TECHNICAL RESEARCH REPORT

Experiments in Hybrid Networking with the Advanced Communications Technology Satellite

by D. Friedman, S. Gupta, C. Zhang, A. Ephremides

CSHCN T.R. 96-8
(ISR T.R. 96-26)

CENTER FOR SATELLITE AND HYBRID COMMUNICATION NETWORKS

EXPERIMENTS IN HYBRID NETWORKING WITH THE ADVANCED COMMUNICATIONS TECHNOLOGY SATELLITE

Daniel Friedman, Sonjai Gupta, Chuanguo Zhang and Anthony Ephremides

Center for Satellite & Hybrid Communication Networks
Institute for Systems Research
University of Maryland
College Park, MD 20742
(301) 405-3641

Abstract

This paper describes experiments conducted over ACTS that were motivated by the commercial potential of low-cost receive-only satellite terminals operating in a hybrid network environment. The first experiment tested highly adaptive methods of satellite bandwidth allocation in an integrated voice-data service environment. The second involved comparison of FEC and ARQ methods of error control for satellite communication with emphasis on the advantage that a hybrid architecture provides, especially in the case of multicasts. Finally, the third experiment demonstrated hybrid access to databases through the use of Mosaic and compared the performance of internetworking protocols for interconnecting LANs via satellite.

Introduction

The Center for Satellite and Hybrid Communication Networks (CSHCN) at the University of Maryland was founded to explore and develop the commercial possibilities of satellite and combined satellite/terrestrial communication technologies. In many communication applications there is a need to transmit much information primarily in one direction between two points and much less information in the opposite direction. This is typically the case in file transfer and database services. While a two-way satellite channel may be used for such asymmetric applications, it is also possible to support them by means of a combination of a one-way satellite channel for the bulk information transfer and a parallel terrestrial channel for the low-bandwidth portion of the application traffic. In such a hybrid network--with parallel satellite and terrestrial channels--the satellite terminal need not have transmit capability, and thus it can be a much less expensive receive-only terminal. In addition to a cost savings, it is possible that improvement in network performance may also be achieved since the terrestrially-carried traffic does notsuffer the high propagation delay incurred through satellite links.

These advantages of hybrid networks, coupled with the availability of low-cost, receive-only satellite terminals, suggest commercial potential for hybrid networks. To develop this potential, the CSHCN proposed in 1992 a series of experiments in hybrid interconnection of local area networks (LANs). The scope of the experiments was later expanded to include other facets of hybrid networking, and three experiments were ultimately devised. The first one examined dynamic allocation of satellite bandwidth in response to variations in the amount of traffic to be sent through the satellite. The second experiment investigated error-control schemes for use in both point-to-point and point-to-multipoint hybrid networks. Finally, the third experiment considered using a hybrid network architecture for remote multimedia database access and also compared the performance of some networking protocols in local area network interconnection.

All the experiments were based on a two-node satellite network configuration, shown in Figure 1. One node was at the University of Maryland at College Park, and the other was at the University of Colorado at Boulder. The experiment equipment configuration at Maryland consisted of two Sun workstations connected through high-speed serial interfaces to the FRACS, and the FRACS was in turn connected to the ACTS T1 VSAT. The FRACS-to-VSAT connection consisted of a T1 connection for traffic, and an RS-232 connection for control messages (such as call setup and bandwidth allocation commands). The workstations ran CSHCN-developed software to implement the bandwidth allocation algorithm, the error-control schemes and the multimedia database server. SunLink Frame Relay software was used for the frame relay connections between the workstations and the FRACS.

A similar arrangement was used at the University of Colorado, except for one important difference. The Colorado T1 VSAT was not collocated with the FRACS but instead was on the premises of the National Telecommunications and Information Administration (NTIA), also in Boulder. An optical link was used for the T1 traffic connection between the University of Colorado and NTIA sites, and a modem link was used for the FRACS control messages.

Figure 1: Configuration for ACTS experiment.

