

Experimental Investigation of the Effects of Delay and Error Rate on Throughput and Performance of ATM Satellite Links Serving LAN Islands

Sufian Yousef and Caroline Strange

Anglia Polytechnic University, Victoria Road South, Chelmsford, Essex CM1 1LL, UK

Abstract. Satellite networks, offering broad geographical coverage and deployment, appear to be an attractive option, provided a number of difficulties deriving from the nature of satellite systems can be overcome. A group of experiments has been performed using Marconi Research Centre ATM Test-bed applying the NetPIPE benchmarking tool to investigate these difficulties such as the influence of the round-trip-delay and the error rate on the throughput, delay variations and performance of LAN islands characterised by their backward low bit rate (2.048 Mbps) over ATM satellite networks. Consequently, some characteristics are defined and some conclusions regarding the future performance of TCP/IP over satellite ATM networks, are drawn.

1 Introduction

ATM is rapidly emerging as the backbone of the future information networks. The ATM protocol is based on international standards and is designed to support multimedia information services. A combination of flexible bandwidth allocation, statistical multiplexing, priority queuing, and multicasting make ATM potentially more bandwidth-efficient than conventional switching technologies, despite the additional overhead of 9.4 percent for the 5-byte header per 53-byte ATM cell. Flexible bandwidth allocation is a key reason. The ability to allocate any and all of the available bandwidth on demand can outweigh the reduction in peak throughput. The peak throughput is sacrificed in favour of flexibility, with the goal of greater overall bandwidth efficiency and the ability to accommodate multimedia information services on demand. Additionally, the ability to multicast information avoids unnecessary transmission of duplicate information, such as copies of e-mail messages or video transmissions, which may be especially relevant over bandwidth-constrained satellite links.

In near future, ATM is expected to coexist with conventional networks based on IP. Logically, IP over ATM (IP/ATM) is one way to integrate ATM with legacy networks. The integration of ATM LANs into WANs via satellite links provides obvious advantages in terms of the ability to extend communications to geographically distant or remote locations, but it also poses challenges due to the longer delay, delay jitter, and potentially higher error rates. Efficient error control tailored to the ATM satellite protocol stack, is required. Interface and rate conversion

is required since such LAN-satellite links do not support rates as high as those of terrestrial fibre optic links.

Satellite systems have several inherent constraints. The resources of the satellite communication network, especially the satellite and the Earth station, are expensive; they must be used efficiently. The large delays in geo-stationary Earth orbit (GEO) systems and delay variations in low Earth orbit (LEO) systems affect both real-time and non-real-time applications. In an acknowledgement and time-out-based congestion control mechanism such as TCP, performance is inherently related to the delay-bandwidth product of the connection. Additionally, TCP round Trip Time (RTT) measurements are sensitive to delay variations that may cause false time-outs and retransmissions.

Therefore, the congestion control issues for broadband satellite networks are to some extent different from those of low-delay terrestrial networks. Both interoperability issues as well as performance issues need to be addressed before transport layer protocols like TCP can satisfactorily work over long-latency satellite ATM networks.

Broadband networking via LAN-SAT would place significant demands on already limited spectral resources at 14/12 GHz. More bandwidth is available at extremely high frequency (EHF) Ka-band. However, losses due to rain and atmospheric attenuation can be considerable at EHF, where link availability and recovery from outage may place limitations on the services that can be supported.

Broadcast or multicast of broadband video transmission to a user group is much more efficient than sending each user in the group a separate copy. A multicast transmission on the forward link can be coupled with a low-rate return link from individual users, which may be employed for error recovery of data transmission. IP multicasting over ATM networks is more practical at present. In the future, multicasting may be supported by IP version 6 over ATM.

TDMA was the chosen access scheme for satellite links. Stations transmit traffic bursts that are synchronised so that they occupy non-overlapping time slots which are organised within periodic frames. Each station can extend its traffic from the down-link.

