NASA's Broadband Satellite Networking Research

William D. Ivancic, NASA

David Brooks and Brian Frantz, Sterling Software

Doug Hoder, NASA

Dan Shell, Cisco Systems Inc.

David Beering, Infinite Global Infrastructures, LLC

ABSTRACT ATM is currently considered the primary WAN technology with Internet protocols providing the routing and transport requirements. Another WAN technology being considered is packet over SONET. Using the Advanced Communication Technology Satellite, NASA has demonstrated and is experimenting with these technologies. This article summarizes some of the major completed and ongoing experiments and demonstrations performed using commercial standard protocols such as ATM and TCP/IP over broadband satellite networks.

The Space Program Office and Communication Technology Division of the NASA Glenn Research Center have been working with the United States satellite communication industry over the past 16 years to develop advanced communication and networking technologies to improve commercial satellite communications. With the recent explosion of the Internet and the enormous business opportunities available to communication system providers, great interest has developed in using commercial protocols over satellite networks. NASA Glenn has been addressing five major areas of concern regarding broadband satellite communications using commercial protocols: quality of service (QoS), interoperability, routing, traffic management, and protocol bandwidth efficiency.

The first section of this article will briefly highlight various protocols used in the demonstrations and experiments and describe some of the important protocol characteristics relative to satellite communications. The second section will discuss completed and ongoing research that addresses three of the five major areas of concern: QoS, interoperability, and protocol efficiency.

PROTOCOL CHARACTERISTICS

The experiments presented utilize asynchronous transfer mode (ATM), Transmission Control Protocol/Internet Protocol (TCP/IP), and Moving Picture Expert Group version 2 (MPEG-2) as the primary data delivery protocols. The following section briefly describes some important characteristics of these protocols.

ATM STANDARDS

ATM is a protocol originally conceived by the telecommunication industry to handle multimedia traffic over wide area networks (WANs). The protocol is a connection-oriented cell-switched protocol developed for fiber optic systems having near-error-free performance characteristics. The protocol encompasses features of both the data link and routing layers, layers 2 and 3 of the open systems interconnection (OSI) model.

The main features of ATM are guaranteed QoS, ease of switching, and multimedia compliance. The ATM cell structure of 53 bytes was designed for high-bandwidth implementations with low switching and routing overhead. The ATM cell consists of a 5-byte header and 48-byte payload. The 5-byte header of each cell contains all the information necessary for the end-to-end delivery of the cell. If the header of the cell is corrupted beyond the corrective abilities of the included 1-byte cyclic redundancy check (CRC), the cell is discarded. The small header and limited routing rules allow the switching to be performed in high-speed hardware. The small size of the ATM cell — 53 bytes — allows small internal buffers to be used in the switches, keeping cell delay and jitter to a minimum.

ATM requires that the order of ATM cells be maintained and the link be nearly error-free. These requirements directly affect the design of satellite and wireless networks that wish to utilize ATM.

TRANSMISSION CONTROL PROTOCOL

TCP/IP is the primary protocol suite used today over the Internet. TCP is the data transport protocol from the TCP/IP suite for reliable data transmission. TCP is a *reliable protocol*.[1] It was designed to be robust and work in a variety of environments including satellite networks. TCP requires feedback to acknowledge successful data reception. Because of this feedback control, most *general implementations of TCP*[2] perform inefficiently in networks that have large delays relative to the available transmission bandwidth (known as the *bandwidth-delay product*). Many tuning parameters are available to enhance the performance of TCP, including segment size, timers, and window sizes. Embedded in the TCP implementation are numerous congestion avoidance algorithms[3] such as slow start, selective retransmission, and selective acknowledgment, which generally improve performance in most shared networks such as the Internet.

The maximum theoretical throughput for TCP is given in Eq. 1. T_{Max} is the maximum theoretical throughput. *RTT* is

[1] *A reliable protocol is one that guarantees accurate delivery of information.*

[2] *As of 1999, most implementations of TCP — such as those found in personal computers — do not have the necessary extensions to allow for efficient transmission over networks that have a large bandwidth-delay product (RFC 1323).*

[3] *Not all of these congestion control algorithms are implemented or enabled in all versions of TCP code.*

0163-6804/99/$10.00 © 1999 IEEE

[illegible]	[illegible]	[illegible]	[illegible]
16	32,768	262,144	262.14
32	65,536	524,288	524.29
1024	2,097,152	16,777,216	16,777.22

Notes:
Throughput = Receive buffer/RTT
kbyte = 1024 bytes
RTT = 500 ms

Table 1. *Maximum throughput vs. receive buffer for a geostationary satellite.*

[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]
(2) T3	39.883	2,424.79
(3) OC3	135.102	8,213.89
(3) OC12	541.966	32,950.27

Notes:
Receive buffer = bandwidth * delay
kbyte = 1024 bytes
RTT = 500 ms
(1) Assumes MTU of 1500 bytes
(2) Assumes DS3 using classical IP over ATM, MTU 9180 bytes
(3) Assumes SONET using classical IP over ATM, MTU 9180 bytes

Table 2. *Receive buffer requirements for a geostationary satellite.*

the round-trip time — approximately 500 ms for a geostationary satellite. *Buff* is the end host computer's receive buffer size, which corresponds to the TCP window size. Tables 1 and 2 are generated from Eq. 1 and show some typical throughput values for a geostationary satellite. Table 1 shows the maximum throughput that can be achieved for a given receive buffer size. Table 2 shows the required receive buffer size for typical WAN links. Current commercial TCP implementations such as those found in Microsoft Windows are set for an 8-kbyte buffer. Thus, the maximum throughput one could obtain through a geostationary satellite is 131.072 kb/s — approximately that of a 128 kb/s modem. Note that for high-bandwidth connections such as an OC-12 (622 Mb/s) connection, a single user requires extremely large buffers at the receiver.

$$T_{max} = \frac{Buff}{RTT} \qquad (1)$$

User Datagram Protocol

The User Datagram Protocol (UDP) is a datagram-oriented transport layer protocol. UDP is an unreliable protocol; it does not have feedback control, and is thus insensitive to delay. However, many reliable applications have been written using UDP as the transport mechanism. These reliable applications may be sensitive to delay.

