Performance Evaluation of TCP Extensions on ATM over High Bandwidth Delay Product Networks

Charalambos P. Charalambous, Sprint

Victor S. Frost and Joseph B. Evans, The University of Kansas

ABSTRACT Practical experiments in a satellite network environment assist in the design and understanding of future global networks. This article describes the practical experiences gained from TCP/IP on ATM networks over a high-speed satellite link and presents performance comparison studies of such networks with the same host/traffic configurations over local area and wide area networks. These comparison studies on the LAN, WAN, and satellite environments increase our understanding of the behavior of high-bandwidth networks. NASA's Advanced Communications Technology Satellite (ACTS), with its special characteristics and high data rate satellite channels, and the ACTS ATM Internetwork (AAI) were used in these experiments to deliver broadband traffic. Network performance tests were carried out using application-level software (Netspec) on SONET OC-3 (155.52 Mb/s) satellite links. Finally, in this article we experimentally study the performance, efficiency, fairness, and aggressiveness of TCP Reno, TCP New Reno, and TCP SACK end hosts on ATM networks over high BDP networks.

Communication satellites with their broadcast characteristics provide an effective and useful platform for establishing links between areas that are inaccessible by terrestrial communication facilities. Advanced satellite systems may compete under certain circumstances with terrestrial fiber optic networks in terms of high transfer rates and very low bit error rates (BERs), and will become significant players in the global information infrastructure (GII) in the future. Asynchronous Transfer Mode (ATM) cell switching technology offers the flexibility to handle and transfer advanced broadband services (voice, data, video) and integration of those in the same network at high transfer rates. The combination of ATM and satellite technologies with the widely used Transmission Control Protocol/Internet Protocol (TCP/IP) suite forms an internetwork architecture that has the potential to provide seamless networking. Such networks will enable the transfer of data and multimedia communications on the same network, everywhere in the world. The well-known drawbacks of this architecture are that in geosynchronous earth orbit (GEO) satellite systems, performance is affected by the inherent latency due to the speed of light impairment and the distance of the satellite from the earth's surface, as well as the probability of bit errors on the satellite links.

This article presents results from performance experiments conducted as a part of the Advanced Communications Technology Satellite (ACTS) ATM Internetworking (AAI) project [1]. The goal of this study was to experimentally test and evaluate the performance of different TCP implementations, which include the necessary extensions for high-performance networking [2], over high bandwidth delay product (BDP) channels, including high-data-rate (HDR) geosynchronous earth orbit (GEO) satellite channels. For a review of TCP over ATM/satellite environments the reader is referred to [3]. The latter environment represents a typical high BDP network infrastructure. The AAI testbed [1] provided the wide-area ATM connectivity for this study. The AAI testbed includes several Department of Defense (DoD) high-performance computing centers, and connections to the MAGIC and ATDnet gigabit testbeds. Our experiments focus on performance measurements on OC-3 and OC-12 ATM using TCP/IP hosts with the necessary extensions for high performance [2, 4, 5] over LAN, WAN, and ACTS HDR channels. The three TCP implementations evaluated are TCP Reno [5], TCP New Reno [6], and TCP Selective Acknowledgments (SACK) [4].

Even though the experiments reported here are for a satellite system, these results are representative of expected performance for future high-speed terrestrial networks. This statement is valid for our work, since the satellite links in our testbeds exhibited very low BER (10^{-12}, 10^{-13}) similar to that of fiber networks. Therefore, all losses observed in these experiments were due to network congestion; hence, loss patterns can be considered similar to those of terrestrial high-speed fiber networks. Host and switch interface rates are rapidly increasing; OC-48c (2.4 Gb/s) interfaces are available today. The future will bring even higher rates. Thus, results from OC-3c networks over satellite channels can be used as an indicator of the performance of such future high BDP terrestrial systems.

BANDWIDTH DELAY PRODUCT

Two fundamental properties of networks are:

- Delay, the two-way latency for information to propagate from the sending node to the receiving node and commonly called *round-trip time* (RTT)
- Bandwidth, the number of bits that can be transmitted in a certain period of time

BDP is the product of the above two characteristics (i.e., the number of bits the network can hold). In transport/data link layers, the BDP represents the maximum amount of allowed unacknowledged data outstanding at any moment on the network, keeping the link or pipeline full. The BDP is also an indication of the upper limit of the buffering requirements the end hosts must have in order to be able to obtain optimal performance.

TRANSMISSION CONTROL PROTOCOL

TCP is a sophisticated transport protocol that offers connection-oriented and reliable byte stream service. It is also

This research is partially supported by DARPA under prime contract DABT63-94-C-0068.

extremely flexible in that almost any underlying technology can be used to transfer TCP/IP traffic [7], such as ATM, which is used in the network architecture considered here. It is an end-to-end protocol with error, flow, and congestion control functions. Detailed information on TCP can be found in [7, 8]. Its basic implementation [9] is unsuitable for high-speed and high-delay networks; therefore, modifications and additions were added to enhance the performance of the protocol over such networks [2, 4, 5]. A major requirement is the support of large flow control window sizes. This is one of the reasons for the variety of TCP versions available today. Here, three TCP protocols are considered.

TCP Reno

The major characteristics of this transport protocol are slow start, congestion avoidance, fast retransmit, fast recovery [5], and support for large windows [2] and delayed acknowledgments [10].

