Traffic Management for TCP/IP over Satellite ATM Networks

Rohit Goyal, Raj Jain, Mukul Goyal, Sonia Fahmy, and Bobby Vandalore
The Ohio State University

Sastri Kota, Lockheed Martin Telecommunications

ABSTRACT

Several Ka-band satellite systems have been proposed that will use ATM technology to seamlessly transport Internet traffic. The ATM UBR, GFR, and ABR service categories have been designed for data. However, several studies have reported poor TCP performance over satellite ATM networks. We describe techniques to improve TCP performance over satellite ATM networks. We first discuss the various design options available for TCP end systems, IP-ATM edge devices, as well as ATM switches for long-latency connections. We discuss buffer management policies, guaranteed rate services, and the virtual source/virtual destination option in ATM. We present a comparison of ATM service categories for TCP transport over satellite links. The main goal of this article is to discuss design and performance issues for the transport of TCP over UBR, GFR, and ABR services for satellite ATM networks.

Asynchronous transfer mode (ATM) technology is expected to provide quality of service (QoS)-based networks that support voice, video, and data applications. ATM was originally designed for fiber-based terrestrial networks that exhibit low latencies and error rates. With the increasing demand for electronic connectivity across the world, satellite networks play an indispensable role in the deployment of global networks. Ka-band satellites using the gigahertz frequency spectrum can reach user terminals across most of the populated world. ATM-based satellite networks can effectively provide real-time as well as non-real-time communications services to remote areas.

However, satellite systems have several inherent constraints. The resources of the satellite communication network, especially the satellite and the Earth station, are expensive and typically have low redundancy; these must be robust and be used efficiently. The large delays in geostationary Earth orbit (GEO) systems and delay variations in low Earth orbit (LEO) systems affect both real-time and non-real-time applications. In an acknowledgment- and timeout-based congestion control mechanism (like TCP), performance is inherently related to the delay-bandwidth product of the connection. Moreover, TCP round-trip time (RTT) measurements are sensitive to delay variations that may cause false timeouts and retransmissions. As a result, the congestion control issues for broadband satellite networks are somewhat different from those of low-latency terrestrial networks. Both interoperability issues as well as performance issues need to be addressed before a transport-layer protocol like TCP can satisfactorily work over long-latency satellite ATM networks.

In this article we describe the various design options for improving the performance of TCP/IP over satellite ATM networks. The next section describes the ATM service categories and options available to TCP/IP traffic. We then describe each ATM design option as well as the TCP mechanism, and evaluate their performance over satellite networks. We conclude with a comparison of ATM service categories for TCP transport over satellite links.

DESIGN ISSUES FOR TCP/IP OVER SATELLITE ATM

Satellite ATM networks can be used to provide broadband access to remote locations, as well as to serve as an alternative to fiber-based backbone networks. In either case, a single satellite is designed to support thousands of Earth terminals. The Earth terminals set up virtual channels (VCs) through the onboard satellite switches to transfer ATM cells among one another. Because of the limited capacity of a satellite switch, each Earth terminal has a limited number of VCs it can use for TCP/IP data transport. In backbone networks, these Earth terminals are IP-ATM edge devices that terminate ATM connections, and route IP traffic in and out of the ATM network. These high-capacity backbone routers must handle thousands of simultaneous IP flows. As a result, the routers must be able to aggregate multiple IP flows onto individual VCs. Flow classification may be done by means of a *QoS manager* that can use IP source–destination address pairs, as well as transport-layer port numbers.[1] The QoS manager can further classify IP packets into flows based on the differentiated services code points in the TOS byte of the IP header.

In addition to flow and VC management, Earth terminals also provide a means for congestion control between the IP and ATM networks. The onboard ATM switches must perform traffic management at the cell and VC levels. In addition, TCP hosts implement various TCP flow and congestion control mechanisms for effective network bandwidth utilization. Figure 1 illustrates a framework for the various design options available to networks and TCP hosts for congestion control. The techniques in the figure can be used to implement various ATM services in the network. Enhancements that perform intelligent buffer management policies at the switches can be developed for unspecified bit rate (UBR) to improve transport-layer throughput and fairness. A policy for selective cell drop based on per-VC accounting can be used to improve fairness.

Providing a minimum guaranteed rate (GR) to UBR traffic has been discussed as a possible candidate to improve TCP performance over UBR. The goal of providing GR is to protect the UBR service category from total bandwidth starvation, and provide a continuous minimum bandwidth guarantee. It has been shown that in the presence of high load of higher-priority constant bit rate (CBR), variable bit rate (VBR), and available bit rate (ABR) traffic, TCP congestion control mechanisms benefit from a minimum GR.

