

Evaluation of SCTP Multistreaming over Wireless/Satellite Links

Mohammed Atiquzzaman
School of Computer Science
University of Oklahoma,
Norman, OK 73019-6151.

William Ivancic
Satellite Networks & Architectures Branch
NASA Glenn Research Center
21000 Brookpark Rd. MS 54-8,
Cleveland, OH 44135.

Abstract—In this paper, we study the impact of multistreaming on the performance of SCTP over satellite networks. We first show that multistreaming results in higher goodput than single streams when the receiver buffer is constrained as in the case of wireless handheld devices. We then demonstrate that the multistreaming feature of SCTP results in reduced buffer requirements at the receiver in the presence of losses in the satellite network. The above advantages makes SCTP an attractive transport protocol for wireless handheld devices.

I. INTRODUCTION

The Stream Control Transmission Protocol (SCTP) [1] is a new end to end message-based reliable transport protocol. Although initially conceived as a reliable transport protocol to carry PSTN signaling messages over IP, it has emerged into a highly reliable and fault-tolerant transport protocol which has the potential to replace TCP in the future.

Multistreaming is one of the most powerful features of SCTP. It allows data from a number of upper layer applications to be multiplexed onto one channel (called association in SCTP) as shown in Fig. 1. Sequencing of data is done within a stream; if a packet belonging to a certain stream is lost, packets (from that stream) following the lost one will be stored in the receiver's stream buffer until the lost packet is retransmitted from the source. Data from other streams can still be passed to upper layer applications. This avoids the head of line blocking found in TCP, where only one stream carries data from all the different upper layer applications.

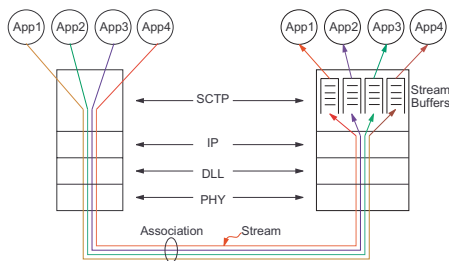


Fig. 1. An SCTP association consisting of four streams carrying data from four upper layer applications.

The HTTP protocol uses TCP as the transport mechanism. The early HTTP protocol used a separate TCP connection to fetch each URL, increasing the load on HTTP servers and causing congestion on the Internet [2]. Future HTTP protocols could use the multistreaming feature of SCTP to allow fast

downloading of web pages by simultaneously downloading different objects over different streams belonging to the *same association*.

The amount of data that can be sent by a TCP (or SCTP) sender depends on the congestion window and the receiver buffer size. When a packet is lost in the network, subsequent packets are queued up in the receiver buffer for resequencing until the lost packet is retransmitted from the source and arrives at the receiver. In the case of limited receiver buffer size, the receiver could run out of buffer space due to head of line blocking resulting in the sender being unable to send data. In the case of multistreaming, some of the streams could be delivering packets to the upper layer when a particular stream is blocked while waiting for a lost packet to arrive. Multistreaming could therefore result in higher throughput and reduced buffer requirements than single streaming.

The *objective* of this paper is to investigate the effect of multistreaming on the goodput of SCTP, and to determine the buffer requirements at a receiver which is connected using a satellite link. The *contributions* of this work can be summarized as:

- Determine the goodput of SCTP over satellite networks for various link error rates and receiver buffer sizes.
- Demonstrate that multistreaming can significantly reduce the receiver buffer size for wireless handheld devices.
- Determine the optimal buffer requirement at the receiver.

The rest of the paper is organized as follows. SCTP congestion control is discussed in Sec. II, followed by definitions and notations used in this paper in Sec. III. The simulation topology and assumptions are described in Sec. IV, and results are presented in Sec. V. Finally, concluding remarks are presented in Sec. VI.