The satellite has a link capacity of 25 Mbps (36 Mhz bandwidth) per transponder at present on Eutelsat II. The satellite link capacity has to be shared by a number of earth stations making a multiple islands interconnection. It is required for the satellite communication to provide the required QOS with an efficient utilisation of the satellite resources. The delay variation is done through the buffering in the ground segment due to the traffic load on the buffer. Most of the delay variation is at Terrestrial Interface Module for ATM (TIM-ATM) buffers. Cell loss occurs when buffer overflows. Ideally controlling the number of applications, traffic load and adequate bandwidth allocation for each application can control cell loss and the delay variation.

The frame duration is designed to be so small to produce reasonable constant delays due to frame length transmission. The RTT specified by the telephone services is at least 400 ms, taking into account the terrestrial (130 ms) and propagation delays (270 ms). The chosen frame size as 20 ms is meant to give acceptable performance. As the satellite header is normally long, the granularity of the capacity assignment to the required bit rate is increased through decreasing the frame length [1].

This scenario is similar to VSAT satellite system typical of corporate networking and rural communications. A large number of small dishes of 3.4 metres diameter stations carrying traffic of bursty nature, need to be served by the satellite. The multiple access is the major problem due to that we deal with various classes of traffic where compatibility of ATM format is needed to be provided (CBR, ABR, UBR, etc).

However, a concept for combining the benefits of ATM with those of SATCOM at Ka-band and error control concepts tailored to ATM SATCOM are proposed by [2]. The point which was not clear is what is sensitive to error rate in TCP/IP transmission and why. Also, the effect of delay on throughput and the effect of various data block sizes when combined to the effect of error rate on delay variations were missing.

This rest of the paper is organised as follows. In section 2, the experiment for the RTT influence on the performance of the low bit-rate islands at 2.048 Mbps using NetPIPE bench marking tool is presented. The effect of delay on throughput is included in section 3. In section 4, the effect of bit error rate on throughput and delay Variations is investigated. The paper is concluded in section 5.

2 Imposing Delay and Error Rate Experiment Using NetPIPE

The aim of this experiment was to impose different RTT values on the low bit-rate TCP and UDP/IP generated traffic on the simulated network to measure their influence on the throughput, block transfer time and its variance at different block sizes and window socket sizes. The results were represented in specific network performance graphs.

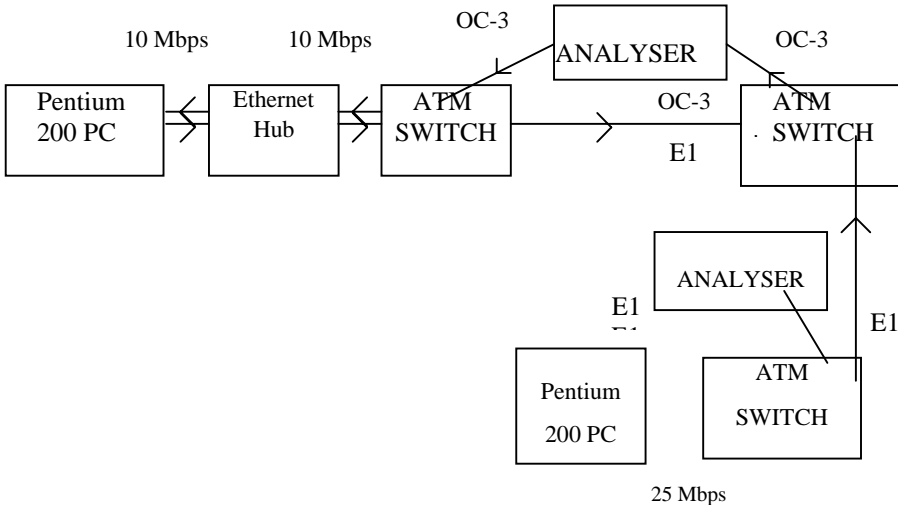


Fig. 1 Set-up for the Simulation of Ground and Satellite Segments

This experiment was performed using NetPIPE benchmark tool [3] for TCP/IP and Netperf [4] for UDP/IP traffic. The transmitting command was originated from a

Pentium 200 MHz PC connected to a LAN emulation Network Interface Card (NIC). The receiving command was originated from another Pentium 200 MHz PC connected to an ATM NIC. The E1 bit-rate was chosen to represent traffic originating from a distributed LAN island communicating with an ATM satellite network similar to Sat-Com network.

