

Web access over a multiple access channel: Evaluations and improvements

Xiaoming Zhou, Nelson X. Liu and John S. Baras
Institute of System Research
University of Maryland at College Park, MD, 20742

Abstract—We evaluate web browsing performance over a multiple access satellite channel for three different MAC layer protocols. A web user behavior model is used to generate realistic source traffic. A new transport layer protocol called RWBP is proposed to solve the TCP problems inside satellite networks. RWBP uses per-flow queuing, round robin scheduling and receiver window backpressure for congestion control. We compare its performance with TCP over the three multiple access protocols. Our simulation results show that RWBP can reduce the multiple access channel load and at the same time improve the web page response time when driven by realistic web traffic. For the MAC layer protocols, combined free demand assignment multiple access (CFDAMA) always performs better than the other two protocols.

I. INTRODUCTION

We consider a star network shown in figure 1, the clients download web pages from the Internet through the forward channel from the satellite gateway to the very small aperture terminals (VSATs). The metrics, which are important to web users, are the web object and web page response time. The transport layer protocol used for web browsing is TCP and there are at least two aspects in TCP which increases the end-to-end response time dramatically.

First is the slow start algorithm in TCP congestion control. When a TCP connection is set up, it first enters the slow start phase with its initial congestion window size set to one segment. Since most of the web file sizes are small, they usually complete the transmissions in the slow start phase before they could enter the congestion avoidance phase. Hence, the response time increases dramatically due to the long propagation delay in the satellite networks.

Second is the congestion losses at the satellite gateway. TCP detects packet loss by three duplicate acknowledgements. If the loss can be recovered by the fast retransmission algorithm, it will increase the response time by at least one round trip time. For a small web file, the congestion window may be so small that it doesn't have enough data packets to trigger the receiver to generate three duplicate acknowledgements. Under such circumstance, timer will become the last resort for loss recovery which increases the response time even more.

Besides the congestion control in the forward channel, the reverse channel from the VSATs to the satellite gateway poses another bandwidth allocation problem. The reverse channel is a multiple access channel (figure 1) and a MAC scheme is needed to allocate the reverse channel bandwidth efficiently




Fig. 1. Direct to user satellite networks (star topology)

so that it will not become the bottleneck of the end-to-end performance.

According to the control policy, MAC schemes can be classified into static or dynamic MAC [10]. Fixed assignment falls into the static MAC. Fixed assignment includes TDMA, FDMA and CDMA. The fixed assignment is suitable for the gateways with relatively smooth aggregate traffic arrival patterns. For the VSATs with very bursty traffic load, the efficiency of fixed assignment is low. Dynamic MAC includes random access and reservation-based MAC. Random access obtains reasonable throughput only at low load and offers little quality of service guarantees such as delay, jitter etc. Reservation based MAC can achieve very high efficiency. However the reservation phase takes a significant amount of time in the satellite network because of the large propagation delay. In figure 1, the reverse channel bandwidth controller is located at the network operation center (NOC), therefore the reservation phase increases the packet delay by two hops (more than 0.5 sec).

In this paper, we evaluate web performance over the multiple access satellite channel. We address this problem from three layers. First, at the application layer, a web user behavior model is developed to generate realistic source traffic; Second, a new transport layer protocol is proposed which can solve the problems of TCP inside satellite networks; Third, three MAC layer protocols are evaluated with the given web traffic characteristics and the transport layer protocols. The MAC protocols we evaluate are very similar to those implemented in practice. The contribution of this paper is that we are the first to address this problem in a systematic manner rather than focus on one or two layers. Our goal is to analyze web browsing performance over satellite and to improve its performance by designing appropriate protocols.

The rest of this paper is organized as follows. Section II

relates our work to previous research. Section III describes the web user behavior model. Section IV presents the new transport layer protocol (RWBP). Section V describes the MAC protocols we evaluate. Section VI gives the simulation results. Finally, Section VII concludes this paper.