---

[1] *TCP/UDP port numbers are accessible only if end-to-end security protocols are not used.*

Guaranteed frame rate (GFR) has recently been proposed in the ATM Forum as an enhancement to the UBR service category. GFR will provide a minimum rate guarantee to VCs at the frame level. The GFR service also allows for the fair usage of any extra network bandwidth. GFR is likely to be used by applications that can neither specify the traffic parameters needed for a VBR VC, nor have capability for ABR (for rate-based feedback control). Current internetworking applications fall into this category, and are not designed to run over QoS-based networks. Routers separated by satellite ATM networks can use the GFR service to establish VCs between one another. GFR can be implemented using per-VC queuing or buffer management.

The ABR service category is another option to implement TCP/IP over ATM. This service category is specified by a peak cell rate (PCR) and a minimum cell rate (MCR) which is guaranteed by the network. ABR connections use a rate-based closed-loop end-to-end feedback control mechanism for congestion control. The network tries to maintain a low cell loss ratio by changing the allowed cell rate (ACR) at which a source can send. Switches can also use the virtual source/virtual destination (VS/VD) feature to segment the ABR control loop into smaller loops. Studies have indicated that ABR with VS/VD can effectively reduce the buffer requirement for TCP over ATM, especially for long delay paths. ABR can be implemented using the feedback control mechanisms in Fig. 1.

In addition to network-based drop policies, end-to-end flow control and congestion control policies can be effective in improving TCP performance over UBR. The fast retransmit and recovery mechanism can be used in addition to slow start and congestion avoidance to quickly recover from isolated segment losses. The selective acknowledgment (SACK) option has been proposed to recover quickly from multiple segment losses. A change to TCP's fast retransmit and recovery has been suggested in [1]. The use of performance-enhancing TCP gateways to improve performance over satellite links has also been proposed in recent studies. The following sections discuss the design and performance issues for TCP over UBR, GFR, and ABR services for satellite networks.

TCP OVER UBR

In its simplest form, an ATM switch implements a tail drop policy for the UBR service category. If cells are dropped, the TCP source loses time waiting for the retransmission timeout. Even though TCP congestion mechanisms effectively recover from loss, the link efficiency can be very low, especially for large delay-bandwidth networks. In general, link efficiency typically increases with increasing buffer size. Performance of TCP over UBR can be improved using buffer management policies. In addition, TCP performance is also affected by TCP congestion control mechanisms and TCP parameters such as segment size, timer granularity, receiver window size, slow start threshold, and initial window size.

TCP Reno implements the fast retransmit and recovery algorithms that enable the connection to quickly recover from isolated segment losses. However, fast retransmit and recovery cannot efficiently recover from multiple packet losses within the same window. A modification to Reno is proposed in [1] so that the sender can recover from multiple packet losses without having to time out.
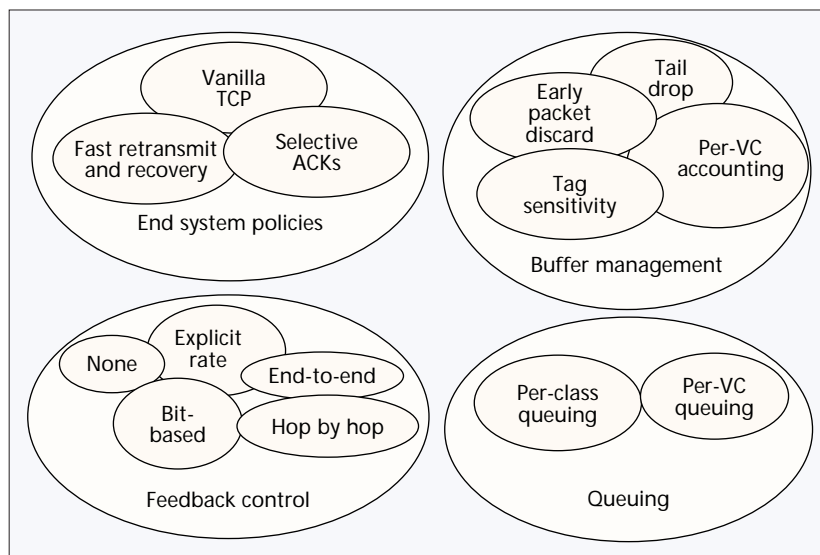
■ **Figure 1.** *Design issues for TCP over ATM.*

