

The following issues are discussed in this article:

- Adaptive multimedia applications can be efficiently transported over ATM ABR service. ABR service has attractive features such as low cell loss, short queuing delays, and closed-loop feedback control.
- ABR service provides MCR guarantees, which can provide an acceptable quality of service for voice and video applications. Modifications are necessary to switch algorithms which assume zero MCR values to provide nonzero MCR guarantees.
- Loss can be minimized for multimedia ABR connections by using a separate queue and reserving part of the link bandwidth for these connections. Switch queue, and hence delay, can be controlled using appropriate queue control functions to scale the available bandwidth estimate.
- Many multimedia applications are multicast. Point-to-multipoint connections need feedback consolidation algorithms at the branch points. Some of these algorithms suffer from the consolidation noise problem. Several schemes have been proposed to solve this problem. Fair bandwidth allocation for multipoint-to-point and multipoint-to-multipoint connections is currently under study.

REFERENCES

[1] T. V. Lakshman, P. P. Mishra, and K. K. Ramakrishnan, "Transporting compressed video over ATM networks with explicit rate feedback control," *Proc. INFOCOM*, 1997.

[2] H. Kanakia, P. P. Mishra, and A. Riebman, "An daptive congestion control scheme for real-time packet video transport," *Proc. ACM SIGCOMM*, pp. 20–31, Sept. 1993.

[3] W. Luo and M. Zarki, "Adapive data partitioning for MPEG-2 video transmission over ATM based networks," *Proc. IEEE Int'l. Conf. Image Processing*, vol. 1, 1995.

[4] R. Aravind, M. Civanlar and A. Reibman, "Packet loss resilience of MPEG-2 scalable video coding algorithms," *IEEE Trans. Circuits and Sys. for Video Tech.*, vol. 6, no. 5, 1996, pp. 426–35.

[5] B. J. Vickers, M. Lee, and T. Suda, "Feedback control mechanisms for real-time multipoint video services," *IEEE JSAC*, vol. 15, no. 2, Apr. 1997.

[6] N. G. Duffield, K. K. Ramakrishnan and A. R. Reibman, "SAVE: An algorithm for smoothed adaptive video over explicit rate networks," *Proc. INFOCOM*, Apr. 1998.

[7] B. Zheng and M. Atiquzzaman, "Traffic management of multimedia over ATM networks," *IEEE Commun. Mag.*, vol. 37, no. 1, Jan. 1999.

[8] B. Vandalore *et al.*, "A definition of general weighted fairness and its support in explicit rate switch algorithms," *Proc. IEEE ICNP*, Oct. 1998, pp. 22–30.

[9] B. Vandalore *et al.*, "Design and analysis of queue control functions for explicit rate switch schemes," *Proc. IEEE ICCCN*, Oct. 1998, pp. 780–86.

[10] S. Fahmy *et al.*, "Feedback consolidation algorithms for ABR point-to-multipoint connections," *Proc. IEEE INFOCOM*, vol. 3, Mar. 1998, pp. 1004–13.

[11] S. Fahmy *et al.*, "Fairness for ABR multipoint-to-point connections," ATM Forum/97-0832, Sept. 1997.

BIOGRAPHIES

BOBBY VANDALORE [StM] (vandalor@cis.ohio-state.edu) received his B.Tech degree in 1993 from the Indian Institute of Technology, Madras, and his M.S. degree in 1995 from The Ohio State University, both in computer science. He is currently a Ph.D. candidate at The Ohio State University. His main research interests are in the areas of multimedia communications, traffic management, and performance analysis. He is the author of several papers and ATM Forum contributions. He is a student member of the ACM.

SONIA FAHMY [StM] (fahmy@cis.ohio-state.edu) received her M.S. degree in computer science in 1996 from The Ohio State University, where she is currently a Ph.D. candidate. Her main research interests are in the areas of multipoint communication, traffic management, and performance analysis. She is the author of several papers and ATM Forum contributions. She is a student member of the ACM and the IEEE Computer Society.

RAJ JAIN [F] (jain@cis.ohio-state.edu) is an active member of ATM Forum Traffic Management group and has influenced its direction considerably. He is a Fellow of ACM, and serves on the editorial boards of *Computer Networks and ISDN Systems*, *Computer Communications* (U.K.), and the *Journal of High Speed Networks*. He is the author of two popular books: *FDDI Handbook: High Speed Networking Using Fiber and Other Media* (Addison-Wesley) and *The Art of Computer Systems Performance Analysis* (Wiley). His publications and ATM Forum contributions can be found at http://www.cis.ohio-state.edu/~jain.

ROHIT GOYAL (goyal@cis.ohio-state.edu) is a Ph.D.. candidate with the Department of Computer and Information Science at The Ohio State University, Columbus. He received his B.S. in computer science from Denison University, Granville. His work has been published in several papers and ATM Forum contributions. His other interests include distributed systems, artificial intelligence, and performance analysis.

MUKUL GOYAL (mukul@cis.ohio-state.edu) is currently a Ph.D. student at the CIS Deprtment of The Ohio State University. He received his B.Tech. in electronics and communication engineering from Regional Engineering College, Kurukshetra, India, in 1995. He worked as a software engineer at Siemens Communication Software Ltd., Bangalore, India, from 1995 to 1997.