

Broadband Satellite Network: TCP/IP Performance Analysis

Sastri Kota, Mukul Goyal, Rohit Goyal and Raj Jain

Abstract

A number of satellite communication systems have been proposed using geosynchronous (GEO) satellites, as well as low earth orbit (LEO) constellations operating in the Ka-band and above. At these frequencies satellite networks are able to provide broadband services requiring wider bandwidth than the current services at C or Ku-band. As a consequence, some of the new services gaining momentum include mobile services, private intranets and high data rate internet access carried over integrated satellite-fiber ATM networks. Several performance issues need to be addressed before a transport layer protocol, like TCP can satisfactorily work over satellite ATM for large delay-bandwidth networks. In this paper, we discuss some of the architectural options and challenges for broadband satellite ATM networks. The performance results of TCP enhancements for Unspecified Bit Rate over ATM (ATM-UBR+) for large bandwidth-delay environments with various end system policies and drop policies for GEO satellite configurations for several buffer sizes are presented.

Keywords: Broadband Satellite, ATM, UBR, TCP, Performance Analysis

1. Introduction

The rapid globalization of the telecommunications industry and the exponential growth of the Internet is placing severe demands on global telecommunications. Satisfying this demand is one of the greatest challenges before telecommunications industry in 21st century. Satellite communication networks can be an integral part of the newly emerging national and global information infrastructures (NII and GII).

Satellite communication offers a number of advantages over traditional terrestrial point-to-point networks. Satellite networks can cover wide geographic areas and can interconnect remote terrestrial networks ("islands"). In case of damaged terrestrial networks, satellite links provide an alternative. Satellites have a natural broadcast capability and thus facilitate multicast communication. Finally, satellite links can provide bandwidth on demand by using Demand Assignment Multiple Access (DAMA) techniques.

The growing congestion of the C and Ku bands have increased the interest of satellite system

developers in the Ka-band. Several factors influence the development of broadband satellite networks at Ka-band frequencies:


- *Adaptive Power Control and Adaptive Coding*: Adaptive power control and adaptive coding technologies have been developed for improved performance, mitigating propagation error impacts on system performance at Ka-band.
- *High Data Rate*: A large bandwidth allocation to geosynchronous fixed satellite services (GSO FSS) and non-geosynchronous fixed satellite services (NGSO FSS) makes high data rate services feasible over Ka-band systems.
- *Advanced Technology*: Development of low noise transistors operating in the 20 GHz band and high power transistors operating in the 30 GHz band have influenced the development of low cost earth terminals. Space qualified higher efficiency traveling-wave tubes (TWTAs) and ASICs development have improved the processing power. Improved satellite bus designs with efficient solar arrays and higher efficiency electric propulsion methods resulted in cost effective launch vehicles.

2. BROADBAND SATELLITE NETWORK

There are several options that drive the broadband satellite network architecture [8]:

- (GSO) Geosynchronous Orbits versus (NGSO) Non-Geosynchronous Orbits (e.g., LEOs, MEOs)
- No onboard processing or switching
- Onboard processing with ground ATM switching or "ATM like," cell or fast packet switching
- Onboard processing and onboard ATM or "ATM like" fast cell/packet switching

However, most of the next generation broadband satellite systems have in common features like onboard processing, ATM or "ATM-like" fast packet switching, terminals, gateways, common protocol standards, and inter-satellite links [8]. Figure 1 illustrates a broadband satellite network architecture represented by a ground segment, a space segment, and a network control segment. *The ground segment* consists of terminals and gateways (GWs) which may be further connected to other legacy public and/or private networks. *The Network Control Station* (NCS) performs various management and resource allocation functions for the satellite media. Inter-satellite crosslinks in the *space segment* provide seamless global connectivity via the satellite constellation. The network allows the transmission of ATM cells over satellite, multiplexes and demultiplexes ATM cell streams for uplinks, downlinks, and interfaces to interconnect ATM networks as well as legacy LANs.



The gateway support several protocol standards such as ATM User Network Interface (ATM-UNI), Frame Relay UNI (FR-UNI), Narrow-band Integrated Digital Network (N-ISDN), and Transmission Control Protocol/Internet Protocol (TCP/IP). The gateway interface unit provides external network connectivity. The number and placement of these gateways in both GEO and MEO systems depend on the traffic demand, performance requirements, and other international regulatory issues. The *user Terminals Interface Unit (TIU)* supports several protocol standards adapting to the satellite network interface. It includes the physical layer functionalities such as channel coding, modulation/demodulation, and other RF functions. The *space segment* consists of either a GEO or MEO constellation depending on the system design. Within payloads full onboard processing and ATM or "ATM-like" switching is assumed.