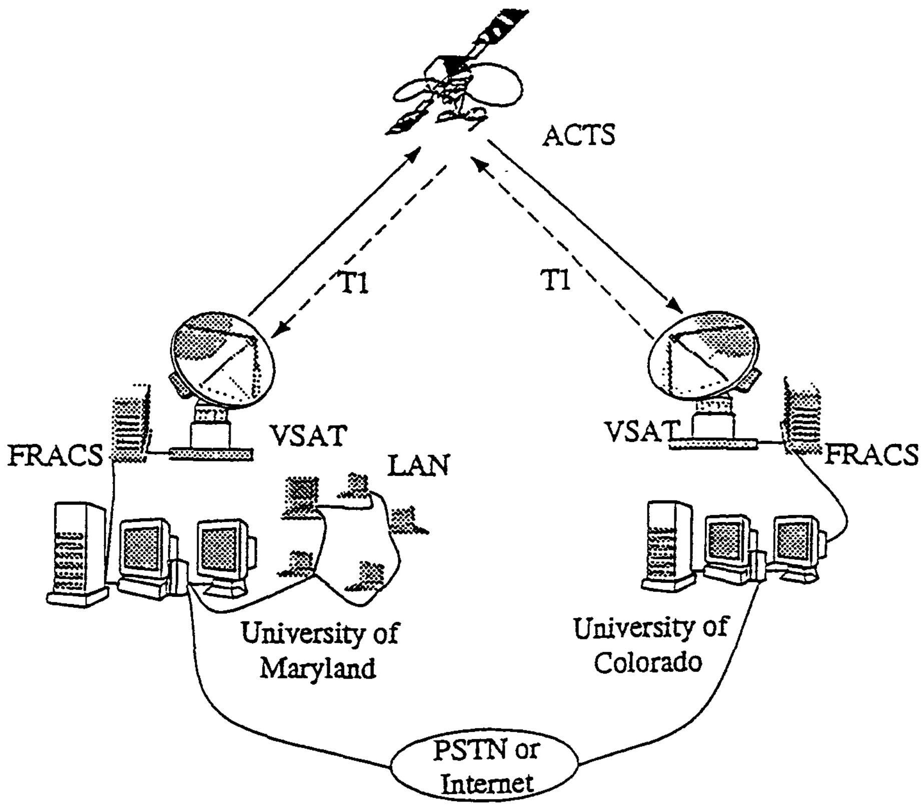

Experiment 1: Dynamic Bandwidth Allocation

The purpose of the first experiment was to find effective methods to rapidly adapt the link bandwidth to fluctuating traffic levels, both for data traffic and for mixed-media traffic such as packetized voice and data. This adjustment of bandwidth was accomplished using a feature of the ACTS system, the ability to establish circuits of different bandwidths, in multiples of 64 kbit/s channels. The rationale behind the experiment is to request, and use, the minimum necessary amount of bandwidth that will permit achievement of satisfactory performance levels, and release it when not needed. Since the performance criteria are several and antagonistic to each other, there is a need for fine-tuned trade-offs amongst them.

The bandwidth allocation algorithms investigated for data traffic were a rate-based algorithm implemented in the FRACS (Frame Relay Access Control Switch) and a threshold-based algorithm of our own design. For mid-media (voice and data), an algorithm that generalizes the concept of thresholds and is based on so-called switch-functions was tested, both for a single station and in a setting of two stations competing for bandwidth. The FRACS was used to signal the ACTS system for requesting or releasing channels either according to its own algorithm, or upon instruction by our experiment software, for our algorithms.

The transmitter and the receiver may either be located at separate nodes, or they may be collocated at the same node, but on separate computers. It was important to run the transmitter and receiver software packages on separate workstations, in order to avoid any delays caused by the sharing of the computer processor. All the workstations involved in the experiment were isolated from the network to prevent unauthorized users from stealing processor time.