TCP New Reno

TCP New Reno was proposed by [6] and requires changes only to the sender-side protocol suite. This modification attempts to overcome TCP Reno's fast retransmit and fast recovery problems to recover from multiple segment losses within one window of transmitted data. Studies on TCP Reno [11, 12] have shown that when more than one loss occurs within one window of data, fast retransmit and fast recovery will be triggered several times in one RTT, resulting in reduction of the congestion window several times, and then linear growth or, eventually, slow start. This leads to throughput reduction. TCP New Reno avoids the above pitfall by extending the fast retransmit phase. In case of only one segment loss, TCP New Reno acts exactly the same as TCP Reno.

TCP SACK

TCP SACK is specified in [4], and requires changes to both the sender and receiver protocol suites. It was designed to achieve better channel utilization than earlier TCP protocols since it is able to recover from multiple segment losses in one RTT. The receiver also provides enough information about segment losses in a window of data; therefore, the sender knows exactly which segments were lost, and retransmits them in one RTT. SACK compatibility is verified during TCP three-way handshake connection establishment with SYN segments that contain the Sack-Permitted option. TCP SACK still contains the congestion principles and mechanisms of TCP in order not to be aggressive and degrade the network performance. Once more, in cases of only one segment loss, the reaction of TCP SACK to congestion is the same as that of TCP Reno.

TCP/IP on ATM Networks

ATM is a scalable cell switching and multiplexing technology that was chosen by the International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) to be the transport technology for the Broadband Integrated Services Digital Network (B-ISDN). Studies [1, 13] have shown that TCP with the necessary extensions for high performance [2, 5] can achieve maximum performance over ATM. The maximum obtainable user bandwidth of TCP over OC-3 ATM is about 134 Mb/s [1, 14]

In our implementations, ATM is transported by Syn-

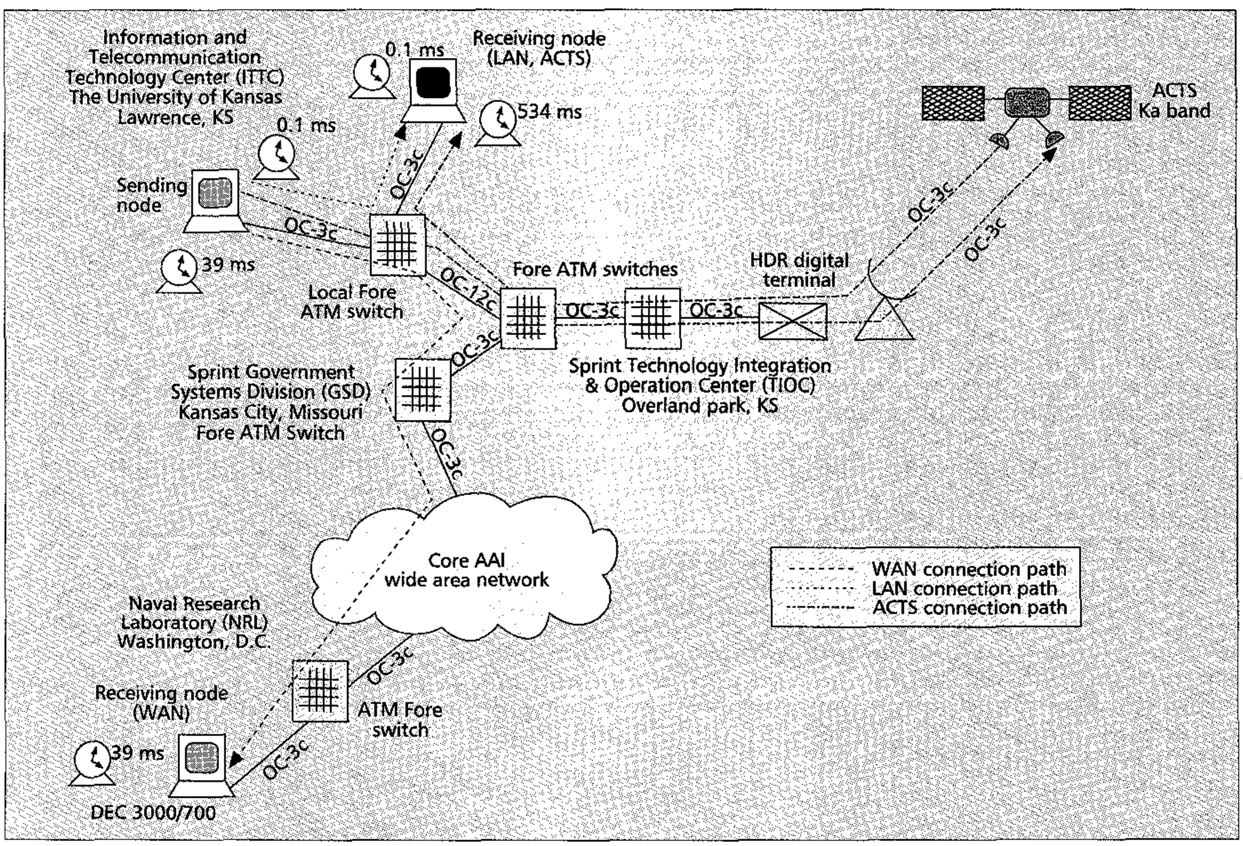

■ **Figure 1.** *Congestion-free configuration of TCP/IP on ATM over KU LAN, AAI WAN, and ACTS.*

chronous Optical Network (SONET) protocols permitting irregular ATM cell arrivals and transporting them via its Synchronous Payload Envelope (SPE). We are also using ATM adaptation layer 5 (AAL 5) in our network architecture, which offers unreliable data transfer services with error detection. All experiments were run over ATM unspecified bit rate (UBR) services. When ATM UBR services are used for a connection, the only operation the switch performs is to drop cells when there is buffer overflow, according to a specific drop policy [12]; thus, all the congestion control is done by the higher protocols. In our case, control is done by TCP.