II. RELATED WORK

Based on the observation that TCP slow start yields poor performance for short and bursty web traffic in term of bandwidth utilization and page response time, Padmanabhan [13] proposes a fast start algorithm to replace the slow start algorithm in TCP. In fast start, the sender caches network parameters so that it can avoid paying slow start penalty for each web page. In Henderson's experiments [8], he assumes no packet loss and fixed RTTs so the results he gets are the response time for the best case. He investigates both end-to-end TCP and connection splitting based schemes. Akyildiz etc propose a transport layer protocol called TCP Peach [1] for satellite networks. TCP Peach replaces the slow start and fast recovery algorithms in TCP Reno with sudden start and rapid recovery. TCP Peach is shown to perform better than TCP Reno in term of average response time of web pages. However the above related work focuses on TCP layer. The satellite channel is abstracted as a point-to-point link with large delay and the MAC problem over the reverse channel hasn't been addressed.

Connors [6] proposes a response initiated multiple access (RIMA) protocol for web access over a star satellite network. The basic idea is based on the following observation: when a large packet is received, with high probability the receiver will send an acknowledgement; when a small packet is received, with reasonable probability the receiver will send a data packet. The satellite acts as a scheduler. It checks the packet size and port number to find out how many slots needed by the receiver and try to allocate enough slots to the receiver. Choi [7] adopts a MAC protocol from a previous research in the context of hybrid fiber coax (HFC) networks. By grouping and piggy backing MAC requests, the MAC protocol can mitigate the performance degradation caused by large propagation delay. Connors and Choi focus on the MAC layer and use end-to-end TCP for their simulations while the TCP over satellite problems have not been addressed.

Mhatre [12] considers both the transport layer and the MAC layer problems. However the STP [8] adopted in his scheme still uses TCP congestion control algorithms which well known is not efficient and effective for the satellite connections. He over-provisions the forward channel bandwidth and the congestion over this channel at the gateway hasn't been considered. In his simulations, the web user's behavior is simplified to always request objects of fixed size 4KB from a server.

III. WEB USER BEHAVIOR MODEL

Web traffic has become the largest component of the Internet traffic. A good web traffic model is essential for simulations and experiments to investigate end-to-end performance such

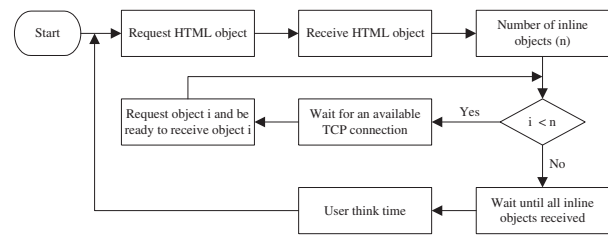


Fig. 2. Web user behavior model

as object response time and page response time. There are several empirical web user behavior models in the literature [11][5][2][14], these models are based on the traffic traces collected either in a local area network or in a wide area network. The elements of a HTTP model are: 1) HTTP request length; 2) HTTP reply length; 3) number of inline objects per page; 4) user think time between retrieval of two successive pages. Mah and Smith [11][14] argue that it is not sufficient to simply transmit data into the network according to these traffic models. This is because these application-dependent but network-independent models should be layered over the transport layer protocol so that the sizes and timing of the packets can be modeled accurately.

We will use the the model shown in figure 2 to generate the web traffic at the application layer for our experiments. First the web browser requests the HTML object. Once the HTML object is received, the browser figures out how many inline objects are in the page and begins to request the inline objects. After all the inline objects are received, the user 'thinks' for some time and starts to retrieve another web page. This model can model HTTP 1.0 with one or multiple TCP connections as well as HTTP 1.1 with a persistent connection. If there is a TCP connection available whether it is one of the parallel connections in HTTP 1.0 or the persistent connection in HTTP 1.1, the browser can send a new request through that connection. The HTTP request length [11] will be modeled by a bimodal distribution with one large peak occurring around 250 bytes and another, a smaller one around 1KB. The reply file sizes will be modeled by Pareto distribution with $k = 1KB$ and $\alpha = 1.1$. The number of inline objects per page will be modeled by a Gamma distribution with mean of 5.55 and standard deviation of 1.14. The user think time will be modeled also by Pareto distribution with $k = 5sec$ and $\alpha = 1.5$.