The software for the transmitter and the receiver was developed in-house. The transmitter software consisted of a source, a "packetizer" to append index numbers to the data packets, a "transmit-logger" to store the relevant statistics, a "transmit-buffer" for the data packets arriving from the source, a "threshold-monitor" to implement the appropriate bandwidth allocation algorithm, and a transmitter to send out the packets at the appropriate rate. The statistics stored by the "transmit-logger" included the times of packet arrivals into the "transmit-buffer", the times when bandwidth requests were placed, the times when bandwidth changes actually took place, etc..

A problem presented by the ACTS system for this experiment was the significant delay in allocating the requested bandwidth. The encountered delay for allocating bandwidth using the ACTS system varied between two and 15 seconds and was not fixed, but changed as a function of the overall load on the ACTS network. To simulate a system in which the delay is either fixed and/or smaller, and hence to assess the effect of such delay on performance, a scheme was devised and tested for pre-allocating, through request from the ACTS system, all necessary satellite channels, but using only those required according to our dynamic bandwidth allocation algorithms. Under such a scenario, the amount of delay used was completely under our control.

Results

This experiment has not yet been completed. However, a few early results are available and summarized below.

Each run of the experiment involved sending approximately 4 million bits of data. This data was organized in 4000 packets of 126 bytes each. It was decided to send 4000 packets, in order to get reproducible results, as determined by earlier tests. Information about the send and receive times of the data packets was stored in the transmit and receive loggers, and the data packets were destroyed once they reached the receiver. For the case of p[re-allocated channels and data-traffic only, the rate-based algorithm (FRACS's algorithm) had a lower average per packet end-to-end delay, but the threshold based algorithm used a lower number of channels. Further tests are being conducted to allow a fairer comparison of the two algorithms.

Some of the results obtained are shown in Table I. The buffer size used was 100 packets. An MMPP source was used. The four values shown indicate performance corresponding to different settings of the threshold values.

The rate-based algorithm was also run, both with and without memory. The results are shown in Table II. The three values shown indicate performance for the rate-based algorithm of the FRACS for three different internal configurations.

As can be seen from these numbers, the threshold-based algorithm suffers greater delay because it uses bandwidth more economically. The FRACS algorithm tends to request additional channels earlier than the threshold-based one and thus pays a higher penalty in bandwidth usage. This phenomenon is common in any service system. The relatively large difference in delay performance is due to the fact that the offered input load rate is close to the average channel usage for the parsimonious algorithm and, thus, tends to load the system to near-capacity levels. Additional tests in progress are expected to show a much smaller discrepancy in these values.

Table I: Threshold-based algorithm results.

Delay (seconds)	Average Channel Usage	Offered Load Rate (kbit/s)
0.837	1.82	58
0.883	1.81	62
0.866	1.78	76
1.878	1.19	70

Table II: Rate-based algorithm results.

Delay (seconds)	Average Channel Usage	Offered Load Rate (kbit/s)
0.346	1.98	57
0.386	2.34	82
0.414	1.98	82

Implementation Considerations

In order to accurately measure the end-to-end delays for data packets, it was necessary for the clocks on the transmitting and receiving computers to be synchronized to within a few milliseconds. To achieve such synchronization, and to combat natural drifting of the clocks, XNTP (Network Time Protocol, Version 3) was used. The Maryland computers were synchronized to stratum--2 time servers at the University of Maryland campus, while the Colorado computer was synchronized to a stratum 1 server in Boulder, Colorado. Synchronization of Maryland and Colorado computers to within a few milliseconds was thus achieved.

In order to do real-time channel allocation, requests were forwarded to the FRACS, which sent out the appropriate signals to the ACTS system. However, the combination of FRACS signaling and the Call Manager software of the T1 VSAT was unable to handle frequent requests. When the transmitter and the receiver stations were collocated at the same node, only one outstanding request was allowed at any given point in time. Thus, parts of the experiment had to be redesigned, and their testing is under way.

When using our bandwidth allocation algorithms, it was necessary to poll the FRACS to

Table III: Residual bit error rates.