For our experiment UDP is mainly used as a testing tool to force congestion into the network.

MPEG-2 Audio/Video Compression and Transportation

The MPEG-2 transport stream is a complicated multiplexing protocol that allows multiple programs of video, audio, mixed video and audio, and user specific data to be transmitted in a single stream. MPEG-2 is an unreliable transport protocol and is therefore insensitive to delay. However, a reliable transport protocol can be encapsulated into an MPEG-2 data stream. This encapsulated reliable protocol would be delay-sensitive.

The transport stream is composed of 188-byte packets containing program specific information such as the program association table (PAT), program map tables (PMTs), conditional access tables (CATs), network information table (NIT), program clock reference (PCR), and program element stream (PES) packets. The PES packets contain the element stream data, as well as the program time stamp (PTS) indicating the time that a presentation unit is presented in the system target decoder, and the display time stamp (DTS) indicating the time that an access unit is decoded in the system target decoder [1].

The MPEG-2 transport stream can be segmented into and placed into ATM cells using either ATM application layer 1 or 5 (AAL1/5). For our experiments, the AAL5 segmentation was utilized (Fig. 1) [2].

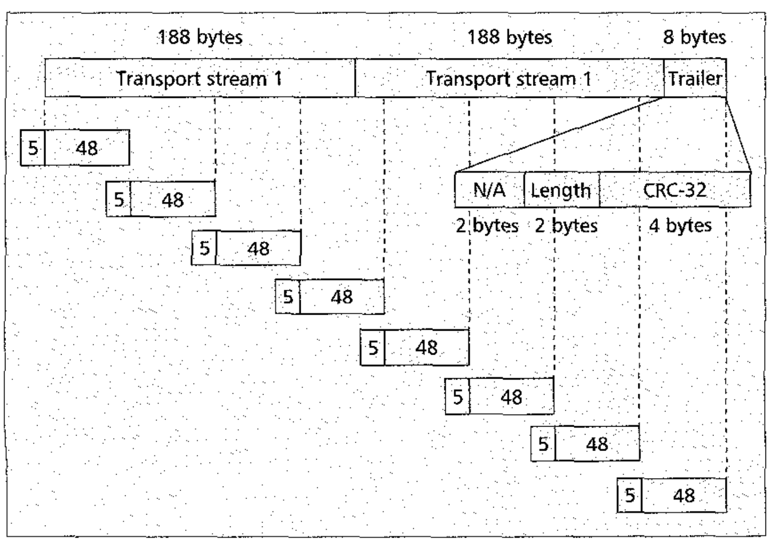

Figure 1. *AAL-5 common part convergence sublayer.*

Quality of Service

Quality of service is application-dependent. Voice quality may be acceptable with bit error rates (BERs) of 10^{-6} or higher. Medical imaging may find 10^{-7} acceptable, since the high-level protocols will resolve any problems resulting from non-error-free transmission. In addition, one has to consider the error distribution characteristic as well as average error rate. For instance, for the same average error rate, burst errors have less effect than random errors on TCP. Thus, the question to be resolved is "What link quality will the satellites need to provide in order to be globally interoperable with terrestrial systems?"

The current version of International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) Draft Recommendation I.356, "B-ISDN ATM Layer Cell Transfer Performance," provides the QoS class definitions and end-to-end network performance objectives. These objectives are given, for each performance parameter, as "upper bounds" (worst-case values) that need to be met on a virtual channel (VC) or virtual path (VP) for the duration of the connection. Many of these parameters are still being debated, in particular the Class I "stringent class" QoS parameters.

In an attempt to address some of these issues and relate the I.356 objective requirements with satellite performance characteristics, NASA Glenn Research Center has performed experiments using MPEG-2 compressed digital video over ATM to determine the link quality that satellites must provide for Class I stringent class services over ATM. In addition, NASA is currently performing ATM QoS research to determine the link qualities necessary for acceptable performance using TCP over ATM.

MPEG-2 over ATM over Satellite

We performed laboratory experiments using MPEG-2 (AAL5) over ATM over an emulated satellite link. The purpose of these experiments was to determine the free-space link quality necessary to transmit high-quality multimedia information using the ATM protocol. The compressed video standard, MPEG-2, was chosen as the baseline application in order to stress the overall link quality. All equipment and protocols are directly traceable to international specifications for clarity, consistency, and repeatability by other researchers.

