

Experimental and Simulation Performance Results of TCP/IP over High-Speed ATM over ACTS *

Charalambous P. Charalambos, Georgios Y. Lazarou, Victor S. Frost
Joseph Evans, Roelof Jonkman

Information and Telecommunication Technology Center
Department of Electrical Engineering & Computer Science
The University of Kansas
Lawrence, KS 66045
E-mail: frost@ittc.ukans.edu

Abstract

To assist in the design and understanding of future global networks, this paper describes the practical and simulation experiences gained from a TCP/IP on ATM network over a high speed satellite link and presents performance comparison studies of such networks with the same host/traffic configurations over local area (LAN) and wide area (WAN) networks. It was found that the satellite systems deliver performance similar to the terrestrial networks regardless their path latencies in cases where the communication channels exhibit a low bit error rate (BER). 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 (ttcp, Netspec) on OC-3 and OC-12 ATM satellite links.

1: Introduction

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 will compete with terrestrial fiber optic networks in terms of high transfer rates and very low bit error rates (BER), 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 deliver advanced integrated broadband services at high transfer rates. The combination of ATM and satellite technologies, using the Transmission Control Protocol/Internet Protocol (TCP/IP) to form an internetwork architecture, has the potential to provide seamless networking.

This paper presents results from experiments conducted as a part of the ACTS ATM Internetworking (AAI) project. The AAI testbed provides wide area ATM connectivity to several DoD High Performance Centers, and the MAGIC and ATNet testbeds. Our experiments focus on performance measurements on OC-3 and OC-12 ATM using standard TCP/IP hosts over LAN, WAN, and ACTS high data rate (HDR) channels. These experiments show that the performance of

TCP/IP on ATM networks over NASA's ACTS, with its special characteristics and low BER, are comparable with similar architectures on terrestrial fiber networks. Also, simulations modeling these experiments were constructed using the BONEs DESIGNER software [20]. Performance predictions based on these models were validated by the experiments and thus showed that these are appropriate tools to gain additional understanding of network behavior.

The rest of this paper is organized as follows: Section 2 provides background on TCP, TCP/ATM, and ACTS; Section 3 presents the experimental scenarios; Section 4 presents the experimental results; Section 5 analyzes the results and discusses TCP pitfalls; Section 6 presents the simulation results; Section 7 states the conclusions of our work.

2: Background Information

2.1: Transmission Control Protocol

TCP is the reliable, connection-oriented, end-to-end error, flow and congestion control protocol in the transport layer of the Internet suite. It was designed to work independently of lower layer implementations for data transport, such as ATM. Its basic implementation [18] is unsuitable for high speed and high delay networks, and therefore modifications and additions were added to enhance the performance of the protocol over such networks [11, 15, 19]. This is one of the reasons for the variety of TCP versions available today. *TCP Reno* was used in the hosts under test in our experiments. The major characteristics of these transport protocols are slow start, congestion avoidance, fast retransmit, fast recovery, and support for large windows. More details on these TCP extensions can be found in [5, 11, 19].

2.2: TCP/IP on ATM networks

ATM is a scalable cell switching and multiplexing technology that was chosen by ITU-T to be the transport technology for the Broadband Integrated Services Digital Network (B-ISDN). Studies [1, 3, 7] have shown that TCP (with the enhancements discussed in the section above) can achieve maximum performance over ATM.

In our implementations, ATM is transported by Synchronous Optical Network (SONET) protocols [21], permitting irregular ATM cell arrivals and transporting them via its Synchronous Payload Envelope (SPE). We are also using AAL 5 (ATM Adaptation Layer 5) in our network architec-

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

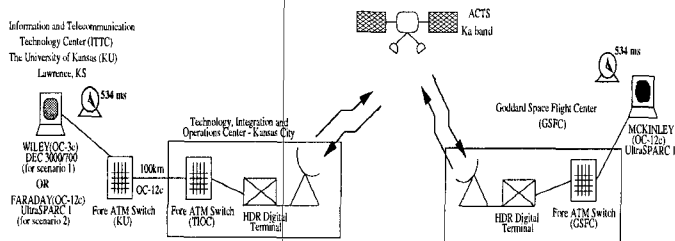


Figure 1. Diagram of scenario 1 and scenario 2; Congestion free TCP/IP over ATM over ACTS networking, for SONET OC-3c and OC-12c rates. NOTE: All links are assumed to be OC-3c for scenario 1 and OC-12c for scenario 2, unless otherwise noted.

ture, which offers unreliable data transfer services with error detection.

2.3: ACTS

ATM transported over SONET systems is rapidly emerging as the transport mechanism for future high speed networks [3]. Broadband communication satellite systems provide an effective platform for world wide communications. Thus, ATM/SONET over satellite channels is the next step towards the implementation of high speed global networks [3].