II. CONGESTION CONTROL OF SCTP

SCTP's congestion control [3] is based on the well proven rate-adaptive window-based congestion control scheme of TCP, viz. Slow Start, Congestion Avoidance, Fast Retransmit and Fast Recovery [4], [5] and Selective Acknowledgement [6], [7]. This ensures fairness for both protocols as they work together in the Internet. SCTP provides reliable transmission, and detects lost, reordered, duplicated or corrupt packets. It provides reliability by retransmitting lost or corrupt packets, and implements flow control by the receiver's advertised window size.

III. DEFINITIONS AND NOTATIONS

We define the following terms.

- *Goodput* is defined as the total number of packets (without considering the retransmitted packets) reaching the destination during the simulation period.

This work was supported by NASA grant NAG3-2528 and NAG3-2922.




Fig. 2. Network simulation topology.

- *Optimal receiver buffer size* is the smallest amount of receiver buffer for which the SCTP goodput is independent of the receiver buffer size.

We use the following notations in this paper.

- s = number of streams per SCTP association.
- r_i = transmission rate of link i in Mbps.
- d_i = propagation delay of link i in msec.
- ϵ_i = error rate of link i . It is the probability that a packet is lost in the link due to errors.
- B = Receiver buffer size in bytes.
- a_rwnd = Advertised receiver window size in bytes.
- $cwnd$ = Congestion window size in bytes.

IV. SIMULATION SETUP AND ASSUMPTIONS

To evaluate the performance of SCTP multistreaming over a satellite link, the simulations were carried out using the *ns* simulator [8] with an SCTP patch from the University of Delaware. The network topology for the simulation is shown in Figure 2, where an SCTP source sends one way traffic to an SCTP sink through a router, possibly located in the satellite. $L_i, 1 \leq i \leq 2$ are links whose transmission rates, propagation delays and random error rates are expressed by the tuple (r_i, d_i, ϵ_i) . In our simulation, we assume that L1 and L2 are terrestrial and satellite links, respectively. We make the following assumptions regarding the simulation setup:

- Data transfers are long. The SCTP source has an infinite supply of data which is being transferred to the destination using *ftp*.
- The terrestrial link is error free, i.e. $\epsilon_1 = 0$.
- Packets are of fixed length equal to one MTU.
- The upper layer at the destination is always ready to accept data.

V. RESULTS

We first present the results for the performance of SCTP, both in the absence and presence of errors. We then show the improvement of SCTP's goodput due to multistreaming, for various error rates and receiver buffer sizes. Finally, we determine the optimal receiver buffer size for different error rates and number of streams. Throughout this study, we have used $r_1 = 5, r_2 = 10, d_1 = d_2 = 130$, and $\epsilon_1 = 0$. The values of r_1 and r_2 were chosen to avoid any congestion loss in the network, thereby leaving all possible losses due to satellite link errors.

A. Effect of Zero Errors

Fig. 3 shows SCTP goodput as a function of the receiver buffer size for $s = 4$ and $\epsilon_2 = 0$ representing zero packet loss due to errors. In the absence of packet losses, there is no blocking at the destination. $cwnd$ initially increases and then becomes constant when it reaches B . After this point, the source sends a burst of $\lfloor \frac{B}{MTU} \rfloor$ packets every RTT as seen in Fig. 4 which plots the TSNs of the packets leaving and Cumulative TSN of the ACKs arriving at the source as a function of time.




Fig. 3. Goodput as a function of receiver buffer size for $\epsilon_2 = 0$ and $s = 4$.

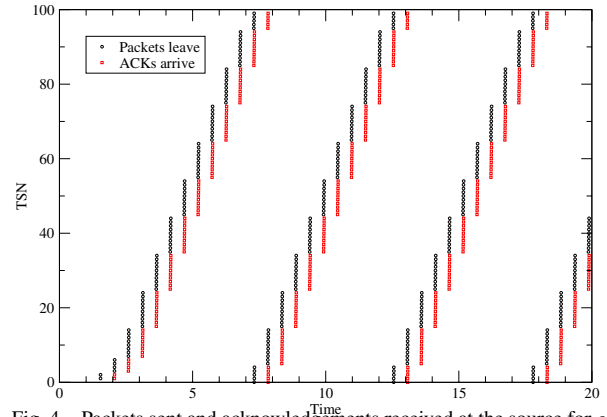


Fig. 4. Packets sent and acknowledgements received at the source for $\epsilon_2 = 0, s = 4$ and $B = 15K$.