TCP with selective acknowledgments (SACK TCP) is designed to efficiently recover from multiple segment losses. With SACK, the sender can recover from multiple dropped segments in about one round-trip. Comparisons of TCP and drop policies for persistent traffic over satellite ATM are presented in [2]. The studies show that in low-delay networks, the effect of network-based buffer management policies is very important and can dominate the effect of SACK. The throughput improvement provided by SACK is very significant for long-latency connections. When the propagation delay is large, timeout results in the loss of a significant amount of time during slow start from a window of one segment. Reno TCP (with fast retransmit and recovery) results in the worst performance (for multiple packet losses) because timeout occurs at a much lower window than in Vanilla TCP. With SACK TCP, a timeout is avoided most of the time, and recovery is complete within a small number of round-trips. For lower-delay satellite networks (LEOs), both NewReno and SACK TCPs provide high throughput, but as the latency increases, SACK significantly outperforms NewReno, Reno, and Vanilla.[2]

UBR+: ENHANCEMENTS TO UBR

Recent research has focused on fair buffer management for best-effort network traffic. In these proposals, packets are dropped when the buffer occupancy exceeds a certain threshold. Most buffer management schemes improve the efficiency of TCP over UBR. However, only some of the schemes affect the fairness properties of TCP over UBR. The proposals for buffer management can be classified into four groups based on whether they maintain multiple buffer occupancies (multiple accounting, MA) or a single global buffer occupancy (single accounting, SA), and whether they use multiple discard thresholds (multiple thresholds, MT) or a single global discard threshold (single threshold, ST). Table 1 lists the four classes of buffer management schemes and examples of schemes for these classes. The schemes are briefly discussed below.

The SA schemes maintain a single count of the number of cells currently in the buffer. The MA schemes classify the traffic into several classes and maintain a separate count for the number of cells in the buffer for each class. Typically, each class corresponds to a single connection, and these schemes maintain

---

[2] *Vanilla TCP refers to TCP with only slow start and congestion avoidance.*

per-connection occupancies. In cases where the number of connections far exceeds the buffer size, the added overhead of per-connection accounting may be very expensive. In this case, a set of active connections can be defined as those connections with at least one cell in the buffer, and only the buffer occupancies of active connections need to be maintained.

ST schemes compare the buffer occupancy(ies) with a single threshold and drop packets when the buffer occupancy exceeds the threshold. Multiple thresholds can be maintained corresponding to classes or connections, or to provide differentiated services. Several modifications to this drop behavior can be implemented, including averaging buffer occupancies, static versus dynamic thresholds, deterministic versus probabilistic discards, and discard levels based on packet tags. Examples of packet tags are the cell loss priority (CLP) bit in ATM cells or the TOS octet in the IP header of the Internet Engineering Task Force's (IETF's) differentiated services architecture.

The SA-ST schemes include early packet discard (EPD), partial packet discard (PPD), and random early detection (RED). EPD and PPD improve network efficiency because they minimize the transmission of partial packets by the network. Since they do not discriminate between connections in dropping packets, these schemes are unfair in allocating bandwidth to competing connections [3]. RED maintains a global threshold for the average queue. When the average queue exceeds this threshold, RED drops packets probabilistically.

However, it has been shown in [4] that RED cannot always guarantee equal bandwidth sharing. The article also contains a proposal for flow random early drop (FRED). FRED maintains per-connection buffer occupancies and drops packets probabilistically if the per-connection occupancy exceeds the average queue length. In addition, FRED ensures that each connection has at least a minimum number of packets in the queue. FRED can be classified as one that maintains per-connection queue lengths, but has a global threshold (MA–ST).

The selective drop (SD) [2] and fair buffer allocation (FBA) schemes are MA–ST schemes proposed for the ATM UBR service category. These schemes use per-connection accounting to maintain the current buffer utilization of each UBR VC. A fair allocation is calculated for each VC, and during congestion (indicated when the total buffer occupancy exceeds a threshold), if the VC's buffer occupancy exceeds its fair allocation, its subsequent incoming packet is dropped. Both SD and FBA improve both the fairness and efficiency of TCP over UBR. This is because cells from overloading connections are dropped in preference to underloading ones.

The virtual queuing (VQ) [5] scheme achieves equal buffer allocation by emulating on a single FIFO queue, a per-VC queued round-robin server. At each cell transmit time, a per-VC variable (γ_i) is decremented in a round-robin manner, and is incremented whenever a cell of that VC is admitted in the buffer. When γ_i exceeds a fixed threshold, incoming packets of the ith VC are dropped. An enhancement called *dynamic EPD* changes the above drop threshold to include only those sessions that are sending less than their fair shares.