The testing set-up shown in Fig. 1 simulates a ground and satellite network segments. The generated TCP, UDP/IP traffic originated in a 25 Mbps bit-rate. Then, it is directed towards the ATM switch [5] in 2 Mbps bit-rate. The satellite link is simulated by generating signal towards the Analyser where a chosen delay is imposed on an OC3 (155 Mbps) bit-rate and then towards the receiving ATM switch (Earth Station). This traffic was then forwarded towards an Ethernet hub on an E1 (2.048 Mbps) bit-rate and then to a receiving PC.

The receiving side of traffic was forwarded from a PC towards the ATM switch via the Ethernet hub and then towards the ATM switch on OC3 bit-rate. From this ATM switch it headed towards the other ATM switch on an E1 (2.048 Mbps) bit-rate through the analyser which was used to impose the testing required delay on the E1 bit-rate. Traffic is transferred back from the ATM switch to the PC on a 25 Mbps bit-rate.

The testing procedure requires the use of NetPIPE and Netperf benchmarks as well as the analyser.

The experiment is repeated several times such that the total time for the experiment is far greater than timer resolution. NetPIPE uses a ping-pong transfer for each block size. This forces the network to transmit just the data block without streaming other data blocks in with the message. The result is the transfer time of a single block, thus providing the information necessary to know which block size is best, or what is the throughput given a block of size k . a sample of the result files is shown in Table. 1.

Table 1. A sample of TCP over ATM Satellite Network Performance at Default Socket Size.

Number of bytes	Time to transfer block in s	Throughput in Mbps	Variance
1	0.2045	0.000037	0.000001
128	0.2052	0.0.0048	0.000003
512	0.2054	0.019	0.000004
1024	0.211	0.038	0.000004
2048	0.214	0.073	0.000006
4096	0.253	0.123	0.00218
6144	0.320	0.147	0.00255
8192	0.653	0.096	0.004443
12288	0.719	0.131	0.00438
16384	0.951	0.132	0.016791
24576	1.257	0.149190	0.012120

Due to the high delay imposed (400 ms RTT) NetPIPE was unable to measure block sizes > 24 Kbytes. Measurements were taken on the end-to-end configuration. Traffic of 2.048 Mbps was chosen from Ethernet-to-Ethernet which represented LAN islands interconnected via satellite systems using ATM.

The imposed RTT delay is 400 ms on the satellite connection. Other delay values such as 100, 50, 20, 10, 1ms are imposed to simulate respectively Medium Earth Orbit (MEO), Low Earth Orbit (LEO), terrestrial wireless links, mobile links and optical fibre links.

Delay Variation is an important factor in satellite communication. It depends on the data block sizes and the retransmissions of the Ethernet LAN. This delay variation is found empirically to be equal (TDMA Frame Time / 2). But in this report we managed to measure the delay variation accurately through the measured variance of the block transfer time by using NetPIPE.

The Request and Response Time = Time to generate Request + Time to process Request + RTT.

The Buffer Requirement (BB) in the satellite link applying ATM cells is computed as follows: $BB = (\text{Trans speed} - \text{Link speed}) * (\text{Burst size} / (53 * \text{Trans speed}))$.

Throughput is restricted by the wait for the acknowledgements. The window size of the protocol can be used to adjust the amount of data to be sent before waiting for the acknowledgement. The following empirical formula estimates the throughput of burst traffic based on connection-oriented protocols with acknowledgements:

$\text{Throughput} = \text{Window size} / (\text{Transmission rate} * \text{RTT})$.

