

Evaluation of SCTP for Space Networks

**Shaojian Fu, Mohammed Atiquzzaman,
William Ivancic**

TR-OU-TNRL-05-101

January 2005



Telecommunication & Network Research Lab

School of Computer Science

THE UNIVERSITY OF OKLAHOMA

200 Felgar Street, Room 159, Norman, Oklahoma 73019-6151
(405)-325-4042, atiq@ou.edu, www.cs.ou.edu/~atiq

Evaluation of SCTP for Space 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. SCTP is based on the TCP protocol, but incorporates a number of advanced and unique features which are not available in TCP. Although, the suitability of TCP over satellite networks has been widely studied, the suitability of SCTP over satellite networks remains to be evaluated. The objective of this paper is to investigate the suitability of SCTP for data communications over satellite networks. We focus on the advanced features of SCTP that enhance its performance in satellite networks. 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

Satellite networks have a large coverage area, and currently provide television, radio, telephony, and navigation services. Satellites are expected to play a significant role in the future global Internet to provide broadband data services. Currently, two types of satellites, viz. Geostationary Earth Orbit (GEO) and Low Earth Orbit (LEO) satellites, are mostly used for the above applications. Traditionally, GEO satellites have been used to provide a *bent-pipe* transmission channel, where all packets received on an up-link are transparently piped to the corresponding down-link, i.e. a GEO satellite is merely a physical layer repeater in space, which is invisible to the routing protocols. To increase system capacity and reduce end-to-end delay, newer satellites are increasingly adopting a regenerative paradigm where the satellites have on-board switching and routing units [1]. This is also consistent with the current efforts of the National Aeronautics and Space Administration [2] and the European Space Agency [3].

The work reported in this paper was funded by National Aeronautics and Space Administration (NASA) grant nos. NAG3-2528 and NAG 3-2922.

The long propagation delay of GEO satellites make them less desirable for real-time applications, such as voice communications. The concept of LEO satellite constellation was introduced in 1990s to provide satellite services at a lower orbit by utilizing a large number of satellites than a GEO constellation. The advantages of LEO over GEO include lower link propagation delay, reduced free space attenuation, lower power consumption for user terminals, and higher spectrum efficiency due to frequency reuse. However, these advantages come at the cost of a large number of satellites required to be launched and maintained (even though a LEO satellite is less expensive than a GEO one). Additionally, mobility management issues, arising due to the non-stationary nature of LEO satellites with respect to the Earth, have to be considered.

When satellites are used for data communications, the application throughput depends, to a large extent, on the throughput of the transport protocol. Certain characteristics of satellite networks, such as long propagation delay of GEO links, high Bit Error Rate (BER) due to channel fading, and frequent handovers of LEO satellites, present a challenge in the design of transport protocols. Although TCP is the dominant transport protocol in the IP protocol suite, it was not initially designed for long bandwidth delay product networks, such as satellite networks which are characterized by long propagation delays and corruption losses due to wireless links. Consequently, enhancements to improve the performance of TCP over satellite networks have been proposed [4], [5], [6].

The recent increase in interest in transmitting Voice over IP (VoIP) networks [7] has led to the development by the Internet Engineering Task Force (IETF) of a new transport layer protocol, called Stream Control Transmission Protocol (SCTP) [8], for the IP protocol suite. Although, the initial objective of developing SCTP was to provide a robust protocol for the transport of VoIP signalling messages over an IP network, later developments have also made it useful as a transport protocol 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 (such as window based congestion control, error detection and retransmission, etc.) of TCP and its enhancements [4] for satellite networks, that made TCP a success during the explosive growth of the Internet. The implementation of some of the enhancements in SCTP are, however, different from their corresponding implementations in TCP. SCTP also incorporated several unique features, such as multistreaming and multihoming (discussed in Sec. IV), that are not available in TCP. Some of these unique features may also help SCTP in achieving a better performance than TCP in satellite networks.

There has been work done in the last few 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. [9], and has been shown that SCTP traffic has the same impact on the congestion control decision of TCP connections as normal TCP traffic. Jungmaier et al. [10] investigated the effect of SCTP multi-homing in the recovery of SS7 network linkset failures. They have shown that multihoming feature of SCTP can help the endpoints to detect link failures earlier than traditional approach and is transparent to upper-layer applications. Caro et al. [11] show the ability of SCTP multi-streaming in reducing the latency of streaming multimedia in high-loss environments. Their experimental results show that multistreaming results in a slower degradation in network throughput as the loss rate increases. Moreover, user satisfaction is increased with the improved multimedia quality provided by this feature.

In the wireless networking area, the performance of SCTP in mobile network [12], [13] and wireless multi-hop networks [14] have been studied. In [12], Fu et al. showed that, like TCP, SCTP also suffered a spurious timeout after a long sudden delay that happens frequently in mobile networks. However, SCTP SACK could be used to eliminate the ensuing spurious fast retransmission after the spurious timeout. The performance of SCTP in Mobile-IP was investigated by Fu et al. [13], and it was shown that the support of a large number of SCTP GapACK blocks in its SACK chunks can expedite the error discovery and lost packet retransmission, and result in better performance than TCP-Reno and TCP-SACK. In [14], Ye et al. have shown that the throughput of an SCTP association degrades when the number of hops between the sender and receiver increases, mainly

due to the hidden node and exposed node problems. A new scheme is proposed by Fu et al. to support low latency, low packet loss mobility called Seamless IP diversity based Generalized Mobility Architecture (SIGMA) [15], which utilizes the multihoming feature of SCTP to achieve lower handover latency, lower packet loss rate, higher throughput than MIPv6 enhancements. In [16], SIGMA is also applied to space networks to support inter-satellite handovers in LEO/MEO constellations. Excellent surveys summarizing results of recent studies on SCTP can be found in [17], [18], [19]. Despite the considerable research on the effectiveness of various SCTP features in terrestrial data networks, *the authors are not aware of any in-depth study to investigate the suitability of SCTP for data communication over satellite networks.*

The *primary objective* of this paper is to evaluate and highlight those advanced features of SCTP that make it suitable for data communication over satellite networks. This paper *differs* from previous work on SCTP in the sense that it investigates, evaluates, and recommends SCTP features that can be exploited to increase SCTP's performance over satellite networks. The results of the evaluation, and the recommendations provided in this paper, can be used to enhance the performance of SCTP for data communications over satellite networks.