Since the above MA–ST schemes compare the per-connection queue lengths (or virtual variables with equal weights) with a global threshold, they can only guarantee equal buffer occupancy (and thus throughput) to the competing connections. These schemes do not allow for specifying a guaranteed rate for connections or groups of connections. Moreover, in their present forms, they cannot support packet discard levels based on tagging.

Another enhancement to VQ, called MCR scheduling [6], proposes the emulation of a weighted scheduler to provide MCR guarantees to ATM connections. In this scheme, a per-VC weighted variable (W_i) is updated in proportion to the VC's MCR, and compared with a global threshold. Reference [7] proposes a combination of a packet marking engine (PME) and an enhanced RED scheme based on per-connection accounting and multiple thresholds (MA–MT). PME+ERED is designed for IP networks, and can provide loose rate guarantees to connections. The PME measures per-connection bandwidths and probabilistically marks packets if the measured bandwidths are lower than the target bandwidths (MT). High-priority packets are marked, low-priority packets unmarked. The ERED mechanism is similar to RED except that the probability of discarding marked packets is lower that that of discarding unmarked packets.

The DFBA scheme [8] proposed for the ATM GFR service provides MCR guarantees for VCs carrying multiple TCP connections. DFBA maintains high and low target buffer occupancy levels for each VC, and performs probabilistic drop based on a VC's buffer occupancy and its target thresholds. The scheme gives priority to CLP = 0 cells over CLP = 1 cells.

A simple SA–MT scheme can be designed that implements multiple thresholds based on packet discard levels. When the global queue length (SA) exceeds the first threshold, packets with the lowest discard level are dropped. When the queue length exceeds the next threshold, packets from the lowest and the next discard level are dropped. This process continues until EPD/PPD is performed on all packets.

Buffer management class	Examples	Threshold type	Drop type	Tag-sensitive?	Fairness
SA–ST	EPD, PPD	Static	Deterministic	No	None
	RED	Static	Probabilistic	No	Equal allocation in limited cases
MA–ST	FRED	Dynamic	Probabilistic	No	Equal allocation
	SD, FBA	Dynamic	Deterministic	No	Equal allocation
	VQ+Dynamic EPD	Dynamic	Deterministic	No	Equal allocation
MA–MT	PME+ERED	Static	Probabilistic	Yes	MCR guarantee
	DFBA	Dynamic	Probabilistic	Yes	MCR guarantee
	VQ+MCR scheduling	Dynamic	Deterministic	No	MCR guarantee
SA–MT	Priority drop	Static	Deterministic	Yes	—

■ Table 1. *Classification of buffer managment schemes.*

As discussed in the previous section, for satellite ATM networks TCP congestion control mechanisms have more effect on TCP throughput than ATM buffer management policies. However, these drop policies are necessary to provide fair allocation of link capacity, to provide differentiated services based on discard levels, and to provide MCR guarantees to low-priority VCs. The GFR service described in the next section makes extensive use of the intelligent buffer management policies described here.

GUARANTEED FRAME RATE

GFR service guarantee requires the specification of an MCR and a maximum frame size (MFS) for each VC. If the user sends packets (or frames) of size at most MFS at a rate less than the MCR, all the packets are expected to be delivered by the network with low loss. If the user sends packets at a rate higher than the MCR, it should still receive at least the minimum rate. The minimum rate is guaranteed to the untagged (CLP = 0) frames of the connection. In addition, a connection sending in excess of the minimum rate should receive a fair share of any unused network capacity. The exact specification of the fair share has been left unspecified by the ATM Forum.

There are three basic design options that can be used by the network to provide the per-VC minimum rate guarantees for GFR; these are discussed below.

TAGGING

Network-based tagging (or policing) can be used as a means of marking nonconforming packets before they enter the network. This form of tagging is usually performed when the connection enters the network. Figure 2 shows the role of network-based tagging in providing a minimum-rate service in a network. Network-based tagging on a per-VC level requires some per-VC state information to be maintained by the network and increases the complexity of the network element. Tagging can isolate conforming and nonconforming traffic of each VC so that other rate-enforcing mechanisms can use this information to schedule the conforming traffic in preference to nonconforming traffic.

BUFFER MANAGEMENT

Buffer management is typically performed by a network element (e.g., a switch or a router) to control the number of packets entering its buffers. In a shared buffer environment, where multiple VCs share common buffer space, per-VC buffer management can control the buffer occupancies of individual VCs. Figure 2 shows the role of buffer management in the connection path. The DFBA scheme can be used by the onboard ATM network to provide minimum-rate guarantees to GFR VCs.