For example: Suppose that the transmission rate is 2.048 Mbps and RTT is 0.5 second. The window size = Transmission rate * RTT = 2.048 Mbps * 0.5 = 1.024 Mbps = 128 k bytes. This means that the throughput is 100% if we ignore the overheads. Throughput is directly proportional to the window size. If throughput is measured as 0.8 Mbps, the Acknowledgement window size is 400 k bits which is approximately 52 k bytes. In our measurements in Marconi Research Centre (MRC) the default socket size on NetPIPE was 8192 bytes. This is why the maximum measured throughput was 8.192/128. To convert the measured throughput to any other socket value we need to multiply it with X/128 where X is the socket size in k bytes. All measurements in this experiment have taken into account a default error rate of 1×10^{-12} .

3 The Effect of Delay on Throughput

In Fig. 1, the satellite network signature graph is shown. This graph shows the network latency when 1 byte is transferred and the effect of the buffer on the throughput when the transfer time increases. It is apparent that at small transfer times (0.33 s), both buffer lengths have the same throughput. Beyond the 0.33 s transfer time point in the graph, throughput increases for the larger buffer size. This shows how larger buffer sizes are important in improving throughput at larger message sizes and consequently larger transfer times.

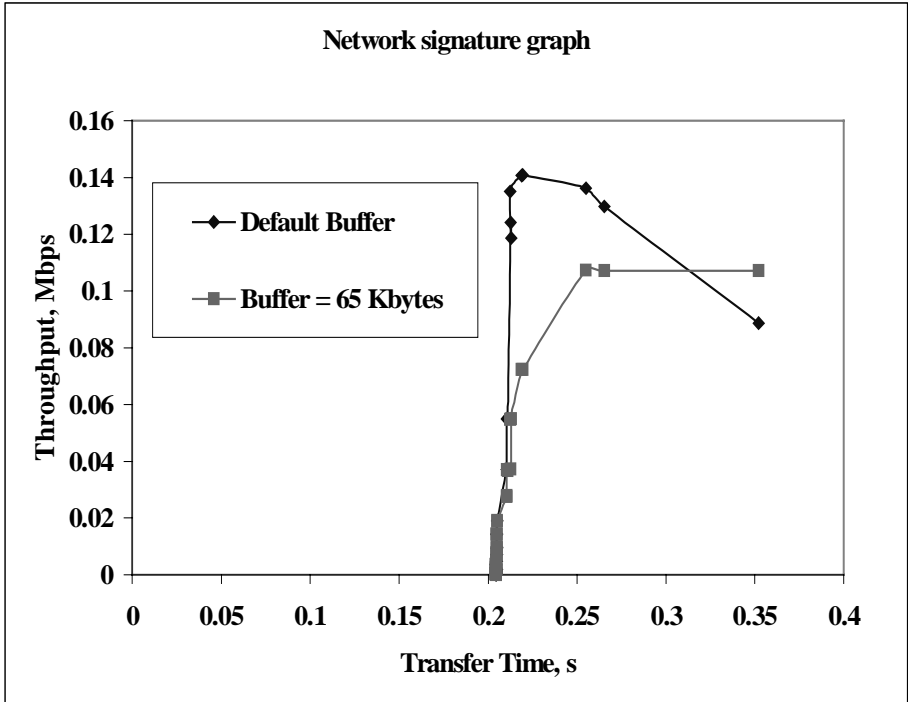


Fig. 1. Satellite Network Signature Graph

In Fig. 2 the effect of the round trip time (RTT) is shown. Normally increasing the block size increases the throughput. But the slope of throughput versus block size is different at different RTT values. The maximum throughput can be achieved when RTT is around 1 to 2 ms (short haul radio links). The 20 ms delay is typical of ATM wireless links where throughput is reasonably good. At Low Orbit military satellite links where delay is 40 to 50 ms, throughput declines abruptly but the maximum block size (65 k bytes) is still appears in the output if required to be processed which makes the Low Orbit satellites unique in this characteristic if compared with other satellite types. At medium orbit satellite links where RTT is around 200 ms throughput drops down and hence other factors such as the buffer size and the socket size has to be taken into account to compensate for this throughput decline. However, the maximum block size to be transferred is upper bounded at around 49 k bytes if such data block size has to be processed. At the traditional GEO high orbit satellite (34000 km) where the RTT value is ≥ 400 ms, the throughput is very small comparatively and the maximum data block size in bytes is limited to 16 k bytes if it is to be processed.