We divide the evaluation of the suitability of SCTP features over satellite networks into two parts:

- Evaluation of *standard* SCTP features. These are features which are either available in TCP, or have been proposed as enhancement to improve TCP's performance in satellite networks (Sec. III);
- Evaluation of *unique* SCTP features. These are features which are not available in TCP, but those that might help SCTP in achieving a high throughput in satellite networks (Secs. IV and V).

In addition to evaluating the suitability of SCTP for satellite networks, a *secondary objective* of this paper is to make recommendations regarding the use of SCTP features for enhancing its 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 are as follows:

- *Investigate* the suitability of SCTP for satellite networks;
- *Evaluate* the performance of the unique features of SCTP over satellite networks;
- Provide *recommendations* on the use of 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. Standard and unique SCTP features that make it suitable for satellite communications are discussed in Secs. III and IV, respectively. Simulation results on the performance of the unique SCTP features over satellite constellations is presented in Sec. 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 [4], [6]. Because SCTP and TCP use similar congestion control, retransmission, and round trip time estimation algorithms, the characteristics have many similar effects on the two protocols. The following are the satellite link characteristics which are of interest in this paper.

- *Long propagation delay*: The propagation delay between an earth station and a GEO satellite is around 120-140ms (milliseconds), which means that it takes the sender a long time to probe the network capacity and detect possible loss of segments, resulting in expensive satellite bandwidth being 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 [20].
- *Errors due to propagation corruption and handovers*: The frequent fading of satellite links results in a low signal-to-noise ratio (SNR) and consequently a high Bit Error Rate (BER) during free space propagation. The GSL handovers in LEO constellations will also contribute to the burst errors observed by the endpoints. These errors will cause TCP and SCTP senders to activate congestion control mechanisms, and reduce their transmission rates unnecessarily.
- *Variable Round Trip Time and Link handovers*: The ground stations in LEO satellite system generally experience a handover interval of only a few minutes between two satellites. Propagation delay between ground and LEO varies rapidly as a satellite approaches and leaves a ground station. During the handover, packets can experience a much higher RTT than during normal periods. Transport layer protocols, like TCP and SCTP, depend on accurate RTT estimation to perform congestion control; too

frequent RTT change may cause problems for TCP RTO calculation algorithms.

In Secs. III and IV, we discuss the SCTP features that can be used to alleviate the effect of the above satellite link characteristics on the transport layer performance.

III. STANDARD FEATURES OF SCTP FOR SPACE ENVIRONMENT

A number of TCP's built-in features and enhancements have been recommended for use in satellite networks [4]. These features, which are also available in SCTP (we call them *standard* SCTP features¹), are known to improve TCP performance in satellite networks, and hence are also helpful for SCTP in satellite networks. However, as described below, implementation of some of these standard SCTP features are different from their implementation in TCP.

A. Support for Path MTU Discovery

Path MTU discovery employs a "DF" bit in IP header and "fragmentation needed and DF set" ICMP message to discover the most appropriate Path MTU to be used for data transmission [21]. As in TCP, Path MTU discovery provides SCTP with information about the largest possible segment that will not cause packet fragmentation at intermediate routers. 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 Path MTU (PMTU) as the smallest MTU discovered for all destination IP addresses. A large segment size can reduce packet overhead, and enable an SCTP sender to increase the congestion window (in terms of bytes) rapidly. PMTU discovery is, therefore, recommended for enabling transfer of large SCTP segments over satellite networks.

B. Congestion Control Mechanisms

Like TCP, SCTP also uses Slow Start and Congestion Avoidance algorithms [8] to probe the available capacity of a network. These algorithms force a sender to wait for ACKs before sending new data in order to prevent congestion collapse. Given the long propagation delay of satellite links, the satellite link bandwidth is not utilized efficiently when the sender is going through these algorithms. To ensure that the sender throttles back under adverse network conditions, thus allowing the network to quickly recover from congestions, the algorithms are

¹The *unique* features of SCTP will be described in Sec. IV

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			

necessary in a shared network, especially when a satellite network becomes part of the Internet. These algorithms are, therefore, recommended for use in SCTP over satellite networks.

C. SCTP Selective Acknowledgment

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 [8].

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's *Options* field is limited to four. If SACK is used together with the timestamp option (requires 12 bytes), the maximum number of blocks is reduced to three.

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 four bytes. Therefore, the maximum number of blocks allowed is 16,380.

Use of SACK allows robust reaction in case of multiple losses from a single window of data [22]. This avoids a time-consuming slow start stage after multiple segment losses in a satellite environment and saves network bandwidth.

Satellite links have high BER and require a large transmission window to utilize the satellite network bandwidth (see Sec. III-E). This translates to a higher probability of multiple non-consecutive segment losses in a single window of data transmission. The number of available Gap Blocks of three or four in TCP may not be sufficient for reporting all the lost segments. 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. SCTP allows more Gap Blocks, thereby making it more robust to multiple losses in a window of data caused by satellite link corruption errors.