SCHEDULING

Figure 2 illustrates the position of scheduling in providing rate guarantees. While tagging and buffer management control the entry of packets into a network element, queuing strategies determine how packets are scheduled onto the next hop. FIFO queuing cannot isolate packets from various VCs (or groups of VCs) at the egress of the queue. Per-VC queuing, on the other hand, maintains a separate queue for each VC (or group of VCs) in the buffer. A scheduling mechanism can select between the queues at each scheduling time. However, scheduling adds the cost of per-VC queuing and the service discipline. For a simple service like GFR, this additional cost may be undesirable for an onboard switch.
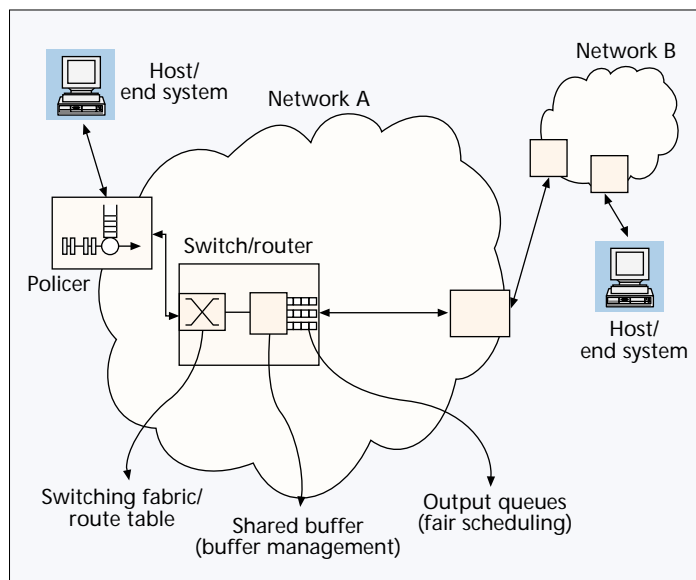
■ **Figure 2**. *Buffering, scheduling, and policing in the network.*

ABR OVER SATELLITE

Reference [9] provides a comprehensive study of TCP performance over ABR service. We discuss a key ABR feature, VS/VD, and highlight its relevance to long delay paths. Most of the discussion assumes that the switches implement a rate-based switch algorithm like ERICA+. Credit-based congestion control for satellite networks has also been suggested. However, in this article we focus on rate-based control as specified in the ATM standards.

In long-latency satellite configurations, the feedback delay is the dominant factor in determining the maximum queue length. A feedback delay of 10 ms corresponds to about 3670 cells (at OC-3) of queue for TCP over ERICA+, while a feedback delay of 550 ms corresponds to 201,850 cells. This indicates that satellite switches need to provide at least one feedback delay worth of buffering to avoid loss on these high-delay paths. A point to consider is that these large queues should not be seen in downstream workgroup or WAN switches, because they will not provide as much buffering. Satellite switches can isolate downstream switches from such large queues by implementing the VS/VD option.

Reference [10] has examined some basic issues in designing VS/VD feedback control mechanisms. VS/VD can effectively isolate nodes in different VS/VD loops. As a result, the buffer requirements of a node are bound by the feedback delay-bandwidth product of the upstream VS/VD loop. VS/VD helps to reduce the buffer requirements of terrestrial switches connected to satellite gateways. Figure 3 illustrates the results of a simulation experiment showing the effect of VS/VD on the buffer requirements of terrestrial switch S. In the figure the link between S and the end host is the bottleneck link. The feedback delay-bandwidth product of the satellite hop is about 160,000 cells, and dominates the feedback delay-bandwidth product of the terrestrial hop (about 3000 cells). Without VS/VD, terrestrial switch S, a bottleneck, must buffer cells of up to the feedback delay-bandwidth product of the entire control loop (including the satellite hop). With a VS/VD loop between the satellite and the terrestrial switch, the queue accumulation due to the satellite feedback delay is confined to the satellite switch. The terrestrial switch only buffers cells that are accumulated due to the feedback delay of the terrestrial link to the satellite switch.

