

Improved Explicit Congestion Notification for Satellite Networks

Arjan Durresi^{*a}, Mukundan Sridharan^{*a}, Chunlei Liu^{*a}, Raj Jain^{**b}

^aComputer and Information Science Department, The Ohio State University

^bNayna Networks, Inc.

ABSTRACT

Due to the fundamental satellite system characteristics such as global coverage, broadcast nature, and bandwidth on demand, satellite systems are excellent candidates for providing high data rate Internet access and global connectivity accommodating the a wide variety of applications. Provisioning of quality of service (QoS) within the advanced satellite system is the critical requirement. Congestion remains the main obstacle to Quality of Service (QoS) on the Internet.

Explicit Congestion Notification (ECN) is the only mechanism that delivers explicit congestion signals to the source. So improving the ECN feedback is essential for the future data and satellite networks and their QoS guarantees. In this paper we present a new traffic management scheme based on an enhanced ECN mechanism. In particular we used the mark-front strategy instead of the mark-tail one in ECN. The main advantage of mark-front strategy is that it sends faster feedback information about the congestion and consequently enables faster reaction from the source.

Our analysis and simulation results show that the mark-front strategy of ECN provides better link-efficiency, fairness among users, lesser buffer requirements and much less losses than the mark-tail strategy.

Keywords: Mark-front ECN, Improved ECN, Congestion control in Satellite Networks, Satellite based Internet

1. INTRODUCTION

The rapid globalization of the telecommunications industry and the exponential growth of the Internet are increasing the demands on global telecommunications. This demand is further increased by the convergence of computing and communications and by the increasing new applications such as Web surfing, desktop and video conferencing. Satisfying this requirement is one of the greatest challenges before telecommunications industry in the 21st century. Satellite communication networks can be an integral part of the newly emerging national and global information infrastructures. [1]

Satellite networks offer global coverage, broadcast capabilities, flexibility in bandwidth allocation, support for mobility and short time for implementing services even in areas with little infrastructure. All these qualities make satellite networks a good candidate for the future Internet infrastructure as broadband access network, high speed backbone and combination of both of them. However, to meet this goal, provisioning of quality-of-service (QoS) within the advanced satellite network systems is the critical requirement. Congestion remains the main obstacle to Quality of Service (QoS) on the Internet. Congestion is a critical problem especially in satellite networks, where TCP congestion control performance is affected by intrinsic satellite link characteristics such as latency, bandwidth, packet loss due to congestion, and losses due to transmission errors links [2].

Despite the fact that a number of schemes have been proposed for network congestion control in satellite networks, the search for new schemes continues [4-19]. One of most promising schemes to improve TCP congestion control is Explicit Congestion Notification (ECN) [4]. The Explicit Congestion Notification ECN, motivated by the DECBit scheme, provides a mechanism for intermediate routers to send early congestion feedback to the source before actual

* durresi@cis.ohio-state.edu; phone: 1 614 688 5610; fax: 1 614 292 2911; <http://www.cis.ohio-state.edu>; Computer and Information Science Department, The Ohio State University; 2015 Neil Ave., Columbus, OH, 43210, USA

** raj@nayna.com; phone: 1 408 956 8000X309, <http://www.nayna.com>, Nayna. Inc.; 157 Topaz St. Milpitas, CA 95035, USA

packet losses happen. The routers monitor their queue length. If the queue length exceeds a threshold, the router marks the Congestion Experienced bit in the IP header. Upon the reception of a marked packet, the receiver marks the ECN-Echo bit in the TCP header of the acknowledgment to send the congestion feedback back to the source. In this way, ECN delivers an even faster congestion feedback explicitly set by the routers. ECN is the only mechanism that delivers explicit congestion signals to the source. So improving the ECN feedback is essential for the future data and satellite networks and their QoS guarantees [3]

In most ECN implementations, when congestion happens, the congested router marks the incoming packet. When the buffer is full or when a packet needs to be dropped as in Random Early Detection (RED), some implementations have the "drop from front" option to drop packets from the front of the queue. However, none of these implementations mark the packet from the front of the queue.