D. Large Initial Congestion Window

RFC 2960 [8] 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. If the initial congestion window is one segment, the receiver must wait for the 200ms timer to expire before acknowledging the first received segment. Because SCTP requires the initial $cwnd \leq 2$ segments [8], we would like to recommend two 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. 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 [20] to extend the TCP usable window size to 65535×2^{14} bytes, and has been recommended for use in satellite communication [4].

The *Advertised Receiver Window Credit* field in the SCTP SACK header (Table I) has a length of 32 bits. It 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 satellite environments. The implicit support of large receiver window size in SCTP, therefore, makes it suitable for satellite networks.

IV. UNIQUE FEATURES OF SCTP FOR SPACE ENVIRONMENT

In this section, we describe the *unique* SCTP features, i.e. SCTP features that are not available in TCP. These include multihoming, multistreaming, byte counting, and explicit support for ECN. The performance impact of some of these unique SCTP features on data communication over satellite networks will be evaluated by simulation in detail in Sec. V. Fig. 1 shows an SCTP association with multihoming and multistreaming with three streams in the association.

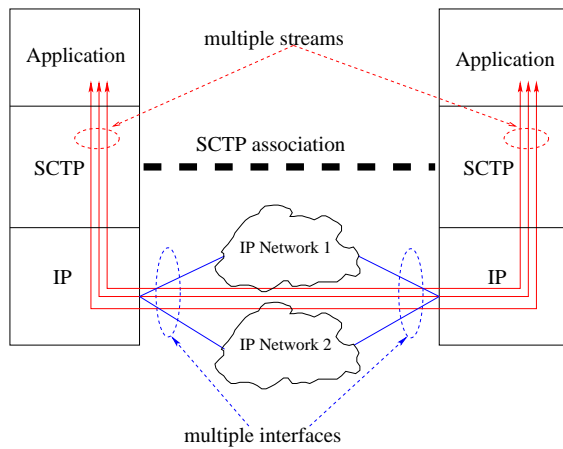


Fig. 1. Schematic view of an SCTP association.

A. Multihoming

Multihoming allows an association (in SCTP terminology, **“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). This built-in support of SCTP for multihomed endpoints can increase the reliability of high-availability applications by switching over data communication to the secondary link when the primary link fails. Retransmission of lost packets is done over the secondary address transparently by the transport protocol without involvement of the application layer. This increases the reliability of retransmitted packets and simplifies the design of applications.

An example use of SCTP multihoming is shown in Fig. 2, where two endpoints are connected through 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 or blackouts periods when the primary satellite is cut out of 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. Multihoming can thus make a satellite network highly reliable and fault tolerant.

B. Multistreaming

Multistreaming allows data from an upper layer application to be split into multiple *streams* in an association as shown in Fig. 3. Sequencing of data within a stream is maintained; 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 segment is retransmitted from the source and received at the receiver. However, data from other streams can

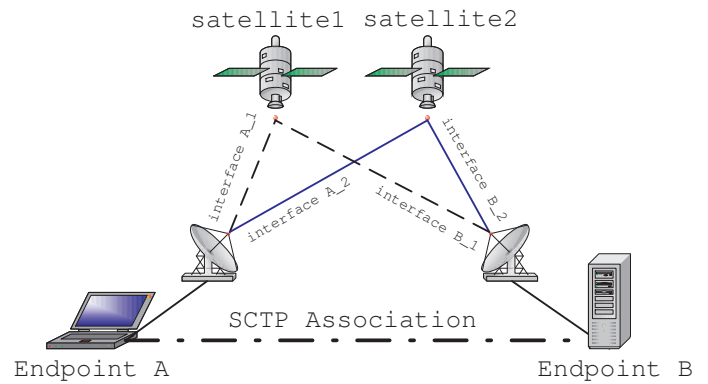


Fig. 2. An SCTP association with multihomed endpoints.

still be delivered to upper layer applications when they arrive at the destination. This avoids the head of line (HOL) blocking [8] found in TCP, where a single stream carries all the data from upper layer applications. The use of multistreaming is especially beneficial in error-prone environments, like satellite communications. It can reduce the HOL blocking at the receiver, and reduce the receiver buffer requirement.

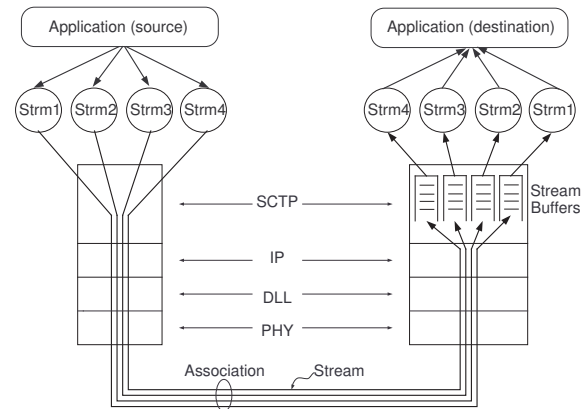


Fig. 3. An SCTP association consisting of four streams carrying data from four upper layer applications.

Effects of multistreaming and HOL blocking are illustrated in Fig. 4, where an SCTP association consisting of four streams is shown. Segments are identified by Stream Sequence Numbers (SSN) [8] which are unique within a stream, although different streams can have the same SSN. In the figure, SSN 11 in stream number 1 has been delivered to the upper layer application, and SSN 9 of the second stream is lost in the network; SSNs 10, 11, 12 are, therefore, queued in the buffer of the second stream while waiting for retransmitted SSN 9 to arrive. Arriving SSN 13 at stream number 2 will also be queued. Similarly SSN 4 of stream 3 is lost during transmission, resulting in the blocking of SSNs 5, 6, and 7. For stream

number 4, SSN 21 is being delivered to the upper layer application, while arriving SSN 23 will be queued in the buffer because of missing SSN 22. Note that when SSN 12 arrives at the buffer of stream 1, it can be delivered immediately even if the other streams are blocked. This illustrates that segments arriving on stream 1 can still be delivered to the upper layer application, although streams 2 and 3 are (and stream 4 will be) blocked because of lost segments.