[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Throughput (average Mb/s)	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Throughput (maximum Mb/s)	126.955	126.109	119.000	129.694	123.272	129.763	126.757
Standard dev. (Mb/s)	3.167	3.567	5.544	2.079	4.189	2.359	4.584

Table 1. *Throughput obtained by TCP extensions on a congestion-free ATM network over LAN, WAN, and ACTS for 10 replications.*

The Satellite System: ACTS

ATM transported over SONET systems is rapidly emerging as the transport mechanism for future high-speed networks [13]. Broadband communication satellite systems provide an effective platform for worldwide communications. NASA's ACTS is an experimental satellite system providing SONET STS-3 (155.52 Mb/s) and SONET STS-12 (622.08 Mb/s) point-to-point and point-to-multipoint services [15]. All the ground earth stations (GES) are equipped with standard SONET OC-3/3c and SONET OC-12/12c interfaces to ensure the interoperability of the satellite network with the terrestrial network. More details on ACTS and GES architectures can be found in [13–15].

System Overview and Tools

The workstations used in our satellite experiments were DEC Alpha 3000/700 with a 225 MHz clock; DEC Alpha 3000/400 with 133 MHz clock; SUN UltraSPARC 1 with 167 MHz clock; and SUN SPARC 20 with 125 MHz clock. Their operating systems supported the TCP Reno, TCP New Reno, and TCP SACK extensions. All workstations had sufficient buffers to accommodate the BDP of the connection and their memory bandwidth, that is, how fast the microprocessor can read/write from/to the main memory, was found to be greater than 400 Mb/s. Thus, theoretically there should not be any problem achieving OC-3c SONET rates. Of course, other host considerations have an impact on performance, like the I/O bus bandwidth (how fast the bus can read/write data from/to the network adapter), ATM drivers' buffering capabilities, and operating system kernel issues.

Therefore, adequate memory bandwidth is a necessary but insufficient condition for determining the ability of the end hosts to achieve the required throughput of the transport network technology. It is, though, a good indication of the capabilities of the end hosts.

The ATM switches used were FORE ASX-1000 and FORE ASX-200BX models. These switches provide shared buffer space for UBR traffic for each network module, which is allocated per virtual circuit (VC) dynamically on an as needed basis. These switches also support early packet discard (EPD), which in case of congestion discards the entire sequence of ATM cells belonging to a single packet, thereby not loading the link with unnecessary cells that will be retransmitted by TCP (in the packet level).

Beyond the basic program *ping* used to determine the two-way latency between the two end hosts, a network performance evaluation tool had to be used to test the network performance. For this purpose we used *Netspec* [1, 16], which provides a block-structured language for specifying repeatable distributed network experiments and can support both TCP and UDP, as well as application-layer traffic shaping, which will be discussed later.

Methodology and Results

Experimental Parameters

All the experimental scenarios were carried out with the default MTU size for classical IP over ATM AAL5 networks, which is 9180 bytes. This configuration results in a TCP maximum segment size (MSS) of 9140 bytes. The most important parameter that had to be calculated for the experiments was the BDP of the connection, which was the upper limit of the TCP window, or equivalently the send and receive socket buffer sizes. This parameter was passed to the operating system kernel via the setsockopt system-level function by the application-level testing software, Netspec. The required BDP value for optimal performance of TCP on OC-3 ATM/SONET over ACTS with RTT of 534 ms is approximately 8.53 Mbytes. In our experimental scenarios we used 10-Mbyte window sizes. The performance metrics in our investigation are throughput, channel utilization (or protocol efficiency), and protocol fairness using the fairness index described in [12].

Congestion-Free Performance and Results

The experimental configuration of the congestion-free scenario is shown in Fig. 1. In this scenario one-to-one memory transfers at full workstation speeds are carried out using Netspec over the University of Kansas (KU) ATM network, AAI WAN, and ACTS system at OC-3c rates. The three different network environment paths are represented by dotted, dashed, and solid lines, respectively. The RTT of each connection is noted by clocks next to the workstations. The purpose of this scenario is to note the maximum TCP over ATM throughput that we can obtain on a SONET OC-3 satellite link with the current generation of workstations, as well as to compare the performance between LANs, WANs, and satellite systems. All three TCP protocols were tested over the LAN and satellite environments, while only TCP Reno was tested over the WAN environment. Congestion is not present in this case, since the maximum rate of the sender machine is the same as the rate of the ATM switch interfaces or the satellite link.