We have studied the mark-front strategy in satellite networks. When a packet is sent from a router, the router checks whether its queue length is greater than the pre-determined threshold. If yes, the packet is marked and sent to the next router. The mark-front strategy differs from the current mark-tail policy in two ways. First, the router marks the packet in the front of the queue and not the incoming packet, so the congestion signal does not undergo the queuing delay as the data packets. Second, the router marks the packet at the time when it is sent, and not at the time when the packet is received. In this way, a more up-to-date congestion feedback is given to the source. In this paper we extend the ECN mark-front strategy presented in our team's previous work [20].

Our study finds that, by providing faster congestion feedback, mark-front strategy reduces the buffer size requirement at the routers; it avoids packet losses and thus improves the link efficiency when the buffer size in routers is limited. Our simulations also show that mark-front strategy improves the fairness among users.

The paper is organized as follows. In Section 2 we discuss the use of satellite networks in the future Internet infrastructure. TCP congestion control issues are presented in Section 3. In Section 4 we describe the mark-front strategy. The simulation results that show the advantages of mark-front strategy are presented in Section 5. In Section 6 we present the conclusions of this study.

2. SATELLITE-BASED INTERNET

Newly proposed broadband satellite systems promise to offer service comparable to current broadband terrestrial solutions. The increasing worldwide demand for more bandwidth and Internet access has created an extremely lucrative market for telecommunications system operators and service providers. Consequently, a large number of new and existing satellite operators have proposed ambitious broadband satellite networks utilizing recent advances in satellite technology and spectrum allocation. These proposed broadband satellite networks plan to change the way that users communicate with each other by offering direct broadband satellite access to every potential user for virtually the same price, regardless of location.

In recent years a significant investment has been made in the planning and development of broadband satellite networks. Interest in Ka-band satellite systems has dramatically increased, with over 450 satellite applications filed with the ITU. In the U.S., there are currently 13 Geostationary Satellite Orbit (GSO) civilian Ka-band systems licensed by the Federal Communications Commission (FCC), comprising a total of 73 satellites. Two Non-Geostationary Orbit (NGSO) Ka-band systems, comprising another 351 satellites, have also been licensed. Eleven additional GSO, four NGSO, and one hybrid system Ka-band application for license and 16 Q/V-band applications have been filed with the FCC [1]. The main advantages and disadvantages of GSO versus non-GSO architectures have been discussed in [2].

Some of the main opportunities for satellite networks in the Internet infrastructure are: Internet Access Service, Internet Backbone, caching, IP multicast, and interconnection of LAN/WANs.

Despite the widespread deployment of terrestrial and wireless broadband capacity, there is still a need for Internet access services in many parts of the world, especially in areas of low subscriber density and with little telecommunication infrastructure. The rapid deployment of satellites and their large coverage areas open the opportunities for to satellite

solutions in Internet access. Web access is already being provided on a small scale to residential users and business users by a number of satellite operators and service providers through mainly Ku-band capacity. The future broadband satellite market offers the potential to expand on this initial market entry service to offer a less expensive, higher-bandwidth solution, capable of data rates higher than current cable modem and ADSL implementations, from 1.5 Mbps to 45 Mbps.

The Internet today can be characterized as U.S.-centric, with the majority of Web content residing on U.S. servers. Even though the share of U.S. Internet hosts is predicted to fall relative to the rest of the world, global ISPs are increasingly demanding access to the U.S. A possible future application for Internet backbone services via satellite may be Intra-regional backbone services, especially over large landmasses. The rise in announced regional GSO satellite systems point to this trend in backbone connectivity. The asymmetry in bandwidth demand may favor the flexibility on bandwidth allocation of satellite networks. The satellite backbone service also can provide a good alternative to congested public telecommunication backbone networks.

Caching may be one of the most important applications for the growth of Internet traffic over satellite. Today, only 1,000 Web sites represent over 80 percent of the most frequently accessed content on the Internet. These Web pages would be transmitted at regular intervals to designated server destinations. Content is then stored in the servers for local users to access. Navigating the Internet through cached content relieves the need to retrieve Internet data from the source, thus reducing delays and congestion inherent with accessing popular Internet content.