IV. IMPROVE WEB PERFORMANCE BY DESIGNING A NEW TRANSPORT LAYER PROTOCOL

Considering the interoperability issue, we adopts the connection splitting based scheme [15][3][8] which is currently used in the industry, and design a new transport layer protocol for reliable data transfer over the satellite link.

In the network as shown in figure 1, an end-to-end TCP connection is split into three connections at the proxies. The first one is set up between the server and the upstream proxy; the second is from upstream proxy to downstream proxy; and




Fig. 3. Queuing model for the satellite gateway and proxies. Flow control is done between downstream proxy and upstream proxy, and between satellite gateway and upstream proxy. It is also done between link layer and the IP layer, between the IP layer and transport layer, and inside transport layer.

the third is from the downstream proxy to the client. Normal TCP is used for the server-proxy and proxy-client connections. Receiver Window Backpressure Protocol (RWBP) is designed for the proxy-proxy connection to transfer data over the satellite channel.

A. Queuing Model at the Satellite Gateway

The satellite gateway and the VSATs are connected to the local proxies (figure 1) through high-speed links whose bandwidth is much larger than the satellite link bandwidth. Therefore between the upstream and downstream proxies, the satellite link is the bottleneck link. The satellite link is used to transfer TCP traffic as well as multicasting video or audio traffic. At the satellite gateway, we assume that a high priority queue is used for multicasting traffic and a low priority queue is used for TCP traffic. These two queues are link layer queues at the terrestrial-satellite output interface (figure 3).

B. Queuing Model at the Proxies

For a normal router, only those packets waiting for transmission are buffered in the IP output queue. However, the proxies have to buffer the packets waiting for transmission as well as those packets that have been transmitted but not acknowledged. A normal router keeps all the packets in a FIFO queue while the proxies have a queue for each connection. From figure 3, we can see that the input queues at IP layer and link layer should be almost always empty if we assume that the processing rate is not the bottleneck. Therefore the possible queuing points at the proxies are transport layer receive/send buffer, IP output queue and link layer output queue.

C. Receiver Window Backpressure Protocol

TCP uses slow start to probe the bandwidth at the beginning of a connection and uses additive increase and multiplicative decrease (AIMD) congestion avoidance to converge to fairness in a distributed manner. RWBP is based on TCP; however RWBP cancels all the congestion control algorithms in TCP and uses per-flow queuing, round robin scheduling [4] and receiver window backpressure for congestion control (figure 3).

Flow control is done between the downstream proxy and the upstream proxy at the transport layer by using the receiver window (figure 3). For each RWBP connection, the downstream proxy advertises a receiver window based on the available

buffer space for that connection just as in TCP. RWBP does not use window scaling to advertise large windows to upstream proxy because large window scale factor can produce inaccurate values. In RWBP, the 16-bit receiver window field is still used but its unit is maximum segment size rather than byte. Similarly flow control is also done between the satellite gateway and the upstream proxy at the link layer (figure 3). The low priority queue at the satellite gateway advertises a receiver window to the upstream proxy so that the low priority queue will not overflow.