The bit error rate (BER) on the satellite links was measured to be in the range of 10^{-12} at the time of testing. The RTT between the end hosts was measured by the program ping and found 534 ms on average for the satellite connection, 0.1 ms for the LAN connection, and 39 ms for the WAN connection. The satellite round-trip latency is the sum of the propagation delay over the satellite link, the transmission delay of ATM cells, the ATM segmentation and reassembly (SAR) delay, the processing delay within

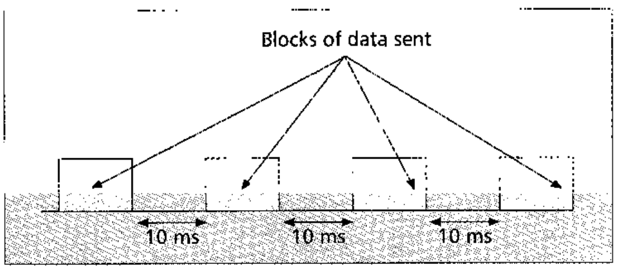

■ **Figure 2.** *The ON-OFF source scheme Netspec uses to shape traffic.*

the TCP/IP stack, and the propagation delay through the terrestrial fiber networks used to connect the hosts under test with the ground stations. Of these factors, the satellite propagation time is by far the dominant one. Table 1 shows the performance results of the three transport protocols for this scenario. The results indicate that TCP hosts, which support the necessary extensions, on ATM/SONET networks over LANs, WANs, and satellite environments with very low BER and high-speed channels, are similar to each other, independent of the large differences in path latencies which they exhibit.

We also used this scenario to test the performance of the three protocols under a noisy satellite connection. We used an OC-12 satellite link with a BER of 8.9 * 10^{-4} and performed the same one-to-one transfers as above. The throughput obtained on average by TCP Reno was 1.235 Mb/s, by TCP New Reno 1.527 Mb/s, and by TCP SACK 2.059 Mb/s. These results confirm that none of the TCP extension implementations can efficiently operate on satellite links with high probabilities of bit error.

Congested Network Performance and Results

When different OC-3c traffic sources are competing for the same OC-3c link, the offered traffic is too high for the switch buffers to handle, resulting in congestion and cell overflow. As a result of this, cells will be dropped and TCP retransmissions will occur. A TCP packet contains about 192 ATM cells. In the experiments conducted here one cell loss will result in the loss of the whole packet. Under these conditions, the goodput will be dramatically decreased. To study this scenario one can look at Fig. 1 and just imagine three DEC 3000/700 workstations at KU instead of one, transmitting at full speed to the network. The transmitting workstations inject data into the link at a faster rate than the switches can forward, and therefore congestion is present. The purpose of this experimental scenario is to note the maximum throughput we can obtain on a SONET OC-3 LAN, WAN, and satellite environment under congestion conditions using the three TCP extension implementations. In addition, we also used application-level traffic shaping and examined the effect of such traffic shaping on TCP performance. Netspec [16], was used to shape the traffic. Shaping was accomplished by specifying the size of the data blocks written to the network at a predefined time interval in milliseconds (10 ms in our experiments). By doing so, a specific source rate can be specified or traffic-shaped. Figure 2 shows the ON-OFF source scheme that Netspec uses to shape the traffic.

Figure 3 shows the graph of the offered load versus aggregate throughput obtained when TCP Reno is used in these experiments. One can observe from these results that the maximum value of TCP Reno throughput in all environments under test was obtained when the aggregate offered load was 105 Mb/s. Above that load, the switch buffers overflowed, cell losses occurred, and the throughput dropped sharply. It also continued to drop while we were increasing the offered application load. The peak throughput under congestion and traffic shaping for the LAN was 103.023 Mb/s, for the WAN 99.265 Mb/s, and for the satellite environment 90.342 Mb/s. When all three sources transmit simultaneously as fast as possible (full rate) (i.e., no application layer pacing), the obtainable TCP Reno throughput is much lower, and was measured as 91.887 Mb/s, 68.646 Mb/s, and 54.105 Mb/s for the LAN, WAN, and satellite environments, respectively. This is shown in the graph of Fig. 3 by the three horizontal lines used as reference for performance comparison to the scenario where traffic shaping was used.

Similar results were obtained with TCP New Reno and TCP SACK, but with higher peaks (thus better channel utilization) in all tested environments. Figure 4 shows the measured performance comparisons among TCP Reno, TCP New Reno, and TCP SACK over the satellite channel with and without traffic shaping. When unregulated full-rate traffic is considered, TCP SACK delivers 60 percent efficiency, TCP New Reno 48 percent, and TCP Reno 40 percent. When traffic shaping is used in the sending nodes, TCP SACK obtains peak throughput of 103.923 Mb/s, TCP New Reno 91.875