The testbed setup is shown in Fig. 2. This testbed is extremely flexible and can easily be expanded to run over actual satellite and terrestrial links, as we have direct ATM connections to NASA's Advanced Communication Technology Satellite (ACTS) high-data-rate (HDR) terminal and the NASA Research and Education Network (NREN). Some of the major features of this testbed are embedded in the equipment and its integration into the system. The video switch provides a mechanism for placing markers in the recorded video so that we can identify which test configurations were set for each test. The Broadband Networks MPEG-2 encoder has three MPEG-2 transport stream channels output through either RS-422 or TAXI ports as well as an OC3 optical ATM connection. This aids in monitoring the encoded video before it passes through the test network. Errors can be generated using three different pieces of equipment. Errors can be generated in the IF link using the HP 3708A noise test set, or digitally using the Adtech SX/14 or HP Impairment Module. ATM link statistics can be obtained using either the Adtech AX/4000 or the Hewlett Packard Broadband Series Test System (BSTS). Generally, we chose to use the AX/4000 because it provides cell error ratio (CER) and cell loss ratio (CLR) simultaneously. The Tektronix MTS 100 allows capture or playback of MPEG-2 transport streams. The MST 100 is also

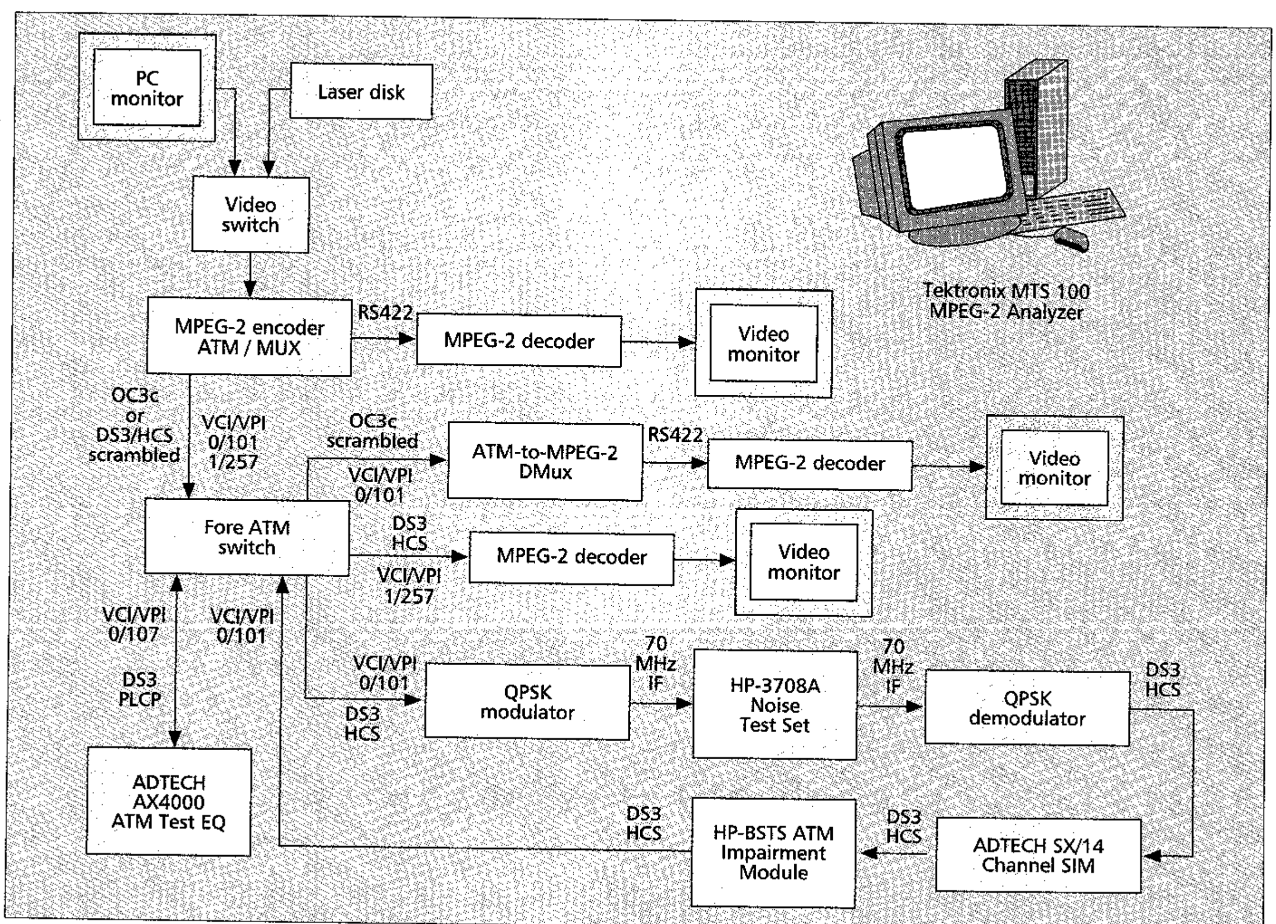

■ **Figure 2.** *MPEG-2 over ATM QoS test configuration.*

used to analyze MPEG-2 transport streams offline. *A laser disk is used for* the video source in order to repeat specific, short video segments without degrading the video, as would occur if we used videotape.

The video takes the following path through the ATM network. It is first encoded and passed onto the ATM switch — the Broadband Networks encoder can provide up to three channels of ATM-encapsulated MPEG-2 video. The video then passes through a combination of video impairment equipment including the satellite modem, the Adtech SX/14, and the HP BSTS impairment module. The video returns to the ATM switch and is forwarded to either the Stellar 1000 set-top box decoder, which has a DS3 interface built in, or to the Broadband Networks ATM-to-MPEG-2 demultiplexer and on to the Broadband Networks decoder. The Adtech AX/4000 produces ATM test cells, which are simultaneously passed through the impairment path and returned for analysis.

[illegible]	[illegible]
Decoder type	Broadband Networks VF-1000D, Stellar 1000
Channel characteristics	
•BER	10^{-5} – 10^{-9}
•Error distribution	Binomial, burst (modem), payload only, header cell loss only
•Modulation format	QPSK 3/4 conv. code, QPSK 3/4 conv. and R/S code, 8-PSK with RS code[1]
AAL type	AAL5, AAL1[1]

[1] Not performed during this study period.