In addition, flow control is done between the transport layer and the IP layer, and between the IP and the link layer. At the upstream proxy, a round-robin scheduler can send a packet for a RWBP connection only if its advertised receiver window is open and there is at least one packet buffer space available at the IP output queue. When there is no packets can be sent or the available advertised receiver window size is zero, the scheduler goes on to serve the next connection. When the IP layer output queue sends packets to the link layer, it has to make sure that the link layer queue is not going to be overflowed. This allows the link layer congestion backpressure to propagate to IP layer and then to transport layer. Inside the transport layer, when packets are moved from upstream connection receive buffer to the downstream send buffer, flow control is performed so that the send buffer will not overflow. This way the congestion is back pressured to the receive buffer of the upstream connection and a smaller receive window is going to be sent to the source. Finally the congestion is back pressured to the source. When the traffic load decreases, more buffer space will be available so a larger advertised receiver window will be sent to the source and the source can speed up. If some connections are bottlenecked upstream or are idle because the application layer does not have data to send, the scheduler can send packets from other connections and high satellite link efficiency can be achieved. The round-robin scheduler does not take into account the packet sizes. Connections with larger packet sizes can get more bandwidth than those with smaller packet sizes. This problem can be solved by a more sophisticated scheduler and is left as future work.

The above flow control scheme can guarantee that there is no buffer overflow in the downstream proxy queues, in the upstream proxy queues or in the low priority queue at the satellite gateway. Therefore RWBP eliminates congestion losses inside the satellite network. RWBP also includes a new error recovery scheme for link layer corruption and a buffer management scheme, please refer to [16] for the detail description of RWBP.

V. MULTIPLE ACCESS PROTOCOLS IN THE REVERSE CHANNEL

Most of the time, the terminal users browse web or download files from the Internet so the traffic for the reverse channel mainly contains TCP acknowledgement packets and HTTP request packets (figure 4).




Fig. 4. Reverse channel traffic characteristics for web browsing

Random access such as slotted aloha can achieve small delay at light traffic loads. When the system is heavily loaded, the packet delay will increase dramatically. Therefore random access is ruled out of our consideration. We will use fixed assignment TDMA as a baseline MAC protocol to compare with other two MAC protocols described in the following.

A. Combined Free Demand Assignment Multiple Access

Combined free demand assignment multiple access (CFDAMA) was proposed as a MAC protocol for the satellite reverse channel [9]. CFDAMA first allocates reverse channel bandwidth to the VSATs on a demand basis as in a reservation based scheme. However when there is no demand, the scheduler allocates the free bandwidth to the VSATs in a round-robin manner. When the channel is lightly loaded, the probability of a VSAT obtains free assignment is high and small packet delay can be achieved. While the reverse channel is heavily loaded, it behaves like a reservation scheme.

B. Aloha Periodic Stream

The aloha periodic stream scheme tries to explore the regeneration cycles of the reverse channel traffic as shown in figure 4. If the VSAT hasn't transmitted traffic for a period of time, it will be in the idle state. When the VSAT becomes active, an Aloha request will be sent to the NOC. After the Aloha request is received, the VSAT is assigned periodic bandwidth based on the number of active terminals and their requests. If the persistent backlog packets of a terminal exceed some threshold during the active period, additional bandwidth is requested by piggybacking the request in the data packets. Additional bandwidth is provided until the maximum is attained or the backlog is decreasing. The bandwidth is given an inactivity timeout values. If no packet arrives from the VSAT during the timeout period, the bandwidth assigned to the user will be released.

VI. PERFORMANCE EVALUATION

In this section, we evaluate the performance of web browsing over a satellite multiple access channel with OPNET. The metrics we are interested in are web object and web page response time. In this paper, we consider web traffic only. For the performance evaluation of mixed traffic with both TCP and high priority multicasting traffic, please refer to [16].