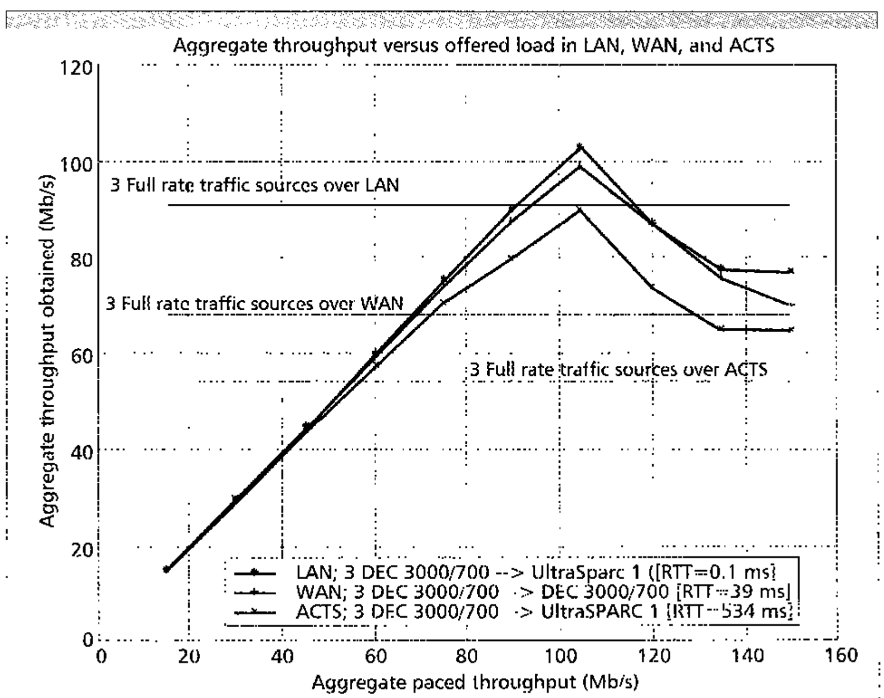

■ **Figure 3.** *Throughput obtained vs. offered load with TCP Reno OC-3c connections over the KU LAN, AAI WAN, and ACTS.*

Mb/s, and TCP Reno 90.342 Mb/s. These results indicate that TCP SACK is more efficient than the others since it has the ability to obtain higher throughput and thus better channel utilization.

Under these conditions, TCP SACK performs better than TCP New Reno or TCP Reno. The cost of using TCP SACK is the protocol implementation complexity, since changes on both sender and receiver stacks are required, while the protocol overhead is negligible compared to the performance improvement. The complexity is not a significant issue with the current generation of workstations; note that over 100 Mb/s throughput to the application layer was experimentally observed. More details on the protocol overheads can be found in [4].

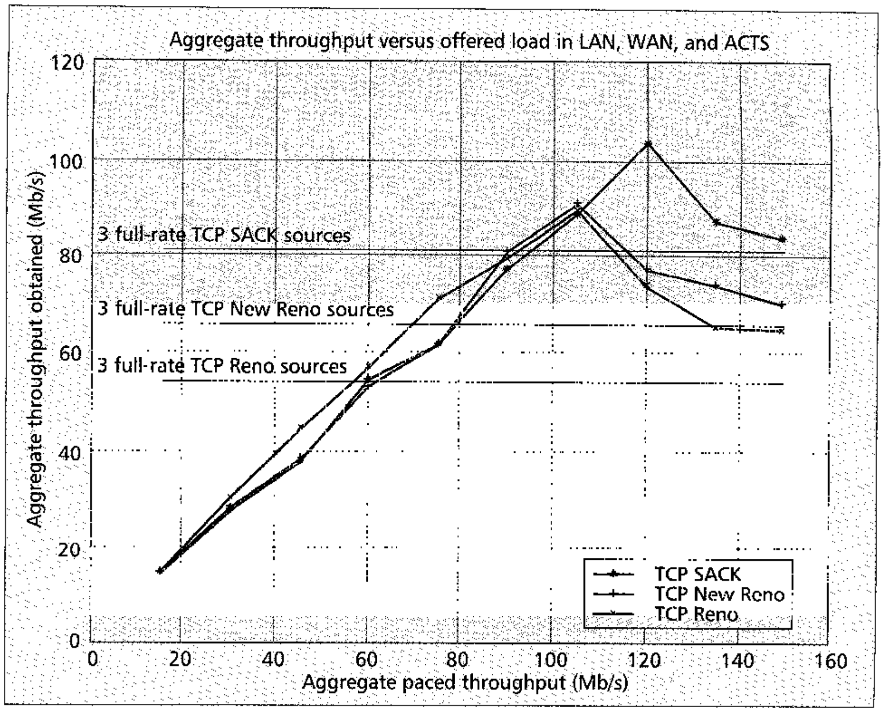

■ **Figure 4.** *Performance comparisons of TCP Reno, TCP New Reno, and TCP SACK over the satellite channel.*

Evaluation and Fairness

To further evaluate the TCP extensions on the TCP/ATM over satellite network, another metric beyond throughput and efficiency has to be taken into consideration: the fairness of the protocols. To study fairness we used the fairness index described in [12], which is the ratio of the square of the sum of all individual connections' throughput obtained over the product of the number of connections competing for the link times the sum of the square of all individual connections' throughput. This index indicates whether workstations with the same TCP protocol competing for the bottleneck link can about equally share the available channel bandwidth. A protocol is ideally fair if the fairness index equals unity. This scenario was also used to measure the channel utilization achieved by each protocol on stressed ATM satellite networks. Figure 5 shows the network configuration of this scenario. In this experiment Netspec generated 15 simultaneous traffic flows, each delivering 15 Mb/s load controlled by the TCP window flow control mechanism. The data traffic passes through the satellite system and is received by another workstation on the KU ATM LAN. All three protocols were tested using this testbed. The BER on the satellite link during testing was $1.1 * 10^{-12}$.

