

# Performance Analysis of TCP-AFC for Satellite-based Networks

N. J. Kothari  
Dharmsinh Desai University,  
Nadiad, Gujarat, India

Bhavika Gambhava  
Dharmsinh Desai University,  
Nadiad, Gujarat, India

Dr. K. S. Dasgupta  
Space Application Centre,  
Ahmedabad, India

## ABSTRACT

Satellite has been identified as a potential candidate to meet the explosive Internet demand and to evolve the global Internet services. With a view to combat channel errors predominant in satellite based networks, TCP with Adaptive Flow Control and *Delayed Fast Recovery* (TCP-AFC) has been designed to identify a random loss with the help of selective acknowledgments. TCP-AFC has demonstrated significant performance enhancement in error prone environments through simulations and experiments on an active emulated network. In order to substantiate the improvement, we investigate the performance of TCP-AFC in a real environment consisting of a Ku band satellite link, which is more susceptible to atmospheric conditions. This paper focuses on evaluation of TCP-AFC in real life situations having appreciable channel noise and delay. Results of the extensive experiments conducted on a test bed consisting of a symmetric GEO satellite link for different channel conditions, data volume, data type and data traffic are presented in this paper. Analysis of the results reconfirms the compatibility of TCP-AFC in a heterogeneous network besides the performance improvement over a dedicated satellite link.

## General Terms

Performance, Experimentation, Verification.

## Keywords

TCP-AFC, SACK\_OK, cwnd, Cd, BER

## 1. INTRODUCTION

Broadband satellite networks are becoming an important niche of Internet particularly in the areas, which are underserved by terrestrial infrastructure. Technological advances have enabled new systems that combine broadband data rates with small terminals, thereby providing more affordable 'last-mile' network access to home and small business users worldwide [1]. However, the success of satellite networks hinges on the ability of the underlying protocols to function efficiently in the environment that is characterized by longer propagation delays (for GEO satellites) and poorer BER. As reported previously, these factors degrade the performance of TCP [2] particularly [3].

In order to improve the performance against above factors, several mitigation techniques have been suggested over standard TCP versions like New Reno [4] and SACK TCP [5]. The schemes proposed to improve the flow using the bandwidth estimation algorithms [6] [7] [8] may exhibit aggressive behaviour resulting

from over estimation of the available bandwidth [8] [9]. It is often not possible to draw sound conclusions from network delay measurements [10]. Additionally, the performance of current TCP protocols suffers in satellite networks due to long propagation delays. A recent congestion control scheme, TCP-Peach with two new algorithms, *Sudden Start* and *Rapid Recovery* is introduced for satellite networks [11]. It requires dummy segments to probe the availability of resources and routers to implement a priority mechanism at the IP layer.

On the other hand, the algorithms exploiting implicit indications are independent of any critical parameter values or the router reconfiguration. Therefore, the existing network can use network accessories without any modification in the subsystems. Adaptive Flow Control with *delayed fast recovery* in TCP (TCP-AFC), which exploits implicit indications in SACK information has exhibited significant performance improvement in wireless environments through simulations [12] [13]. Moreover, it has been successfully tested for its compatibility with the existing terrestrial Internet [14]. This paper focuses on validation of the performance of TCP-AFC in a satellite-based network with the help of results of extensive experiments conducted over a set up consisting of a real Ku band geostationary satellite link.

Section 2 briefly describes the algorithm of Adaptive Flow Control with *Delayed fast recovery* proposed in TCP-AFC. The test bed consisting of a real satellite link set up for the performance evaluation is explained in section 3. We analyze the results of experiments conducted over the satellite subnetwork for different network parameters in section 4.

## 2. ADAPTIVE FLOW CONTROL WITH DELAYED FAST RECOVERY

A simple and effective approach towards identifying a packet loss resulting from random channel errors is proposed previously [12] by using SACK information [15] to keep track of network condition related to congestion. The approach decouples loss recovery and rate adjustment conditionally when a packet loss is detected at TCP sender. Consequently, a corrupted or lost data packet does not trigger congestion control necessarily, which otherwise reduces the sending rate upon packet corruption [16]. In view of the fact that, SACK information has remained underutilized [17], the proposed TCP-AFC uses an algorithm that introduces a new flag SACK\_OK on sender side implementation to determine TCP's response to a packet loss. The flag indicates continuous increment in sequence number of the packets selectively acknowledged by receiver. The algorithm sets SACK\_OK to true on detection of the first loss and resets to false

on a subsequent loss within the same window. The *delayed fast recovery* algorithm [12] assumes the first loss in a window as a random loss and delays *fast recovery* while the flag is true. The *fast recovery* is invoked later following a high to low transition of SACK\_OK. Since all except the first packet are selectively acknowledged, there is no sign of significant congestion in the network demanding flow control.