NASA's ACTS is a satellite system providing SONET STS-3 (155.52 Mbps) and SONET STS-12 (622.08 Mbps) point-to-point and point-to-multipoint services, with possible direct support of STS-1 (51.84 Mbps) with the necessary multiplexing equipment in the ground stations [16]. 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 [4, 5, 10, 16].

3: Experimental Scenarios

3.1: System Overview

The workstations used in the satellite experiments were DEC Alpha 3000/700 with 225 MHz clock; SUN UltraSPARC 1 with 167 MHz clock; and SUN SPARC 20 with 125 MHz clock. Their operating systems supported the TCP Reno extensions. The ATM switches used were FORE ASX-1000 and FORE ASX-200BX models. These switches provide a shared buffer space of 8192 cells for Unspecified Bit Rate (UBR) traffic for each network module (four ports for SONET OC-3c or one port for SONET OC-12c). The UBR buffer space is allocated per virtual circuit (VC) dynamically on an as needed basis [17]. 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).

3.2: Description of scenarios

Scenario 1 is illustrated in Figure 1. In this scenario an OC-3c equipped Alpha workstation at KU transmits to an OC-12c equipped UltraSPARC workstation at GSFC (Goddard Space Flight Center) via an OC-3 ACTS link. The purpose of this experimental scenario is to note the maximum TCP over ATM throughput that we can obtain on a SONET OC-3 satellite link. 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.

Scenario 2 is illustrated in Figure 1 as well. In this scenario an OC-12c equipped UltraSPARC workstation at KU

transmits to an OC-12c equipped UltraSPARC workstation at GSFC via an OC-12 ACTS link. The purpose of this experimental scenario is to note the maximum throughput that we can obtain on a SONET OC-12 satellite link. The two hosts are configured as above, and congestion is not present.

Scenario 3 is illustrated in Figure 2. In this scenario congestion is present. Three OC-3c equipped workstations at

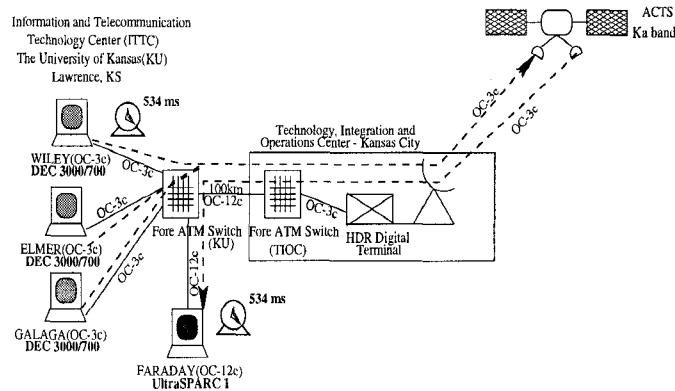


Figure 2. Diagram of scenario 3; TCP/IP over OC-3c ATM over ACTS networking under congestion conditions.

KU transmit to an OC-12c workstation at KU via an ACTS OC-3 link in loop-back mode. The transmitting stations are injecting data into the link at a faster rate than the link can handle, and therefore congestion is present. The purpose of this experimental scenario is to note the maximum throughput that we can obtain on a SONET OC-3 satellite link under congestion conditions. All the stations are configured as in scenario 1.

Scenario 4 is illustrated in Figure 3, where performance measurements are taken from the KU ATM LAN and the AAI WAN, for performance comparisons between LANs, WANs, and the satellite system of scenario 1. This test was carried out under conditions of no congestion. All the hosts are configured as in scenario 1.

Scenario 5 is illustrated in Figure 4, where performance measurements are taken from the KU ATM LAN and the AAI WAN, for performance comparisons under congestion conditions with the LAN, WAN, and the satellite system of scenario 3. Congestion is present in this case since all three workstations shown in the diagram transmit at a faster rate than the OC-3c link can handle. All the workstations are configured as in scenario 1.

3.3: Performance of TCP/IP over ATM/SONET over ACTS

The physical layer of our implementations is based on 155.52 Mbps SONET OC-3c and 622.08 Mbps SONET OC-12c interfaces. Classical IP over ATM is implemented according to RFC-1577 [13], where IP datagrams are encapsulated using IEEE 802.2 LLC/SNAP and segmented into ATM cells using AAL 5 (48 bytes). The ATM layer adds another 5 bytes of header information. The default Maximum Transmission Unit (MTU) size for Classical IP over ATM networks is 9180 bytes [2], with a SNAP header of 8 bytes, [9, 13] and a maximum AAL 5 Protocol Data Unit (PDU) of 65535 bytes [21]. TCP and IP add another 20 bytes of header information each. That leads to a user rate of 134.513 Mbps and 538.053 Mbs for OC-3c and OC-12c interfaces respectively [5].