Interconnectivity to the external private or public networks is possible with the support of the standard protocol. For the satellite ATM case, the signaling protocols based on ITU-T Q.2931 can be used when necessary. For other networks, the common channel signaling protocol, e.g., Signaling System No. 7 (SS7), can be used. The other interconnection interfaces between public and private ATM networks are the ATM Inter-Network Interface (AINI), the Public User Network Interface (PUNI) or the Private Network-Network Interface (PNNI), and the default interface between two public ATM networks, namely, the B-ISDN Inter Carrier Interface (B-ICI). However, these interfaces require further modifications to suit the satellite interface unit development. There is a definite need for an integrated satellite-ATM network infrastructure and standards for interfaces and protocols are in development process.

Effective traffic management and media access protocols constitute main challenges for successful deployment of Satellite ATM networks. Limited bandwidth available on satellite links make it necessary to use DAMA techniques in order to support multimedia applications [7]. Congestion control is an essential part of traffic management. ATM-ABR service uses Explicit Rate Congestion Control where feedback from the network contains the explicit rate at which sources should send data.

However, this scheme needs to be analysed in terms of the end-to-end delay requirements for satellite-ATM networks. In the long propagation delay satellite configurations, the feedback delay is the dominant factor in determining the maximum queue length. A feedback delay of 10 ms corresponds to about 3670 cells of queue for TCP over ERICA, while a feedback delay of 550 ms corresponds to 201,850 cells. Satellite switches can isolate downstream switches from such large queues by implementing Virtual Source/Virtual Destination (VS/VD) options [4].

3. TCP/IP TRAFFIC TRANSPORT OVER SATELLITE ATM

TCP/IP is the most popular network protocol suite and hence it is important to study how well these protocols perform on long delay satellite links. The main issue affecting the performance of TCP/IP over satellite links is very large feedback delay compared to terrestrial links. The inherent congestion control mechanism of TCP causes source data rate to reduce rapidly to very low levels with even a few packet loss in a window of data. The increase in data rate is controlled by ACKs received by the source. Large feedback delay implies a proportional delay in using the satellite link efficiently again. Consequently, a number of TCP enhancements (NewReno, SACK) have been proposed that avoid multiple reductions in source data rate when only a few packets are lost [2,9]. The enhancements in end-to-end TCP protocol are called *End System Policies*.

Satellite ATM link performance can also be improved by using *intelligent switch policies*. The Early Packet Discard policy [10] maintains a threshold R , in the switch buffer. When the buffer occupancy exceeds R , then all new incoming packets are dropped. Partially received packets are accepted if possible. The Selective Drop policy [3] uses per-VC accounting, i.e., keeps track of current buffer utilisation of each active UBR VC. A UBR VC is called "active" if it has at least one cell currently buffered in the switch. The total buffer occupancy, X , is allowed to grow until it reaches a threshold R , maintained as a fraction of the buffer capacity K . A fair allocation is calculated for each active VC, and if the VC's buffer occupancy X_i exceeds its fair allocation, its subsequent incoming packet is dropped. Mathematically, in the Selective Drop scheme, an active VC's entire packet is dropped if

$$(X > R) \text{ AND } (X_i > Z \cdot X/N_a)$$

where N_a is the number of active VCs and Z is another threshold parameter ($0 < Z \leq 1$) used to scale the effective drop threshold.

4. End-System Policy Vs Switch Policy for Satellite-ATM

[5] discusses the relative impact of end system policies (TCP flavors: Vanilla, Fast Retransmit Recovery/Reno, NewReno, SACK), switch drop policies (Early Packet Drop and Selective Drop) and switch buffer sizes (0.5 RTT, 1 RTT, 2 RTT) on the performance of MEO and GEO links satellite UBR+ links for Internet traffic. The same issues have been studied earlier for *persistent/infinite* TCP traffic in [3]. Both studies establish that for long delay satellite links, end system policies are far more effective than switch policies in ensuring good performance.