Table 3. *Test variables.*

A series of tests were run to baseline the MPEG-2 video in an error-free and errored environment for both transmission of MPEG-2 transport streams directly over an emulated satellite channel, and for MPEG-2 transport streams over AAL5 over various emulated satellite channels. The variables shown in Table 3 were introduced into the test in a systematic and controlled manner in order to determine which parameters are affected by errored channels. The systematic reduction of variables was necessary in order to reduce the number of permutations necessary for complete and accurate results.

The following tests were performed:
- Encoding rate testing
- MPEG-2 transport stream with errors
- MPEG-2 over ATM with errors
- MPEG-2 over ATM over satellite channel (emulated)
- MPEG-2 over ATM channel characteristics

Experimental results using MPEG-2 compressed digital video indicate that the MPEG-2 decoders temporarily lose synchronization at CLRs and CERs on the order of 1.0E–7 and 4.0E–7, respectively [2]. As an example, the movie "Apollo 13" was videotaped to get some idea of the long-term QoS requirements. An EF-Data SDM 9000 45 Mb/s quaternary phase shift keying (QPSK) modem was configured to utilize 3/4 convolutional coding (without the Reed-Solomon code activated) at an E_b/N_0 of 8.0 dB, which corresponds to a 2.0E–8 BER, 5.0E–7 CER, and 4.0E–8 CLR. In two hours and 20 minutes, there were at least 12 identifiable errors, block errors, or loss of synchronization. Video was also recorded using combined 3/4 convolutional coding and Reed-Solomon block coding, resulting in a BER, CLR, and CER of approximately 1.0E–8, 4.5E–7, and 5.5E–7, respectively. Block errors and synchronization losses averaged approximately four to six per half-hour. Thus, one would expect the CLR and CER requirement to be at least 1.0E–8 and 1.0E–7 or better depending on the customer's acceptance level.

These experimental results have been integrated into ITU — Radiocommunication Standardization Sector (ITU-R) Recommendation S.atm, dealing with broadband integrated services digital network (B-ISDN) ATM performance requirements, and S.atm-av, dealing with B-ISDN ATM availability requirements. The QoS contributions have also been submitted to the T1A1.3 United States standards group which works on ATM performance issues and is responsible for U.S. contributions to the ITU-T on ITU-T Recommendation I.356, "B-ISDN ATM Layer Cell Transfer Performance" [3].

TCP OVER ATM QOS

There is a great amount of interest in understanding the ATM QoS requirements for services such as large file transfers that utilize TCP, particularly over networks with long delay-bandwidth products, because TCP requires acknowledgments for received packets and retransmission for lost packets. NASA is currently performing ATM QoS testing using TCP. In addition to documenting the performance of TCP in an error-prone environment for various error distributions, we are investigating the relationship of the ATM CERs and CLRs for various link error distributions and error rates.

Figure 3 provides a detailed block diagram of the test configuration; note the similarity to Fig. 2. For experiments that do not utilize the satellite modem, the physical layer can be OC3c using the SONET physical layer protocol or DS3 using clear channel (HCS) or PLPC physical layer protocols. For

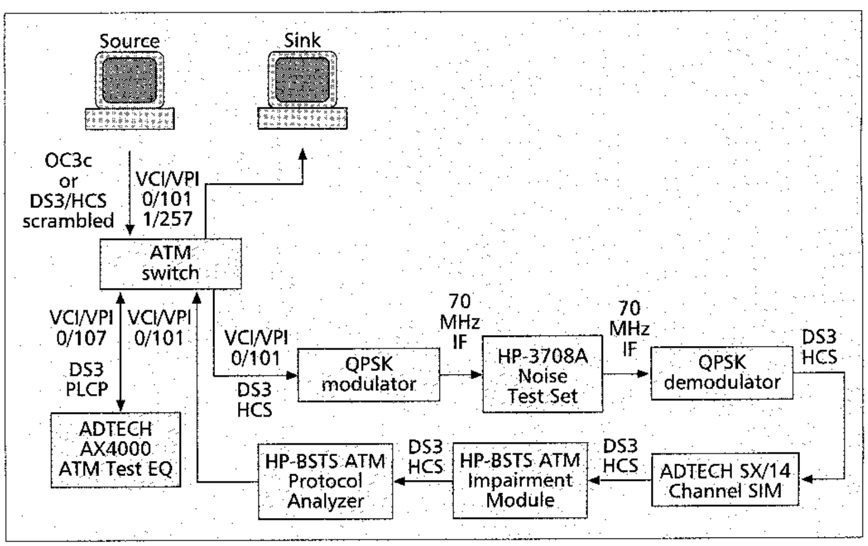

Figure 3. *ATM QoS testing using TCP.*

experiments using the satellite modem and noise test set, only DS3 using HCS or PLPC physical layer protocols can be accommodated due to our modem limitations — 45 Mb/s. Delay and errors can be digitally inserted into the link using either or both the Adtech SX/14 and the HP-BSTS Impairment Module. Analog noise can be inserted into the IF link using the HP-3708A Noise Test Set. The HP-BSTS Protocol Analyzer Module is used to monitor cell loss and AAL5 errors. The Adtech AX4000 is used to verify that all error and delay link parameters have been properly set up prior to performing an experiment.

These experiments were ongoing as of April 1999. The results should be available by August 1999.

