



Communications
Research Centre
Centre de recherches
sur les communications

**ACCORD Broadband ATM Satellite Experiment
(BASE)
*-protocols characterization***

by

**Isabelle Labbé, Louis Gravel, Gérard Nourry,
Corey Pike, John Butterworth, Gretchen Bivens,
Sean Lanigan and Brian Spink**

19980602 116

CRC REPORT NO. 97-008

December 1997
Ottawa



Industry
Canada

Industrie
Canada

The work described in this document was sponsored by the Defence R&D
Branch of the Department of National Defence under work unit 5CB11.

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

Canada

**ACCORD Broadband ATM Satellite Experiment
(BASE)**

-protocols characterization

by

**Isabelle Labbé(DRX), Louis Gravel(DRX),
Gérard Nourry(DRX), Corey Pike(DSST)
and John Butterworth(DSAP)**
Communications Research Centre

and

**Gretchen Bivens(Rome Lab), Sean Lanigan(Rome Research)
and Brian Spink(Rome Lab)**
Rome Laboratory

The work described in this document was sponsored by the Defence R&D Branch of the Department of National Defence under work unit 5CB11.

INDUSTRY CANADA
CRC REPORT NO. 97-008

Canada

December 1997
Ottawa

ABSTRACT

A Broadband ATM satellite Experiment (BASE) was carried out between the Communications Research Centre, Canada and Rome Laboratory, USA during the period of May 95 to March 96. The objectives of the experiment was to perform a series of tests to characterize ATM over broadband satellite bearers.

*This document reports on the second phase of the experiment that was conducted from December 95 to March 96 over the NASA Advanced Communication Technology Satellite (ACTS) satellite (Ka-band). Tests consisted in throughput measurements of TCP over a broadband ATM satellite link as a function of Bit Error Rate (BER). The throughput performance of TCP/IP was very poor due mostly to the (bandwidth * delay) product being limited by the TCP window size implementation. The results are presented along with the influence of others factors such as the Maximum Transmission Unit (MTU) size and packet size. Finally, a performance comparison with a special implementation of TCP for network with large delays and high bandwidth (TCP for Long Fat Networks - TCP-LFN) is presented.*

RÉSUMÉ

Une expérience de satellite MTA à large bande a été effectuée conjointement par le Centre de Recherches sur les Communications (CRC) du Canada et Rome Laboratory des Etats-Unis, de mai 1995 à mars 1996. L'objectif était de procéder à une série de tests afin de caractériser les liens satellites MTA à large bande.

*Dans le présent document, il est question de la deuxième phase de l'expérience, effectuée de décembre 1995 à mars 1996 en utilisant le satellite ACTS (bande Ka) de la NASA. Les tests ont consisté à prendre des mesures de capacité du protocole TCP sur un lien satellite MTA à large bande en fonction du taux d'erreur sur les bits. La performance de capacité du protocole TCP/IP a été médiocre surtout à cause du produit (largeur de bande * délai), qui était limité par l'implantation de la dimension de la fenêtre TCP. Dans ce document, les résultats obtenus sont présentés en considérant aussi l'influence qu'on put avoir d'autres facteurs tels la dimension du MTU et la dimension du paquet. Enfin, on présente une comparaison de performance avec une implantation spéciale de TCP (TCP-LFN) pour réseaux à large délai.*

EXECUTIVE SUMMARY

ATM is becoming an important network backbone technology. For many years to come, ATM will have to be interfaced with legacy network as well as carry the traffic originating from various types of networks.

The Transmission Control Protocol (TCP) and Internet Protocol (IP) form a protocol pair which is widely used in both civilian and military networks. Much work has been performed on issues related to interfacing TCP/IP and ATM, however, little information is available on the characterization of TCP over ATM broadband satellite links.

Results obtained from the characterization of ATM over satellite links (described in the companion reports) have shown that the Cell Loss Ratio (CLR) is not negligible over degraded links. Responsibility to recover from cell loss is then left to the higher layers (network, transport, or application) which, in order to provide reliable services, will often need to use retransmission schemes. Satellite links introduce significant delays which could, in those situations, have a serious effect on application performance. It is thus important to study how well high layer protocols can be interfaced to an ATM satellite link.