The ACTS network, from the end-user point of view, provides the functions of terrestrial SONET-based fiber networks. Therefore, the throughput for standard SONET OC-3c and SONET OC-12c should be achievable over the satellite network.

4: Experimental Results

All the experimental scenarios were carried out with the default MTU size for Classical IP over ATM networks, which is 9180 bytes [2]. This configuration results in a TCP Maximum Segment Size (MSS) of 9140 bytes. The most impor-

tant parameter that had to be calculated for the experiments 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 network performance tools (Netspec, *ttcp*) that were used to test the network capacity. *Netspec* and *ttcp* were the performance evaluation tools used throughout the experiments. More details on these tools can be found in [5, 12].

4.1: Results from scenario 1

In this scenario, as shown in Figure 1, an OC-3 satellite link was established between the TIOC (Technology Integration and Operations Center) and GSFC. The BER of the satellite links was measured to be in the range of 10^{-11} at the TIOC and 10^{-13} at GSFC throughout the experiments, which is close to the specifications of terrestrial fiber networks. The round trip time (RTT), that is the path latency between the two hosts, was measured by the program *ping* to be 534 ms on average. For comparison, let us note here that the average RTT in the local ATM network, shown in Figure 3, is only 0.1 ms, while the average RTT in the AAI WAN, shown in the same figure, is 39 ms. 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 reassemble (SAR) delay, the processing delay within 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.

The send and receive buffer sizes to fully utilize the OC-3 ATM link for optimal performance, were calculated at 9 MB. In our experiments we used a transmit and receive window size of 10 MB. We obtained high throughput results, averaged in ten trials to 109.717 Mbps with a standard deviation (σ) of 5.544 Mbps.

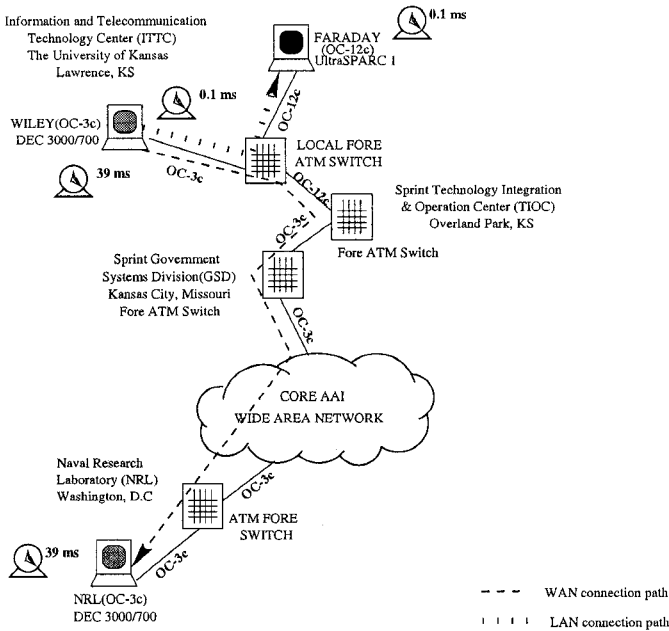


Figure 3. Diagram of scenario 4; Congestion free, TCP/IP over OC-3c ATM in the local ATM network (LAN) and the AAI WAN.

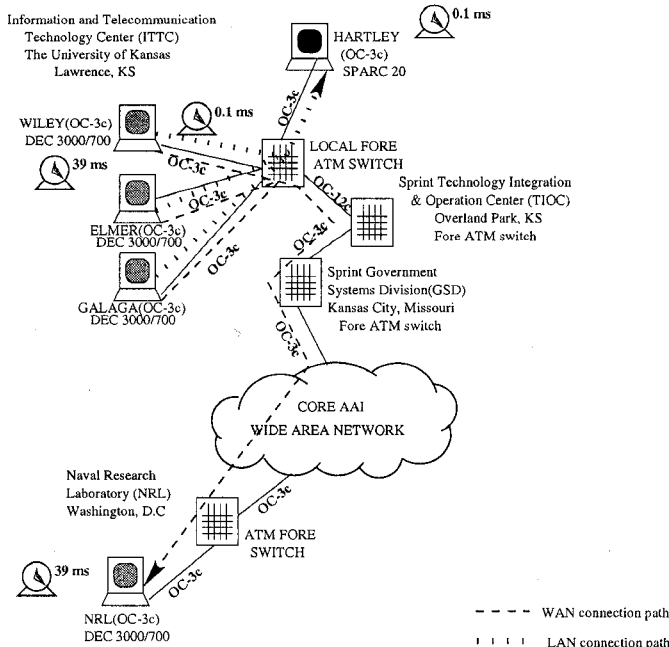


Figure 4. Diagram of scenario 5; TCP/IP over OC-3c ATM in the local ATM network (LAN) and the AAI WAN, under congestion conditions.