The results obtained are as follows: TCP Reno obtained an aggregate throughput of 44.621 Mb/s, achieving 33 percent channel utilization and a fairness index of 0.89; TCP New Reno obtained an aggregate throughput of 86.651 Mb/s, achieving 64 percent channel utilization and a fairness

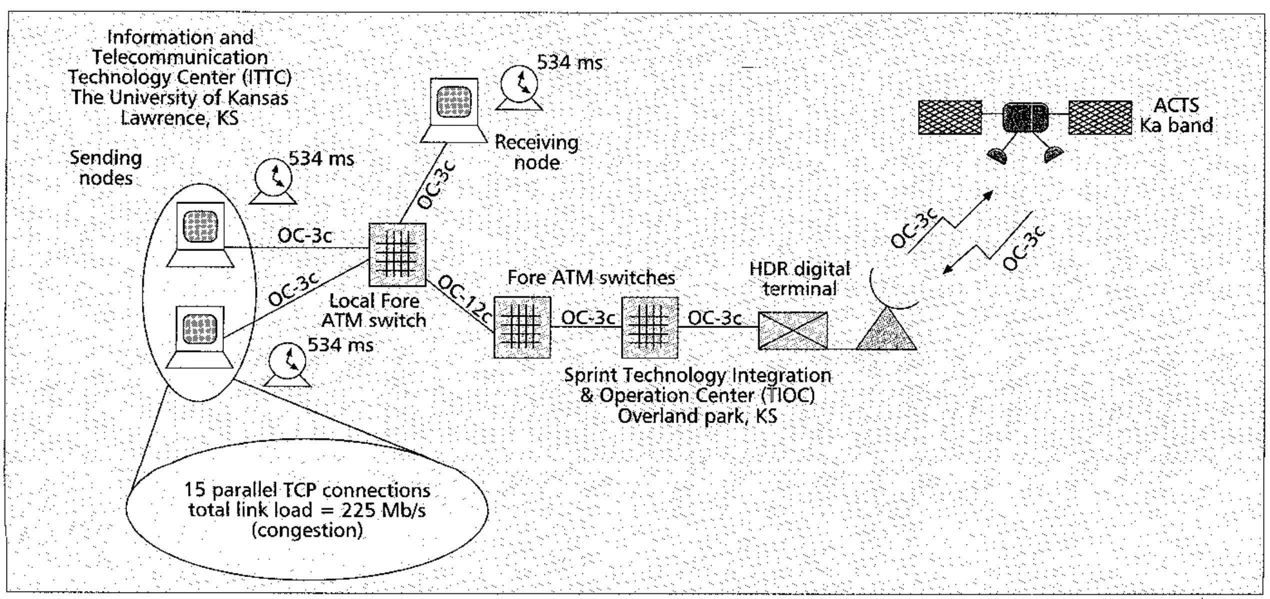

■ **Figure 5.** *Network configuration of 15 parallel TCP connections, each one generating 15 Mb/s of traffic to the network.*

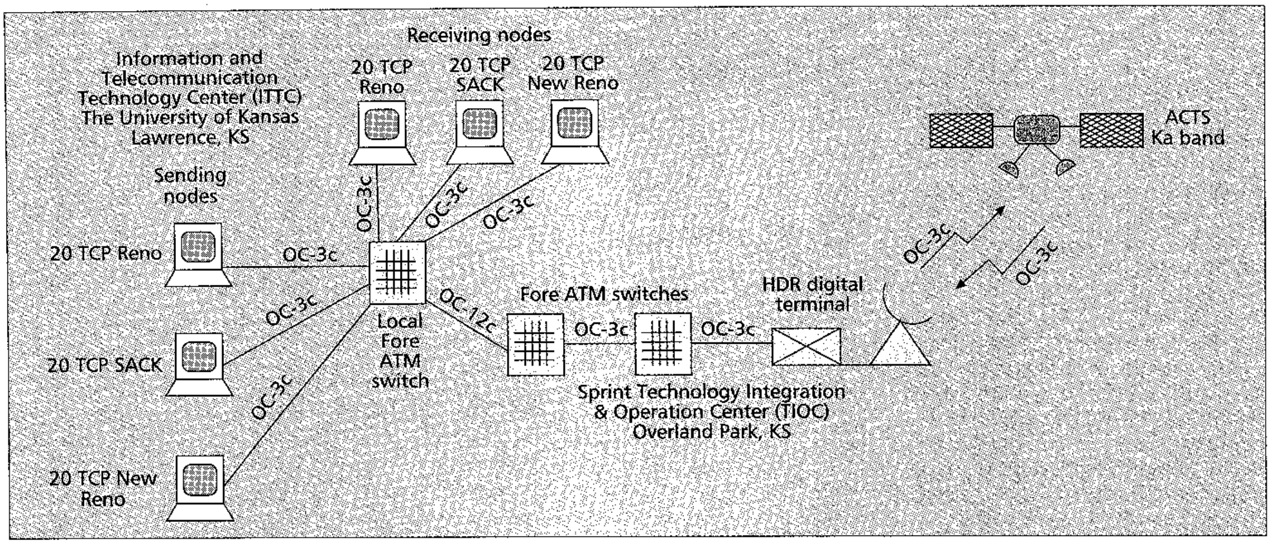

■ **Figure 6.** *Network configuration of 60 parallel TCP connections (20 connections of each protocol) competing for the stressed ATM satellite link. Each connection generates 7.5 Mb/s of traffic, contributing to a total network load of 450 Mb/s. Each protocol source transmits to the corresponding protocol sink.*

index of 0.96; TCP SACK obtained an aggregate throughput of 113.318 Mb/s, achieving 84 percent channel utilization and a fairness index of 0.84. These results indicate that TCP SACK with the selective acknowledgment scheme and the ability to retransmit multiple segments in one RTT is more efficient than the other two protocols but is the least fair, while TCP New Reno with the ability of recovering from N losses in N RTTs is the most fair compared to both TCP Reno and TCP SACK. In this case one TCP implementation was used each time for all connections. Thus, the fairness of TCP SACK in our testing environment was found to be a little lower. This is due to the fact that connections which do not time out (connections with retransmitted packets lost, timeout) are not going through slow start, hence utilizing more link bandwidth as also discussed and observed based on simulations in [17].