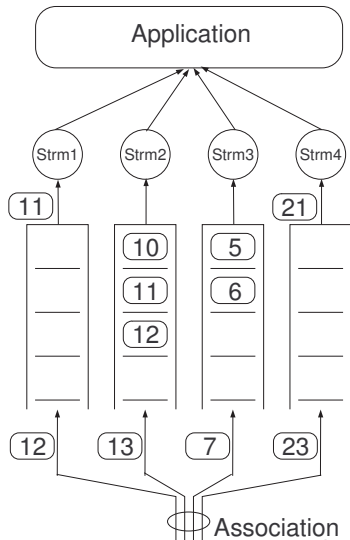


Fig. 4. Illustration showing head of line blocking of individual streams at the receiver.

C. Byte Counting in Acknowledgments

As discussed in Sec. III-D, delayed SACK is recommended by RFC 2960. *Byte counting* algorithm increases the *cwnd* by the number of bytes acknowledged by SACK segments instead of by the number of SACKs. 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 *cwnd* by the number of acknowledgments received by the sender, delayed SACK in TCP increases the time required by the sender to increase *cwnd* during Slow Start.

SCTP limits *cwnd* increase to one PMTU per SACK; we call this Byte Counting Limit (BCL). The benefit of byte counting is impaired when the total number of bytes acknowledged by a single SACK exceeds one PMTU. This effect can be illustrated by Fig. 5, where the PMTU is 1500 bytes. For the SCTP association in Fig. 5(a), the SACK chunk acknowledges 1072 bytes (less than PMTU), so *cwnd* increases by two segments. As a comparison, in Fig. 5(b), the SACK chunk acknowledges

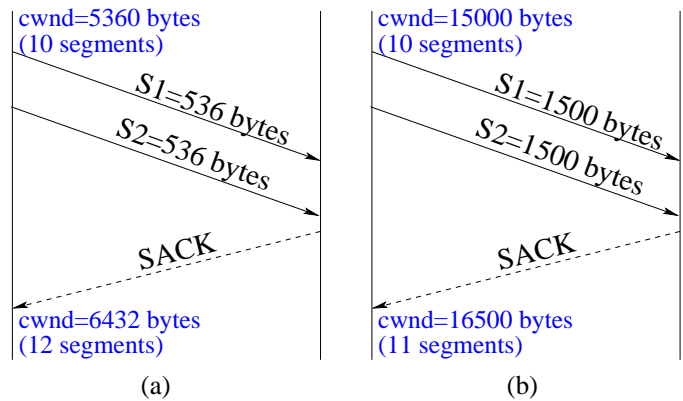


Fig. 5. SCTP byte counting with BCL=1 and 2 respectively.

3000 bytes, but *cwnd* can only be increased by one segment. We, therefore, recommend increasing BCL to two PMTU in order to speedup the slow start phase when delayed SACK is used.

D. Explicit Congestion Notification (ECN)

SCTP defines ECN capable TLV (the optional parameters in an SCTP chunk use the Type-Length-Value (TLV) format [8]) in both INIT and INIT-ACK chunks which are exchanged between endpoints during association setup. When an endpoint initiates 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 in the association. Once ECN enabled, detecting and responding to congestion in SCTP are almost similar to those defined in [23]. 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 *cwnd* has been reduced.

Due to the high BER of satellite links (as compared to terrestrial 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 congestion state to the endpoints [23]. 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, a sender can utilize the feedback from a receiver to differentiate corruption losses from congestion drops. When it is determined that a segment loss is due to corruption during transmission over satellite links, the sender can avoid unnecessary reductions of the

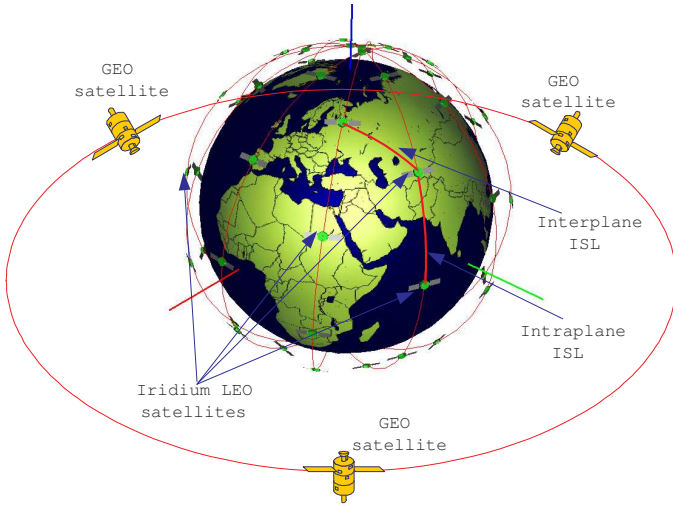


Fig. 6. Mixed constellation of Iridium and GEO.

congestion window, which is an important advantage in long delay satellite networks.

V. PERFORMANCE EVALUATION OF SCTP FOR SATELLITE NETWORKS

In Sec. IV, we described a number of *unique* features of SCTP that can be used to enhance transport layer throughput in satellite networks. In this section, we use discrete event simulation (*ns-2*) to study the effect of those unique features of SCTP on its performance over satellite networks. We use end-to-end *throughput*, defined as the number of useful bits delivered per second to upper layer applications at the destination endpoint, as the measure of transport layer performance.