As part of the TTCP ACCORD project, the Communications Research Centre (CRC)/Defense Research Establishment Ottawa (DREO), Canada and Rome Laboratory, USA have agreed to perform a series of tests to characterize broadband ATM satellite bearers.

The Broadband ATM satellite experiment (BASE) was divided into two phases. The objective of the first phase was to characterize broadband ATM channels and embedded sub-channels over Ku- and Ka-band satellites. The second phase aimed at characterizing the performance of standard networking protocols such as TCP/IP over broadband ATM Ka-band satellite links.

The first phase was conducted in May 95 over the AnikE satellite (Ku-band) and from July 95 to September 95 over the NASA ACTS satellite (Ka-band). Tests consisted in measurements of ATM Quality of Service (QoS) parameters such as Cell Loss Ratio (CLR) and Cell Error Ratio (CER) as a function of Bit Error Rate (BER). Details of this work can be found in documents CRC-RP-97-007 and CRC-TN-97-008 respectively.

This document presents the second phase of the experiment: the networking protocol characterization performed from December 95 to March 96 over the NASA ACTS satellite.

Table of Contents

	ABSTRACT/RESUME	iii
	EXECUTIVE SUMMARY	v
	LIST of FIGURES/LIST of TABLES	ix
1.0	INTRODUCTION	1
2.0	BACKGROUND.....	1
3.0	EXPERIMENTAL SETUP	2
4.0	TRANSMISSION CONTROL PROTOCOL	4
4.1	TCP Overview	4
4.2	Transmission Scheme.....	4
4.3	Performance Related Issues.....	6
4.3.1	Protocols Definition	6
4.3.2	Protocol Implementation	8
4.4	TCP/IP Tests	8
4.5	TCP Results	9
4.5.1	TCP results over fibres.....	9
4.5.2	TCP results over satellite	10
4.5.3	TCP-LFN over satellite.....	11
5.0	UDP	11
5.1	Overview of the UDP protocol.....	11
5.2	UDP format	12
5.3	UDP Issue	12
5.4	UDP tests	12
5.5	UDP results.....	13
6.0	CONCLUSIONS	13
7.0	REFERENCES	14

List of Figures

1	Experimental Satellite Configuration	3
2	Data Flow	5
3	Protocol Stack.....	6
4	Protocol Overhead	7
5	Experimental Laboratory Configuration.....	9
6	UDP Datagram format.....	12
7	Application throughput with TCP over fibre	16
8	Application throughput with TCP over satellite	16
9	Application throughput with TCP over degraded satellite link	17
10	Application throughput with TCP-LFN over degraded satellite link	17
11	Application throughput with UDP over fibers and satellite links.....	18

List of Tables

1	Bandwidth Available to Layers -DS3 link.....	6
---	--	---

1.0 INTRODUCTION

As part of the TTCP ACCORD project, the Communications Research Centre (CRC) from Canada and Rome Laboratory from the USA have agreed to perform a series of tests related to Asynchronous Transfer Mode (ATM) over broadband (DS3) satellite bearers. These tests were conducted in two phases.

The objective of the first phase was to conduct Quality of Service (QoS) performance measurements of ATM over a Ku- and Ka-band satellite links. The first portion of the first phase lasted 10 days and was completed in May 1995 using the Anik-E satellite. A report [1] describes the experiment and gives the preliminary results.

The second portion of the first phase was completed during the period extending from July to September 1995 over the ACTS satellite. A report [2] gives the results that were obtained.

This document reports on the second part which consisted in characterizing the performance of standard transport and network protocols such as TCP/IP and UDP/IP over broadband ATM satellite links. This was completed during the period of December 1995 to March 1996 over the ACTS Satellite (Ka-band). This report is organised as follows: Section 2 gives the background. Section 3 describes the experimental configuration. Section 4 provides the TCP overview, tests, and results. Section 5 addresses UDP, while Section 6 supplies conclusions.

