

SCTP over Satellite Networks

Shaojian Fu

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—The Stream Control Transmission Protocol (SCTP) has recently been standardized as a new transport layer protocol in the IP protocol suite. In addition to the core features of TCP, SCTP incorporates a number of advanced and unique features which are not available in TCP. The objective of this paper is to investigate the suitability of SCTP for data communications over satellite links. We describe SCTP features that allow SCTP to better utilize the bandwidth of satellite networks, while at the same time avoiding congestion collapse in a shared network. Finally, we provide recommendations on the use of SCTP over satellite networks.

Keywords: Stream Control Transmission Protocol, Satellite networks, Transport protocols, Next Generation Networks.

I. INTRODUCTION

Recent interest in transmitting voice over IP networks [1] has led IETF to develop a new transport layer protocol, called Stream Control Transmission Protocol (SCTP) [2], for the IP protocol suite. Although, the initial aim of SCTP was to provide a robust protocol for the transport of signalling messages over an IP network, later developments have made it also useful for a wider range of applications, resulting in moving the standardization work of SCTP from SIGTRAN to the Transport Area Working Group (TSVWG) of IETF in February 2001.

SCTP is a reliable network-friendly transport protocol which can co-exist with TCP in the same network. The design of SCTP absorbed many strengths and features of TCP (such as window based congestion control, error detection and retransmission, etc.), that made TCP a success during the explosive growth of the Internet. Moreover, SCTP incorporated several unique features, such as multistreaming and multihoming (discussed in Sec. III), that are not available in TCP.

Satellite links are an indispensable part of the global Internet to provide broadband data, television, telephony, and navigation services, etc. Although, TCP is the dominant transport protocol in the IP protocol suite, it was not initially designed for long bandwidth product networks, such as satellite networks which are characterized by long propagation delays and corruption losses due to wireless links. Consequently, a number of enhancements to TCP have been proposed to enhance its performance over satellite networks [3], [4], [5]. Although, a number of those TCP enhancements may have been incorporated in SCTP, the implementation of some of them are different from the way they are implemented in TCP. Some of the features

which are unique to SCTP can also help SCTP to achieve a better performance than TCP in satellite environments.

Despite the extensive research on TCP performance and enhancements over satellite links, *the authors are not aware of any such study to investigate the suitability of SCTP for data transmission over satellite networks.* The goal of this paper is to evaluate and highlight those SCTP features that make it suitable for satellite networks. Specifically, we divide the evaluation into two parts (Secs. IV and V):

- SCTP features which *currently exist* in TCP or its enhancements for satellite networks;
- *Unique* SCTP features that help SCTP to achieve a high throughput in satellite networks.

There has been work done in the last couple of years in evaluating the performance of many aspects of SCTP. For example, the co-existence of SCTP and TCP in the Internet has been studied by Jungmaier et.al. [6]. Use of multistreaming and multihoming to reduce latency and fault tolerance of data transmission in highly lossy environments has been reported by Jungmaier et.al. [6] and Conrad et.al. [7]. The effect of SCTP multihoming was investigated in high-availability environments to achieve fast recovery over fault conditions [8]. In the wireless networking area, the performance of SCTP in Mobile network [9], [10] and wireless multi-hop networks [11] have been studied. This paper *differs* from previous work in the sense that it investigates and evaluates the SCTP features that can be exploited to increase SCTP's performance over satellite networks, while at the same time using its advanced congestion control algorithms (Sec. III-A) to prevent congestion collapse in the Internet. The results and recommendations provided in this paper can be used to increase SCTP throughput over satellite networks.

The *objective* of this paper is to highlight and evaluate the suitability of SCTP for satellite networks, and make recommendations regarding the use of its features for enhancing the transport layer performance over satellite networks. Such recommendations could possibly be incorporated into the SCTP protocol, which is still in its early stages of development.

The *contributions* of this paper can be summarized as follows:

- Provide insights into the suitability of SCTP over satellite links;
- Highlight the different features of SCTP which will help it to achieve the performance of TCP with enhancements in satellite environments;
- Determine the effects of the unique features of SCTP in

TABLE I
SCTP SACK CHUNK FORMAT.