Evaluation and Aggressiveness

The last set of experiments was aimed at measuring the aggressiveness of each protocol relative to the other two on a congested network. This scenario could be assumed to exist during the transition period where various end hosts on the network would use different TCP protocols. The network configuration of this scenario is shown in Fig. 6. Twenty TCP Reno connections, 20 TCP New Reno connections, and 20 TCP SACK connections (a total of 60 connections were generated using Netspec) simultaneously compete for the bottleneck link. Each connection generates 7.5 Mb/s of traffic load contributing to a total link load of 450 Mb/s. The BER on the satellite link during testing was $2.3 * 10^{-12}$.

From this scenario we obtained the following results: TCP Reno obtained an aggregate throughput of 35.558 Mb/s with an average per-connection throughput of 1.777 Mb/s and a standard deviation of 0.271 Mb/s; TCP New Reno obtained an aggregate throughput of 31.300 Mb/s with an average per-connection throughput of 1.565 Mb/s and a standard deviation of 0.181 Mb/s; TCP SACK obtained an aggregate throughput of 50.716 Mb/s with an average per-connection throughput of 2.536 Mb/s and a standard deviation of 0.270 Mb/s. These results indicate that when the three protocols compete simultaneously on a stressed congested ATM network, TCP SACK is slightly more aggressive than TCP Reno and TCP New Reno. This kind of behavior is also discussed in [11, 18]. The same scenario was carried out under no congestion. In this case each source was generating 2 Mb/s contributing to a total network load of 120 Mb/s. In this case TCP Reno obtained an aggregate throughput of 38.207 Mb/s with an average per-connection throughput of 1.910 Mb/s, TCP New Reno obtained an aggregate throughput of 36.920 Mb/s with an average per-connection throughput of 1.846 Mb/s, and TCP SACK obtained an aggregate throughput of 37.146 Mb/s with an average per-connection throughput of 1.857 Mb/s. Thus, when there is no congestion on the network, the three protocols are equally (approximately) sharing the available link bandwidth.

Conclusions

In this article we experimentally evaluated the performance, efficiency, fairness, and aggressiveness of TCP Reno, TCP New Reno, and TCP SACK end hosts on ATM networks over high BDP networks. We also compared throughput obtained from experimental scenarios over the local ATM network, AAI WAN, and ACTS. The results of this study indicate:

- In high BDP networks, performance is affected by the inherent latency due to the speed of light impairment and the distance between the end hosts, as well as the probability of bit errors on the specific link. However, TCP over ATM, with different enhancements to TCP (TCP Reno, TCP New Reno, TCP SACK) can achieve high throughput over satellite links.
- Throughput results of TCP extensions for high-performance end hosts on ATM/SONET networks over LANs, WANs, and satellite environments with very low BER and high-speed channels (e.g., ACTS) are similar to each other independent of the large differences in path latencies they exhibit.
- In cases where noisy (with high BER) high-speed satellite links are established, throughput obtained by TCP Reno, TCP New Reno, and TCP SACK end systems is very low, showing that none of the tested TCP protocols are able to recover under these conditions as expected.
- When traffic sources are competing for the same link and no traffic shaping is used, throughput drops dramatically. Throughput degradation is much worse in the high BDP environments as expected, because of the larger time needed to recover from losses. TCP SACK is more efficient than TCP New Reno and TCP Reno since it is able

to achieve better channel utilization. The worst channel utilization is achieved by TCP Reno.

- Traffic shaping at the user layer will help to achieve higher peak values of throughput compared with unregulated transmissions from the traffic sources under congestion conditions.
- When multiple simultaneous connections with the same TCP protocol are competing for the bottleneck link, TCP New Reno is the most fair, while TCP SACK is the least fair but the most efficient.
- Under a hybrid environment with all three TCP protocols competing for the bottleneck link, TCP SACK was found to be slightly more aggressive than TCP New Reno and TCP Reno. Unfortunately, we cannot generalize on this observation, since all connections where generated by only three physical workstations. Therefore, it is hard to say whether operating system issues had some effect on the results. Further work is needed to address this issue. Under the same environment but with no congestion present, all protocols share the available bandwidth about equally.
- The experimental results reported here confirm performance predictions based on simulation studies [11, 17, 18].

Further investigation with these protocols must be carried out with different traffic flows (MPEG, FTP, Web, videoconference, voice) in order to study the effect of TCP extensions on delay-sensitive traffic on stressed ATM high-BDP networks.

References

[1] L. DaSilva *et al.*, "ATM WAN Performance Tools, Experiments, and Results," *IEEE Commun. Mag.*, Aug. 1997.

[2] V. Jacobson, R. Braden, and D. Borman, "TCP Extensions for High Performance," RFC 1323, May 1992.

[3] S. Kota, R. Goyal, and R. Jain, "Satellite ATM Network Architectural Considerations and TCP/IP performance," *Proc. 3rd Ka Band Utilization Conf.*, Sorrento, Italy, Sept. 1997.

[4] M. Mathis *et al.*, "TCP Selective Acknowledgment Options," Oct. 1996, RFC 2018.