2.0 BACKGROUND

ATM is becoming an important network backbone technology. For many years to come, ATM will have to be interfaced with legacy network as well as carry the traffic originating from various types of networks.

The Transmission Control Protocol (TCP) [3] over Internet Protocol (IP) form a protocol pair which is widely used in both civilian and military networks. Much work [4], [5], [6] has been performed on issues related to interfacing TCP/IP and ATM. However, little information is available on the characterization of TCP over ATM broadband satellite links.

Results obtained from the characterization of ATM over satellite links have shown that the Cell Loss Ratio (CLR) is not negligible over degraded links [1][2]. Responsibility to recover from cell loss is then left to the higher layers (network, transport, or application) which, in order to provide reliable services, will often need to use retransmission schemes. Satellite links introduce significant delays which could, in those situations, have a serious

effect on application performance. It is thus important to study how well standard protocol pairs will perform when interfaced to an ATM satellite link.

3.0 EXPERIMENTAL SETUP

The hardware configuration for this experiment can be seen in Figure 1. It was kept as simple as possible to minimize potential source of problems that could affect the measurements. Proper action was taken to ensure that the workstations operated in a stand alone mode to prevent external traffic from degrading the system performance (e.g. by introducing congestion, by taking CPU time, etc...). Also, the ATM switches were isolated, i.e. no traffic other than the experimental traffic could travel through.

The CRC configuration included the following pieces of equipment:


1. SPARC 10 Workstation: The workstation used the SunOS 4.1.4 operating system. It was equipped with an ATM SBA-200 Network Interface Card (NIC) from FORE Systems. It was configured to limit the outgoing traffic to a peak rate of 40 Mbps in order to prevent buffer overflow at the ATM switch.
2. ATM switch: The ATM switch was the ASX-200 from FORE Systems. This device was required to interconnect the workstation (OC-3) and the satellite modem (DS-3).
3. Fibre Loop Converter: Converted DS-3 coaxial signal to DS-3 fibre-optic signal or vice versa. Conversion was required because of the half kilometre distance between the ATM equipment (switch) location and the satellite ground station (EF Data Modem).
4. EF Data Modem: Performed the modulation/demodulation as well as the error detection/correction scheme for satellite communication (45 Mbps).

The Rome Lab configuration included the following pieces of equipment:

1. SPARC 10 Workstation: Same as described for CRC configuration with the exception that the operating system was SunOS 4.1.3.
2. ATM switch: The ATM switch was a GTE SPANet. This device was required to interconnect the workstation (TAXI 140 Mbps) and the satellite modem (DS-3).
3. Adtech AX/4000: ATM generator/analyser configured to operate in Line Monitor mode. It was used as a DS-3 signal regenerator to overcome the attenuation problem due to the distance between the ATM equipment location and the satellite ground station (EF Data Modem).

The TCP/UDP throughput measurements were performed using the "nttcp" software tool developed by FORE Systems. This application was specifically written for network bench-

CRC, Canada



Rome Lab, USA




Figure 1: Experimental Satellite Configuration

mark purposes and is public domain shareware. Among the various parameters it allows to be set, the ones that were of particular interest for this experiment are:

1. size of packets written to the socket interface.
2. total number of packet to be sent.
3. Send/Receive buffer size for TCP window.
4. Nagle's algorithm enable/disable (TCP_NODELAY option),

4.0 TRANSMISSION CONTROL PROTOCOL

4.1 TCP Overview

The Transmission Control Protocol (TCP) provides a reliable transfer service over the Internet Protocol (IP). TCP is a connection oriented protocol. This means that a connection must first be established between two End Systems (ES) before any data can be sent across the communication path. The reliability is provided by the implementation of segment acknowledgements and retransmissions. TCP guarantees delivery of data between machines without duplication, misordering, or packet loss. The transmission of data is performed using a windowing scheme (sliding window). It allows for end-to-end flow control by having the receiver advertise the available space in its window [7].