Noise Effect Bit Error Rates	Plain Text	BCH (15, 11)	BCH (15, 7)	Golay (23, 12)
3.16×10^{-3}	2.042×10^{-3}	9.990×10^{-5}	1.180×10^{-5}	1.300×10^{-6}
10^{-3}	8.163×10^{-4}	1.160×10^{-5}	1.100×10^{-6}	0
3.16×10^{-4}	2.877×10^{-4}	1.100×10^{-6}	0	0
10^{-4}	9.970×10^{-5}	0	0	0
3.16×10^{-5}	3.210×10^{-5}	2.000×10^{-7}	0	0
10^{-5}	1.010×10^{-5}	0	0	0
0	0	0	0	0

determine when bandwidth changes occurred, since the FRACS was not capable of reporting this information on its own. This polling was conducted by a portion of the transmitter software and thus it ran on the same workstation, competing with the date-transmission functions for processor time, and thus slowing the traffic generation and transmission processes. On the other hand, no such polling was needed while running the rate-based algorithm, since the FRACS itself ran this algorithm and kept a log of bandwidth changes. However, for the purposes of a fair comparison, FRACS polling was conducted in all configurations of this experiment.

Experiment 2: Error Control Schemes for Satellite and Hybrid Networks

As mentioned earlier, this program of experiments was motivated by the commercial potential of hybrid networks. Information transferred using a hybrid network must be protected against errors just as in a satellite network. However, the availability in a hybrid network of a terrestrial link with less propagation delay than the satellite channel presents additional problems as well as possibilities for error control, particularly in the case of automatic-repeat-request (ARQ) schemes. The second experiment explored such possibilities. Furthermore, a satellite is an excellent means for point-to-multipoint communication. Hence this experiment also investigated ARQ error control schemes for multicast communication in hybrid networks.

This experiment was additionally motivated by the possibility of improving throughput by sending ARQ acknowledgments terrestrially instead of by satellite, thus avoiding the satellite propagation delay for the acknowledgments. Throughput might be increased yet more by retransmitting packets (as may be necessary) terrestrially instead of by satellite. Calculations showed that judiciously using the hybrid configuration could indeed increase the throughput in some cases, sometimes significantly.

Data was sent from one station to another via satellite. Before transmission, an error control protocol was applied by the transmitter to protect the data as it traveled through the satellite channel. Since the ACTS channel, when used with a T1 VSAT, typically exhibits a bit error rate of 10^{-7} or less, it was necessary to inject artificially-produced noise in order to study error control schemes. An independent, identically-distributed model of bit inversion was used for the artificial noise. After corruption by artificial noise, the receiver applied the error control protocol to the to correct any errors which may have developed in the data. The data was then stored for later comparison with the original data to determine a residual bit error rate.

Both forward error correction (FEC) and ARQ error control schemes were investigated. For all the error control methods tested, a continuous source of traffic was used. The FRACS was used as an interface between workstations and VSATs. The Internet and a 14.4 kbit/s telephone modem connection were each employed as the terrestrial link for hybrid operation. The parameters of interest were throughput, end-to-end delay, and the residual error rate of the data received. Not only were the results from FEC testing compared with each other, as were the point-to-point ARQ results, but the results of these two parts were compared as well.

For FEC testing, the data was sent via satellite from one workstation at Maryland to another beside it. For point-to-point ARQ testing, a terrestrial link was sometimes required, so using two adjacent workstations would yield an unrealistically small terrestrial link propagation delay. Hence the Colorado station was used for transmitting to the Maryland station during such testing. During point-to-two point ARQ testing, a workstation at Maryland transmitted to another beside it and to one in Colorado, with a software delay used for terrestrially-carried messages from the Maryland workstation.