Closely related to caching described above, IP multicasting has become an extremely popular topic in the Internet community because it addresses both the emerging "push" applications popular on the Internet and the congestion they cause. Satellite networks with their intrinsic broadcast capability are a good solution for multicast transmissions.

As globalization expands into regions of the world where the telecommunications infrastructure is inadequate for broadband applications, the need for high-speed reliable connectivity between the regional and central offices becomes evident. Satellites, due to their ability to offer global coverage, are poised to gain a foothold in this market.

3. TCP CONGESTION CONTROL IN SATELLITE NETWORKS

The delays in GSO systems and delay variations in NGSO systems affect both real-time and non-real-time applications in an acknowledgement- and time-out-based congestion control mechanism. As a result, the congestion control issues for broadband satellite networks are somewhat different from those of lower-latency terrestrial networks [21-25].

The main characteristics of the end-to-end path that affects transport protocols are latency, bandwidth, packet loss due to congestion, and losses due to transmission errors links [21-25]. If part of a path includes a satellite channel, these parameters can vary substantially from those of terrestrial networks. Especially latency in satellite networks will impact the TCP congestion control mechanism.

Among the three components of latency propagation delay, transmission delay and queuing delay, propagation delay is the dominant part in the broadband satellite links. The RTT between terminals using GSO satellite links could be around 600 ms. TCP uses acknowledgment feedback to probe the network and achieve congestion control and reliable delivery. The slow feedback due to the large RTT as in all feedback control systems will weakness the TCP congestion control and consequently will decrease the link utilization. In case of LEO satellites the propagation delay will vary and the connection path may change and the queuing delay may be significant. Large variations of RTT may lead to false timeouts and retransmissions.

Both methods employed in TCP/IP to deliver congestion signals, timeout and loss detection use packet loss as congestion signal. Packet losses not only increase the traffic in the network, but also add large transfer delay. Especially in satellite networks it is not convenient to use loss detection as congestion signals. In satellite networks, dropped packets have already consumed expensive network bandwidth and will contribute in aggravating the consequences of high latency in the TCP congestion control scheme.

Recognizing the need for a more direct feedback of congestion information, the Internet Engineering Task Force (IETF) has come up with Explicit Congestion Notification (ECN) method for IP routers [4, 5]. ECN is much more powerful than the simple packet drop indication used by existing routers and is more suitable for high distance-bandwidth networks.

In order to manage traffic, IP routers need to inform the sources about their load levels so that the sources can increase or decrease their traffic to match the available capacity. In the simplest case, the feedback can consist of dropping packets as is currently done in the IP routers. The next step is to send feedback in the network layer header. The Internet Engineering Task Force (IETF) has recently introduced with RFC 2481 [4] the "Explicit Congestion Notification (ECN)." Two bits in the IP header have been reserved for this purpose. One of these bits is used to indicate congestion in the router. The routers start marking this bit using a RED-like algorithm based on the average queue length. The second bit is used by the sources to indicate whether the flow is ECN capable. Newer implementation of TCP will be ECN capable and will react to ECN marking in a manner similar to that of packet drop. The receivers echo the ECN bit back to the source in TCP header of the returning ACKs. Sources respond to ECN once per round trip time (RTT). To guard against loss of ACKs, receiver continues to set ECN-Echo bit in subsequent ACKs (even if further packets do not have CE bit set) until it receives a packet with Congestion Window Reduced (CWR) bit set. A source, after responding to congestion indication by halving the congestion window (CWND), sets CWR bit in next packet sent (in order to inform receiver about action taken in response to congestion). After receiving a packet with CWR bit set, receiver does not set ECN-Echo bit in ACKs until it gets another packet with CE bit set.

The two major advantages of the use ECN are: in case of not very high level congestion the packets are not dropped and second it could provide provides a more detailed information about congestion as will be seen in our simulation study later. Both these ECN features are very suited to satellite networks. It is very important to avoid the dropping of packets that already have consumed the critical resources of satellite long links. Also because satellite links can be seen as feedback control system with a long feedback time, it is very important to have the most accurate information about the congestion situation, in order provide the best reaction and not to spend several long RTTs to adjust the sender reaction.