A Comparison of ATM Service Categories

Existing and proposed ATM standards provide several options for TCP/IP data transport over a satellite ATM network. The three service categories — ABR, UBR, and GFR — and their various implementation options present a cost-performance trade-off for TCP/IP over ATM. A comparison of the service categories can be based on the following factors:

• Implementation complexity
• Buffering requirements for switches and ATM end systems
• Network bandwidth utilization
• Bandwidth allocation (fairness and MCR guarantees)

Higher complexity arises from resource allocation algorithms for connection admission control (CAC) and usage parameter control (UPC), as well as from sophisticated queuing and feedback control mechanisms. While UPC is performed at the entrance of the ATM network to control the rate of packets entering the network, CAC is performed during connection establishment by each network element. UBR is the least complex service category because it does not require any CAC or UPC. Typical UBR switches are expected to have a single queue for all UBR VCs. Buffer management in switches can vary from a simple tail drop to the more complex per-VC accounting-based algorithms such as FBA. An MCR guarantee to the UBR service would require a scheduling algorithm that prevents the starvation of the UBR queue. The GFR service could be implemented by either a single queue using a DFBA like mechanism or per-VC queues and scheduling. The ABR service can be implemented with a single ABR queue in the switch. The VS/VD option requires the use of per-VC queu-

ing and increases the implementation complexity of ABR. The CAC requirements for GFR and ABR are similar. However, the tagging option, CLP conformance and MFS conformance tests in GFR add complexity to the UPC function.

The additional complexity in ABR feedback control presents a trade-off with ABR buffer requirements. Network buffering is lower for ABR than for UBR or GFR. In addition, ABR has controlled buffer requirements that depend on the bandwidth-delay product of the ABR feedback loop. At the edge of the ATM network, network feedback can provide information for buffer dimensioning. Large buffers in edge routers can be used when the ABR network is temporarily congested. In the case of UBR and GFR, edge devices do not have network congestion information and simply send the data into the ATM network as fast as they can. As a result, extra buffers at the edge of the network do not help for UBR or GFR. This is an important consideration for large-delay-bandwidth satellite networks. With ABR, satellite gateways (routers at the edges of a satellite ATM network) can buffer large amounts of data, while the buffer requirements of the onboard ATM switches can be minimized. The buffer requirements with UBR/GFR are reversed for the gateways and onboard switches.

The ABR service can make effective use of available network capacity by providing feedback to the sources. Edge devices with buffered data can fill up the bandwidth within one feedback cycle of the bandwidth becoming available. This feedback cycle is large for satellite networks. With UBR and GFR, available bandwidth can be filled up immediately by edge devices that buffer data. However, the edge devices have no control of sending rate, and data is likely to be dropped
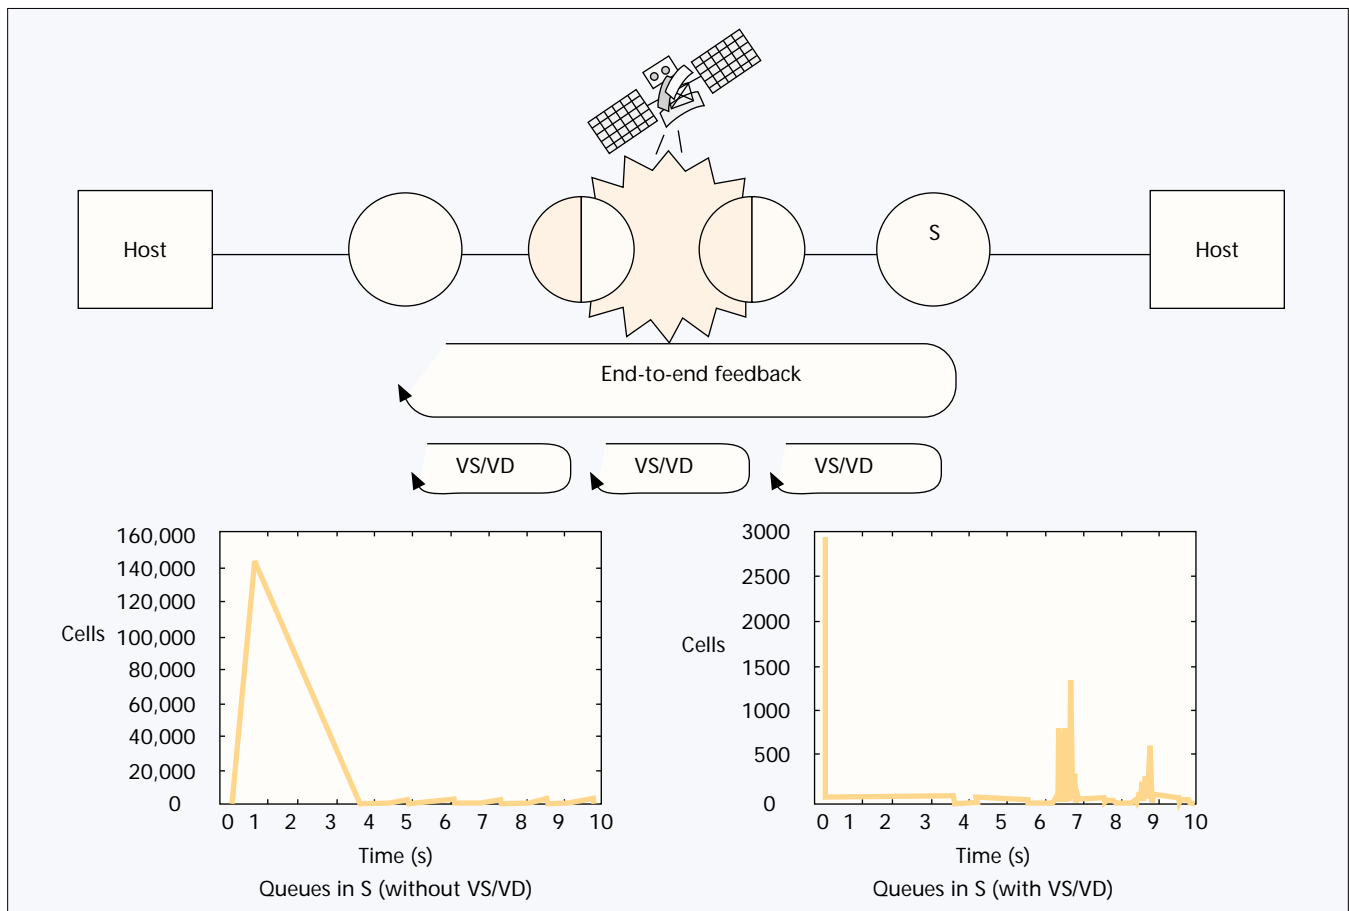
■ Figure 3. *The VS/BD option in ATM-ABR.*