To make meaningful statistical inferences from the data, about 100 error events (corrupted bits) were deemed minimally necessary. Hence at least 10 million bits would have to be sent through the channel when operating the noise effects at the lowest bit error rate (10^{-5}). A more meaningful approach, which was adopted, was to send at least 10 million information bits from one end of the system to the other. Hence, in all testing, at least 10 million information bits were transferred.

Figure 2: Remote multimedia database access with a hybrid network.

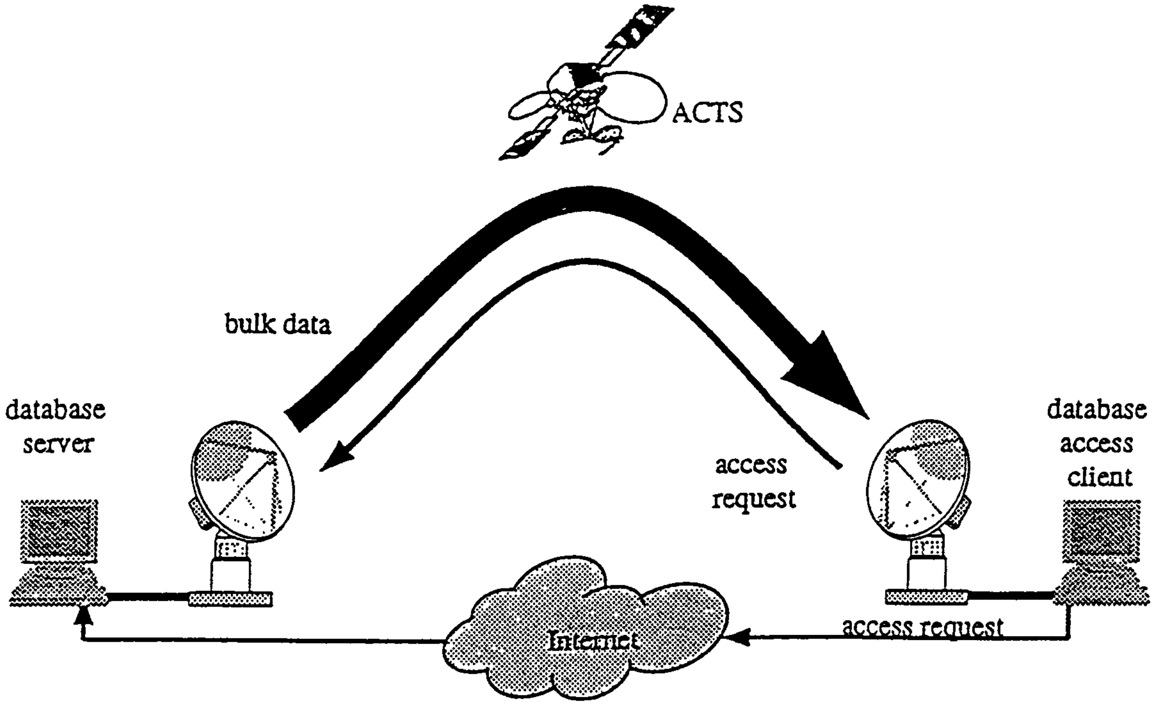

Results

The FEC portion of the experiment has been nearly completed, while the ARQ portions are not yet complete. The residual bit error rates achieved by sending 10 million encoded bits through the satellite channel, including the noise effects generator, are given in Table III. These results will develop their value when compared with the results of the ARQ portion of the experiment.

One result already available from the ARQ portion of the experiment is that the 14.4 kbit/s modems introduce significant delay in transporting bits from one end of a telephone link to another. This delay, even after subtracting propagation delays, is about 140 ms in each direction. This delay is speculated to be due to compression/decompression of the data and the modems' trellis-coded modulation scheme, particularly the Viterbi decoding. The precise causes for this delay have not yet been determined. As the delay is a significant fraction of the time required to send a bit between two points via geostationary satellite, it is possible little throughput improvement over a satellite network might be achieved by using such a modem link as the terrestrial link in a hybrid network. This finding does not, however, diminish the significant cost savings achievable with the hybrid architecture instead of a pure-satellite architecture.