Simulation Configuration and Experiments




Figure 2 shows the configuration used in all simulations. The configuration consists of 100 WWW clients being served by 100 WWW servers, one server for each client. Both WWW clients and servers use underlying TCP connections for data transfer. The WWW traffic model used in this study is an extension of that specified in SPECweb96 benchmark. [11] and is based on HTTP/1.1 standard [1]. The switches implement the UBR+ service with optional drop policies described before.

Configuration Parameters

Links connecting server/client TCPs to switches have a bandwidth of 155.52 Mbps (149.76 Mbps after SONET overhead), and a one way delay of 5 microseconds. The link connecting the two switches simulates MEO and GEO link respectively and has a bandwidth of 45Mbps (T3). The corresponding one-way link delays are 100 ms and 275 ms respectively. Since the propagation delay on the links connecting client/server TCPs to switches is negligible compared to the delay on the inter-switch link, the round trip times (RTTs) due to propagation delay are 200 ms and 550 ms for MEO and GEO respectively. All simulations run for 100 secs. TCP maximum segment size (MSS) is set to 9180. TCP timer granularity is set to 100 ms. Using window scaling option, TCP maximum receiver window size is set to 2,097,120 and 4,194,240 bytes for MEO and GEO links respectively. The value of maximum receiver window is set so that it is greater than RTT-bandwidth product of the path. The TCP delay ACK timer is NOT set. Segments are ACKed as soon as they are received. The drop threshold R is 0.8 for both switch drop policies - EPD and SD. For SD simulations, threshold Z also has a value 0.8. We use three different values of buffer sizes corresponding to 0.5 RTT, 1 RTT and 2 RTT - bandwidth products of the end-to-end TCP connections for each of the propagation delays. The performance is measured in terms of the efficiency of link usage, i.e., the ratio of total throughput of all connections and the maximum possible throughput on the link.

Simulation Analysis Technique

We analyze the effects of 3 factors - TCP flavor, buffer size and drop policy - in determining the efficiency and for MEO and GEO links. The values a factor can take are called 'levels' of the factor. For example, EPD and SD are two levels of the factor 'Drop Policy'. The analysis consists of the calculating the following terms. A detailed description of analysis procedure is available in [6,5].

- **Overall mean:** This consists of the calculation of the overall mean 'Y' of the result (efficiency or fairness).
- **Total variation:** This represents the variation in the result values (efficiency or fairness) around the overall mean 'Y'. *The goal of the analysis to calculate, how much of this variation can be explained by each factor and the interactions between factors.*
- **Main effects:** These are the individual contributions of a level of a factor to the overall result. A particular main effect is associated with a level of a factor, and indicates how much variation around the overall mean is caused by the level. We calculate the main effects of 4 TCP flavors, 3

buffer sizes, and 2 drop policies.

- **First order interactions:** These are the interaction between levels of two factors. In our experiments, there are first order interactions between each TCP flavor and buffer size, between each drop policy and TCP flavor, and between each buffer size and drop policy.
- **Allocation of variation:** This is used to explain how much each factor contributes to the total variation.

Simulation Results

Following observations can be made about MEO and GEO links from Tables 1-2. TCP flavor explains 56.75% for MEO and 69.16% for GEO of the efficiency variation and hence is the major factor in deciding efficiency value. SACK results in substantially better efficiency than other TCP flavors. Thus, for long delay satellite links, SACK is the best choice in spite of complexity of its implementation. Buffer size explains 21.73% for MEO and 13.65% for GEO of the variation and interaction between buffer size and TCP flavors explains 13.42 for MEO and 7.54% for GEO of the variation. Efficiency values are largely unaffected as we increase buffer size from 0.5 RTT to 1 RTT. There is a marginal improvement in performance as buffer size is increased further to 2 RTT. Vanilla and Reno show substantial efficiency gains as buffer size is increased from 1 RTT to 2 RTT. Note that a buffer size of 0.5 RTT is sufficient for SACK. Further increase in buffer size brings very little performance improvement for SACK. Drop policy (EPD or Selective Drop) does not have an impact on efficiency as indicated by negligible allocation of variation to drop policy. From the observations above, it can be concluded that SACK with 0.5 RTT buffer is the optimal choice for MEO and GEO links with either of EPD and SD as switch drop policy.