during congestion. This data must be retransmitted by TCP, and can result in inefficient use of the satellite capacity.

In addition to efficient network utilization, a satellite ATM network must also fairly allocate network bandwidth to the competing VCs. While Vanilla UBR has no mechanism for fair bandwidth allocation, UBR or GFR with buffer management can provide per-VC fairness. ABR provides fairness by per-VC rate allocation. A typical satellite ATM network will carry multiple TCP connections over a single VC. In ABR, most losses are in the routers at the edges of the network, and there routers can perform fair buffer management to ensure IP-level fairness. In UBR and GFR, on the other hand, most losses due to congestion are in the satellite ATM network, where there is no knowledge of the individual IP flows. In this case, fairness can only be provided at the VC level.

CONCLUDING REMARKS

Several issues arise in optimizing the performance of TCP when ATM is deployed over satellite links. This article emphasizes that both TCP mechanisms as well as ATM mechanisms should be used to improve TCP performance over long-delay ATM networks. ATM technology provides at least three service categories for data: UBR, ABR, and GFR. Each of these can be implemented by a number of mechanisms, including:

• UBR with intelligent buffer management
• UBR with guaranteed rate
• ABR with network feedback
• ABR with virtual source/virtual destination

In addition, TCP provides several congestion control mechanisms, including:

• Vanilla TCP with slow start and congestion avoidance
• TCP Reno with fast retransmit and recovery
• TCP New Reno
• TCP with selective acknowledgments (SACK)

It has been shown that Vanilla TCP over the UBR service category achieves low throughput and high unfairness over satellite networks. This is because during packet loss, TCP loses time waiting for its coarse granularity retransmission timeout. In the presence of bursty packet losses, fast retransmit and recovery (without SACK) further hurts TCP performance over UBR for long delay-bandwidth product networks.

Frame-level discard policies such as EPD improve throughput significantly over cell-level discard policies. However, fairness is not guaranteed unless intelligent buffer management using per-VC accounting is used. Throughput increases further with more aggressive New Reno and SACK. SACK gives the best performance in terms of throughput. It has been found that for long delay paths, the throughput improvement due to SACK is more than that from discard policies and buffer management. Using guaranteed rates (GR, GFR) helps in the presence of high loads of higher-priority traffic such as CBR or VBR traffic.

For TCP over ABR, VS/VD can be used to isolate long-delay segments from terrestrial segments, which helps in efficiently sizing buffers in routers and ATM switches. As a result, terrestrial switches only need to have buffers proportional to the bandwidth-delay products of the terrestrial segment of the TCP path. Switches connected to the satellite VS/VD loops must have buffers proportional to the satellite delay-bandwidth products.