4.2: Results from scenario 2

In this experimental scenario, as shown in Figure 1, an OC-12 satellite link was established between the TIOC and GSFC. In this case, the minimum window size required for maximum throughput is about 34 MB.

Unfortunately, due to technical problems in the OC-12 digital terminal at the GSFC ground station at the time of the experiment, we managed to run only two successful tests with a 10 MB window size in the sender and receiver hosts. The theoretical throughput that we should obtain with the 10 MB window size on an OC-12 link is 157.088 Mbps. Due to a relatively high BER (10^{-8} at TIOC and 10^{-9} at GSFC) reported on the satellite links, we observed 121.835 Mbps and 127.870 Mbps instead.

4.3: Results from scenario 3

When different OC-3c traffic sources are competing for the same OC-3c link, the offered traffic is too high for the switches to handle, resulting in cell overflow. As a result of this, cells will be dropped and TCP retransmissions will occur. A TCP packet includes about 192 ATM cells. One cell loss will result in the loss of the rest of the cells in that packet. Under these conditions, the goodput will be dramatically decreased.

In the case shown in Figure 2, the satellite network was tested under congestion conditions. Three OC-3c sources were transmitting over an OC-3 satellite link. This will result in congestion and the throughput will be decreased. Netspec with constant rate traffic was used to gradually increase

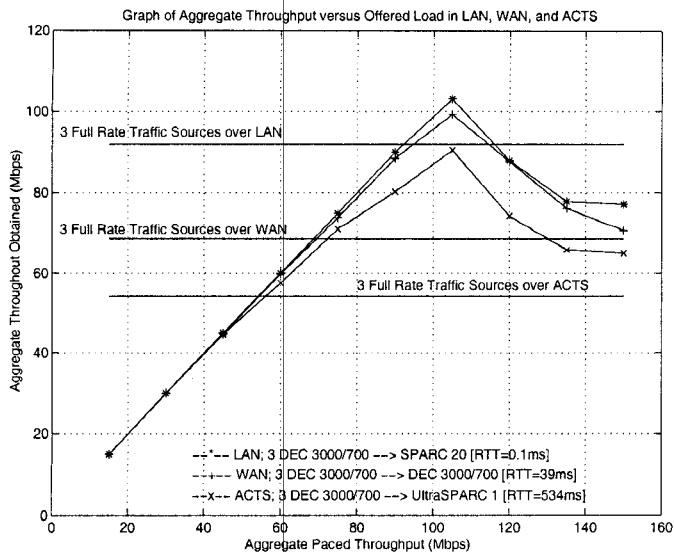


Figure 5. Graph illustrating throughput obtained versus offered load with three TCP connections over LAN, WAN and satellite environments.

the offered load, and measurements were taken for the aggregate throughput of the source machines. The available throughput was measured when the source machines were transmitting full rate traffic. The results from this experiment are shown in Figure 5, and compared with similar conditions over LANs and WANs.

4.4: Results from scenario 4

In this experimental scenario, as shown in Figure 3, measurements were taken in the local ATM network and in the AAI WAN on OC-3c connections under no congestion conditions, in order to compare these results with the ones obtained from the satellite network of scenario 1. The hosts participating in the LAN and WAN environments were configured with the default IP over ATM settings, like the hosts in the satellite experiment of scenario 1, and the tests were run with a window size of 800 kB, which guarantees maximum throughput. Table 1 shows the average and maximum throughput obtained from the LAN, WAN and satellite network experiments. It is obvious from this table that the throughput results obtained in TCP/IP on ATM networks over terrestrial LANs and WANs, and satellite links with low BER, are very close regardless of the large difference in their network path latencies.

4.5: Results from scenario 5

In the experimental scenario shown in Figure 4, measurements were taken in the local ATM network and in the AAI WAN on OC-3c connections under congestion conditions in order to compare these results with the ones obtained from the satellite network of scenario 3. All the hosts participating in the LAN and WAN environments were configured as in scenario 3 discussed above. Figure 5 shows the offered load versus aggregate throughput obtained in the experiments under congestion conditions over LANs, WANs, and satellite environments. One can observe from these results that the maximum value of throughput in all environments under test was obtained when the aggregated offered load was 105 Mbps. 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 load. The peak throughput under congestion for the LAN was 103.023 Mbps, for the WAN 99.265 Mbps, and

	LAN	WAN	ACTS
Throughput(Mbps) (Average)	119.838	114.058	109.717
Throughput(Mbps) (Maximum)	126.955	126.109	119.000
σ (Mbps)	6.398	7.079	5.544
RTT(ms)	0.1	39	534

Table 1. Table showing the average and maximum throughput, standard deviation, and round trip time on OC-3c connections over a LAN, WAN, and ACTS environments in ten trials.

for the satellite environment 90.342 Mbps.