When factors like convergent nature of upstream traffic and limited wireless bandwidth cause congestion, the conditional nature of the algorithm reacts as it immediately triggers *fast recovery* after detecting another loss. Thus, the performance of TCP-AFC is enhanced over wireless links prone to channel errors and protected on conventional wired networks. In order to utilize the available network resources optimally, the algorithm is extended with a technique for estimation of *cwnd* without sacrificing its suitability in mixed environments [13]. The extension calculates the gap in sequence numbers of the first two packet losses. This gap is an indication of the degree of congestion and is referred as Congestion distance, *Cd*. Therefore, when *Cd* is larger than *cwnd*/2, a fair rise in *cwnd* beyond the conventional value is suggested such that the network is not over burdened. The rate adjustment is based on the following expression.

$$cwnd_{n+1} = \frac{cwnd_n}{2} + R_{adj} \quad \dots \dots \dots (1)$$

Where  $R_{adj}$  = Rate Adjustment

$$= \frac{Cd^4}{2 \times cwnd_n^3} \quad \text{if } Cd > \frac{cwnd_n}{2}$$

$$= 0 \quad \text{if } Cd \leq \frac{cwnd_n}{2}$$

The component  $R_{adj}$  of  $cwnd_{n+1}$  produced by (1) can be anywhere between 0 to  $cwnd_n/2$ . If *Cd* is less than or equal to *cwnd*/2, the conventional approach is followed to ensure that the performance does not deteriorate on wired networks where congestion induced losses are common. Conventional loss recovery mechanism is adopted for the packet indicated to be lost after first three *dupacks* or a partial acknowledgment [4].


### 3. SATELLITE BASED TEST SETUP

Results of simulations and subsequently, experiments on an emulated network have demonstrated that TCP-AFC yields substantial performance improvement over a point-to-point satellite connection without hampering normal TCP performance on a conventional network [13][14]. A variety of experiments using multiple test methods generally provides a more compelling argument in favor of a change to TCP [18]. Since the results obtained from performance investigations involving actual system elements are always trustworthy, the observations from the tests conducted over a real satellite system will be a strong evidence in favour or against of a proposal to modify the existing scheme considering that a Ku band link is more susceptible to atmospheric conditions [19].

Figure 1 shows the satellite sub system set up at SAC<sup>∇</sup>. Important specifications of the setup are listed in Table 1. Though recent advances in error control codes have increased immunity of satellite channels, even mild error rates on high-speed links can have a crippling effect on throughput [20]. A variable attenuator on the transmitter side changed the error rate over the satellite link. The transmitter was constrained by maximum data rate of 2 Mbps. The transceiver connecting the source and destination through a satellite link in Figure 1 contains a modem and up/down converters. During the experiments, we changed various network conditions like channel noise, data rate, and network traffic to examine the performance of TCP-AFC. Performance of TCP-AFC is evaluated in two phases: (a) Point-to-point network (b) Shared network with concurrent (i) TCP and (ii) multimedia traffic.

**Table 1**  
**Specifications of the Test Bed of a Satellite Network**

Satellite	Geostationary / GSAT-3
Uplink Frequency (F/R)	14062 MHz / 14066 MHz
Downlink Frequency (F/R)	11012 MHz / 11016 MHz
Modulation type	QPSK
Forward Error Correction	½ Rate
BER block size	$1 \times 10^{-6}$ , $1 \times 10^{-5}$
Data Rates (Synchronous)	256, 512, 1024, 2048 (in kbps)
System Software	Linux Kernel 2.4.19
Audio/Video Transcoding	64 / 512, 64 / 128 (in kbps)




**Figure 1. Test bed for Satellite-based Subnetwork (Congestion Free Network)**

#### 3.1 Point-to-Point Network

Data was transferred over a satellite subnetwork as shown in Figure 1. In absence of traffic sharing the satellite link, there was no congestion in the network. Therefore, the measurements from these experiments indicate the performance for the losses due to channel errors only. Different file sizes – from small to large (5, 10, 20 MB) were used to observe initial and sustained response to data traffic for different parameter settings.