The goodput therefore, depends directly on the receiver buffer size as seen in Fig. 3. The y-axis shows $(TSN \bmod 100)$.

In the absence of packet losses, there is no blocking at the destination buffer. The goodput is therefore independent of the number of streams, i.e. although Fig. 3 is shown for $s = 4$, it applies for any value of s .

B. Effect of Errors

Fig. 5 shows the goodput as a function of the receiver buffer size for $s = 4$ and $\epsilon_2 = 0.01, 0.03$ and 0.05 . For a particular error rate, the goodput initially increases as B increases indicating that the goodput is constrained by the receiver buffer size. This is evident from Figure 6 which shows $cwnd$ and a_rwnd for $s = 4, \epsilon_2 = 0.01$ and $B = 15K$. Receiver buffer size of $15K$ was chosen to make the goodput of the connection constrained by the receiver buffer size as seen by frequent dropping of a_rwnd below one MTU which in turn restricts the increase of $cwnd$ beyond $15K$.

As B increases (Fig. 5), a point is reached after which any further increase of B does not have any effect on the goodput; at this point, the goodput is limited by the congestion control mechanism of SCTP invoked by packet losses due to link errors. Further increase in goodput can only be achieved by lowering the link error rate as seen in Fig. 5. The fact that the goodput is limited by $cwnd$ for large values of B is evident from Fig. 7 which shows $cwnd$ and a_rwnd for $s = 4, \epsilon_2 = 0.01$ and $B = 35K$. a_rwnd falls below one MTU only once, while at other times, the decrease of $cwnd$ (which is a measure of goodput) is governed by SCTP's congestion control




Fig. 5. Goodput as function of receiver buffer size for different error rates and $s = 4$.




Fig. 7. Advertised receiver window and congestion window with $B = 35K$, $\epsilon_2 = 0.01$, and $s = 4$.




Fig. 6. Advertised receiver window and congestion window with $B = 15K$, $\epsilon_2 = 0.01$, and $s = 4$.




Fig. 8. Packets sent and acknowledgements received at the source for $\epsilon_2 = 0.01$, $s = 4$ and $B = 35K$.

mechanisms (slow start, congestion avoidance, fast retransmit, etc.) after packet drops. Our observation that the goodput at the flat region of the curves in Fig. 5 is limited by the congestion control mechanism of SCTP is validated by the fact that a_rwnd never fell below the MTU in the flat region of the goodput curve; i.e. the receiver buffer size was not a limiting factor in the goodput.

Fig. 8 shows packets sent and ACKs received at the source. We can see long delays in retransmitting lost packets while waiting for DUP ACKs resulting from the lost packets. These long delays and the drop of $cwnd$ results in reduced goodput when receiver buffer is not a constraint.

C. Effect of Multistreaming

We have shown in the previous sections that the receiver buffer size could be a limiting factor in the case of small buffers. In this section, we show how multistreaming *reduces the receiver buffer requirements and increases the goodput* by avoiding HOL blocking at the receiver.

Figure 9 shows the goodput of an SCTP connection using one and four streams for error rates of 1% and 5% in the bottleneck link. The goodput of SCTP with four streams per connection is better than one stream per connection for small receiver buffer sizes. This is because, for small receiver buffer sizes and a single-stream per association, the sender is often prevented from sending packets because of lack of receiver buffer space arising due to HOL blocking. However, in the case of multiple streams per association, some of the streams may be able to pass packets to the upper layers even though other streams are

blocked due to lost packets belonging to those streams. Fewer packets are dropped when the error rate is low. This results in higher goodput due to a smaller chance of blocking at the receiver.

Figure 10 compares the goodput of SCTP for one and four streams as a function of the error rate for $B = 15K$. The higher goodput of multistreaming is shown for various error rates. As an example, four streams per association results in about 10% increase in throughput over one stream at 3% error rate. A small receiver buffer size (15K) was chosen to demonstrate the *advantage of multistreaming when the buffer size is limited*.