4.2 Transmission Scheme

The basic flow of information between two hosts using TCP as the transfer mechanism is shown in Figure 2 [8]. The application sends data to the communication protocols through a socket interface using the "write" system call. The data is copied to a buffer list and ends up in the TCP Send Buffer Window, as long as there is space available. If there is no more space, the application waits for the tcp_output to receive an acknowledgement which will free space in the Send Buffer Window. TCP then makes a packet and sends it to the IP Layer. The packet will include a portion of the data which will vary in size from 1 byte to the Maximum Segment Size (MSS). IP appends its header and sends it to the Device Driver Layer. If the IP packet size is greater than the Message Transfer Unit (MTU) size, the packet will be segmented. The Device Driver Layer will use the IP packet and "stuff" it into whatever protocol is used at the link layer. For ATM specific, it will be the sub-layer convergence protocol, which is, in this case, an AAL5 frame. ATM cells are then sent over the physical medium. At the receiving end, the reverse process is applied and as the TCP segments are being put together, acknowledgements are sent back to the sender, advertising the space available in the Receive Buffer.




Figure 2: Data Flow

Figure 3 illustrates the protocol stack utilized during the experiment.




Figure 3: Protocol Stack

4.3 Performance Related Issues

Performance issues for TCP over ATM can be classified into two categories:

- how the protocols are defined in the specifications.
- how the protocols are implemented.

4.3.1 Protocol Definitions

At each layer, packets are encapsulated with specific protocol overhead (in the form of headers/trailers) which varies in size and in information content. Figure 4 depicts the overhead added at each layer. The total amount of Protocol overhead can become at times significant and limit the application performance [9]. For a DS3 link (45Mbps), Table 1 shows the amount of bandwidth that remains after each successive layer has claimed its share of the overhead. Each row in the table shows how much bandwidth is available to the indicated protocol layer.

	Without PLCP ^a (Mbps)	With PLCP (Mbps)
Line Rate	44.736	44.736
DS3	44.209	40.72
ATM	40.04	36.88
AAL 5	37.5	33.81
IP	37.4	33.7
TCP	37.1	33.4
Application	36.8	33.1

TABLE 1. Bandwidth Available to Layers - DS3 link

a. Physical Layer Convergence Protocol




Figure 4: Protocol Overhead

As can be seen from Table 1, both PLCP and ATM reduce significantly the throughput available. Although the line rate is 44.736 Mbps the actual rate available to the application itself is closer to 33Mbps after all the overhead has been added.

4.3.2 Protocol Implementation

The following elements will have an impact on the application throughput results:

- a) TCP Send and Receive Window Size. For most implementations the TCP header uses a 16 bit field to report the receive window size to the sender. The largest window size will then be 2^{16} bytes, or 65536 bytes. The practical number, as mentioned in "Unix Network Programming" [10] and tested in our laboratory, is closer to 52224 bytes (51 Kbytes). The application can only send segments up to one window without receiving an acknowledgement. This is usually not an issue for a network with little delay. The situation becomes different for satellite communications (1 hop) where it takes around 0.25 seconds for the data to reach its destination and the same amount of time to receive the acknowledgement. In this case, this means that two windows, or 102 kbytes, would be transmitted every second which is far from the available bandwidth. TCP LFN (Long Fat Network) which implements TCP extensions for high performance [11] provides a solution to this problem.
- b) Maximum Segment Size (MSS) and Maximum Transfer Unit (MTU). The MTU size has a direct influence on the MSS. Comer [7] mentions that TCP usually computes a MSS so that the IP packet will match the network MTU for end systems on the same network. If the IP packet is bigger than the MTU size, this will result in segmentation which means more overhead. This overhead, however, is negligible compared to the ATM and DS3 overhead.
- c) Buffer Management Strategy and Data Copy Algorithm at the Socket Layer Interface. This is implementation dependant and has been discussed in [12]
- d) Timers. The round trip time (RTT) is the basis to calculate the retransmission timeout (RTO). An accurate dynamic determination of RTO is essential to good TCP performance.
- e) Implementation of Segmentation and Reassembly (SAR) Sublayer on ATM Network Interface Card. The SAR will usually be implemented with a certain amount of buffer space to help in the reconstruction of the AAL frame.