A. Satellite Constellations

We consider two types of satellite constellations in our simulation: a GEO constellation proposed by the Clarke model [24], and a LEO constellation called Iridium [25], [26]. The GEO constellation resides at an altitude of 35786 km, and each satellite has on-board processing capability to route the packets. We choose Iridium as the LEO constellation in this paper because it is the first operational LEO system that provides truly global coverage. Fig. 6 shows the satellites and their orbits in both a GEO and Iridium LEO constellation.

The Iridium constellation consists of 66 satellites, grouped into 6 *planes* with each plane have 11 satellites. Each satellite has four 25Mbps inter-satellite links (ISL), which operate in the frequency range of 22.55 to 23.55 GHz. Two of the ISLs (called *intraplane* ISL) connect a satellite to its adjacent satellites in the same plane, and the

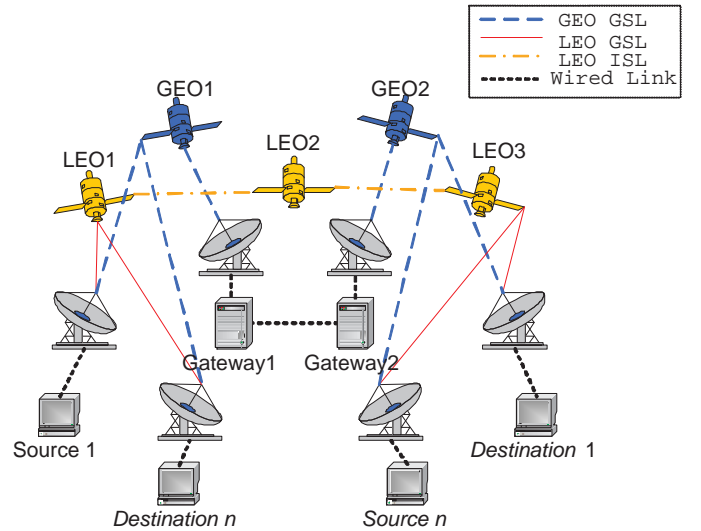


Fig. 7. Simulation topology.

other two ISLs (*interplane* ISL) connect it to the satellites in the neighboring co-rotating planes. The interplane ISLs is temporarily deactivated near the poles because of antenna limitations in tracking these ISLs in polar areas [27]. Each earth endpoint can be connected to a GEO and/or a LEO satellite through ground-to-satellite link (GSL). In the case of connection to LEO, the GSL links experience periodical handovers to accommodate the relative movement of the LEO satellites and the Earth.

B. Simulation Setup and Parameters

The orbit and link characteristics of the GEO and LEO satellite constellations used in our simulation are shown in Tables II and III, respectively, and SCTP protocol parameter values are summarized in Table IV. The simulations were repeated a number of times by varying a number of the parameters in Tables II, III, and IV, while choosing new set of positions for the earth endpoints for every simulation run. For ease of understanding of the interconnections between the senders, receivers and the satellites, a partial network topology consisting of two GEO satellites and three LEO satellites of the constellation for a particular simulation run is shown in Fig. 7. In the topology, each GEO satellite is connected to a ground station using a GSL, and the ground stations are interconnected using a terrestrial mesh network.

In the simulation, SCTP associations between a number of sender-receivers pairs are setup between randomly chosen earth endpoints within the latitude range between -40° and 70° . This selection of latitude range is based on the statistics that over 99% percent of the world population resides in this range of latitudes [27].

TABLE II

ORBIT AND LINK CHARACTERISTICS OF GEO CONSTELLATION.

Number of satellites	3
Altitude	35786.1 km
Longitude	$-90^\circ, 30^\circ, 150^\circ$
Period time	24 hours
GSL link bandwidth	2Mbps
Bandwidth between gateways	100Mbps
Link queue size	50
Path BER	10^{-4} to 10^{-9}

TABLE III

ORBIT AND LINK CHARACTERISTICS OF IRIDIUM CONSTELLATION.

Number of planes	6
Number of satellites/plane	11
Altitude	780.0 km
Geometry	polar orbits at 86.4° incl.
Period time	100.4 minutes
Interplane separation	31.6°
Minimum elevation angle	8.2°
ISLs per satellite	4
GSL link bandwidth	1.5Mbps
ISL link bandwidth	25Mbps
Link queue size	50
Path BER	10^{-4} to 10^{-9}

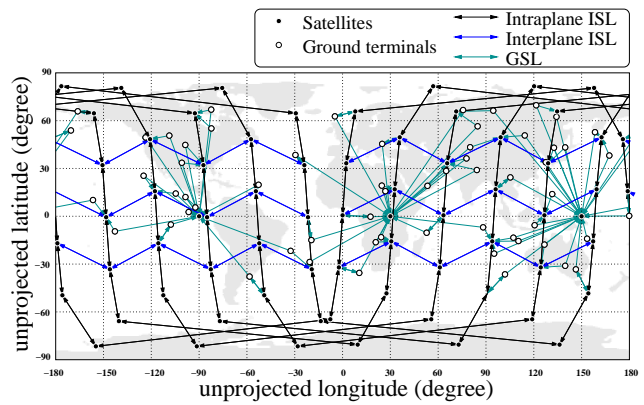
A GEO GSL link is setup between the ground station of an endpoint and its nearest (in terms of longitude) GEO satellite. For LEO GSL links, the simulator automatically connects a ground station to its nearest (depending on latitude and longitude) LEO satellite. A global routing agent (GRA) within the simulator accomplishes the task of connecting endpoints to the LEO satellites. When the network topology changes due to movement of LEO satellites, the GRA recomputes new routing tables at all the nodes (including satellites and ground stations).

Fig. 8 shows a complete snapshot of the ISLs and GSLs in an Iridium-GEO constellation through which

TABLE IV

SIMULATION PARAMETERS FOR THE TOPOLOGY OF FIG. 7.