When all three sources transmit simultaneously as fast as they can (full rate), the obtainable throughput is much lower, and it was measured as 91.887 Mbps, 68.646 Mbps, and 54.105 Mbps for the LAN, WAN and satellite environments, respectively.

5: Analysis of results and TCP pitfalls

The results obtained from the experiments are less than expected, for all three environments. This is due to the protocol implementations under different operating system kernels in the transmitting and receiving hosts, as well as due to the congestion algorithms implemented in the TCP protocol.

The slow-start and congestion avoidance algorithms in the TCP protocol have a negative effect on throughput, especially for high delay networks. It is well known that maximum throughput will be achieved if the send and receive window sizes are set correctly. This is of course true in cases where the increase in the window size takes effect instantaneously and not gradually, as in the case of TCP with slow start and congestion avoidance mechanisms [3]. According to equation (1)¹, and under no congestion or segment losses, with a window size of 10 MB as used in our satellite experiments (thus, about 1148 MSSs of 9140 bytes in one window) and a RTT of 534 ms, it takes 9.813 seconds ($18.377 \times RTT$) to fill the pipe, that is,

$$Slow_Start_Time = RTT \times (1 + \log_{1.5} W) \quad (1)$$

where RTT is the round trip time between the two hosts, and W is the number of segments in the receiver window size.

In our experiments, we were transmitting an average of 1 GB (about $102 \times Window_Size$) of data per trial. Therefore, one window of data was transferred in 9.813 seconds, and after that, each remaining window of data was transferred in one RTT (534 ms); so in our case we transferred 102 windows of data in about $18.377 \times RTT + 101 \times RTT = 119.377 \times RTT$ seconds (63.747 seconds). This total data transfer time of 63.747 seconds is equivalent to $(102 \times Window_Size)/(119.377 \times RTT) = 0.854 \times (Window_Size/RTT)$, and means that there is an 14.5% reduction in throughput, compared with an instantaneous transmission, caused by the TCP slow-start mechanism in our set of satellite experiments even when no congestion or segment losses were present.

When there is a segment loss, TCP assumes that is caused by congestion [6, 18, 19], and therefore the transmitter has to reduce the rate of injecting data into the network. In the congestion avoidance phase, the CWND halves in response to a loss, and recovers by increasing linearly (approximately one segment per RTT) until it reaches its original value. In the satellite experiments we ran using a 10

¹ Valid for TCP systems with delayed acknowledgments.

MB window size and default MTU sizes (9180 bytes), assuming that the TCP sender managed to reach the receiver's window (10MB) without losses, it will then take $571.5 \text{ segments} (5 \text{ MB}) \times \text{RTT} = 305.2 \text{ seconds}$ to fill the pipe using this algorithm. In the fast recovery phase, the CWND halves (plus 3 segments due to the three duplicated ACKS received) and congestion avoidance follows. Thus, for our satellite experiments it will take $(571.5/3) \text{ segments} \times \text{RTT} = 303.5 \text{ seconds}$ to fill the pipe. If the segment loss occurs early in slow start, then it will take hours to fully recover using this algorithm.

TCP Reno with fast retransmit and fast recovery improves performance over the basic TCP implementation, but it exhibits another pitfall. Studies [1, 8] have shown that when more than one loss occurs within one window, fast retransmit and fast recovery will be triggered several times in one RTT, resulting in reduction of the CWND several times and then linear growth. This leads to throughput reduction.

Due to these attributes of TCP, the satellite links must have a very low BER and no losses due to congestion, otherwise the throughput will drop dramatically. In the OC-3 experiments, the satellite links exhibited very low BER, comparable with those in fiber-optic terrestrial networks, therefore lost segments and retransmissions were limited, and throughput was comparable to that obtained over LANs and WANs. In the OC-12 experiments, the BER in the satellite links was relatively high, resulting in a degradation in throughput which can be justified using the same logic as above. Even if we have only one segment loss per window which is detected by fast retransmit, it will take 303.5 seconds (assuming that the sender has already reached the receiver's window) for the TCP end host to utilize the available bandwidth in the satellite environment. Using the same logic, in the AAI WAN this time is reduced to 22.18 seconds and for the local ATM network to 0.057 seconds.