<sup>∇</sup> Space Application Center (SAC), Indian Space Research Organization(ISRO), Ahmedabad.



**Figure 2. Shared Network  
(With & Without Multimedia UDP Traffic)**

### 3.2 Shared Network

TCP performance is sensitive to the presence of other traffic in its path [21]. Therefore, in the second phase, the additional source shown in Figure 2 caused buffer overflow at the modem. In this phase, we focused on one particular aspect of TCP performance: the interaction between TCP-AFC algorithm and congestion-induced losses along an end-to-end path that contains both a satellite channel and a wired network segment. Thus, the performance evaluation is carried out in an environment that suffers from mixed type of losses in presence of different types of concurrent traffics.

## 4. SATELLITE BASED TEST SETUP

In this section, we analyze results of the experiments conducted for different parameters and configurations of a satellite based network described in section 3.

### 4.1 Congestion Free Satellite Subnetwork


The objective of these experiments is to measure the performance of TCP-AFC over a point-to-point connection in congestion free environment of Figure 1. Multiple attempts of the same file transfer were made for each combination of file size (5-20MB), data rate (1,2Mbps) and BER ( $1 \times 10^{-6}$ ,  $1 \times 10^{-5}$ ) to average out the impact of link dynamics. Performance of TCP and TCP-AFC was measured alternately in each attempt of file transfer as shown in Figure 3. TCP-AFC performs better than TCP as observed in Figure 3 for the error rate of  $1 \times 10^{-6}$ . Moreover, the throughput measurements are higher and have less variation in TCP-AFC. This indicates a more stable performance on an error prone channel. As a result, Minimum, Maximum and Average throughputs are higher as shown in Figure 4.

We repeated the test with the BER of  $1 \times 10^{-5}$  over the satellite link. Increase in channel errors caused more packet losses and required frequent loss recovery phases over a TCP connection. Frequent reduction in flow coupled with loss recovery in TCP degraded its throughput severely as seen in Figure 5. The adaptive flow control of TCP-AFC protects its performance from degradation. The difference in performance of TCP and TCP-AFC significantly increased in this situation as compared to the BER  $1 \times 10^{-6}$ .


It is worth to clarify that the points of TCP-AFC and TCP corresponding to a particular iteration in Figure 3 do not mean that the tests were conducted in identical situation due to the dynamic nature of a real environment. In order to provide a fair comparison of performance of TCP and TCP-AFC, we examined

packet dumps obtained during the tests. We identified two samples of the dumps, in which TCP and TCP-AFC suffered nearly same number of packet losses during transfer of the same file. Total number of losses in TCP and TCP-AFC connections were 90 and 93 respectively for 2 Mbps data rate and BER of  $1 \times 10^{-5}$ . A comparison based on these two samples in turn justifies the performance enhancement of TCP-AFC with case of channel errors as illustrated in Figure 6 and Figure 7.


As explained above, the experiment was conducted over a congestion free dedicated satellite subnetwork (Figure 1). Hence, the losses spotted in the dumps (represented by symbol R in Figure 6) are caused by the channel errors. The persistent rise in sequence numbers of TCP-AFC suggests that it avoided rigid flow control, which the conventional TCP unconditionally invoked in case of a packet loss occurring in absence of congestion. The difference in the sequence numbers of TCP and TCP-AFC also reflects in the throughput vs. time plots of the samples in Figure 7.



**Figure 3. Throughput Measurement in a Series of File Transfers (2 Mbps Data Rate,  $1 \times 10^{-6}$  BER, 20 MB File)**



**Figure 4. Performance of 20 MB File Transfers  
(2 Mbps Data Rate &  $1 \times 10^{-6}$  BER)**



**Figure 5. Performance of 20 MB File Transfers  
(2 Mbps Data Rate &  $1 \times 10^{-5}$  BER)**


TCP-AFC outperforms TCP throughout the period of file transfer. The ability of utilizing the available bandwidth enables TCP-AFC to complete the file transfer with greater efficiency (in 245 seconds) whereas TCP requires 554 seconds to transfer the same amount of data. As a result, the aggregate throughput of TCP-AFC was significantly higher (82KBps) than TCP (36KBps).