Traffic type	FTP
Number of associations	10 to 40
Header size	52 bytes
Payload size	512 bytes
Number of streams per association	1 to 4
Receiver buffer size	4 to 40 segments
Byte counting limit	1 PMTU
Initial <i>cwnd</i>	2 segments
Initial <i>ssthresh</i>	20 segments



Credit: background map 1 (outline) provided by the Xerox Parc Map Viewer background maps 2 and 3 (grayscale photo montages) provided as samples by Living earth background map 4 (land masses) courtesy of the footprint generator from L. Wood

Fig. 8. Unprojected map of mixed Iridium/GEO constellation

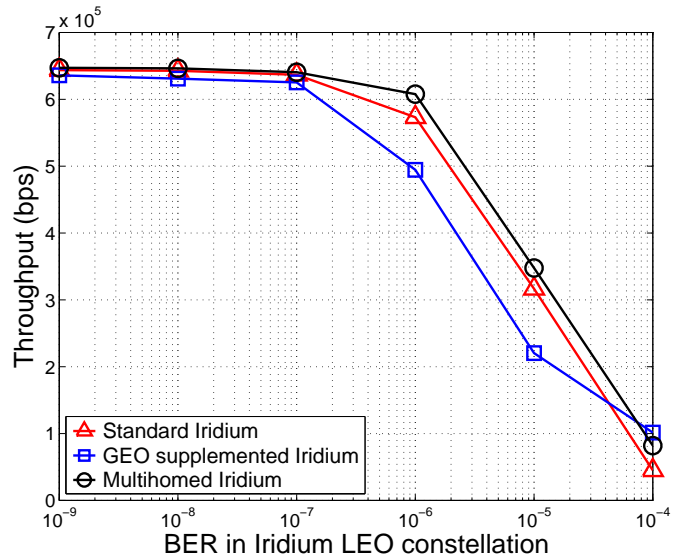


Fig. 9. End to end throughput comparison of standard Iridium and GEO supplemented Iridium for GEO GSL BER of 10^{-6} and 30 SCTP associations.

30 SCTP associations between 30 pairs of end points are setup. The six planes of the Iridium constellation are shown by the nearly-vertical lines (since Iridium's inclination is 86.4°). Each LEO satellite has four ISLs: two intraplane and two interplane. Since the satellites in the two planes near 0° longitude are counter rotating, there is no interplane ISL between the two planes near 0° longitude. Three GEO satellites reside at longitudes of $-90^\circ, 30^\circ,$ and 150° . Since each of the GEO satellites has, on the average, 20 GSLs setup to support 30 SCTP associations requiring 60 GSLs, the GSLs connected to the three GEO satellites appear to be denser as compared to the LEO GSLs.

C. Effect of Multihoming

In this section, we study the effectiveness of SCTP multihoming in improving the end to end throughput in satellite network. We assume that every endpoint is multihomed with two interfaces, with the possibility of connecting it to a GEO and a LEO satellite. We compare three configurations: (1) Standard Iridium (where only Iridium constellation is used) with single homed SCTP associations; (2) GEO supplemented Iridium using multihomed associations, where each earth endpoint is connected to an Iridium and a GEO satellite, and Iridium and GEO GSLs are used as primary and backup paths, respectively; (3) Standard Iridium with multihomed associations, where each earth endpoint is connected to two adjacent Iridium satellites, and these two LEO satellites are used as the primary and backup paths. During our simulations shown in this section, the BER of the GEO GSL links was fixed at 10^{-6} . BER of LEO GSLs and LEO ISLs were varied between 10^{-4} and 10^{-9} . We keeping BER of GEO GSL fixed since we are mainly interested in the effect of SCTP multihoming under different BERs in LEO constellation caused by relative movement and frequent handovers. Also, in this section, the number of streams in each SCTP association is fixed at one to eliminate the effect of multistreaming.

Fig. 9 shows the end to end throughput per association as a function of the BER in Iridium LEO constellation. For low BER, standard Iridium and GEO supplemented Iridium have similar performance. This is because of very few packets being corrupted at low values of BER, resulting in very few retransmissions over GEO. For medium values of BER (between 10^{-7} and 10^{-5}), although a larger number of packets are corrupted and hence retransmitted through the GEO link, the long propagation delay of GEO links slows down the error recovery, resulting in low throughput of GEO supplemented Iridium. For high values of LEO link BER (larger than 10^{-4}), the low link BER of GEO links (fixed at 10^{-6}) compensates for the longer error recovery time, therefore, the GEO supplemented Iridium configuration has the highest performance. In most cases, Standard Iridium with multihomed associations has the highest performance because it takes advantage of both higher reliability of SCTP multihoming and lower delay of LEO GSLs.

To dimension the GEO link bandwidth, we show in Fig. 10 the total bandwidth requirement (for all associations) of a GEO link (when it is used as a backup path) in the GEO supplemented Iridium configuration as a function of LEO link BER. The load on the GEO GSLs were changed by varying the number of associations between