4. MARK-FRONT STRATEGY IN ECN

We have studied and experimented the mark-front strategy in satellite networks. When a packet is sent from a router, the router checks whether its queue length is greater than the pre-determined threshold. If yes, the packet is marked and sent to the next router. The mark-front strategy differs from the current mark-tail policy in two ways. First, the router marks the packet in the front of the queue and not the incoming packet, so the congestion signal does not undergo the queuing delay as the data packets. Second, the router marks the packet at the time when it is sent, and not at the time when the packet is received. In this way, a more up-to-date congestion feedback is given to the source. In this paper we extend the ECN mark-front strategy presented in our team previous work [20].

We make the following assumptions, Figure 1. (1) Receiver windows are large enough so the bottleneck is in the network. (2) Senders always have data to send. (3) There is only one bottleneck link that causes queue buildup. (4) Receivers acknowledge every packet received and there are no delayed acknowledgments. (5) The queue length is measured in packets and all packets have the same size.

In [20] has been estimated the benefits of mark-front compared to -mark tail strategy. In particular was proved that:

1. In a TCP connection with ECN mark-tail congestion control, if the fixed round trip time is r seconds, the bottleneck link rate is d packets per second, and the bottleneck router uses threshold T for congestion detection, then the maximum queue length that can be reached in slow start phase is less than or equal to $2T + rd - l$.
2. In a TCP connection with ECN mark-front congestion control, if the fixed round trip time is r seconds, the bottleneck link rate is d packets per second, and the bottleneck router uses threshold T for congestion detection, then the maximum queue length that can be reached in slow start phase is less than or equal to $T + rd$.
3. In a path with only one connection, the optimal threshold that achieves full link utilization while keeping queuing delay minimal in congestion avoidance phase is $rd - l$. If the threshold is smaller than this value, the link will be

under-utilized. If the threshold is greater than this value, the link can be full utilized, but packets will suffer an unnecessarily large queuing delay.

Combining the results in 1, 2 and 3, we can see that the mark-front strategy reduces the buffer size requirement from about 3rd to 2rd. It also reduces the congestion feedback's delay by one fixed round-trip time. From this conclusion it is clear that applying mark-front ECN in satellite networks will bring substantial advantages, because of high values of r , round trip times, in satellite networks.




Figure 1. ECN mechanism

5. SIMULATION RESULTS

In order to compare the mark-tail ECN with our mark-front ECN scheme, we carried out a set of simulations using the NS simulator [26]. We modified the RED algorithm in ns simulator to deterministically mark the packets when the real queue length exceeds the threshold.

For all our simulations we used the following configuration, shown in Figure 2. A Number of sources $S_1, S_2, S_3, \dots, S_n$ are connected to a router R_1 through 10Mbps, 2ms delay links. Router R_1 is connected to R_2 through a 1.5Mbps, T_S ms delay link. R_2 is connected to R_3 through a 1.5Mbps, T_S ms delay link and a number of destinations $D_1, D_2, D_3, \dots, D_n$ are connected to the router R_3 via 10Mbps 4ms delay links. The link speeds are chosen so that the congestion will happen only between routers R_1 and R_2 where our scheme is tested. This configuration can simulate the case LEO, MEO and GSO satellite networks.

An FTP application runs on each source. Reno TCP is used as the transport agent. (The modifications were made to the Reno TCP). The packet size is 1000 bytes and the acknowledgement size is 40 bytes.

Simulation Scenarios

With the basic configuration described above the following simulation scenarios were used to test our scheme.

1. Two sources with the same RTT.
2. Ten overlapping connections with same RTT, each connection starting 0.3 seconds after the previous one.
3. Ten connections with different RTT. The minimum RTT is 259 ms and for each connection the RTT increases by 10 ms, so RTTs are 269, 279, ... up to 349 ms.