4.4 TCP/IP Tests

In order to study the effect of delay and error rate, two series of tests were conducted:

- a. Lab configuration (Figure 5), negligible delay over fibres.
- b. Satellite configuration (Figure 1), high round trip delay (~525ms).

Both series of tests were performed with and without PLCP framing embedded into DS3 frame structure.




Figure 5: Experimental Laboratory Configuration

4.5 TCP Results

For each configuration (over fibres and over satellite), the application throughput was measured for various MTU sizes and various packet sizes. Over satellite, these experimental measurements were recorded over an error free link ($BER < 10e-10$) as well as over degraded links (i.e. $10e-8 < BER < 10e-5$). The tests were performed with the ATM Adaptation Layer AAL5.

4.5.1 TCP results over fibres

In order to minimize the delay component and to gain some insight into the basic behaviour of TCP/IP over ATM, initial testing was first performed over a simple laboratory setup (Figure 5). The performance of TCP over AAL5 ATM was measured between two Sparc10 workstations. Results of the protocol throughput for various MTU sizes is shown in Figure 7. The MTU size determines the TCP Maximum Segment Size (MSS) that will be assigned at connection time. The MSS is normally set to the network MTU minus the size of the TCP/IP header which is 40 bytes. The TCP segments may thus vary in size from 1 up to MSS bytes.

According to Figure 7, it is found that the application throughput is better for a greater MTU size (9188 bytes). Given some combined fixed overhead (TCP, IP, and AAL), a larger MTU size incurs a smaller percentage of overhead and thus, an improved throughput. Therefore, under error free and minimal delay conditions, a large transmission unit is desirable.

The previous performance measurements were recorded for two cases: with and without PLCP framing embedded into the DS3 frame structure. Results were found to be in close agreement with the theory. While the available bandwidth to the application layer was cal-

culated to be 33.1 Mbps with PLCP and 36.8 Mbps without PLCP (ref. Table 1), the experimental results gave respectively 33.3 Mbps and 35.7 Mbps. Since PLCP framing introduces significant overhead and is furthermore not recommended to be used over burst error links [1], the satellite tests were conducted without PLCP framing.

4.5.2 TCP results over satellite

TCP performance depends not only upon the transfer rate itself, but rather upon the product of the transfer rate and the round-trip delay(RTT). This (bandwidth * delay) product measures the amount of data that would fill the “pipe”. In our case, for a satellite channel with a typical round-trip delay of 0.5 sec, the TCP window required to fill the DS3 “pipe” would be:

$$(BW * RTT) = \text{TCP window size}$$

$$(45 \text{ Mbits/sec} * 0.5 \text{ sec}) = 2810 \text{ kbytes}$$

However, Sun’s implementation limits the TCP window size to 51 kbytes [10]. The measured round-trip delay of the satellite channel was approximately 0.525 second, hence the maximum practical throughput which could be achieved would be:

$$(51 \text{ kbytes} * 8 \text{ bits/bytes}) / 0.525 \text{ sec} = 796 \text{ kbps}$$

Figure 8 shows the throughput results obtained for various MTU sizes over an error-free ($BER < 10e-10$) broadband satellite link. The maximum experimental throughput obtained was 770 kbps which is in close agreement with the theory. However, this result was obtained with the smallest MTU size (1500 bytes) as opposed to the results obtained in the lab which gave the highest throughput when the MTU size was the largest (9188). The combination of a large delay and a limited window size results in transmission idling time, i.e. in time spent waiting for an acknowledgment. This situation seems to favour the use of smaller MTU sizes. The buffering management between layers, the data copy strategy and the TCP window update/acknowledgment algorithm are all complex elements that are implementation-dependant. Also, when transmission idling time is involved, their combined effects on the overhead and protocol mechanism would require further study [12].