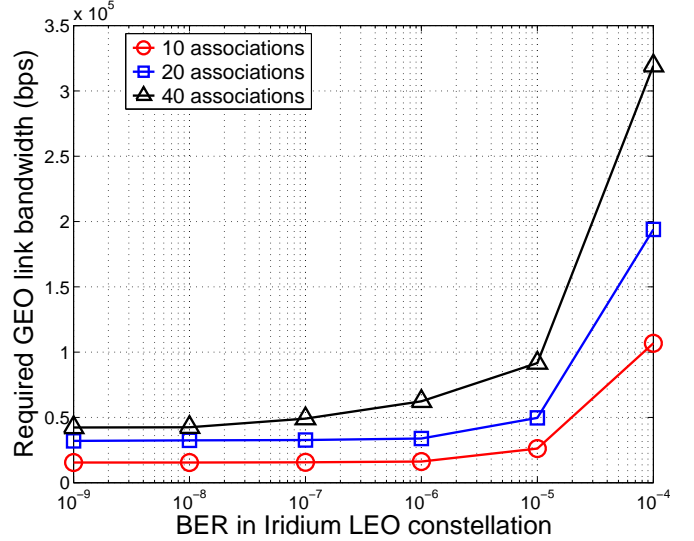


Fig. 10. Requirement on GEO link bandwidth when used as backup path

10 and 40. As expected, higher LEO link BER results in larger number of retransmitted packets due to larger number of corrupted packets. The larger number of retransmitted packets result in large bandwidth requirement of GEO GSLs. Considering the highest value of BER (10^{-4}) and dividing the total bandwidth by the number of association, the maximum bandwidth requirement of a GEO GSL is found to be about 10Kbps per association. This bandwidth requirement on GEO GSL links is not high, even for bandwidth hungry FTP traffic, allowing GEO satellites as backup links to be a cost effective approach to increase the end to end throughput for high LEO link BERs.

D. Effect of Multistreaming

To study the effect of SCTP multistreaming (described in Sec. IV-B) on buffer size requirement at the receiver endpoint, and end to end throughput over GEO and LEO satellite networks (having different characteristics as shown in Tables II and III), the number of streams per association was varied between one to four. We simulated multistreaming over GEO and LEO constellations separately using only one interface (either GEO or LEO) of each endpoint.

Multistreaming can be used to alleviate the head-of-line (HOL) blocking resulting from TCP's strict byte-order delivery (see Sec. IV). Each stream is kind of a "sub-flow" within the overall data flow, where the delivery of packets in a sub-flow are independent of other sub-flows. Under error-prone satellite link conditions and limited receiver buffer size, multistreaming can significantly reduce the receiver buffer size requirements




Fig. 11. The effect of multistreaming on Iridium throughput as a function of receiver buffer size.

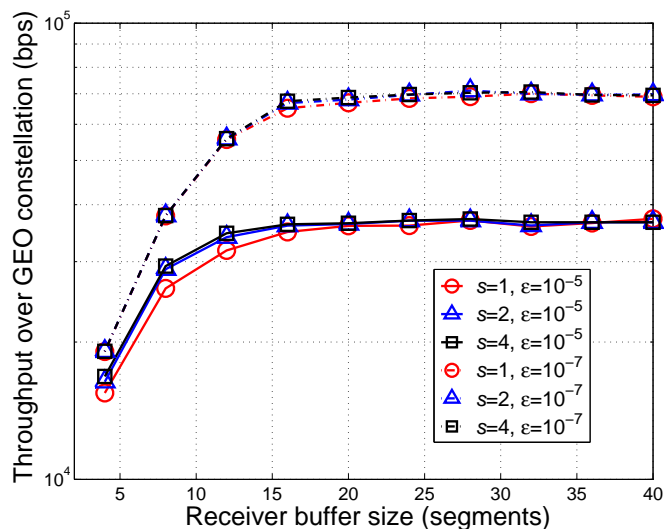


Fig. 12. The effect of multistreaming on GEO throughput as a function of receiver buffer size.

and increase end to end throughput. This is illustrated in Figs. 11 and 12 which show the end to end throughput per association as a function of the receiver buffer size. s is the number of streams and ϵ is the BER of Iridium or GEO links. We can see that when the BER is low ($< 10^{-7}$), multistreaming does not have much impact on the end to end throughput. For higher BER ($\geq 10^{-5}$) with limited receiver buffer, the HOL blocking will result in buffer overflow at the receiver and reduction in throughput.

Figs. 13 and 14 show the end to end throughput as a function of BER for LEO and GEO satellites, respectively. We also varied the number of streams and buffer

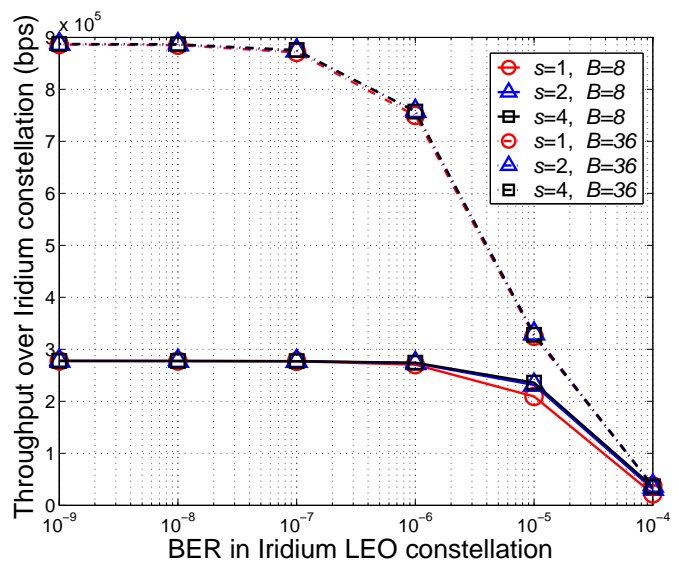


Fig. 13. The effect of multistreaming on throughput as a function of Iridium BER.