Metrics

We used four metrics to compare the two strategies. The first metric is the *Buffer size requirement* for zero loss congestion control. This is the maximum queue size that can be built up at the router in the slow start phase before the congestion signal takes effect at the congested router. If the buffer size is greater or equal to this value, no packet loss will happen. This metric is measured as the maximum queue length in the entire simulation.

The second metric, *Link efficiency*, is calculated from the number of acknowledged packets (not counting the retransmissions) divided by the possible number of packets that can be transmitted during the simulated time. Because of the slow start phase and possible link idling after the window reduction, the link efficiency is always smaller than 1. Link efficiency should be measured with long simulation time to minimize the effect of the initial transient state. We tried different simulation times from 5 seconds to 100 seconds. The results for 10 seconds show the essential features of the strategy, without much difference from the results for 100 seconds. So the simulation results presented in this paper are based on 10-second simulations.

The third metric, number of *Dropped Packet*, counts the number of packet dropped during the simulation time. Because the high cost of satellite links, it is desirable not to drop packets that already have used satellite bandwidth.

The fourth metric, *Fairness index*, is calculated according to the formula in [27]. If n connections share the bandwidth, and x_i is the number of acknowledged packets of connection i , then the fairness index is calculated as:

$$Fairness = \frac{(\sum x_i)^2}{n \sum x_i^2}$$

and unfairness = 1 - fairness.

Buffer Size Requirement

Figure 3 shows the buffer size requirement for mark-tail and mark-front for ten connections with same RTT with $TS = 59ms$, which simulates LEO connections. The mark-front provides better results than tail-mark. So using mark-front less buffering is required in order to avoid losses.

In Figure 4a there are compared the results of Buffer Requirement for ten overlapping sources with same RTT and $TS = 259$. In Figure 4b there are compared the results of Buffer Requirement for ten sources with different RTT and $TS = 259$ ms. Both experiments simulate GSO or hybrid satellite networks. Again the Buffer Requirement is lower in the case of front-mark strategy.

In all experiments mark-front strategy requires less buffers in order not to have losses. This result represents two practical advantages of mark-front strategy. First less buffers means less expensive devices and what is more important, less queuing delay in the routers. So mark-front strategy adds less delay to the already high satellite latency.




Figure 2. Simulation configuration




Figure 3 Ten connections with same RTT with TS = 59 ms

Link Efficiency

Figure 5 shows the link efficiency for two scenarios. In both cases, the efficiency increases with the threshold, until the threshold is about rd, where the link reaches almost full utilization. Small threshold results in low link utilization because it generates congestion signals even when the router is not really congested. Unnecessary window reduction actions taken by the source lead to link idling.

In Figure 5a there are compared the Link Efficiency obtained with two sources, same RTTs, TS = 59ms. This experiments simulates LEO or hybrid satellite networks. Mark-front has better link efficiency because when congestion happens, this strategy provides a faster congestion feedback than mark-tail. Faster feedback prevents the source from sending more packets that will be dropped at the congested router.

In Figure 5b there are compared the Link Efficiency obtained with two sources, different RTTs, TS = 259ms. This experiments simulates GSO or hybrid satellite networks. Again mark-front performs better than mark-tail.

Dropped Packets

In Figures 6 and 7 there are compared the dropped packets using mark-front and mark-tail in different conditions. We would like to stress that the number of Dropped Packets is an important metric especially in satellite networks. The goal is to minimize the dropping of packets that have already have consumed expensive satellite bandwidth.

Figure 6a shows Dropped Packets in case of ten overlapping sources, same RTTs, TS = 59ms and Figure 6b shows Dropped Packets for two sources, same RTTs, TS = 59ms. In both experiments, which simulate LEO or hybrid satellite

networks, mark-front result in much less dropped packets. This is explained by the faster feedback, which enables the sources to change the rate faster and to avoid losses.

Figure 7 shows Dropped Packets for $TS = 259ms$, which could be the case of GSO or hybrid satellite networks. In Figure 7a there are ten overlapping sources, same RTTs, in Figure 7b there are two sources, same RTTs, and in Figure 7c there are ten sources, different RTTs. In all the experiments mark-front results in less packet dropped. In important point to be stressed is that the case with different RTT, which is a more realistic experiment, provides the best results.