0	7	15	31
Type=3	Chunk Flags		Chunk length
Cumulative TSN Ack			
Advertised Receiver Window Credit			
Number of GapAck Block		Number of Dup TSN	
Gap Ack Block #1 Start		Gap Ack Block #1 End	
.....			
Gap Ack Block #N Start		Gap Ack Block #N End	
Duplicate TSN 1			
.....			
Duplicate TSN X			

improving its performance over satellite links;

- Provide recommendations on using SCTP over satellite networks.

The rest of the paper is organized as follows. In Sec. II, the characteristics of satellite links and their effects on the performance of transport layer protocols are described. The SCTP features that exist in TCP and some unique SCTP features are discussed in Sec. III. The suitability of SCTP for satellite communications is presented in Secs. IV and V. Recommendations on using SCTP over satellite networks, and our conclusions from this research are presented in Sec. VI.

II. EFFECTS OF SATELLITE LINK CHARACTERISTICS ON TRANSPORT PROTOCOLS

A number of satellite link characteristics, which are different from terrestrial links, may limit the performance of transport protocols over satellite networks [3]. The characteristics have similar effects on SCTP and TCP, since these two protocols use similar congestion control, retransmission, and round trip time estimation algorithms. The characteristics, described below, form the basis of various features that we will discuss later in Secs. IV and V.

- *Long propagation delay*: The propagation delay between an earth station and a Geostationary Earth Orbiting (GEO) satellite is around 120-140ms (milliseconds), which means it takes the sender long time to probe the network capacity and detect the possible loss of segments, and the expensive satellite bandwidth is wasted.
- *Large delay-bandwidth product*: The GEO satellite link is a typical case of the Long Fat Pipe (LFP), which features a large delay bandwidth product. For example, the DS1-speed GEO channel has a 96500-byte size pipe. The fundamental performance problems with the current TCP over LFN links were discussed in [12].
- *Corruption loss during transmission*: The large transmission distance of satellite links results in a low signal-to-noise ratio (SNR) and consequently a high Bit Error Rate (BER). The errors loss will cause the TCP and SCTP senders to reduce their transmission rates unnecessarily.

III. FEATURES OF SCTP

A number of TCP's built-in features and later enhancements have been recommended for use in satellite networks [3]. In this section, we first describe those TCP features which are also

available in SCTP, followed by unique features of SCTP that can benefit communication over satellite networks (Sec. III-B).

A. Features of SCTP that are common with TCP

The following features of TCP and its enhancements have been recommended for satellite networks. These features are, therefore, also helpful for SCTP over satellite networks. However, the implementation of some of those features in SCTP are different from TCP as described below.

- *Slow Start and Congestion Avoidance*: Like TCP, SCTP also uses Slow Start and Congestion Avoidance algorithms [2] to probe the available capacity of the network. These algorithms force the sender to wait for ACKs before sending new data in order to prevent congestion collapse. Given the long propagation delay of the satellite link, the channel bandwidth is not utilized efficiently when the sender is going through these algorithms.
- *Fast Retransmit and Fast Recovery*: Like TCP, SCTP also incorporates a Fast Retransmit algorithm based on Selective Acknowledgement (SACK) gap reports [13]. This mechanism speeds up loss detection in satellite links, and therefore, increases network resource utilization. One of the major differences between SCTP and TCP is that SCTP doesn't have an explicit Fast Recovery phase, but achieves this automatically with the use of SACK.
- *Path MTU discovery*: Like in TCP, Path MTU discovery provides SCTP with the information about the largest possible segment size that will not cause packet fragmentation at intermediate routers. However, SCTP has a slightly different support for path MTU discovery as compared to TCP, as will be discussed in Sec. IV-A.
- *SCTP SACK*: Unlike TCP, the use of SACK is mandatory in SCTP. In SCTP, all data are carried in a structure called "chunk" which is fully described by the *Chunk type*, *Chunk flags*, *Chunk length*, and *Chunk data* fields as shown in Table I for an SCTP SACK chunk [2]. For TCP, the length of the *Options* field is limited to 40 bytes. A SACK option consisting of n blocks will have a length of $8 \times n + 2$ bytes. Therefore, the maximum number of SACK gap blocks in TCP is limited to 4. If SACK is used together with the timestamp option (requires 12 bytes), the maximum number of blocks is reduced to 3. Compared to TCP, SCTP allows more gap blocks in its SACK chunk. The total available chunk space, as determined by the *Chunk Length* field (Table I) is 2^{16} bytes. Subtracting the space used by first 16 bytes, the maximum space available for gap blocks is $2^{16} - 16$, with each block requiring 4 bytes. Therefore, the maximum number of blocks allowed is 16380. The effect of this difference in the number of SCTP and TCP SACK blocks on satellite networks will be discussed in Sec. IV-C.