In order to study the performance of TCP over degraded links, throughput measurements were recorded for BER ranging from $10e-10$ up to $10e-5$. Figure 9 shows the results obtained for three different MTU sizes. The break point appears to be located where the BER is approximately $10E-7$. For higher BERs, results drastically deteriorate. If a cell is in error or a cell is lost, the AAL at the receiving end is not able to perform proper reassembly of the data which results in the retransmission of the entire TCP segment. Many cells are therefore retransmitted unnecessarily. Thus, as the BER increases, the application throughput performance decreases rapidly.

All of the previous satellite measurements were recorded with Nagle’s algorithm turned off. Nagle’s algorithm was introduced as a solution to unnecessary header and processing overhead resulting from small TCP segments. It prevents sending TCP segments that are

smaller than the assigned MSS if any previously transmitted data on the connection remains unacknowledged. Because of the large delay involved with satellite communications, it is thus preferable, to increase the efficiency, to operate with Nagle's algorithm turned off. Under BSD Unix, the default TCP configuration uses Nagles' algorithm. However, it can be turned off by the `TCP_NODELAY` socket option.

4.5.3 TCP-LFN over satellite

For comparison purposes, the same tests were performed with a special implementation of TCP for Long Fat Networks (TCP-LFN)¹. TCP-LFN implements RFC 1323 features [11], which includes window scaling, time stamps, and protection against wrapped sequence numbers. The performance results can be seen in Figure 10.

By making use of the window scaling option, the TCP send/receive window size was set to the "optimal" size (~ 3000 kbytes) and the application throughput results increased to 30 Mbps for a BER < 10E-10. This experimental result was to be expected since the transmission idling time is now practically non-existent. As was the case with the laboratory measurements, the channel is now filled with data and the highest throughput is obtained with the largest MTU size (9188).

Results show a break point being now located around a BER value of 10E-8. This earlier break point can be attributed to the fact that the "pipe" is now constantly filled with user data (cells carrying actual information) and thus more vulnerable to burst errors. In the normal implementation of TCP, the idle time between transmissions increases the number of empty cells and thus reduces the probability of burst errors hitting user cells. However, a comparison between Figures 10 and 9 show that even over degraded links (BER of 10E-6), TCP-LFN performance remains higher (1.3 Mbps) than TCP over an error-free link (0.770 Mbps).

5.0 UDP

5.1 Overview of the UDP protocol

The User Datagram Protocol (UDP) is a connectionless transport protocol that provides a primary mechanism to application programs to send datagrams to other application programs. Like TCP, it uses the Internet Protocol (IP) to transport a message however, it does not use acknowledgements to ensure that messages arrive at the destination, it does not order incoming messages and it does not provide feedback to control the rate at which information flows between machines (end-to-end flow control). Thus, UDP messages can be lost, duplicated, or arrive out-of-sequence. Furthermore, packets can arrive faster than the recipient can process them, causing buffer overflows. To summarize, UDP provides an unreliable connectionless delivery service between end-to-end users.

1. The Solaris implementation of TCP-LFN was used for this part of the experiment.

Therefore, it is up to the Application program to take full responsibility for handling problems of reliability. Despite all this, UDP can still be advantageous for real-time applications such as voice and video, for which delay is critical. As long as the packet loss remains under a reasonable ratio, UDP, for those specific applications, will perform much better than TCP.

5.2 UDP format

The UDP datagram consists of two parts: a UDP header and a UDP data area. As Figure 6 shows, the header is divided into four 16-bit fields that specify the port from which the message is sent, the port to which the message is going, the message length and a checksum.

UDP SOURCE PORT	UDP DESTINATION PORT
UDP MSG LENGTH	UDP CHECKSUM
DATA	
...	

Figure 6: UDP Datagram format

As described for TCP (Section 4.2), the Application program exchanges information with the UDP layer through a socket interface. UDP Datagrams are then encapsulated/decapsulated in/out of an IP packet for transmission/reception across/from the media (such as ethernet or ATM).