Table 1 Simulation Results for MEO and GEO Links

Drop Policy	TCP Flavor	Buffer=0.5RTT		Buffer=1RTT		Buffer=2RTT	
		Efficiency		Efficiency		Efficiency	
		MEO	GEO	MEO	GEO	MEO	GEO
EPD	Vanilla	0.848	0.791	0.879	0.792	0.899	0.848
	Reno	0.894	0.805	0.903	0.817	0.909	0.874
	NewReno	0.903	0.866	0.910	0.859	0.912	0.845

	SACK	0.908	0.902	0.912	0.909	0.916	0.921
SD	Vanilla	0.836	0.808	0.872	0.816	0.901	0.868
	Reno	0.876	0.810	0.898	0.781	0.902	0.863
	NewReno	0.892	0.790	0.892	0.832	0.898	0.851
	SACK	0.917	0.918	0.926	0.916	0.937	0.921

Table 2 Allocation of Variation for MEO and GEO Efficiency Values

Component	Sum of Squares		%age of Variation	
	Efficiency		Efficiency	
	MEO	GEO	MEO	GEO
Individual Values	19.3453	17.3948		
Overall Mean	19.3334	17.3451		
Total Variation	0.0119	0.0497	100	100
Main Effects:				

TCP Flavor	0.0067	0.0344	56.75	69.16
Buffer Size	0.0026	0.0068	21.73	13.65
Drop Policy	0.0001	0.0001	0.80	0.25
First-order Interactions:				
TCP Flavor-Buffer Size	0.0016	0.0037	13.42	7.54
TCP Flavor-Drop Policy	0.0007	0.0025	6.11	4.96
Buffer Size-Drop Policy	0.0001	0.0002	0.53	0.41

5. CONCLUSIONS

Broadband satellites networks are the new generation communication satellite systems that will use onboard processing and ATM and/or "ATM-like" switching to provide two-way communications. The proposed satellite or broadband satellite systems operate at Ka-band and above frequencies. Several technical challenges and issues, e.g., traffic management, Quality of Service (QoS) assurance, interoperability, efficient protocols, and standards. In this paper, we analysed design parameters based on end policies and switch parameters for efficient satellite ATM networks. In summary, as delay increases, the gains of end system policies are more important than the gains of drop policies and large buffers.

6. References

1. R. Fielding, J. Gettys, J. Mogul, H. Frystyk, T. Berners-Lee, "Hypertext Transfer Protocol -HTTP/1.1", RFC 2068, January 1997.
2. S. Floyd, T. Henderson, "The NewReno Modification to TCP's Fast Recovery Algorithm," Internet Draft, November 1998, Available from <ftp://ftp.ietf.org/internet-drafts/drafts-ietf-tcpimpl-newreno-00.txt>
3. Rohit Goyal, Raj Jain, Shivkumar Kalyanaraman, Sonia Fahmy, Bobby Vandalore, "Improving the Performance of TCP over the ATM-UBR Service," *Computer Communications*, Vol 21/10, 1998.

4. R. Goyal, R. Jain, M. Goyal, S. Fahmy, B. Vandalore, and S. Kota, "Traffic Management for TCP/IP over Satellite ATM Networks," *IEEE Communications Magazine*, March 1999, Vol. 37, No. 3, pp. 56-61
5. Mukul Goyal, Rohit Goyal, Raj Jain, B. Vandalore, S. Fahmy, T. vonDeak, K. Bhasin, N. Butts, S. Kota, "Performance Analysis of TCP Enhancements for WWW Traffic using UBR+ with Limited Bufers over Satellite Links", ATM_Forum/98-0876R1, December 1998, <http://www1.cse.wustl.edu/~jain/atmf/a98-0876.htm>
6. R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Simulation, and Modeling*, John Wiley & Sons Inc., 1991.
7. S. Kota, J. Kallaus, H. Huey, and D. Lucantoni, "Demand Assignment Multiple Access (DAMA) for Multimedia Services – Performance Results," *Proc. MILCOM'97*, Monterey, CA, 1997.
8. Sastri Kota, "Satellite ATM Networks: Architectural Issues and Challenges," *Proc. Conf. on Satellite Networks: Architectures, Applications and Technologies*, NASA Lewis Research Center, Cleveland, pp. 443-457, June 2-6, 1998.
9. M. Mathis, J. Madhavi, S. Floyd, A. Romanow, "TCP Selective Acknowledgment Options," RFC 2018, October 1996.
10. A. Romanow, S. Floyd, "Dynamics of TCP Traffic over ATM Networks", *IEEE JSAC*, May 1995.
11. SPEC, "An Explanation of the SPECweb96 Benchmark," Available from <http://www.specbench.org/osg/web96/webpaper.html>