[5] W. R. Stevens, "TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms," Jan. 1997, RFC 2001.

[6] J. C. Hoe, "Improving the Start-up Behavior of a Congestion Control Scheme for TCP," *Proc. ACM SIGCOMM '96*, Aug. 1996.

[7] D. E. Comer, *Internetworking with TCP/IP, Volume I, Principles, Protocols, and Architecture*, 3rd ed., Prentice Hall, 1995.

[8] W. Richard Stevens, *TCP/IP Illustrated, Vol. I: The Protocols*, Addison Wesley, 1994.

[9] J. Postel, "Transmission Control Protocol," RFC 793, Sept. 1981.

[10] R. Braden, "Requirements for Internet Hosts — Communication Layers," RFC 1122, Oct. 1989.

[11] K. Fall and S. Floyd, "Simulation-based Comparisons of Tahoe, Reno, and SACK TCP," *Comp. Commun. Rev.*, July 1996.

[12] R. Goyal *et al.*, "UBR+: Improving Performance of TCP over ATM-UBR service," *Proc. ICC '97*, Montreal, Canada, June 1997.

[13] S. M. Bajaj *et al.*, "Performance Characterization of TCP/IP-On-ATM over an ATM/SONET High Data Rate ACTS Channel," *Proc. 16th AIAA Int'l. Commun. Satellite Sys. Conf.*, Feb. 1996.

[14] C. P. Charalambous *et al.*, "Experimental and Simulation Performance Results of TCP/IP over High-Speed ATM over ACTS," *Proc. 1998 IEEE ICC*, Atlanta, GA, June 1998.

[15] D. Hoder and B. Kearny, "Design and Performance of the ACTS Gigabit Satellite Network High Data-Rate Ground Station," *Proc. 16th AIAA Int'l. Commun. Satellite Sys. Conf.*, Feb. 1996.

[16] R. Jonkman, "Netspec: A Network Performance Evaluation and Experimentation Tool," available at http://www.ittc.ukans.edu/netspec

[17] R. Goyal, "SACK TCP/UBR+ over Terrestrial and Satellite Networks," ATM Forum rep. ATM_Forum/97-0423, Apr. 1997.

[18] R. Bruyeron, B. Hemon, and L. Zhang, "Experimentations with TCP Selective Acknowledgments," *Comp. Commun. Rev.*, vol. 28, no. 2, Apr. 1998.

Biographies

CHARALAMBOS P. CHARALAMBOUS [M] (ccharala@ittc.ukans.edu) received his diploma of technician engineer from the Higher Technical Institute, Nicosia, Cyprus, in 1994, his B.Sc. in electrical engineering from the University of Colorado, Denver, in 1996, and his M.Sc. in electrical engineering with emphasis in telecommunication systems from the University of Kansas, Lawrence, in 1998. In the summer of 1998 he joined Sprint as a network engineer in the Broadband Switch Engineering group. His research interests include high-speed ATM terrestrial and satellite networks and transport protocols. More information about the author can be found at www.ittc.ukans.edu/~ccharala

VICTOR S. FROST [F] (frost@ittc.ukans.edu) is currently acting director of the University of Kansas Telecommunications and Information Technology Center. He has been involved in research for numerous corporations, including Sprint, NCR, BNR, Telesat Canada, AT&T, McDonnel Douglas, DEC, and COMDISCO Systems. His research has also been sponsored by government agencies, including, NSF, ARPA, Rome Labs, and NASA. From 1987 to 1996 he was director of the Telecommunications and Information Sciences Laboratory. He has published over 45 journal articles and has had leadership roles and participated in numerous communications and networking conferences and workshops. His current research interest is in the areas of integrated communication networks, high-speed networks, communications system analysis, and simulation. He is currently involved in research on the MAGIC and AAI high-speed wide area testbeds. He received a Presidential Young Investigator Award from the National Science Foundation in 1984, an Air Force Summer Faculty Fellowship, a Ralph R. Teetor Educational Award from the Society of Automotive Engineers, and the Miller Professional Development Awards for Engineering Research and Service in 1986 and 1991, respectively. He is a member of Eta Kappa Nu and Tau Beta Pi. He served as Chair of the Kansas City section of the IEEE Communications Society from June 1991 to December 1992. He received B.S., M.S., and Ph.D. degrees from the University of Kansas, Lawrence in 1977, 1978, and 1982, respectively. In 1982 he joined the faculty of the University of Kansas, where he is currently a Dan F. Servey Distinguished Professor of Electrical Engineering and Computer Science.

JOSEPH B. EVANS (evans@ittc.ukans.edu) received a B.S.E.E. degree from Lafayette College in 1983, and M.S.E., M.A., and Ph.D. degrees from Princeton University in 1984, 1986, and 1989, respectively. He has been at the University of Kansas since 1989, where he is now an associate professor of electrical engineering and computer science. He has been involved in major high-performance networking testbeds such as the MAGIC gigabit testbed and the ACTS ATM Internetwork, and broadband wireless efforts such as the Rapidly Deployable Radio Network (RDRN) mobile wireless ATM project. He has published over 60 journal and conference works. He has been a researcher at Olivetti & Oracle Research Laboratory, Cambridge University Computer Laboratory, USAF Rome Laboratories, and AT&T Bell Laboratories. His research interests include high-speed (gigabit) networks, mobile and wireless networking, and system implementations.