To support our findings and inferences, we calculated throughput at different stages of the same file transfer for different volumes of data as shown in Table 2. Throughput of TCP and TCP-AFC are found to be close to each other in the initial phase of the file transfer. However, as packet losses start occurring frequently, the difference in the performance also becomes noticeable. TCP-AFC is relatively immune to losses than TCP as viewed in Table 2. Since every loss recovery attempt during *slow start* reduces *ssthresh* by half, *congestion avoidance* begins sooner [2]. This earlier initiation of *congestion avoidance* restricts the growth of *cwnd* and prior termination of *slow start* degrades the performance of TCP further. However, TCP-AFC attempts to identify such losses and decouples loss recovery from congestion control.


We validate this further with the help of Figure 8. As discussed in section 2, TCP-AFC behavior differs only after detection of a packet loss. In Figure 8, the performance of TCP and TCP-AFC

**Table 2**  
**Performance Comparison at Different Stages**  
**of the File Transfer**

Volume of Data Transfer	Throughput (Bytes/sec)		Total Losses	
	TCP	TCP-AFC	TCP	TCP-AFC
0.5 MB	74341.79	78112.54	2	4
1 MB	94513.89	109082.8	6	5
2 MB	56827.51	138250	11	6
5 MB	49865.72	147836.2	18	19
10 MB	37525.68	99671.12	48	40
15 MB	36900.66	86212.58	66	68
20 MB	36081.95	81554.97	90	93



**Figure 6. Sequence No. Vs Time (20 MB File Transfer)**



**Figure 7. Throughput Vs Time (20 MB File Transfer)**



**Figure 8 Throughput vs. Time Plots**  
**(First 1 M bytes in 20 MB File Transfer)**

remains same until the occurrence of the first loss in TCP-AFC at 4 sec. However, this loss does not severely affect the performance of TCP-AFC, whereas, TCP throughput is limited after the first loss occurring around 4.5 sec. and the subsequent losses than after. Observations from other file transfers (5,10MB) are not presented here due to space constraints. However, 20MB file transfer covers all phases in a sustained data transfer session.

## 4.2 Congested Network

Heterogeneity is a key property that makes it difficult to model and simulate Internet. It ranges from the individual links that carry the network's traffic, to the protocols that interoperate over the links, to the "mix" of different applications used at a site, to the levels of congestion seen on different links [22]. The success of the network relies on successful deployment and operation of protocol mechanisms that mitigate the problems posed by this heterogeneous environment. Therefore, we consider specifically the transport protocol performance that a satellite-based user is likely to encounter when the connection traverses the wired Internet, in contrast to looking at transport connections in dedicated satellite environments as discussed in subsection 4.1.



Background traffic was added in the topology of Figure 2 first with another file transfer and then a multimedia transfer. Since the background traffic was delivered in the same direction as the foreground transfer, a bottleneck was created at the ingress of the satellite network. Performance of TCP-AFC is evaluated in presence of competing traffic that causes packet drops due to buffer overflow at the modem. On the satellite link, packets are corrupted by channel errors additionally. Thus, we report the tests conducted over a satellite network, which suffered from losses due to congestion and corruption, in this section.

#### 4.2.1 Concurrent TCP Traffic

We conducted the experiments over this mixed environment in two parts. In the first part, we generated concurrent traffic of another 20 MB file transfer from the transmitter 2 of Figure 2. The objective of this experiment was to examine the performance of TCP-AFC when it competes with concurrent TCP traffics. Transmitter 1 transferred the file with TCP-AFC and TCP alternately. However, transmitter 2 used only TCP implementation in all tests to deliver the background traffic. This allowed us to evaluate TCP-AFC performance in comparison with TCP in the same network conditions.

The transmitter power was raised to create a clear sky condition for the satellite link. In this situation, the network suffered from packet losses due to buffer overflow at the modem i.e. congestion. TCP-AFC is expected to follow conventional TCP while facing severe congestion. Figure 9 illustrates the performance of TCP and TCP-AFC for different data rates and reconfirms that TCP-AFC retains conventional TCP behavior in absence of channel error even when the network is congested.