The advertised receiver window for one and four streams is shown in Fig. 11 for $\epsilon_2 = 5\%$ and $B = 15K$. It is seen that, for a single stream, a_rwnd frequently falls below 1500 bytes (one MTU) thereby restricting the sender from sending packets. In the case of four streams, the advertised receiver buffer size falling below one MTU is less frequent than the case of a single stream. This results in higher goodput for the case of multistreaming. Note that one MTU is the minimum size required by the sender to send data.

D. Optimal Receiver Buffer Size

It has been shown in Sec. V-B that as the receiver buffer size is increased, the goodput of SCTP in the presence of errors




Fig. 9. Goodput as a function of receiver buffer size with one and four streams for various error rates.




Fig. 10. Goodput as function of error for one and four streams for $B = 15K$.

(see Fig. 5) becomes independent of the receiver buffer size. In this section, we determine the *optimal receiver buffer size* (see Sec. III for definition) for various error rates. Receiver buffer size larger than the optimal receiver buffer size does not contribute to increasing the goodput, and hence is wasted. The optimal receiver buffer size can be used to dimension the buffer size at the receiver for portable mobile devices (such as PDAs and next generation wireless phones) where it is desirable to reduce the memory size because of weight and power restrictions.

Figure 12 shows the optimal receiver buffer size as a function of the error rate for one and four streams. The buffer size required for multistreaming is significantly lower than for single streams. For example, four streams result in a saving of about 35% of receiver buffer space over a single stream for 1% error rate. Similar saving is also evident for other error rates. The optimal receiver buffer size was determined as the receiver buffer size for which the goodput becomes constant (i.e. independent of the receiver buffer size) in Fig. 5.

Figure 12 also shows that the optimal receiver buffer size is smaller for higher error rates. This is consistent with our observation in Sec. V-B that *as the error rate increases, the goodput is increasingly dominated by the congestion control mechanism, and to a lesser extent by the receiver buffer size.*

VI. CONCLUSIONS

We have shown that multistreaming improves the goodput of SCTP for the case of limited receiver buffer size in the presence of errors in satellite links. With the proliferation of handheld devices operating over wireless (or satellite/radio) links having




Fig. 11. Advertised receiver window for one and four streams. $\epsilon_2 = 0.05$, $B = 15K$.




Fig. 12. Optimal receiver buffer size vs. error rate for one and four streams.

limited receiver buffer capacity, SCTP will perform better than TCP (which uses a single stream) for such devices. We have also demonstrated that multistreaming can reduce the buffer requirements at the receiver, and hence would be attractive for wireless handheld devices where the reduction of weight and power requirements is of utmost importance.

REFERENCES

- [1] R. Stewart, Q. Xie, K. Morneault, and C. Sharp et. al., "Stream control transmission protocol." RFC 2960, Oct 2000.
- [2] Joe Touch, John Heidemann, and Katia Obraczka, "Analysis of HTTP performance." <http://www.isi.edu/lam/publications/http-perf/index.html>, August 1996.
- [3] R. Brennan and T. Curran, "SCTP congestion control: Initial simulation studies," *International Teletraffic Congress (ITC 17)*, Brazil, 2001.
- [4] M. Allman, V. Paxson, and W. Stevens, "TCP congestion control." RFC 2581, April 1999.
- [5] R. Alangir, M. Atiqzaman, and W. Ivancic, "Impact of retransmission mechanisms on the performance of SCTP and TCP in a satellite network," *NASA Earth Science Technology Conference*, Pasadena, CA, June 2002.
- [6] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow, "TCP selective acknowledgement options." RFC 2018, Oct 1996.
- [7] S. Fu and M. Atiqzaman, "Improving end-to-end throughput of mobile IP using SCTP," *2003 Workshop on High Performance Switching and Routing*, Torino, Italy, pp. 171–176, June 24–28, 2003.
- [8] "ns-2 network simulator." www.isi.edu/nsnam/ns/.