The experiments conducted under congestion conditions show that the common TCP congestion algorithms are not efficient for TCP over WANs and satellites, and result in throughput degradation in common cases. TCP end traffic sources competing for the same link will cause throughput to drop because of switch buffer overflows and cell losses. Traffic control is essential under these circumstances to prevent losses. Increasing the offered load above a certain level will cause the throughput to drop sharply. In the satellite environment, throughput drops faster than in the LAN or the WAN environments. This is because of the TCP mechanism, where bandwidth is wasted due to the long time needed to reach the receiver advertised window size after a segment loss occurs. For the same reason, when no traffic shaping is present, and all sources inject data in the network simultaneously as fast as they can, the aggregated throughput obtained over ACTS is lower than that obtained over the WAN or LAN.

6: Simulation of TCP over satellite networks

The creation of simulation models is usually the only means of predicting and evaluating the performance of high speed networks, since mathematical models of such networks are not yet feasible [14]. Simulation is an excellent way of investigating and understanding the behavior of the network architecture under test, as well as an efficient methodology for observing the effects of network parameter changes on overall performance.

In this section we investigate simulation models for TCP/IP end hosts over satellite environments and we use measurements to validate our simulation results. The sim-

System Parameter	Value
TCP segment size	9140 bytes
MTU size	9180 bytes
TCP processing time	62 μs
TCP buffer size	10 MB
Slow-Timer period	0.5 s
Fast-Timer period	0.2 s
Minimum RTO	1.0 s
Switch processing delay	6.0 μs
Satellite propagation delay	125 ms per link
OC-3c Link speed	149.760 Mbps
OC-12c Link speed	599.040 Mbps

Table 2. Simulation parameters used in our models.

ulation software used for our simulations is *BONeS DESIGNER* [20].

6.1: Simulation model primitives and parameters

The core of the simulation models developed is the TCP primitive module. We used the TCP primitive module developed by researchers at the University of Kansas, as described in [14], which is based on the 4.3 BSD Reno version of TCP and supports the functions covered in Section 2. In our simulation models, the network (IP) and physical (SONET) layers are not included. The impact of these layers is captured by accounting for their information overhead. Also, due to the long run-time of the ATM segmentation and reassemble process in the simulated satellite environment, an ATM model was not used in the simulation models of scenarios 1 and 2. In these models, we are actually simulating TCP over the satellite path and it is to be shown that the measurements obtained by experiments using TCP/IP on ATM/SONET networks over satellites can validate the simulation results of TCP over satellite models. In the simulation model of scenario 3, where congestion with many cell losses and TCP retransmissions is present, the ATM module is required for more accurate results. The simulation system parameters used in our models are shown in Table 2 and explained in [5, 14].

6.2: Simulation Results

• Simulation of scenario 1, OC-3 satellite connectivity

The block diagram of the model simulating scenario 1 is shown in Figure 6 and the results are shown in Table 3. The data flows from end host wiley (KU) to end host mckinley (GSFC). The whole path through the satellite is bidirectional to allow ACKs from the receiver to the transmitter. The switches, High Data Rate Terminal (HDRT), and ACTS blocks are modeled by a FIFO (First In First Out) queue and a server. The link propagation delay is represented by the Link blocks (OC-3 or OC-12) and the OC-3 or OC-12 service times are included in the output port of every TCP end host and switch. The total path round trip delay through all the blocks is 534 ms, which is the round trip delay found in the real scenarios. The switch buffer sizes were set to a large value, since there is no congestion and thus their values are irrelevant to performance. The error blocks in the downlink paths of ACTS are random switches from the BONeS library and allow data to pass through with a certain probability (passed as a parameter). This parameter was calculated by equation (2),

$$P_{\text{error-free-packets}} = (1 - \text{BER})^{(\text{Packet_Size})_{\text{bits}}} \quad (2)$$

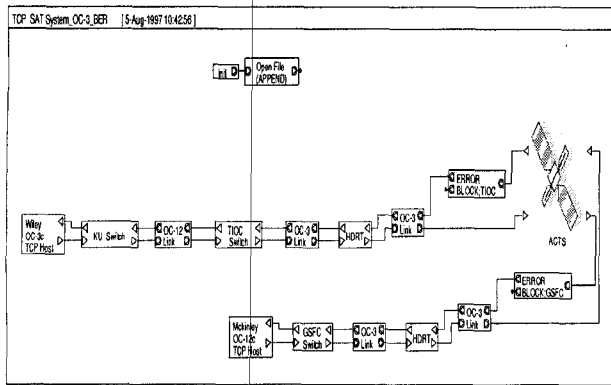


Figure 6. Simulation model for experimental scenario 1, OC-3 satellite connectivity.

where $P_{error_free_packets}$ is the probability with which packets are passing through without any errors, BER is the bit error rate for each link as given by the Network Management Terminal (NMT) during the experiments, and Packet_Size is the MTU used, 9180 bytes multiplied by 8 bits.