The experiment was repeated with channel errors to investigate the performance of TCP-AFC in an environment that suffered from both types of losses: losses due to congestion at the modem and corruption on the satellite link. Figure 10 summarizes the performance of TCP and TCP-AFC for different BER values and data rate of the satellite link with reference to the average of throughputs measured. It is evident from Figure 10 that performance of TCP-AFC improves significantly over TCP for the BER of  $1 \times 10^{-5}$  in comparison with the BER of  $1 \times 10^{-6}$  with either data rate, 2 Mbps and 1 Mbps for the reason discussed earlier. It can also be observed that TCP-AFC almost achieves the performance of a wired connection when the measured BER is moderate i.e.  $1 \times 10^{-6}$ .




Figure 9. Performance in Absence of Errors (with TCP Traffic)




Figure 10. Performance Comparison (for Different Error Rate & Data Rate)

#### 4.2.2 Concurrent TCP and Multimedia Traffic

It has been pointed out earlier that TCP performance degrades in presence of UDP, which is unresponsive to congestion [21]. The objective of the experiments discussed in this part is to show the improved performance of TCP-AFC in presence of UDP based multimedia traffic from transmitter 2 in the topology of Figure 2. The multimedia traffic of an MP4 transfer was generated at different rates of transcoding; 512 kbps and 128kbps for video whereas 64kbps was selected for audio. The multimedia nature of the traffic also appeared bursty some times.

Like the previous phase, the data rate of satellite link was varied from 2Mbps to 1Mbps to produce different congestion conditions in the test environment. However, the change in data rate did not affect the bandwidth occupied by UDP carrying multimedia traffic. The throughput of UDP transfer was observed to be 83KBps and 29KBps corresponding to 512kbps and 128kbps video rates respectively. These different video rates would also vary the degree of congestion in the network along with the data rates of the satellite link. Therefore, the FTP connection depending on TCP or TCP-AFC was required to efficiently utilize the bandwidth left unused by UDP traffic.

The efficiency of TCP and TCP-AFC in this kind of real environment is quantified in Figure 11 and Figure 12 as a percentage of the maximum performance. The maximum throughput is obtained in the tests conducted with clear sky condition of the satellite link. It is evident from the Figure 11 and Figure 12 that both schemes suffer from under utilization of the available bandwidth in presence of poor BER i.e.  $1 \times 10^{-5}$ . However, bandwidth underutilization in TCP is noticeable in most of the cases. This weakness is more apparent in Figure 12, when





Figure 11. Bandwidth Utilization (1Mbps Data Rate, 512 Kbps Video Rate)

the channel data rate is increased to 2Mbps. Thus, these experiments reveal that TCP-AFC is more robust while competing with unresponsive UDP flows in error-prone environments.



**Figure 12. Bandwidth Utilization**  
(2Mbps Data Rate, 128 Kbps Video Rate)

## 5. CONCLUSIONS

In order to enhance the performance of TCP in wireless environments in general and satellite networks in particular, TCP with Adaptive Flow Control and Delayed *fast recovery* (TCP-AFC) has been proposed and evaluated through simulations earlier. In this paper, we validate our previous work of TCP-AFC over a satellite link. Performance of TCP-AFC is analyzed in various network conditions of a set up created involving EDUSAT link at Space Application Centre, ISRO, Ahmedabad.

The analysis of results of extensive experiments carried out confirm that performance of TCP-AFC improves significantly (about 138%) over satellite networks without degradation in heterogeneous environment where losses due to congestion are common. In particular, we find that a well-tuned TCP protocol like TCP-AFC can achieve performance comparable to connections that do not traverse a GEO satellite link.

It was noticed that TCP performance degrades by 63.1% of the performance in error free environment, while competing with UDP traffic because of additional losses due to corruption on the satellite link having poor BER. In case of TCP-AFC this degradation was observed to be only 25.2%. Thus, in congestion situations, TCP-AFC was able to double the throughput of satellite connections by avoiding flow control when the losses did not occur due to congestion in the network. Importantly, this aggressive policy for losses due to channel errors does not penalize the connection while suffering from congestion induced losses. Additionally, the measurements also reinforce its compatibility with the existing terrestrial network, as its performance does not deteriorate in a congested network. The performance enhancement of TCP-AFC can be attributed to its ability of discriminating losses and adaptively controlling flow.