B. Unique features of SCTP that are not present in TCP

In this section, we describe those unique features of SCTP that are not available in TCP, but are useful in sending data over satellite networks. The effect of these unique features on data

transmission over satellite networks will be discussed in detail in Sec. V.

Like TCP, SCTP also fits in the transport layer of the Internet protocol stack. Fig. 1 shows an SCTP association, along with multihoming and multistreaming as described below.




Fig. 1. Schematic view of an SCTP association.

- **Multihoming:** Multihoming allows an association (in SCTP, “**association**” represents the communication relationship between endpoints, which is analogous to “**connection**” in TCP) between two endpoints to span across multiple IP addresses (or network interface cards). One of the addresses is designated as the primary, while the other one can be used as backup in the case of failure of the primary address, or when the upper layer application explicitly requests the use of the backup. For example, retransmission of lost packets can be done over the secondary address to increase the reliability of retransmitted packets.
- **Multistreaming:** Multistreaming allows data from an upper layer application to be split into multiple *streams* in an association as shown in Fig. 1. Sequencing of data is maintained within a stream; if a segment belonging to a certain stream is lost, segments (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. However, data from other streams can still be delivered to upper layer applications when they arrive at the destination. This avoids the head of line (HOL) blocking found in TCP [2].
- **Byte counting in acknowledgments:** RFC 2960 [2] recommends that an SCTP receiver should use delayed SACK in acknowledging user data. This requires an acknowledgment to be generated for every second segment received, or within 200 ms of the arrival of any unacknowledged segment. In SCTP, the *byte counting* algorithm increases the *cwnd* by the number of bytes acknowledged by the SACK. Byte counting decouples the increase of *cwnd* from the arrival frequency of the SACKs, and thus overcomes the problem of slow increase of *cwnd* when delayed SACK is used in long propagation delay networks. Note that, because TCP increases the congestion window (*cwnd*) by the number of acknowledgments received by the sender, delayed SACK in TCP increases the time required by the sender to increase the *cwnd* during Slow Start.

- **Explicit Congestion Notification (ECN):** SCTP defines the ECN capable TLV (the optional parameters in a SCTP chunk use the Type-Length-Value (TLV) format [2]) in both INIT and INIT-ACK chunks which are exchanged between endpoints during association setup. When an endpoint initiated a new association, it adds the ECN capable TLV in the INIT chunk. If the peer endpoint responds with the same TLV in the INIT-ACK chunk, ECN is enabled on the association. Once ECN enabled, detecting and responding to congestion in SCTP are almost similar to those defined in [14]. The difference is when the SCTP receiver detects the “*Congestion Experienced*” bit in the IP header of a received segment, it will use an Explicit Congestion Notification Echo (ECNE) chunk to notify the sender about the congestion, and the sender will respond with Congestion Window Reduce (CWR) indicating that the *cwnd* has been reduced.

IV. SUITABILITY OF SCTP FOR SATELLITE COMMUNICATIONS

SCTP is a relatively new transport protocol which is still being developed. It is extremely important to understand the suitability of SCTP for satellite communications, and make necessary improvement to the protocol while it is still in its early stages of development. In this section, we individually describe the features of SCTP that enhance its performance over satellite networks.