Fairness

When mark-tail strategy is used, old connections have the tendency to occupy the buffer and the lock-out new connections. Also connections with smaller RTT have advantages with mark-tail, because they more ACKs and grow faster the congestion window. Mark-front alleviated the discrimination against large RTT by sending faster feedback.

Figure 8a shows the Unfairness for ten overlapping sources, same RTTs, $TS = 59ms$. Mark-front results in better fairness among flows with different RTTs. Figure 8b shows the Unfairness for ten overlapping sources, different RTTs, $TS = 259ms$. Again Mark-front outperforms mark tail. In case of shorter RTT mark-front provides a better improvement than in the case of longer RTT. In the case of higher RTT the queuing delay is less important component of the end-to-end latency.




Figure 4a. Buffer requirement for ten overlapping sources, same RTTs, $TS = 259ms$




Figure 4b. Buffer requirement for ten sources, different RTTs, $TS = 259ms$




Figure 5a. Link Efficiency for two sources, same RTTs, $TS = 59ms$




Figure 5b. Link Efficiency for two sources, different RTTs, $TS = 259ms$




Figure 6a. Dropped Packets for ten overlapping sources, same RTTs, TS = 59ms




Figure 6b. Dropped Packets for two sources, same RTTs, TS = 59ms




Figure 7a. Dropped Packets for ten overlapping sources, same RTTs, TS = 259ms




Figure 7b. Dropped Packets for two sources, same RTTs, TS = 259ms




Figure 7c. Dropped Packets for ten sources, different RTTs, TS = 259ms




Figure 8a. Unfairness for ten overlapping sources, different RTTs, TS = 59ms




Figure 8b. Unfairness for ten overlapping sources, different RTTs, TS = 259ms

6. CONCLUSIONS

In this paper we have studied the mark-front strategy used in Explicit Congestion Notification, applied in satellite networks. The packets are marked at the front of the queue and not in the tail of the queue as is the case of mark-tail strategy. The mark-front strategy send faster the congestion signal back to the source, which results in faster reaction from the source and in improved feedback-based TCP congestion control.

The mark-front strategy reduces the buffer size requirements at the router for the condition without losses. That means less complex devices and less delay, both important in satellite networks with already high latency. By sending faster feedback signaling about the congestion, mark-front enables the source to adjust its window in time to avoid packet losses and link idling, consequently improves the link efficiency. This result is in particular important for satellite networks, where the link bandwidth is more expensive and less available than in terrestrial networks. The mark-front strategy improves the fairness among old and new flows, and helps to alleviate TCP' discrimination against connections with large round trip time. The mark-front strategy reduces significantly the packet losses, which is again a very important metric in satellite networks, where the goal is not to drop packets that already have consumed expensive satellite bandwidth. All our simulation results show the improvements of mark-front strategy.

Based on the results, we conclude that mark-front strategy is an easy to implement enhancement that provides better TCP congestion control and improves important performance metrics in satellite networks.

ACKNOWLEDGEMENTS

This research was supported in part by grants from OAI, Cleveland, Ohio and NSF CISE grant #9980637.

REFERENCES

1. Arjan Durrezi, Sastri Kota, Raj Jain, Mukul Goyal, "Achieving QoS for TCP Traffic in Satellite Networks with Differentiated Services", Accepted for publication in *Space Communications Journal*.
2. Sastri Kota, "Multimedia Satellite Networks: Issues and Challenges," *Proc. SPIE International Symposium on Voice, Video, and Data Communications*, Boston, Nov 1-5, 1998
3. Arjan Durrezi, Mukundan Sridharan, Chunlei Liu, Mukul Goyal, Raj Jain, " Congestion Control using Multilevel Explicit Congestion Notification in Satellite Networks," submitted to ICCCN2001.
4. K. Ramakrishnan and S. Floyd, "A Proposal to add Explicit Congestion Notification (ECN) to IP," *RFC 2481*, January 1999