TCP OVER SATELLITES

There exists a perception that TCP does not work over satellite networks. This perception is unfounded. In fact, TCP was originally designed to be extremely robust since it was designed in the 1970s to work over the ARPANET and to include operation over satellite networks. However, beginning around the fall of 1986, the Internet began showing signs of congestion collapse. To alleviate this problem, congestion control algorithms such as the slow start algorithm were adopted into the TCP standard implementations. These algorithms have been continually enhanced and provide an elegant solution to congestion control in an environment consisting of multifaceted users operating on a variety of interconnected networks, the Internet. Many congestion control algorithms — slow start in particular — may result in inefficient bandwidth utilization for end-to-end communications where a moderate amount of data is being transferred over a link exhibiting large bandwidth-delay characteristics.

Recently, satellite system providers have shown interest in solving TCP efficiency problems associated with long delays and error-prone links. Similarly, the terrestrial community is interested in solving TCP problems over high-bandwidth links, whereas the wireless community is interested in improving TCP performance over bandwidth-constrained error-prone links.

ACTS EXPERIMENT 118X: TCP INTEROPERABILITY

NASA realized that solutions to TCP had already been proposed for most of the problems associated with efficient data transfer over large bandwidth-delay links (which include satellite links). The solutions are detailed in various Internet Engineering Task Force (IETF) requests for comments (RFCs). Unfortunately, most of these solutions had not been tested at high speeds (155+ Mb/s). Therefore, NASA's ACTS experiments program initiated a series of TCP experiments to demonstrate TCP/IP scalability and determine how far the protocol can be optimized over a 622 Mb/s satellite link. These experiments were known as 118i and 118j. During the 118i and 118j experiments, NASA worked closely with SUN Microsystems and FORE Systems to improve the operating system, TCP stacks, and network interface cards and drivers. We were able to obtain instantaneous data throughput rates of greater than 520 Mb/s using TCP over ATM over a 622 Mb/s synchronous optical network (SONET) OC12 link. Following the success of these experiments and the successful government/industry collaboration, a new series of experiments, the 118x experiments, were developed [4]. The objectives of the 118x experiments were:

- To work in partnership with the computer industry to promote the development of interoperable, high-performance TCP/IP implementations across multiple computing/operating platforms
- To work with the satellite industry to answer outstanding questions regarding the use of standard protocols

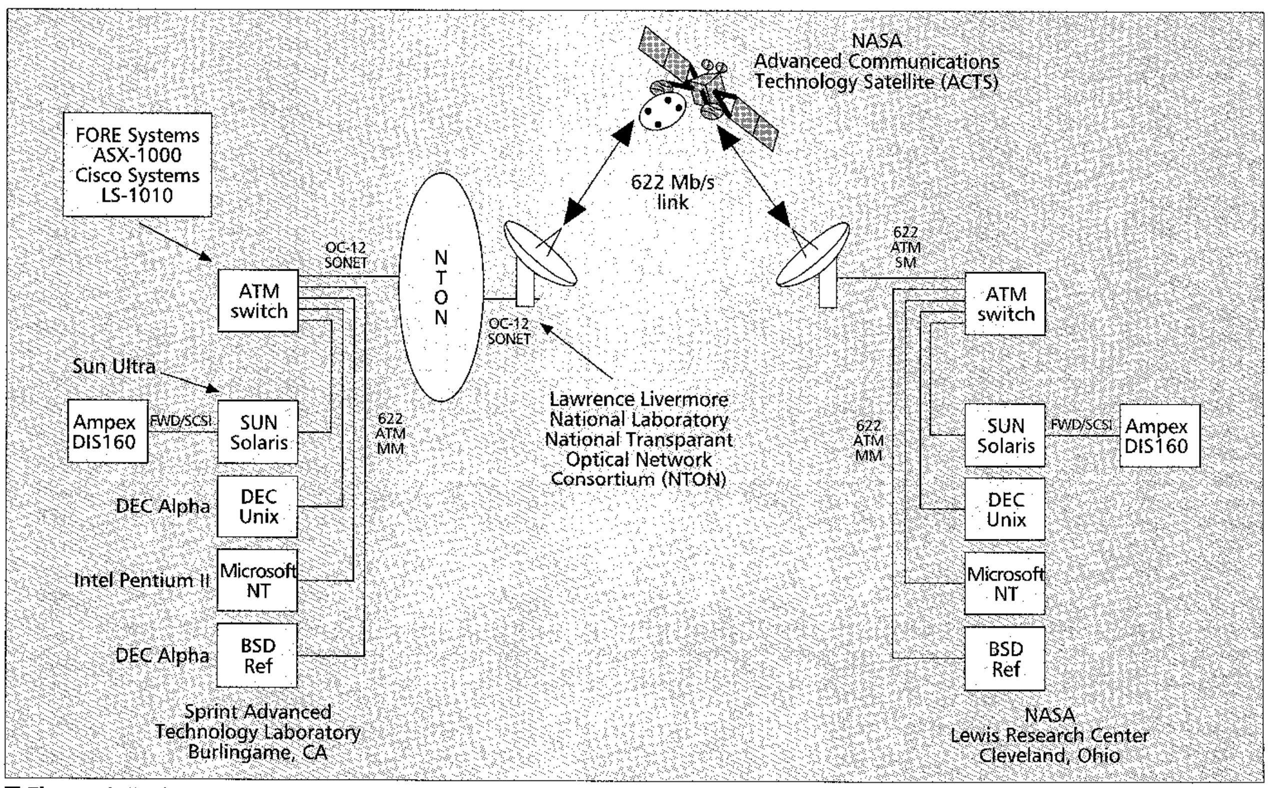

■ **Figure 4.** *End-to-end network layout.*

(TCP/IP and ATM) for the delivery of advanced data services, and for use in spacecraft architectures
- To demonstrate the interoperability of TCP/IP over OC12 ATM over a satellite network in a multivendor environment using ACTS

Figure 4 shows the overall network configuration for the 118x experiments. Two HDR terminals (HDRs) and ACTS provided the satellite link with an effective bidirectional data throughput of 622 Mb/s. The interfaces to the ACTS ground terminals is 622 Mb/s using the SONET physical link protocol. The HDR located at NASA's Glenn Research Center (GRC) in Cleveland, Ohio, was connected directly to a FORE ASX-1000 ATM switch. The workstations and personal computers were also connected directly to this switch. No routers were used for these experiments. The second HDR was located at Lawrence Livermore National Laboratory. The HDR was connected to an ATM switch at Sprint's Advanced Technology Laboratory in Burlingame, California, through the National Transparent Optical Network (NTON) using dense wave-division multiplexing (DWDM) technologies. At Sprint's Advanced Technology Laboratory there was a mirror image of the GRC site. Again, no routers were used at the Sprint site.