A. Support for Path MTU Discovery

SCTP supports Path MTU (PMTU) Discovery. However, its implementation is slightly different from TCP in that an SCTP association may span multiple IP addresses because of multihoming. Consequently, separate path MTU estimates must be maintained for each destination IP address. SCTP defines PMTU as the smallest path MTU discovered for all destination IP addresses. A large segment size can reduce the packet overhead, and enable the SCTP sender to increase the congestion window more rapidly in terms of bytes. PMTU discovery is, therefore, recommended for enabling transfer of large SCTP segments over satellite networks.

B. Congestion Control Mechanisms

Despite the difference between the congestion control mechanisms of SCTP and TCP, both of them have the same objective: to ensure that the sender throttles back under adverse network conditions to recover from congestion quickly. Even though congestion control mechanisms may have a negative effect on the throughput of a SCTP and TCP (see Sec. III-A), the mechanisms are necessary to prevent congestion collapse in a shared network, such as the Internet. These mechanisms are, therefore, recommended for use in SCTP running over satellite networks.

C. SCTP Selective Acknowledgment

Use of SACK allows more robust reaction in the case of multiple losses from a single window of data [13]. This avoids a

time-consuming slow start stage after multiple segment losses in a satellite environment, and thus saves network bandwidth.

Since satellite links feature high BER, and require a large transmission window to utilize the satellite network bandwidth (see Sec. IV-D), there is a higher probability of multiple non-consecutive segment losses in a single window. The number of available Gap Blocks of 3 or 4 in TCP (see Sec. III-A) may not be sufficient for reporting all the segment losses. If all the losses in a single window cannot be reported in a single SACK, the sender has to wait longer to determine all the lost segments. As discussed in Sec. III-A, SCTP allows more Gap Blocks, thereby rendering it more robust in the case of multiple losses in a window of data.

D. Large Receiver Window Support

The length of the *Window* field in the TCP header is only 16 bits, resulting in a maximum window size of 65535 bytes. Because a DS1-speed GEO satellite channel has a 96500-byte size pipe (see Sec. II), TCP cannot fully utilize the channel bandwidth. As a result, the window scaling option was proposed [12] to extend the TCP usable window size to 65535×2^{14} bytes, and has been recommended for use in satellite communication [3].

The *Advertised Receiver Window Credit* field in the SCTP SACK header (Table I) has a length of 32 bits. It thus enables a usable receiver window of up to 2^{32} bytes, compared to 65535×2^{14} bytes in TCP with the Window Scaling option. This inherent large window size of SCTP should be enough for most of satellite environments. The implicit support of large receiver window size in SCTP makes it suitable for satellite networks.

V. EXPLOITING UNIQUE SCTP FEATURES IN SATELLITE NETWORKS

In Sec. IV, we have described the various features of SCTP that can be used to satisfy the requirements of satellite communications, as proposed over the years for TCP over satellite networks. In this section, we describe the *unique and novel features* of SCTP that can be exploited to increase its performance over satellite networks.

A. Use of Multihoming

This built-in support of SCTP for multihomed endpoints can increase the reliability of high-availability applications by transparently switching over data communication to the secondary link when the primary link fails (see Sec. III-B). An example of SCTP multihoming is shown in Fig. 2, where the two endpoints are connected by two satellite links via satellite1 and satellite2. One of the links is designated as the primary, while the other one can be used as backup in the case of failure of the primary address, or blackouts periods when the primary satellite is cut out from communication due to shadowing, satellite handovers, etc. The backup link can also be used when the upper layer application explicitly requests the use of the backup link, or for load balancing among satellites. Multihoming can thus make a satellite network highly reliable and fault tolerant.




Fig. 2. An SCTP association with multihomed endpoints.




Fig. 3. The effect of multistreaming on goodput.

B. Effect of Multistreaming

Multistreaming of SCTP (see Sec. III-B) can be used to alleviate the head-of-line (HOL) blocking effect resulting from TCP's strict byte-order delivery policy. Each stream is kind of a "sub-flow" within the overall data flow, where the delivery of packets in a sub-flow is independent of other sub-flows. We have demonstrated in [15] that, under error-prone satellite link conditions, multistreaming can significantly reduce the receiver buffer size requirements and increase channel goodput when the receiver's buffer is limited. This effect is illustrated in Fig. 3, where s is the number of streams and ϵ is packet error rate. It is seen that for small receiver buffer sizes, multistreaming can increase the SCTP goodput by eliminating HOL blocking in an error-prone satellite environment.