In figure 1, 128 clients download web files from 128 Internet servers. The forward satellite channel bandwidth is 4Mbps and the reverse channel bandwidth is 400kbps. The link bandwidth from each server to the upstream proxy is 4Mbps; and the link

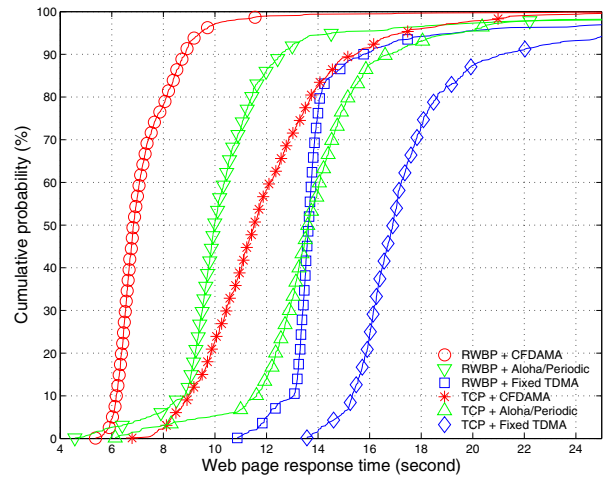


Fig. 5. Web page response time for different transport layer and MAC layer protocols

bandwidth from each downstream proxy to its corresponding client is also 4Mbps. The link delay from each server to the upstream proxy is 40ms and the link delay from each downstream proxy to its corresponding client is 0.01ms. The peak rate in RWBP is set to 900kbps and all the proxy buffer sizes are set to the peak rate satellite delay product. The satellite gateway buffer size is set to the satellite bandwidth delay product. The maximum segment size is 512 bytes. HTTP 1.1 is used at the application layer and the simulation time is 60 minutes.

A. Response Time for Different Transport and MAC Layer Protocols

We evaluate the web access performance for different transport layer protocols and MAC layer protocols by using the realistic web user behavior model described in section III. The results are shown in figure 5. For the transport layer protocols, RWBP always outperforms TCP when used for the satellite connection because RWBP doesn't need to go through the slow start phase and it eliminates congestion losses inside the satellite network. For the MAC layer protocols, fixed assignment gives the worse performance therefore it is not suitable for bursty reverse channel traffic. The aloha periodic stream scheme achieves some statistical multiplexing by only holding the channel for a terminal within a timeout period. However from Figure 4 we can see, during the timeout period if there are no packets arriving at the VSAT, the assigned channel bandwidth to the VSAT is wasted. Actually the period during which the channel release message propagates from NOC to the VSAT cannot be used by any VSAT either. While in CFDAMA, the NOC can assign the free bandwidth to those terminals with more traffic to send and achieve higher efficiency and smaller delay so CFDAMA performs better than the other two MAC protocols. Figure 5 shows RWBP combined with CFDAMA achieves the best web performance.

TABLE I

REVERSE CHANNEL TRAFFIC FOR EACH WEB PAGE WITH DIFFERENT ACKNOWLEDGEMENT FREQUENCY (THE AVERAGE NUMBER OF HTTP GETS PER PAGE IS 5.55 AND THE AVERAGE SIZE OF HTTP GETS IS 390 BYTES)

N	Num of ACKs	Total Traffic (bytes)	Normalized Total Traffic
2	61	$61*40 + 5.55*390 = 4604.5$	1
4	31	$31*40 + 5.55*390 = 3404.5$	0.74
8	16	$16*40 + 5.55*390 = 2804.5$	0.61
16	8	$8*40 + 5.55*390 = 2484.5$	0.54




Fig. 6. Improve response time performance by reducing ACK traffic. One ACK is sent every N data packets are received. CFDAMA is used as the MAC protocol and RWBP is used as the transport layer protocol.

B. Improve Response Time Performance by Reducing Reverse Channel Traffic

Because RWBP does not use acknowledgements to clock out data packets like in TCP, less frequent acknowledgements are needed in the reverse channel. In RWBP, an acknowledgement is sent when every N data packets are received. By increasing N, we can decrease the acknowledgement frequency. According to the web traffic model in section III, each page has 5.55 objects on average and the average size of each object is 11KB, i.e. each page is about 61KB. Because the data packet size is 512 bytes, each page contains 122 data packets and $122/N$ acknowledgements are needed for each page. One HTTP get is needed for each object so totally 5.55 HTTP gets are needed for each page. The average size of HTTP get packets is 390 bytes and the size of each acknowledgement packet is 40 bytes. Table I show the reverse channel traffic load for different acknowledgement frequency. By increase N from two to eight or sixteen, the reverse channel traffic can be reduced by about forty to fifty percent. When the acknowledgement traffic is reduced, smaller delay can be achieved for both ACKs and HTTP requests which leads to improved response time performance. Figure 6 shows the web page response time is reduced with N increased i.e. ACK traffic reduced. The MAC layer protocol used in this