## 6. ACKNOWLEDGMENTS

The authors acknowledge N. G. Vasanthkumar, B. P. Chaniyara and Vishal Agrawal for helping in setting up the test-bed involving real satellite link at Advanced Digital Communication Test Group (ADCTG), Space Application Center, ISRO, Ahmedabad.

## 7. REFERENCES

- [1] T. Henderson, Networking over Next-Generation Satellite Systems, *Ph. D. Thesis, University of California*, 1999.

- [2] J. Postel, Transmission Control Protocol, *RFC 793*, 1981.
- [3] S. Kota, M. Marchese, Quality of Service for Satellite IP Networks, *Inter. Jour. of Satellite Commu. Network*, 2003.
- [4] S. Flyod, T. Henderson, The New Reno Modification To TCP's Fast Recovery Algorithm. *IETF, RCF 2582*, 1999.
- [5] K. Fall, S. Floyd, Comparisons of Tahoe, Reno & SACK TCP, [ftp://ftp.ee.lbl.gov/papers/sacks.ps.Z](http://ftp.ee.lbl.gov/papers/sacks.ps.Z), 1995.
- [6] C. Fu, S. Liew, TCP Veno: TCP Enhancement For Transmission Over Wireless Access Networks, *IEEE Journal On Selected Areas In Comm.*, Vol. 21, No. 2, 2003.
- [7] C. Casetti, M. Gerla, S. Mascolo, M. Sanadidi, R. Wang, TCP Westwood: Bandwidth Estimation for Enhanced Transport Over Wireless Links, *Proc. of Mobicom*, 2001.
- [8] K. Xu et al., Improving TCP Performance In Integrated Wireless Communications Networks, *Computer Networks* 47 (2005) 219–237.
- [9] A. Capone, F. Martignon, Bandwidth Estimates in the TCP Congestion Control Scheme, *Tyrrhenian IWDC 2001*.
- [10] R. Ludwig, A Case for Flow-Adaptive Wireless Links, *University of California, Report UCB/CSD-99-1053*, 1999.
- [11] I. Akyildiz, G. Morabito, S. Palazzo, TCP-Peach: A New Congestion Control Scheme for Satellite IP Networks, *IEEE/ACM Trans. Networking*, vol. 9, no. 3, 2001.
- [12] N. Kothari, K. Dasgupta, Performance Enhancement of SACK TCP for Wireless Network by Delaying Fast Recovery, *WOCN 2006*.
- [13] N. Kothari, B. Gambhava, K. Dasgupta, Adaptive Flow Control: An Extension to Delayed Fast Recovery, *ADCOM 2007*.
- [14] N. Kothari, J. Patel, K. Patel, K. Dasgupta, Performance Evaluation of TCP-AFC in an Emulated Real Network, *ICON 2008*.
- [15] M. Mathis, J. Mahdavi, S. Floyd, A. Romanow, TCP selective acknowledgment options, *RFC 2018*, 1996.
- [16] Y. Wang, H. Kung, Use of TCP Decoupling in Improving TCP Performance over Wireless Networks, *Mobile Networks and Applications*, 2000, pp. 1–20.
- [17] E. Blanton, M. Allman, K. Fall, L. Wang, A Conservative Selective Acknowledgment-based Loss Recovery Algorithm for TCP, 2003.
- [18] M. Allman, A. Falk, On effective Evaluation of TCP, *ACM Computer Communication Review*, 1999.
- [19] M. Richharia, Satellite Communication Systems – Design Principles, 2<sup>nd</sup> Edition, *Macmillan Press*, 1999.
- [20] C. Charalambous, V. Frost, J. Evans, Performance of TCP Extensions on Noisy High BDP Networks, *IEEE Communications Letters*, 1999.
- [21] D. Hong, C. Albuquerque, C. Oliveira, T. Suda, Evaluating the impact of emerging streaming media applications on TCP/IP performance, *IEEE Communication Magazine*, 2001.
- [22] S. Floyd, V. Paxson, Difficulties in Simulating the Internet, *IEEE/ACM Transactions on Networking*, 2001.