Experiment 3: Remote Database Access and Internetworking Protocol Comparison

In the third experiment, the hybrid network mechanism of data access was demonstrated and evaluated against the use of solely satellite connection or solely Internet connection. Two internetworking protocols were applied in the satellite link, namely CCITT's standard X.25 protocol and the frame relay protocol. The logic and the performance of these protocols were studied and compared. We focused on the data link layer functionalities since this is where the key differences of those two protocols reside. In addition, we studied the commonly used transport layer protocol, TCP, over the satellite link; in the process, a problem of using TCP was identified, and a solution was provided, which was to use an extended version of TCP (Jacobson and Dorman 1992). Finally, and most importantly, two emulated LANs were interconnected by the satellite link, and a comparison of the performance of X.25 and frame relay used for their interconnection is being carried out as the culmination of the objectives of this experiment.

Remote Database Access via a Hybrid Network

A common communication need in many applications is to access a computer network with a personal computer for the purpose of transferring

Table IV: Access results.

Internet:

Database Size (bytes)	Delay (seconds)	Throughput (kbyte/s)
5,299,770	170	30
5,299,770	132	40
3,311,616	240	14
103,488	9.5	11

Satellite:

Database Size (bytes)	Delay (seconds)	Throughput (kbyte/s)
103,448	6.5	16
206,976	9.1	22
413,952	14	29
827,904	24	34
1,655,808	44	37
3,311,616	84	38
6,623,232	170	39

Hybrid:

Database Size (bytes)	Delay (seconds)	Throughput (kbyte/s)
103,488	4.2	24
103,488	4.1	24
206,976	5.9	34
206,976	6.3	32
413,952	9.2	44
413,952	8.8	46
827,904	15	53
827,904	16	50
1,655,808	90	18
1,655,808	33	49
1,655,808	29	56
1,655,808	28	57

files and databases. While a telephone modem can be used to connect to the network, often times the low bandwidth of such a connection is insufficient. A leased line of higher bandwidth (56 kbit/s, 1.544 Mbit/s T1, etc.) is an alternative, but an expensive one. A satellite link is also an option, but the required ground terminal can also be expensive. However, typical users of computer networks consume much more information form the network than they send to it. For such users with asymmetric bandwidth requirements the low volume traffic to be sent to the network can be carried over a telephone modem connection instead of the satellite, while the bulk traffic from the network can be carried over a higher-bandwidth satellite channel. The transmit capability of the satellite terminal may so be eliminated, thus significantly reducing the cost of the terminal.

The first part of this experiment demonstrated this concept by remotely accessing multimedia database using a hybrid network (Figure 2). In particular, all short interactive access requests were routed through the Internet, and the bulk information from the database was routed through a T1 satellite link. ASCII text, audio and image files, ranging in size from 50 kilobytes to 6.4 megabytes, were transferred using the "point-and-click" Mosaic user interface. The performance obtained when using only Internet varied with the time of day, with throughputs as low as several kilobytes per second during peak network traffic hours up[to 40 kbyte/s. Operation through the satellite exhibited much less variation, with throughput of about 39 kbyte/s, since the satellite link was a dedicated and directly-connected link. The throughput through satellite usually exceeded that through Internet, when Internet exhibited stable throughput. Table IV provides some of the specific measurement results.

The utilization of the T1 satellite link was also monitored during the satellite and hybrid tests. A utilization of 27% was achieved in the satellite test, but 41% utilization was achieved in the hybrid test. The reason for this is the acknowledgment delay of 290 ms through satellite was reduced to an average of 70 ms when an Internet terrestrial link was used, and this reduction significantly affected the window-based flow control applied by the transmission layer of the application.