experiment is CFDAMA.

VII. CONCLUSION

We evaluate web performance over a multiple access satellite channel with realistic user behavior model. For the MAC layer protocols, CFDAMA performs better than aloha/periodic stream and fixed assignment TDMA. For the transport layer, we design a new protocol called RWBP which does not have slow start and eliminates congestion losses inside satellite networks. We compare its performance with TCP over the three multiple access protocols. Our results show that RWBP can reduce the reverse channel traffic load and at the same time improve the web page response time when driven by realistic web traffic. RWBP combined with CFDAMA gives the best performance among all the combinations.

REFERENCES

- [1] I.F. Akyildiz, G. Morabito, and S. Palazzo. TCP Peach: A new congestion control scheme for satellite IP networks. *IEEE/ACM Trans. on Networking*, 9(3), June 2001.
- [2] P. Badford and M. Crovella. Generating representative Web workloads for network and server performance evaluation. In *ACM SIGMETRICS*, 1998.
- [3] A. Bakre and B. R. Badrinath. Implementation and performance evaluation of indirect TCP. *IEEE Transactions on Computers*, 46(3), March 1997.
- [4] Kongling Chang. IP layer per-flow queuing and credit flow control. Technical report, Ph.D. Thesis Harvard University, Division of engineering and applied science, January 1998.
- [5] Hyoung-Kee Choi and John O. Limb. A Behavior Model of a Web Traffic. In *International Conference of Networking Protocol '99*, September 1999.
- [6] D. P. Connors and G. J. Pottie. Response Initiated Multiple Access (RIMA), a Medium Access Control Protocol for Satellite Channels. In *IEEE GLOBECOM'00*, 2000.
- [7] Hyoung-Kee Choi et al. Interactive web service via satellite to home. *IEEE communication magazine*, March 2001.
- [8] T. R. Henderson and R. H. Katz. Transport protocols for Internet-compatible satellite networks. *IEEE J. Select. Areas Comm.*, 17:326–344, February 1999.
- [9] Tho Le-Ngoc and I. Mohammed Jahangir. Performance Analysis of CFDAMA-PB Protocol for Packet Satellite Communications. *IEEE Transactions on Communications*, 46(9), September 1998.
- [10] Alberto Leon-Garcia and Indra Widjaja. *Communication Networks: Fundamental Concepts and Key Architectures*. McGraw-Hill, 1999.
- [11] Bruce A. Mah. An Empirical Model of HTTP Network Traffic. In *IEEE INFOCOM*, 1997.
- [12] V. Mhatre and C. Rosenberg. Performance improvement of TCP-based applications in a multi-access satellite system. In *IEEE VTC'02 Fall*, September 2002.
- [13] V. Padmanabhan and R. Katz. TCP fast start: A technique for speeding up web transfers. In *IEEE GLOBECOM'98 Internet Mini-Conf.*, 1998.
- [14] F. D. Smith, F. Hernandez, K. Jeffay, and D. Ott. What TCP/IP Protocol Headers Can Tell Us About the Web. In *ACM SIGMETRICS*, 2001.
- [15] Hughes Network System. DirecPC web page. In <http://www.direcpc.com>, December 2002.
- [16] Xiaoming Zhou and J. S. Baras. Congestion management in access networks with long propagation delay. Technical report, <http://www.isr.umd.edu/CSHCN>, July 2003.