We were able to test, to varying degrees, interoperability on four operating systems: Sun's Solaris, Microsoft's NT4 and NT5, Silicon Graphics IRIX, and Compaq's OSF1. Tables 4 and 5 highlight the sustained average throughputs we were able to obtain for symmetric links. Table 4 shows results when operating in a LAN environment with tens of milliseconds of delay. Table 5 shows results when operating over a satellite link with 570 ms delay. The variation of results can be due to TCP and operating system implementations, network interface cards, line drivers, and/or workstation processor speeds and resources. It should be noted that the maximum theoretical throughput for TCP over classical ATM over SONET is approximately 134 Mb/s for a 155 Mb/s link and 537 Mb/s for a 622 Mb/s link when taking ATM and SONET overhead into consideration. These results show the state of the system as of November 1, 1998. One should expect that the systems will become more stable and eventually obtain near theoretical throughput within the next few years, particularly for OC3 links.

We also performed some bulk data transfers over asymmetric links indicative of a relay satellite. Here, the acknowledgment return link bandwidth was constrained in order to determine the ratio of data transmission bandwidth to acknowledgment bandwidth required for bulk data transfers. These tests were only performed using the Solaris operating system. We were able to achieve data throughputs of 317 and 181 Mb/s with acknowledgment links of 2.0 and 1.0 Mb/s, respectively, thus showing a ratio of greater than 150:1 data throughput to acknowledgment for single-user bulk data transfers.

[illegible]	[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]	[illegible]
IRIX	622	155	465[2]
NT4/NT5	622	622	Unstable
Solaris	622	622	473

[1] OSF1 transmitting data, Solaris acknowledging.
[2] Solaris transmitting data, IRIX acknowledging.

Table 4. *Data throughput for TCP over ATM over SONET over a 570 msec delay.*

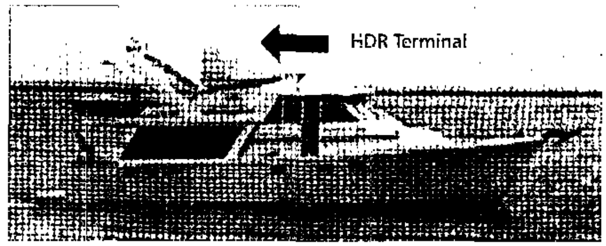

Figure 5. *Shipboard terminal inside the Chicago Breakwall.*

Experiment #149: Shipboard HDR Satellite Communication Demonstration

In October 1998, NASA Glenn Research Center and the Naval Research Laboratory embarked on a collaborative effort designed to investigate the technical feasibility of a bidirectional mobile network operating at 45 Mb/s (22 times faster than the previous record)! This network involved a moving vessel at sea [Fig. 5], a fixed ground station, and access into the terrestrial networks (e.g., the Internet). NASA Glenn Research Center's participation was relevant to the Center's collaboration and technology transfer efforts as well as being extremely relevant to NASA's Human Exploration and Development of Space (HEDS) Enterprise technology thrusts for mission services, high-performance communications, and interoperability. The Naval Research Center's participation was relevant to the Office of Naval Research Communications Technology Program efforts to enhance satellite communication connectivity to naval forces, including investigating those technologies and system attributes pioneered by ACTS that will be commercially available in the future. Specific technical goals included conventional and advanced tracking comparisons, operating tracking schemes in a stressed environment (vessel at sea, inclined orbit satellite), extrapolation of tracking scenarios for use in future low/medium earth orbit (LEO/MEO) satellite communication systems, and the investigation of high-speed data transfer applications. These applications included TCP/IP-based file transfers, interactive and variable TCP/IP-based multimedia, production-quality video, and CD-quality audio occurring simultaneously in a roll-pitch-yaw environment (especially challenging in waters outside the breakwall in Lake Michigan at this time of year). The achievements of this NASA-NRL collaboration were so successful that system modifications and new technical goals are already being considered for additional testing with ACTS [5].

Packet over SONET

NASA Glenn and Cisco Systems are currently performing satellite WAN research using packet-over-SONET (POS) technology. One goal is to compare ATM and POS technologies over a satellite channel to determine the overall improvements in bandwidth utilization obtainable by using POS instead of ATM. A second goal is to determine if the QoS tools available in IP can provide similar performance to that of ATM. NASA mis-

[illegible]	[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]	[illegible]
NT4/NT5	622	622	357
Solaris	622	622	400–500

Table 5. *Data throughput for TCP over ATM over SONET in a LAN environment.*

sion planners will use this information when considering which WAN technologies to utilize on large space platforms such as the Shuttle or International Space Station.