C. Byte counting in delayed Acknowledgment

SCTP limits the congestion window ($cwnd$) increase to one PMTU per SACK, we call this Byte Counting Limit (BCL). When the total number of bytes acknowledged by a single SACK exceeds PMTU, the benefit of byte counting is impaired. This effect can be illustrated by Fig. 4, where the PMTU is 1500 bytes. For the SCTP association in Fig. 4-a, the SACK chunk acknowledges 1072 bytes (less than PMTU), so the $cwnd$ is increased by two segments. As a comparison, in Fig. 4-b, the SACK chunk acknowledges 3000 bytes, but the $cwnd$ can only be increased by one segment. Consequently, we recommend increasing the BCL to 2 PMTU to speedup the slow start phase when delayed SACK is used.

D. Large Initial Congestion Window

As discussed in Sec. V-C, delayed SACK is recommended for use in SCTP. If the initial congestion window is 1 segment,

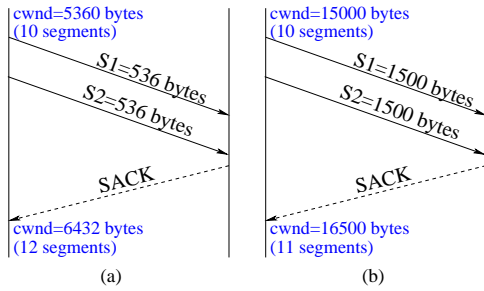


Fig. 4. SCTP byte counting with BCL=1 and 2 respectively.

TABLE II

SUMMARY OF RECOMMENDATIONS FOR SCTP FEATURES WHICH ARE ALSO AVAILABLE IN TCP.

Mechanism	Use	Where
Path MTU Discovery	Recommended	S
Slow Start	Required	S
Congestion Avoidance	Required	S
Fast Retransmit	Recommended	S
Fast Recovery	Implicitly Used	S
SACK	Implicitly Used	S, R
Delayed SACK	Recommended	R
Large Receiver Window	Implicitly Used	S, R

the receiver must wait for the 200ms timer to expire before acknowledging the first received segment. Because the SCTP RFC requires the initial $cwnd \leq 2$ segments, we recommend using 2 segments as the initial value of $cwnd$ for SCTP over satellite links. This will also decrease the time required for the slow start phase by one RTT.

E. Explicit Congestion Notification (ECN)

Due to the relatively high BER of satellite links, determining the exact reason (congestion vs corruption losses) of segment losses can prevent the sender from unnecessarily entering congestion control and thus improve SCTP's throughput. ECN provides a framework that enables the network routers to notify the state of congestion to the endpoints [14]. This mechanism is not a complete solution to the above problem, but helps in increasing the throughput. Due to SCTP's explicit support for ECN (see Sec. III-B), the sender can utilize the feedback from receiver to differentiate corruption losses from congestion drops. When it's determined that the segment loss is due to the corruption losses during transmission, the sender can avoid unnecessary reductions of the congestion window, which is especially useful in long delay satellite networks.

VI. RECOMMENDATIONS AND CONCLUSION

Our recommendations for using SCTP over satellite networks are summarized in Tables II and III. In these tables, the last column denotes network point to implement the mechanism: "S" means the sender, "R" means the receiver, and "S, R" means both sender and receiver. Table II summarizes the recommendations for SCTP features which are also available in TCP, while Table III provides recommendations for the use of the unique features of SCTP, i.e. those that are not available in TCP.

TABLE III

SUMMARY OF RECOMMENDATIONS FOR FEATURES UNIQUE TO SCTP.