The effect of the buffer size on throughput is shown in fig. 3. Throughput increases with the increase in the socket buffer size at different slopes depending on the satellite link RTT. At low orbit satellites of $RTT \leq 100$ ms the throughput increases in a faster slope until the socket size is 32 k bytes. Beyond this socket size throughput becomes nearly constant until it reaches the maximum block size (64 k bytes). At the medium

orbit satellites of $RTT \leq 400$ ms throughput increases quite linearly on the socket size range and it joins the low orbit throughput at the maximum socket size (64 k bytes). At the RTT value of 600 ms, throughput increases linearly but in smaller slope at the whole socket size range. This figure helps the design engineer to allocate a specific socket size to target a specific throughput. This Figure shows that the allocation of the default socket buffer size is impractical when throughput is demanded. Implementing TCP/IP over ABR is shown by the NetPIPE to have a linear relationship at very long RTTs such as 600 ms which is typical of Geo-stationary satellites. This means that the window buffer size on very long RTTs depends on the throughput-delay product only which yields buffer size in bits out of this product: $Mbps \times Time \text{ in } s = M$ buffer size in bits = $M/8$ buffer size in bytes. The formula at $RTT = 600$ ms according to Fig. 3 is: $Buffer \text{ Size} = Throughput \times RTT$. In other words, the equation representing the straight line for this linear relationship is $y = 0.0125 x$ where y is the throughput in Mbps and x is the socket buffer size in k bytes. In LEO satellites this linear relationship is not valid where much smaller socket buffer size is required to achieve the same goal. Although this is a favourable characteristics in Leo satellite, it removes the influence of the socket buffer size in improving the delay variations.

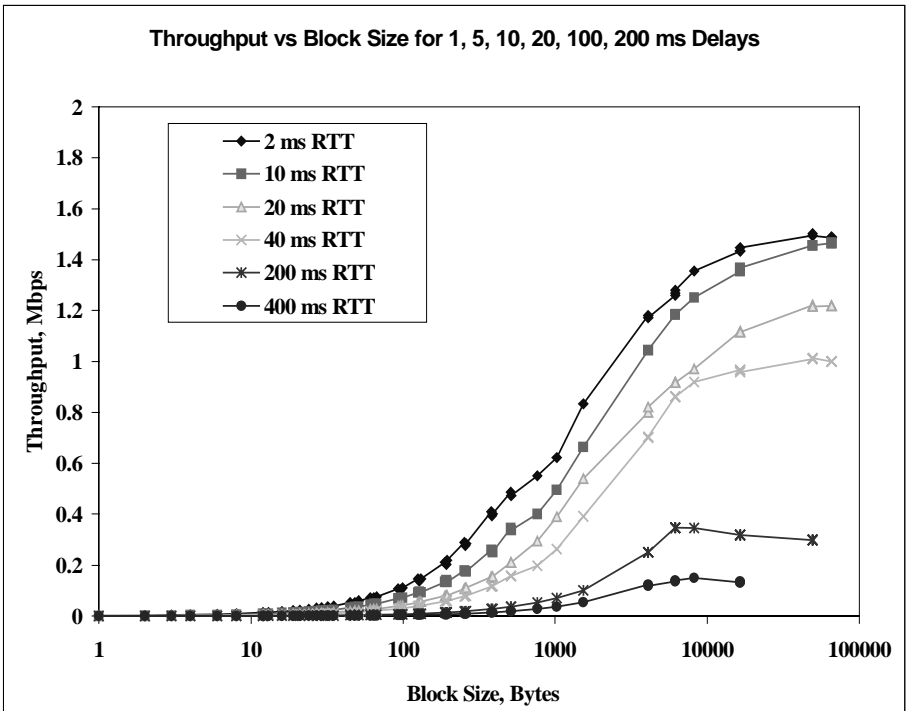


Fig. 2. Block Size in bytes versus Throughput in Mbps at Different RTT Values and at a Default Socket Size.