In parallel with this work, and motivated by the same desire to support asymmetric communication with low-cost receive-only satellite terminals, the CSHCN helped Hughes Network Systems develop a recently-announced product named *DirecPC*. With this product, a user can send an information request to a computer network (in particular, Internet) using a modem and have the information returned to him via a satellite link. To accomplish this, it is necessary the computer accessed over the network return requested information to a satellite uplinking computer instead of to the computer connecting the user's dialup connection to the network. This required special software for encapsulating Internet TCP/IP frames within TCP/IP frames with suitably modified addresses so that the requested information would be routed properly. The CSHCN contribution to this product has been recognized by the University of Maryland Office of Technology Liaison as Outstanding Invention of 1994.

Protocol Comparisons

We first considered the relative merits and performance of the two basic protocols for internetworking, namely X.25 and frame relay. Following that, we considered TCP and extended TCP over satellite and finally we engaged in the comparison of the performance of these protocols for interconnecting two LANs.

The ultimate objective of this experiment is to emulate two LANs interconnected over the satellite link, on which X.25 and frame relay protocols are applied as the internetworking protocol. The performance of particular applications in this interconnected system, over the two different protocols (either X.25 or frame relay), are to be

Figure 3: LAN interconnection.

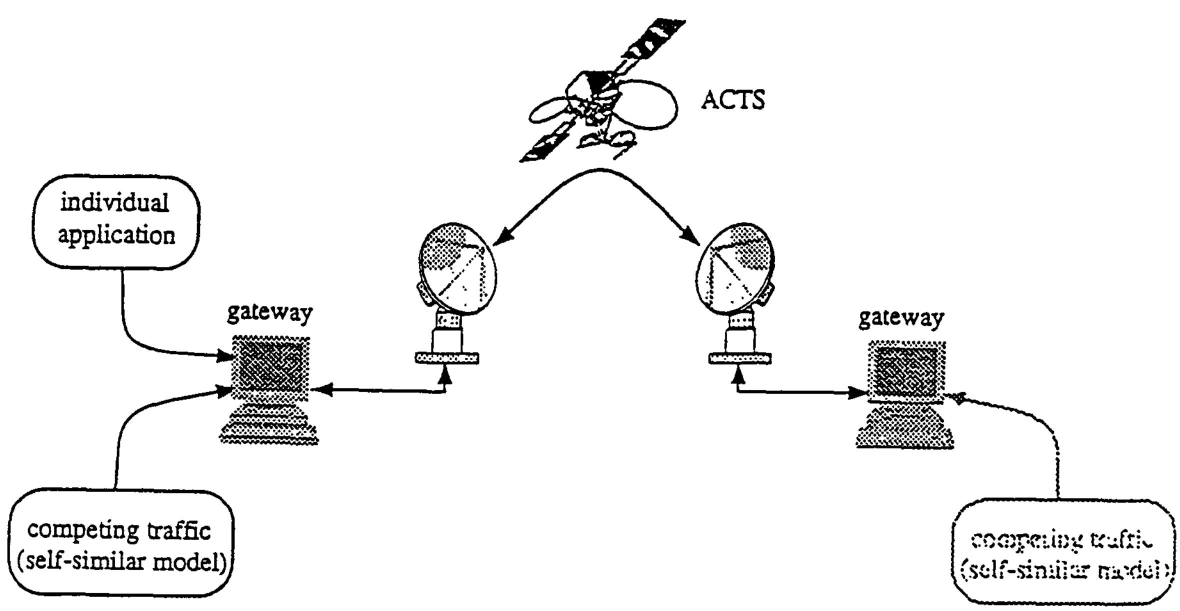

compared. Specifically, the average packet delay of each individual application is the performance criterion.

The emulation of the LANs was composed of two components; the first was emulation of the background traffic of a LAN, and the second was the particular application that we are interested in. The background LAN traffic was emulated by using the so-called Self-Similar model that is based on recent Bellcore studies. It was found to be a more realistic traffic model for LAN traffic. The particular application messages were merged with this emulated LAN traffic and fed into the network gateway, through which they were sent over the satellite link using as internetworking protocol, either X.25 or frame relay (Figure 3).