• **Simulation of scenario 2, OC-12 satellite connectivity**

The block diagram of the model simulating scenario 2 is similar with the one in Figure 6, with the only difference being the satellite links are OC-12. The simulation results are shown in Table 3. The blocks and parameters are the same as in the simulation model of scenario 1, with the only differences being the OC-12 service times in all paths, and the appropriate probability value, calculated by equation (2), for the error blocks in the satellite links.

• **Simulation of scenario 3, OC-3 satellite connectivity under congestion conditions**

The block diagram of the model simulating scenario 3 is shown in Figure 7 and the simulation results are shown in Table 3. This model represents an OC-3 ACTS bent pipe loop-back connection between three TCP/ATM traffic source blocks (no traffic shaping is present) and a block accommodating three TCP/ATM hosts in order to be able to accept data from them. The connection paths between the source and receiver blocks are established through the "KU ATM Switch" block using connection identifiers, which are equivalent to the virtual circuit identifiers (VCI) used in the experiments, and treated at the receiving end with the use of a FIFO queue with a server which allows TCP packets to pass through towards the specified receiver TCP/ATM host at specific time intervals, specified by the TCP processing time. The probability with which cells are passing through is calculated for this model by equation 3, and is passed as a parameter to the error block located on the downlink satellite path,

$$P_{error_free_cells} = (1 - BER)^{(Cell_Size)_{bits}} \quad (3)$$

where $P_{error_free_cells}$ is the probability with which cells are passing through without any errors, BER is the bit error rate for each link as given by NMT during the experiments and Cell_Size is the ATM cell size in bits, which is 424 bits.

The rest of the blocks are the same as in the simulation model of scenario 1, with the only differences being the processing of ATM cells instead of TCP pack-

ets, and the switch buffer size, which was set to 8192 cells per virtual circuit (VC) to model the maximum possible buffer space a VC can have in the FORE ATM switches. Note that the UBR buffer space in the FORE ATM switches is allocated per VC on an as needed basis [17]. For the simulation purposes, setting the switch buffer space to the maximum value of 8192 cells per VC is a close approximation.

7: Conclusions

In this paper we studied the performance of TCP/IP end hosts on ATM networks over high speed satellite links. We also compared throughput obtained from experimental scenarios over the local ATM network, over the AAI WAN, and over ACTS. We did the same using simulation and validated the simulation results with the experimental measurements. The basic results of our study indicate:

- 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. Since we are running TCP over ATM, different enhancements to TCP [11, 19], can achieve high throughput over satellite links.
- Throughput results for TCP/IP hosts on ATM/SONET networks over LANs, WANs and satellite environments with very low BER and high speed channels (like ACTS) are similar to each other independent of the large differences in path latencies they exhibit. Of course this is achievable if TCP supports the necessary extensions, with the most important being the extension of scaling window sizes beyond 64 kB. Even if we assume that the communication channels are ideal, the throughput will be lower than the theoretical values because of the TCP slow-start algorithm.
- In cases where noisy high speed satellite links are established, throughput obtained by TCP Reno end systems will be degraded, because of retransmissions and the fact that bandwidth will be wasted while the source tries to reach the receiver's advertised window size using the fast recovery or congestion avoidance algorithms. A possible solution to this that might increase the performance obtained by TCP over noisy satellite links is to modify TCP to handle retransmissions in a more efficient way. Simulation studies [1, 8], have shown that the implementation of a selective acknowledgment

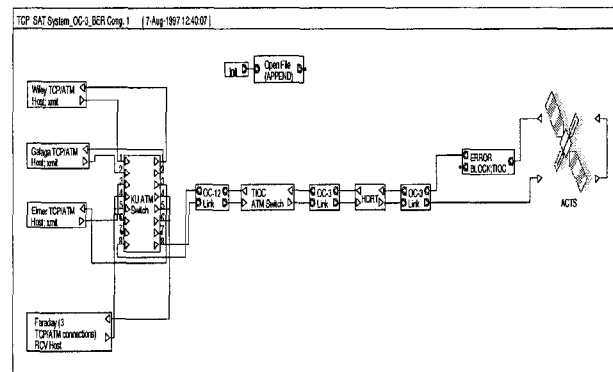


Figure 7. Simulation model for experimental scenario 3, OC-3 satellite connectivity under congestion conditions.

	BER	Experimental Results (Max.Values)	Simulation Results (Max.Values)	% Error
Scenario 1 (OC-3 link) (No ATM simulated)	TIOC/GSFC $10^{-11}/10^{-13}$	119.000Mbps	115.295Mbps	3.1%
Scenario 2 (OC-12 link) (No ATM simulated)	TIOC/GSFC $10^{-8}/10^{-9}$	127.870Mbps	131.360Mbps	2.6%
Scenario 3 (full rate) (OC-3 congestion) (ATM simulated)	TIOC 10^{-11}	54.105Mbps	49.641 Mbps	8.2%

Table 3. Table showing the experimental and simulation results (maximum values obtained over trials) for scenarios 1, 2, and 3, as well as the BER on the satellite links during the experiments.