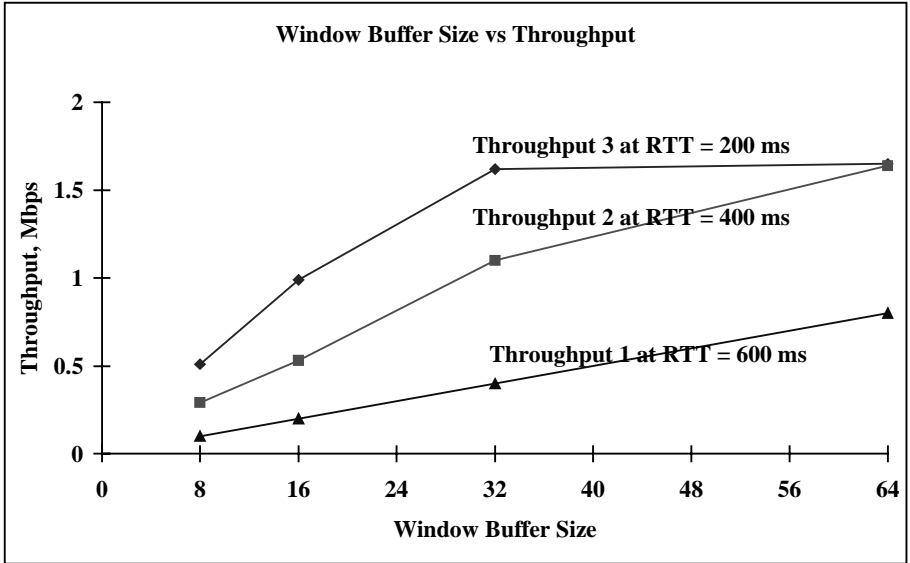


Fig. 3. Window Size in K bytes versus Throughput in Mbps at Different Delay (one direction) Values

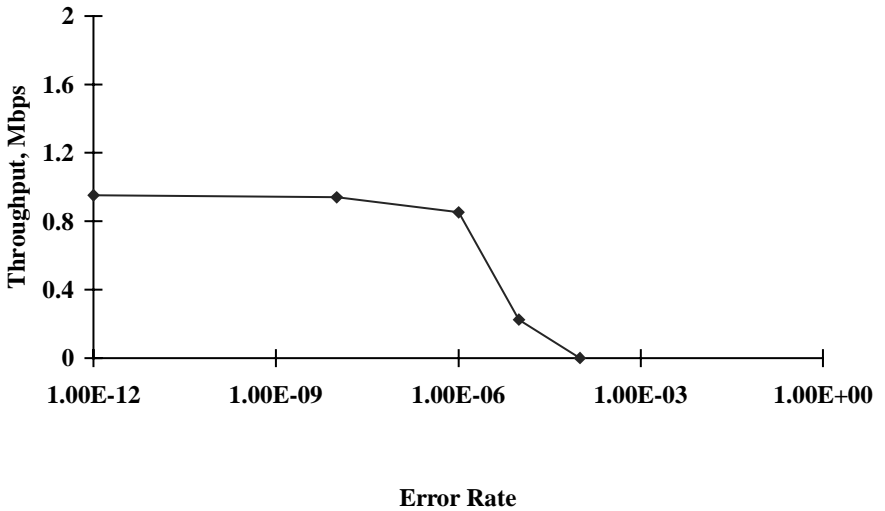


Fig. 4. Error Rate versus Throughput at RTT = 400 ms.