For both ATM over SONET and POS, available information bandwidth is identical. For OC2 (155.520 Mb/s) the information bandwidth is 149.76 Mb/s after the SONET section and path overhead is removed. POS uses Point-to-Point Protocol (PPP) as the data link layer protocol. The overhead for PPP in POS is 10 bytes/IP datagram consisting of the following fields: start flag (one byte), address (one byte), control (one byte), protocol (two bytes), CRC (four bytes), and end flag (one byte). ATM has an initial cell overhead of 5 bytes/48-byte payload. ATM has additional overhead to map IP into ATM. The overhead to map IP to ATM AAL5 requires the following fields: padding (0–47 bytes), common part converge sublayer (CPCS) user–user interface (one byte), common part indicator (one byte), length (two bytes), and CRC (four bytes). The combination of the IP datagram, the eight bytes of CPCS trailer, and the padding bytes must be a multiple of 48 bytes, which is why the padding varies between 0 and 47 bytes. For large IP datagrams, most of the ATM penalty resides in the ATM header inefficiency. For very small IP datagrams, the padding inefficiency also becomes a major factor. In general, we anticipate a bandwidth efficiency improvement of 15–20 percent for POS vs. IP over ATM over SONET depending on the type of data being transmitted.

Figure 6 shows the POS test configuration. UDP and TCP traffic are sent through an IP switch, a router, and an emulated satellite WAN using either ATM or POS. Here we use UDP traffic to congest the buffers in the router in order to remove the interaction of the buffers from the ATM or POS overhead calculations. The Cisco Catalyst 5000 IP switches are configured to provide two virtual LANs. We are then able to transmit two 100 Mb/s data streams into the Cisco 7507 router which enables us to congest the OC3 outputs. Either the ATM or POS interface is configured to be administratively down while the other interface is being tested. This allows us to test either WAN interface without having to physically remove hardware or cabling.

These tests were ongoing as of April 1999. Results should be available by August 1999.

CONCLUSION

This article presented work that has been accomplished and is ongoing regarding ATM and IP technologies including ATM vs. POS and ATM QoS. A summary of ACTS 118x TCP interoperability experiments and the ACTS 149 shipboard experiment is provided. In addition, research is continuing in the area of ATM QoS using TCP and POS. The results of these two investigations should be available by the summer of 1999 at the Satellite Networks and Architectures Branch Web site [6].

This broadband research is being reported to a variety of standards organizations, including the ATM Forum, the Satellite Communication Division of the Telecommunication Industry Association, the ITU, and the IETF. NASA strongly encourages the space communication industry to enhance its participation in these standards organizations.

For the past 20 years NASA Glenn's communication research has been focused on maintaining U.S. preeminence in satellite communication. Recently, numerous filings have been made with the ITU to place into orbit onboard processing communication satellites using both Ku and Ka-bands. These systems include multisatellite constellations such as the Teledesic/Motorola, Hughes Spaceway, and Alcatel

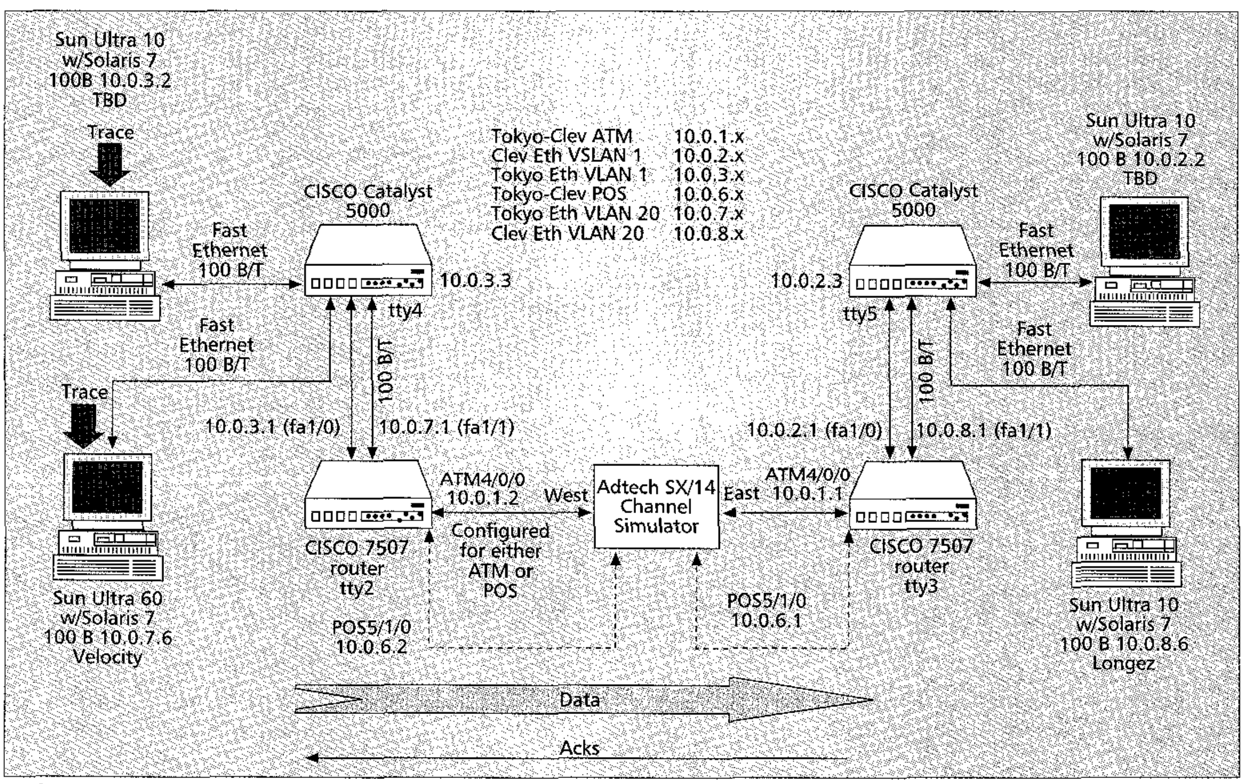

■ **Figure 6.** *Packet-over-SONET test configuration.*

Espace/Loral Space Skybridge systems. In addition, systems such as Iridium and Globalstar either are operational or will be shortly. Thus, it appears that the commercial satellite communication industry has matured to the point where it can carry on without additional government funding for research and technology development. As a result, the NASA Glenn Research Center is shifting its communication technology research and development focus toward NASA-specific mission requirements.