REFERENCES

[1] J. C. Hoe, "Improving the Start-up Behavior of a Congestion Control Scheme for TCP,'' *Proc. SIGCOMM '96*, Aug. 1996.
[2] R. Goyal *et al.*, "TCP Selective Acknowledgments and UBR Drop Policies to Improve ATM-UBR Performance over Terrestrial and Satellite Net-
works," *Proc. ICCCN '97*, Las Vegas, Sept. 1997, pp. 17–27; http://www.cis.ohio-state.edu/~jain/papers/ic3n97.htm.
[3] H. Li *et al.*, "TCP over ABR and UBR Services in ATM,'' *Proc. IPCCC '96*, Mar. 1996.
[4] D. Lin and R. Morris, "Dynamics of Random Early Detection," *Proc. SIGCOMM '97*, 1997.
[5] Y. Wu, K.-Y. Siu, and W. Ren, "Improved Virtual Queuing and Dynamic EPD Techniques for TCP over ATM," *Proc. ICNP '97*, 1997.
[6] K.-Y. Siu, Y. Wu, and W. Ren, "Virtual Queuing Techniques for UBR+ Service in ATM with Fair Access and Minimum Bandwidth Guarantee," *Proc. GLOBECOM '97*, 1997.
[7] W. Feng *et al.*, "Techniques for Eliminating Packet Loss in Congested TCP/IP Networks," Univ. of MI, CSE-TR-349-97, Nov. 1997.
[8] Rohit Goyal *et al.*, "Buffer Management for the GFR Service," ATM_Forum/98-0405, July 1998; http://www.cis.ohio-state.edu/~jain/atmf/a98-0405.htm
[9] S. Kalyanaraman, "Traffic Management for the Available Bit Rate (ABR) Service in Asynchronous Transfer Mode (ATM) Networks," Ph.D. dissertation, Ohio State Univ., 1997.
[10] R. Goyal *et al.*, "Per-VC rate allocation techniques for ABR feedback in VS/VD networks, *Proc. GLOBECOM '98*, Nov.1998; http://www.cis.ohio-state.edu/~jain/papers/globecom98.htm

BIOGRAPHIES

ROHIT GOYAL (goyal@cis.ohio-state.edu) is a Ph.D. candidate in computer and information science at The Ohio State University. His research areas include traffic management, quality of service, and performance analysis for high-speed networks. He is an active participant in the ATM Forum, IETF, and TIA, and has published several conference and journal papers. He received a B.S. in computer science from Denison University, Granville, Ohio, and an M.S. in computer and information science from The Ohio State University.

RAJ JAIN [F] is a professor of computer and information science at The Ohio State University. He is an active member of the 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* (United Kingdom), 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 and can be found at http://www.cis.ohio-state.edu/~jain.

MUKUL GOYAL is currently a Ph.D. student in the CIS Department at 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 in Siemens Communication Software Ltd., Bangalore, India, from 1995 to1997.

SONIA FAHMY [StM] 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.

BOBBY VANDALORE [StM] received his B.Tech. degree in 1993 from 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.

SASTRI KOTA [SM '86] (sastri.kota@lmco.com) is a technical consultant with Lockheed Martin Mission Systems in the areas of broadband networks, commercial and military satellite systems architectures, and standards and regulatory affairs. He has over 25 years' experience in network systems design and analysis, and has held technical and project management positions at SRI International, Ford Aerospace, XEROX, MITRE, and Computer Sciences Corporation. Currently he is very active in the telecommunications and networks standards development organizations and is an ATM Forum ambassador for wireless ATM and the ITU-R Working Party 4B Rapporteur for the new Recommendation "Performance of Satellite Systems Operating over 15 GHz." He has served as technical cochair and assistant technical chair for MILCOM '97 and '90, conference chair for SPIE '90 and '91; and session organizer and chair for MILCOM '97, '96, '95, '94, '93, '92, '90, and '86. He holds a B.Sc. in physics from Andhra University, a B.E. in telecommunications from Birla Institute of Science and Technology, Pilani, an M.E in electronics and communications from the University of Roorkee, and an electrical engineer's degree from Northeastern University, Boston, Massachusetts. He is a senior member of AIAA, member of ACM and SPIE. His research interests include broadband networks, wireless ATM, performance modeling and simulations, multiple access, traffic management, and QoS for satellite ATM networks.