As an extension, a two hop LAN interconnection was tested. Here, on the receiving end (now it is actually the intermediate node), another Self-Similar traffic source was built to represent the LAN traffic, and a portion of this traffic, considered as the inter-LAN traffic which is intended for the third node (ultimate destination), was merged with the incoming traffic from the first gateway. All these traffic streams were sent back through the satellite link to the first gateway, which virtually represented the third gateway. After the particular application reached the final destination, the end-to-end performance of it was logged. In this two-hop transmission over the satellite link, either the X.25 or frame relay was used as the link layer protocol.

This portion of the experiment has not yet been completed.

Conclusion

As of this writing, the experiments have not yet been completed, largely due to uncontrollable factors. The ACTS VSAT suffered several problems, including incorrect installation, cable breakage, two complete failures of its transmit power amplifier and partial failures of its control computer. The weather in Colorado prevented experimentation several times by disrupting the optical link. Conducting experiments with equipment 1500 miles away has presented its own challenges, including operating computers through the unpredictable Internet. Despite these difficulties, work continues in earnest and we expect completions in the near future.

Acknowledgments

This work would not have been possible without the kind and helpful support, including the loan of the T1 VSAT, of ACTS program personnel at the NASA Lewis Research Center. Further, the authors were only part of a larger team which designed, implemented and conducted the experiments. This team included: Dr. Apostolos Traganitis of the University of Crete, Greece; Timothy Kirkwood of the CSHCN; Yannis Konstantopoulos of MCI; Dr. Anil Agarwal, Moorthy Hariharan and Anita Fitzwater of COMSAT Laboratories; Dr. Stanley Bush, Fan Huang, Gil Weiss, Neill Cameron, Wenjin Yang and David Bailey of the University of Colorado at Boulder; Marjorie Weibul of NTIA; and Ruplu Bhattacharya, Samardh Kumar and Anastassios Michail of the University of Maryland.

References

Fisher, W. and K. Meier-Hellstern (1993) "The Markov-Modulated Poisson Process (MMPP) Cookbook," *Performance Evaluation*, 1:18, September 1993.

International Standards Organization, *International Standard ISO 4335, "Information processing Systems - Data communication - High-level data link control procedures - Consolidation of elements of procedures"*.

International Standards Organization, *Proposed Draft International Standard ISO 7776/DAM 2, "Information processing systems - Data communication - High-level data link control procedures - Description of the X.25 LAPB-compatible, DTE data link procedures - DAM 2: Multi-selective Reject Option"*.

Jacobson, R.B.V. and D. Borman (1992) *TCP Extensions for High Performance (RFC 1323)*, May 1992.

Leland, W.E. , M.S. Taqqu, W. Willinger, and D.V. Wilson (1994) "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Trans. Networking*, 2:1-15.

Leland, W.E. , M.S. Taqqu, W. Willinger, and D.V. Wilson (1993) "On the self-similar nature of Ethernet traffic," In *Proc. ACM Sigcomm '93*, 183-193.

Lin, W. and P.R. Kumar, (1994) "Optimal control of a queueing system with two heterogeneous servers, "IEEE*Trans. Automatic Control*, 29:696-703, August 1994.

Rubinovitch, M. (1985) "The slow server problem: A queue with stalling," *J. Appl.* Prob., 22:879-892.

Mills, D.L. (1992) *Network Time Protocol (Version 3) Specification, Implementation and Analysis*

(RFC 1305) University of Delaware, March 1992.

Viniotis, I. and A. Ephremides (1988) "On the optimal dynamic switching of voice and data in communication networks," in *Proc. Computer Network Symposium*, 8-16., Apri l 1988.

Viniotis, I. (1985) "Optimal control of a multiserver queueing system with heterogeneous servers," Master's thesis, University of Maryland, College Park.

Walrand, J. (1994) "A note on 'Optimal control of a control of a queueing system with two heterogeneous servers'," *Systems and Control Letters*, 4:131-134, May 1994.