ACKNOWLEDGMENTS

The authors would like to thank Rich Groh for his support in setting up and maintaining the high-data-rate terminals and his help on the shipboard experiments, and Dave Stewart for his support in maintaining the high-data-rate terminals.

REFERENCES

[1] ISO/IEC 13818, "Information Technology — Generic Coding of Moving Pictures and Associated Audio Information Systems," 1994.
[2] ITU-T Draft H.222.1.
[3] W. D. Ivancic, B. Frantz, and M. Spells, "MPEG-2 Over ATM Over Satellite QoS Experiments: Laboratory Tests," U.S. (T1A1) contribution T1A1.3/97-084 to ITU-R Working Party 4B, Philadelphia, PA, Oct. 1997.
[4] D. E. Brooks *et al.*, "118x Final Report: High-Speed TCP Interoperability Testing Over ACTS," TBP, Mar. 1999; available at http://mrpink.lerc.nasa.gov/papers.html
[5] Shipboard High Data Rate Satcom Experiment #149, http://mrpink.grc.nasa.gov/shipboard/shipboard.html
[6] Satellite Networks and Architectures Branch, http://ctd.grc.nasa.gov/5610/5610.html

BIOGRAPHIES

WILLIAM D. IVANCIC (wivancic@grc.nasa.gov) has been with NASA for more than 16 years. He received his B.S.E.E. and M.S.E.E. form Cleveland State University in 1982 and 1986, respectively. He is a senior research engineer at NASA's Glenn Research Center working in the satellite communications arena. HIs work includes advanced digital and RF design, communications networks, satellite onboard processing, and system integration and testing. In addition, he has been responsible for the development of proof-of-concept onboard processing satellite equipment, ground terminal equipment, and numerous pieces of special test equipment. His recent research has been directed toward evaluating and applying commercial communication protocols such as ATM and TCP/IP to space-based networks.

DAVID BROOKS has been a network engineer with Sterling Software for five years. He received his B.S.E.E. in 1992 and M.S.E.E. in 1997 from Cleveland State University. He has performed extensive research on high-speed TCP/IP and ATM-based terrestrial and satellite networks to support hardware and software application research with NASA's Advanced Communication Technologies Satellite (ACTS) and other long-delay, high-bandwidth environments. His present work includes research and testing of "commercial off the shelf" equipment interoperability in long-delay, high-bandwidth environments. In addition, he is working on advanced communication research for NASA's Advanced Air Transportation Technologies (AATT) project.

BRIAN FRANTZ (bfranz@grc.nasa.gov) has been with Sterling Software for four years as a NASA Support Service contractor. He received his B.S.E.E. in 1992 and M.S.E.E. in 1994 from Cleveland State University while working on a thesis entitled "A reprogrammable FPGA-based ATM traffic generator which emulates an MMPP model." He is a network engineer working in the Telecommunications and Networking and Satellite Communications Branches at NASA Glenn Research Center. His work includes network design, integration, and analysis of high-performance LANs and WANs, and research on the impact of network errors on MPEG-2 video for remote echocardiograph analysis. His recent research has been focused on ATM QoS requirements for TCP/IP over errored LAN, WAN, and satellite environments.

DOUG HODER (dhoder@grc.nasa.gov) has been involved with designing and engineering RF communications systems for the last 17 years. He has worked in the military electronics, television, and cellular phone industries for the Georgia Tech Research Institute, Scientific-Atlanta, and Motorola, Inc. He is the co-inventor of several U.S. and foreign patents. He holds degrees from Cleveland State University (B.E.E., 1983) and the Georgia Institute of Technology (M.S.E.E.,1986). He has been with NASA Lewis since 1990. At NASA he managed the development and operation of the ACTS Gigabit Satellite Network. He is currently developing a wideband satcom system for aeronautical communications, navigation, and surveillance use. His current interests are in the development of wideband hybrid networks and high-performance graphics.

DAN SHELL (dshell@cisco.com) has been with Cisco Systems for five years as a systems engineer in the Federal Division. Prior to working at Cisco he was lead WAN network administrator at BP America and GE Lighting. He has been active in IP over satellite research at NASA Glenn. His field of concentration at Cisco has been IP security, IPdiff serv, and ATM. His recent research efforts have been directed toward voice over IP, weighted RED, and POS using high-delay links and high data rates.

DAVID BEERING (drbeering@sprynet.com) is a principal of Infinite Global Infrastructures, LLC, based in Chicago, Illinois. His current work focuses on the design, integration, and deployment of large-scale high-performance virtual networks. Prior to his current position, he spent five years with AT&T Communications and seven years with Amoco Corporation, both in Chicago. He is presently working extensively with NASA's Glenn Research Center and Johnson Space Center to develop reference architectures, coordinate research, and conduct demonstrations on communications networks using TCP/IP and ATM. Much of his experimental work focuses on NASA's ACTS operating at rates of 155 and 622 Mb/s. He also serves as a network architect for Lockheed Martin Corporation's Consolidated Space Operations Contract (CSOC) team.