4 The Effect of Bit Error Rate on Throughput and Delay Variations

Normally throughput is affected by the delay due to the constant bandwidth-delay product in the absence of bit error rate. However, the error rate experimental influence on throughput of high-orbit (GEO) satellite links is measured and shown in Fig. 4. This Figure shows a gradual decline of throughput due to the error rate on the error rate ranges from 0 to 1×10^{-6} . This influence is extremely higher when the error rate increases from 1×10^{-6} to 1×10^{-4} . Actually, throughput disappears when the error rate is 1×10^{-4} . Therefore, it is important to protect satellite links from error rates $> 1 \times 10^{-6}$ by applying diversity or coding techniques.

In Fig. 5 the influence of error rate on the percentage CDV at default block size and RTT = 400 ms is shown. It shows the exact way the % CDV increases with the error rate. While the error rate range from 1×10^{-12} to 1×10^{-8} has a small effect on the % CDV, the range 1×10^{-8} to 1×10^{-5} has a high slope increase on the % CDV and a drastic slope increase on the error range of 1×10^{-5} to 1×10^{-4} . The lesson learned here is that for the sake of controlling the % CDV, the error rate has to be limited to 1×10^{-8} as a maximum possible error rate.




Fig. 5. Error Rate versus the Percentage Coefficient Delay Variation at RTT = 400 ms and a Default Block Size (8192 bytes).

5 Conclusions

The window buffer size has to be > 65 k bytes to get reasonably acceptable throughput at large delay satellite networks carrying TCP/IP over ATM.

LEO military satellites RTT (40-50) ms is the best to allow the transmission of block sizes as high as 65 k bytes. Geo-stationary satellites limit the maximum block sizes to nearly 16 k bytes only. MEO satellites limit the maximum block sizes to nearly 49 k bytes. Throughput is very low at GEO-stationary satellites and hence other means have to be adopted to improve the throughput such as using selective acknowledgements (SACK Ack.), higher window buffer sizes and lower rates than 2 Mbps in the backward path of the LAN satellite systems.

Implementing TCP/IP by the NetPIPE has shown a linear relationship at very long RTTs such as 600 ms which is typical of Geo-stationary satellites. This is supported by the network graphs where it is apparent that at small transfer times, different buffer lengths have the same throughput. Beyond the intersection transfer time point in the network signature graph, throughput increases for the large buffer sizes as compared to the default buffer size. This shows how larger buffer sizes are important in improving throughput at larger message sizes and consequently larger transfer times. This result simplifies the use of the buffer to improve throughput due to the achieved empirical formula. As a result, switches connected to satellite loops only need to have buffers proportional to the bandwidth-delay products only.

Delay variations depend on the data block sizes, the bit error rate and the RTT. Bit error rate can degrade the throughput and the effect is different at different block sizes and window buffer sizes.

The accepted threshold of the bit error rate on throughput, delay variations and block transfer time is 1×10^{-6} . The superior low percentage coefficient of delay variations (% CDV) can be achieved on error rates of $\leq 1 \times 10^{-8}$.

Diversity as a powerful error control technique, may be employed to improve satellite link availability and reduce power requirements at different rain margins.

In order to minimise the transmission delay effect and support real time applications efficiently, the adjustment of existing protocols or development of new ones is required. The value of a time-out parameter or window size of protocol which requiring acknowledgements of cell arrival might be increased to accommodate the long propagation delay.

References

- 1 M. H. Hadjitheodosiu et al., "End-to-end Performance Analysis for Broadband Island Interconnection via Satellite," RACE 11/R2074 (CATALYST) Project, European Conference of Satellite Communication (ESC-3) Project, European Conference of Satellite Communications (ECSC-3), Nov. 1993, pp. 170-174.
- 2 J. Farserotu and A. Tu, "TCP/IP over Low Rate ATM-SATCOM Links," IEEE MILCOM '96, Mclean, VA.
- 3 NetPIPE, <http://www.ameslab.gov/netpipe>.
- 4 Netperf, <http://www.cup.hp.com>.
- 5 FORE Systems, <http://www.fore.com>.