5. Sally Floyd, "TCP and Explicit Congestion Notification," *Computer Communications Review*, Vol. 24, No. 5, October 1994, pp. 10-23.
6. Sally Floyd and Van Jacobson, "Random Early Detection gateways for Congestion Avoidance," *IEEE/ACM Transactions on Networking*, Vol. 1, No. 4, August 1993, pp. 397-413.
7. Sally Floyd and Van Jacobson, "On Traffic Phase Effects in Packet-Switched Gateways," *Internetworking: Research and Experience*, Vol. 3, No. 3, September 1992, pp. 115-156.
8. Sally Floyd, "Connections with Multiple Congested Gateways in Packet-Switched Networks Part 1: One-way Traffic," *Computer Communications Review*, Vol. 21, No. 5, October 1991, pp. 30-47.
9. David D. Clark and Wenjia Fang, "Explicit allocation of best-effort packet delivery service," *IEEE/ACM Trans. on Networking*, Vol. 6, No. 4, August 1998, pp. 362-373.
10. W. Feng, D. Kandlur, D. Saha and K. Shin, "Blue: A New Class of Active Queue Management Algorithms," University of Michigan, *Technical Report UM-CSE-TR-387-99*, 1999.
11. S. Kalyanaraman, R. Jain, S. Fahmy, R. Goyal and B. Vandalore, "The ERICA Switch Algorithm for ABR Traffic Management in ATM Networks," Submitted to *IEEE/ACM Trans. on Networking*, November 1997, available at <http://www.cis.ohio-state.edu/~jain/papers/erica.htm>
12. Sally Floyd and Kevin Fall, "Promoting the Use of End-to-End Congestion Control in the Internet," *IEEE/ACM Trans. on Networking*, August 1999.
13. Sally Floyd and Kevin Fall, "Router Mechanisms to Support End-to-End Congestion Control," *Unpublished Manuscript*, <http://www-nrg.ee.lbl.gov/floyd/papers.html>
14. M. Mathis, J. Semke, J. Mahdavi and T. Ott, "The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm," *Computer Communications Review*, Vol. 27, No. 3, July 1997, pp. 67-82.
15. D. Chiu and R. Jain, "Analysis of the Increase/Decrease Algorithms for Congestion Avoidance in Computer Networks," *Journal of Computer Networks and ISDN Systems*, Vol. 17, No. 1, June 1989, pp. 1-14.
16. Sally Floyd and Tom Henderson, "The NewReno Modification to TCP's Fast Recovery Algorithm," *Internet Draft - Work in Progress*, February 1999.
17. M. Mathis, J. Mahdavi, S. Floyd and A. Romanow, "TCP Selective Acknowledgment Options," *RFC 2018*, October 1996.
18. K. K. Ramakrishnan and R. Jain, "A binary feedback scheme for congestion avoidance in computer networks," *ACM Transactions on Computer Systems*, Vol. 8, No. 2, May 1990, pp. 158-181.
19. R. Jain, S. Kalyanaraman, R. Viswanathan, "Rate Based Schemes: Mistakes to Avoid." *AF-TM 94-0882*, September 1994.
20. Chunlei Liu and Raj Jain, "Improving Explicit Congestion Notification with the Mark-Front Strategy," *Computer Networks*, 35, 2000, pp. 185-201
21. Nasir Ghani, Sudhir Dixit, "TCP/IP Enhancements for satellite networks," *IEEE Communications Magazine*, pp.64-72, Vol. 37, No 7, July 1999
22. Mark Allman, "TCP byte counting refinements," *ACM Computer Communication Review*, July 1999.
23. Mark Allman, Chris Hayes, Hans Kruse, Shawn Osterman, "TCP Performance over satellite links," *5th International Conference on Telecommunication Systems*, 1997
24. M. Allman, D. Glover, L. Sanchez, "Enhancing TCP over satellite channels using standard mechanisms," *RCF 2488, IETF*, January 1999
25. M Allman editor, "Ongoing TCP research related to satellites," *RFC 2760, IETF*, February 2000
26. NS Simulator, Available from <http://www-mash.cs.berkeley.edu/ns>.
27. Raj Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Simulation and Modeling*, New York, John Wiley and Sons Inc., 1991.