5.3 UDP Issue

As mentioned above, UDP accepts a user packet and sends it to IP after encapsulating it with an UDP header. In most implementations, there is no buffering between UDP and IP. IP segments at the sender and reassembles at the receiver. So for each UDP packet, multiple IP packets may be generated. Therefore, a resource constraint anywhere on the transfer path may cause IP segments to be dropped leading to the loss of entire UDP packets.

In the configuration used for the experiment, loss of data occurs because of the non-availability of the kernel buffers (at the IP level). In order to prevent UDP from sending the data at high rates and thereby minimize buffer overflows, we incorporated within the application software (nttcp) a simple delay between packet writing. The delay value is “tunable” to obtain the optimal throughput without loss of data.

5.4 UDP tests

UDP being unreliable, tests were conducted on error-free links (i.e. BER less than $10e-10$) to minimize as much as possible loss of information. Mainly for comparison purposes with TCP, the following two tests were performed:

- a) Lab configuration (Figure 5), negligible delay over fibres.
- b) Satellite configuration (Figure 1), high round trip delay (~525ms).

5.5 UDP results

For each configuration (over fibres and over satellite), the application throughput was measured for various MTU sizes. Results were recorded only when the nttcp software showed no loss of data i.e. the total amount of bytes received was the same as the total amount of bytes transmitted. This requirement was considered sufficient to allow proper throughput comparison between TCP and UDP.

Figure 11 shows the application throughput using the UDP/IP protocols. The following points are worth mentioning:

1. The results over a satellite channel were identical to the ones obtained over fibre. This was expected since the implementation of the UDP protocol does not include any windowing scheme, therefore the delay component does not have any impact on the performance.
2. There is a small improvement in the application throughput as the MTU size increases. As explained in section 4.5.1, given some combined fixed overhead, a larger MTU size incurs a smaller percentage of overhead. Therefore, a large transmission unit is desirable.
3. The performance of UDP is comparable to the results obtained with TCP-LFN over error-free link. However, the UDP protocol remains unreliable since it does not guarantee the delivery of every packets.

There were no tests performed over degraded links.

6.0 CONCLUSIONS

This paper presented the results of transport protocols used in conjunction with the ATM protocol over a broadband satellite channel. The most important conclusions emerging from the experiment are:

- The throughput performance of TCP/IP over ATM is very poor due mostly to the (bandwidth * delay) product being limited by the TCP window size implementation. When used

over a broadband satellite channel, a modified implementation, such as TCP-LFN, is recommended.

- TCP-LFN performance was found to be highly superior to TCP. Even over highly degraded links (BER of $10E-6$), the TCP-LFN throughput remains higher than any of the TCP results. Clearly, the use of an appropriate window size (i.e. large enough to insure that the pipe is filled), even though more susceptible to burst errors, produces much better results than when transmission idling time is involved.

- A BER better than $10E-7$ (TCP) or $10E-8$ (TCP-LFN) should be maintained in order to avoid significant degradation of the throughput.

- In most cases, a larger MTU size gives better throughput. This can be explained by the ratio of overhead/user data being smaller for larger segments. However, when the "pipe" cannot be filled, the smaller MTU size appears to give the better results. Further studies would be required to explain this situation.

- UDP over (nearly) error-free links gives good performance results but does not provide the reliability required for data transfers. However, for time-critical applications such as voice and video, where a small packet loss ratio can be tolerated, UDP could be used with success.

7.0 REFERENCES

- [1] I. Labbé, et al, "ACCORD Broadband ATM Satellite Experiment (BASE) - DS3 Ku-band channel", CRC-RP-97-007, Ottawa, December 9th, 1997.
- [2] I. Labbé, et al, "ACCORD Broadband ATM Satellite Experiment (BASE) - DS3 Ka-band channel", CRC-TN-97-008, Ottawa, December 9th, 1997.
- [3] Postel, J. "Transmission Control Protocol - DARPA Internet Program Protocol Specification", *RFC 793*, DARPA, September 1981.
- [4] G.J. Armitage, K.M. Adams, "How Inefficient is IP over ATM Anyway", *IEEE Network*, January 1995.
- [5] D.E. Comer, J.C. Lin, "TCP Buffering and Performance Over an ATM Network", Purdue Technical Report CSD-TR 94-026, March 16, 1994.
- [6] A. Romanow, S. Floyd, "Dynamics of TCP Traffic over ATM Networks", Work supported by Contract No DE-AC03-76SF00098 for U.S. Department of Energy.
- [7] D.E. Comer, "*Internetworking with TCP/IP 2nd edition*", Prentice-Hall, 1991.
- [8] A. Mink, Y. Fouquet, S. Wakid, "Hardware Measurement Techniques for High-Speed Networks", National Institute of Standards and Technology.