Mechanism	Use	Where
SCTP Multihoming	Recommended	S, R
SCTP Multistreaming	Recommended	S, R
Byte Counting	Implicitly Used	S, R
Larger Byte Counting limit	Recommended	S
Larger Initial $cwnd$	Recommended	S
ECN	Recommended	S, R

In this paper, we first outlined satellite link characteristics which may limit the performance of transport protocols, followed by mechanisms that can help SCTP to better utilize the bandwidth of satellite environments, while preventing congestion collapse in a shared network. *The authors believe that all these mechanisms should make SCTP very suitable as a transport protocol for satellite networks.*

There are a number of unresolved issues for both TCP and SCTP over satellite links, such as SCTP/IP Header Compression in high BER environment, bias against long-RTT associations during congestion avoidance, and the interaction between SCTP retransmissions and link layer ARQ. They require further research to improve the performance of SCTP over satellite links. Moreover, TCP enhancements, such as Protecting Against Wrapped Sequence (PAWS) numbers and Round Trip Time Measurement (RTTM) [12], require the timestamp option [12] which is not available in SCTP. In order to use these mechanisms in SCTP, new chunk type for timestamp should be considered in future developments of SCTP.

REFERENCES

- [1] M. Hassan, A. Nayandoro, and M. Atiquzzaman, "Internet telephony: services, technical challenges, and products," *IEEE Communications Magazine*, vol. 38, no. 4, pp. 96–103, April 2000.
- [2] R. Stewart and Q. Xie et. al., "Stream control transmission protocol." IETF RFC 2960, October 2000.
- [3] M. Allman, D. Glover, and L. Sanchez, "Enhancing TCP over satellite channels using standard mechanisms." IETF RFC 2488, January 1999.
- [4] N. Ghani and S. Dixit, "TCP/IP Enhancements for Satellite Networks," *IEEE Communication Magazine*, vol. 37, no. 7, pp. 64–72, July 1999.
- [5] C. Partridge and T.J. Shepard, "TCP/IP performance over satellite links," *IEEE Network*, vol. 11, no. 5, pp. 44–49, Sep/Oct 1997.
- [6] A. Jungmaier, "Performance evaluation of the stream control transmission protocol," *Proceedings of the IEEE Conference 2000 on High Performance Switching and Routing*, Heidelberg, Germany, pp. 141–148, June 2000.
- [7] P.T. Conrad, G.J. Heinz, and A.L. Caro et. al., "SCTP in battlefield networks," *IEEE Military Communications Conference (MILCOM)*, McLean, VA, pp. 289–295, October 2001.
- [8] A. Jungmaier, E.P. Rathgeb, and M. Tuxen, "On the use of SCTP in failover-scenarios," *International Conference on Information Systems, Analysis and Synthesis*, Orlando, Florida, pp. 363–368, July 2002.
- [9] S. Fu, M. Atiquzzaman, and W. Ivancic, "Effect of delay spike on SCTP, TCP Reno, and Eifel in a wireless mobile environment," *11th International Conference on Computer Communications and Networks*, Miami, FL, pp. 575–578, October 2002.
- [10] J. Noonan, P. Perry, and J. Murphy, "A study of SCTP services in a Mobile IP network," *IT&T Annual Conference*, WIT, Ireland, October 2002.
- [11] G. Ye, T. Saadawi, and M. Lee, "SCTP congestion control performance in wireless multi-hop networks," *MILCOM2002*, Anaheim, California, pp. 934–939, October 2002.
- [12] V. Jacobson, R. Braden, and D. Borman, "TCP extensions for high performance." IETF RFC 1323, May 1992.
- [13] K. Fall and S. Floyd, "Simulation-based Comparisons of Tahoe, Reno, and SACK TCP," *ACM Computer Communications Review*, vol. 26, no. 3, pp. 5–21, July 1996.
- [14] K. Ramakrishnan and S. Floyd, "A Proposal to add Explicit Congestion Notification (ECN) to IP." IETF RFC 2481, January 1999.

- [15] M. Atiquzzaman and W. Ivancic, "Evaluation of SCTP multistreaming over wireless/satellite links," *12th International Conference on Computer Communications and Networks*, Dallas, Texas, pp. 591–594, October 2003.