(SACK) TCP protocol can improve the performance over TCP Reno.

- When traffic sources are competing for the same link and no traffic shaping is used, throughput drops dramatically. In this case, cells are dropped due to overflowed switch buffers and thus TCP retransmissions occur. This is much worse in the satellite environment, since TCP Reno will decrease the window size when multiple segment drops occur within one window, and fast recovery will not be fast enough due to the time needed for the sender to ramp up and reach the receiver's window size.
- 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.
- Using experimental measurements we validated the simulation models.

Further performance experiments must be carried out to obtain a better understanding of the OC-12 connectivity through ACTS. New TCP implementations must be applied in order to investigate how efficiently they can handle retransmissions caused by congestion or segment losses on noisy high speed satellite channels.

References

- [1] Furquan A. Ansari. Adapting TCP/IP over ATM, *Master of Science Thesis*, University of Kansas, 1996.
- [2] R. Atkinson. Default IP MTU for use over ATM AAL5, May 1994. RFC 1626.
- [3] S.M. Bajaj, C. Brazdziunas, D.E. Brooks, D.F. Daly, S.M. Srinidhi, T. Robe, F. Vakil. Performance Characterization of TCP/IP-On-ATM over an ATM/SONET High Data Rate ACTS Channel. *Proceedings of the 16th AIAA International Communications Satellite Systems Conference, February 1996*.
- [4] M.A. Bergamo. Network Architecture and SONET Services in the NASA/ARPA Gigabit Satellite using NASA's Advanced Communications Technology Satellite (ACTS). *Proceedings of the 15th AIAA International Communications Satellite Systems Conference, February 1994*.
- [5] C.P. Charalambous, G.Y. Lazarou, V.S. Frost, J. Evans, R. Jonkman. Experiments and Simulations of TCP/IP over ATM over a High Data Rate Satellite Channel. *AAI Report, ITTC-FY98-TR-10980-25, August 1997*.
- [6] Douglas E. Comer. *Internetworking with TCP/IP, Volume I, Principles, Protocols, and Architecture*. Prentice Hall, 3rd edition, 1995.
- [7] L. DaSilva, J. B. Evans, D. Niehaus, V. S. Frost, R. Jonkman, B. Lee, G. Lazarou. ATM WAN Performance Tools, Experiments, and Results. *IEEE Communications Magazine*, August 1997.
- [8] Kevin Fall, Sally Floyd. Simulation-based Comparisons of Tahoe, Reno, and SACK TCP. *Computer Communications Review, July 1996*.
- [9] Juha Heinanen. Multiprotocol Encapsulation over ATM Adaptation Layer 5, July 1993. RFC 1483.
- [10] D. Hoder, B. Kearny. Design and Performance of the ACTS Gigabit Satellite Network High Data-Rate Ground Station. *Proceedings of the 16th AIAA International Communications Satellite Systems Conference, February 1996*.
- [11] Van Jacobson, Robert Braden, David Borman. TCP Extensions for High Performance, May 1992. RFC 1323.
- [12] Roelof Jonkman. Netspec: A Network Performance Evaluation and Experimentation Tool. Available at <http://www.ittc.ukans.edu/netspec>.
- [13] M. Laubach. Classical IP and ARP over ATM, January 1994. RFC 1577.
- [14] Georgios Y. Lazarou, Victor S. Frost, Joseph B. Evans, Douglas Niehaus. Using Measurements to Validate Simulation Models of TCP/IP over High Speed ATM Wide Area Networks. In *Proceedings of the 1996 IEEE International Conference on Communications, June 1996*.
- [15] Matthew Mathis, Jamshid Mahdavi, Sally Floyd, Allyn Romanow. TCP Selective Acknowledgment Options, October 1996. RFC 2018.
- [16] NASA System Handbook. *Advanced Communications Technology Satellite*, NASA, Cleveland, 1995.
- [17] Vector General Description Guide. Northern Telecom, Canada, 1996.
- [18] Jon Postel. Transmission Control Protocol, September 1981. RFC 793.
- [19] W. Richard Stevens. TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms, January 1997. RFC 2001.
- [20] Systems & Networks, *BONeS DESIGNER 3.0 Modeling Guide*, Lawrence, KS, 1995.
- [21] Andrew S. Tanenbaum. *Computer Networks, 3rd edition*, Prentice Hall PTR, New Jersey, 1996.