- [9] Cavanaugh, J.D., "Protocol Overhead in IP/ATM Networks", University of Minnesota, 1994.
- [10] Stevens, W. R., "*Unix Network Programming*", Prentice-Hall, 1990.
- [11] Jacobson, V., "TCP Extensions for High Performance", *RFC 1323*, May 1992.
- [12] K. Moldeklev, P. Gunningberg, "How a large ATM MTU causes deadlocks in TCP data transfers", *IEEE/ACM Transactions on networking*, Vol 3. No 4, August 1995.




Figure 7: Application throughput with TCP over fibers




Figure 8: Application throughput with TCP over satellite




Figure 9: Application throughput with TCP over degraded satellite link




Figure 10: Application throughput with TCP-LFN over degraded satellite link




Figure 11: Application throughput with UDP over fibers and satellite links

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM
(highest classification of Title, Abstract, Keywords)

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in Section 8) Communications Research Centre 3701 Carling Av. P.O. box 11490 St. H Ottawa, Ont. K2H 8S2		2. SECURITY CLASSIFICATION (overall security classification of the document including special terms if applicable) UNCLASSIFIED	
3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title) ACCORD Broadband ATM satellite Experiment (BASE) -Protocols characterization (U)			
4. AUTHORS (Last name, first name, Middle initial) Isabelle, Labbe; Louis, Gravel; Gerard, Nourry; Corey, Pike; John, Butterworth; Gretchen, Bivens; Sean, Lanigan; Brian, Spink;			
5. DATE OF PUBLICATION (month and year of publication of document) December 1997		6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc.) 18	6b. NO. OF REFS. (total cited in document) 12
7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.) Final Technical Report (CRC-RP-97-008)			
8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include the address.) DEFENCE RESEARCH ESTABLISHMENT OTTAWA 3701 Carling Av. Ottawa, Ont. K1A 0Z4			
9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant) wk unit# 5cb11		9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written)	
10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document) CRC RPT 97-008		10b. OTHER DOCUMENT NOS. (Any other numbers which may be assigned to this document either by the originator or the sponsor)	
11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification) (x) Unlimited distribution () Distribution limited to defence departments and defence contractors; further distribution only as approved () Distribution limited to defence departments and Canadian defence contractors; further distribution only as approved () Distribution limited to government departments and agencies; further distribution only as approved () Distribution limited to defence departments; further distribution only as approved () Other (please specify)			
12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.) UNLIMITED			

13. **ABSTRACT** (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

A Broadband ATM satellite Experiment (BASE) was carried out between the Communications Research Centre, Canada and Rome Laboratory, USA during the period of May 95 to March 96. The objective of the experiment was to perform a series of tests to characterize broadband ATM satellite bearers.

This document reports on the second phase of the experiment that was conducted from December 95 to March 96 over the NASA ACTS satellite(Ka-band). Tests consisted in throughput measurements of TCP over a broadband ATM satellite link as a function of Bit Error Rate (BER). The throughput performance of TCP/IP was very poor due mostly to the (bandwidth * delay) product being limited by the TCP window size implementation. The results are presented along with the influence of others factors such as the MTU size and packet size. Finally, a performance comparison with a special implementation of TCP for network with large delays (TCP-LFN) is presented

14. **KEYWORDS, DESCRIPTORS or IDENTIFIERS** (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

ATM over satellite
TCP/IP, ATM , satellite
Transport protocols and ATM
Long fat networks