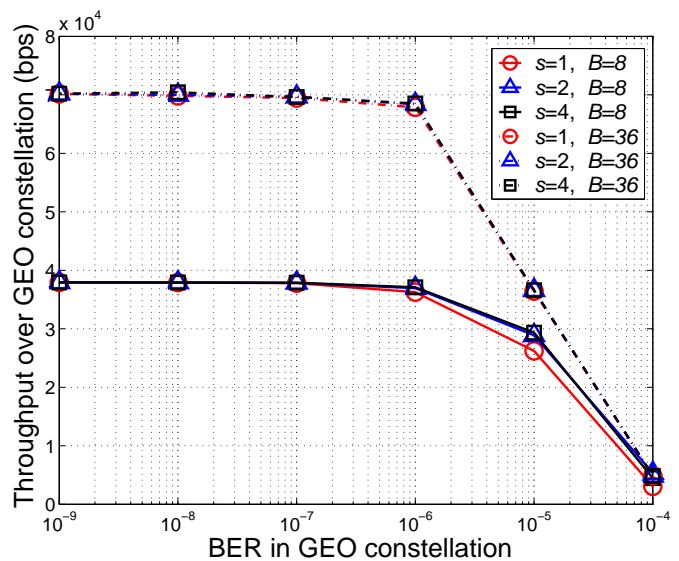


Fig. 14. The effect of multistreaming on throughput as a function of GEO BER.

sizes. We can see that the number of streams has an impact on the end to end throughput for high BER and small receiver buffer size (8 segments). When the buffer size is small, a high BER will result in a high degree of HOL blocking, resulting in higher possibility of buffer overflow, and therefore lower throughput. However, when the receiver buffer size is large, the buffer is sufficient to avoid buffer overflow. The throughput is not noticeably reduced as a result on HOL blocking, and the number of streams has less effect on the end to end throughput.

TABLE V
SUMMARY OF RECOMMENDATIONS FOR STANDARD SCTP
FEATURES.

Feature	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 Initial <i>cwnd</i>	Recommended	S
Large Receiver Window	Implicitly Used	S, R

VI. RECOMMENDATIONS AND CONCLUSION

SCTP is a new transport layer protocol which is being developed by the Internet Engineering Task Force. In this paper, we have studied the suitability of its various features in space environment. In addition to a number of features which are in common with TCP and are known to help transport layer performance in space environment, we have also investigated the suitability of some of its unique features for space environment.

We first outlined satellite link characteristics which may limit the performance of transport protocols, followed by unique SCTP features that may help in better utilizing the bandwidth of satellite networks, while preventing congestion collapse in a shared network. *The authors believe that these features should make SCTP very suitable as a transport protocol for satellite networks.*

With a view to stimulating further research in the area of SCTP in space networks, we summarize our recommendations for the use of SCTP features in space networks in Tables V and VI. In these tables, the last column denotes the point in the network where the feature could be implemented: "S" means the sender, "R" means the receiver, and "S, R" means both sender and receiver. Table V summarizes the recommendations for standard SCTP features (i.e. those which are also available in TCP), while Table VI provides recommendations for use of the unique SCTP features (i.e. those that are not available in TCP).

There are a number of unresolved issues for both TCP and SCTP in a space environment, 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. Such issues require further research to improve the performance of SCTP over satellite links. Moreover, TCP enhancements, such as Protecting Against Wrapped

TABLE VI
SUMMARY OF RECOMMENDATIONS FOR UNIQUE SCTP
FEATURES.

Feature	Use	Where
SCTP Multihoming	Recommended	S, R
SCTP Multistreaming	Recommended	S, R
Byte Counting	Implicitly Used	S, R
Larger Byte Counting limit	Recommended	S
ECN	Recommended	S, R

Sequence (PAWS) numbers and Round Trip Time Measurement (RTTM) [20], require the timestamp option [20] which is not available in SCTP. In order to use these features in SCTP, new chunk type for timestamp should be considered in future developments of SCTP.

REFERENCES

- [1] L. Wood, A. Clerget, and I. Andrikopoulos et. al., "IP routing issues in satellite constellation networks," *International Journal of Satellite Communications*, vol. 19, no. 1, pp. 69–92, January/February 2001.
- [2] K. Bhasin and J. L. Hayden, "Space Internet architectures and technologies for NASA enterprises," *International Journal of Satellite Communications*, vol. 20, no. 5, pp. 311–332, September/October 2002.
- [3] R. Donadio, F. Zeppenfeldt, and S. Pirio et. al. eds., *ESA Telecommunications: IP Networking over Satellite Workshop*. European Space Agency, Noordwijk, Netherlands, May 2004.
- [4] M. Allman, D. Glover, and L. Sanchez, "Enhancing TCP over satellite channels using standard mechanisms." IETF RFC 2488, January 1999.
- [5] N. Ghani and S. Dixit, "TCP/IP Enhancements for Satellite Networks," *IEEE Communication Magazine*, vol. 37, no. 7, pp. 64–72, July 1999.
- [6] C. Partridge and T.J. Shepard, "TCP/IP performance over satellite links," *IEEE Network*, vol. 11, no. 5, pp. 44–49, Sep/Oct 1997.
- [7] 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.
- [8] R. Stewart and Q. Xie et. al., "Stream control transmission protocol." IETF RFC 2960, October 2000.
- [9] 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.
- [10] 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.
- [11] A.L. Caro, P.D. Amer, P.T. Conrad, and G.J. Heinz, "Improving multimedia performance over lossy networks via SCTP," *ATIRP 2001*, College Park, MD, March 2001.
- [12] 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.

- [13] S. Fu and M. Atiquzzaman, "Improving end-to-end throughput of Mobile IP using SCTP," *Workshop on High Performance Switching and Routing*, Torino, Italy, pp. 171–176, June 2003.
- [14] G. Ye, T. Saadawi, and M. Lee, "SCTP congestion control performance in wireless multi-hop networks," *MILCOM2002*, Anaheim, California, pp. 934–939, October 2002.
- [15] S. Fu, L. Ma, M. Atiquzzaman, and Y. Lee, "Architecture and performance of SIGMA: Seamless IP diversity based Generalized Mobility Architecture," *Accepted for publication by ICC*, Seoul, Korea, May 2005.
- [16] S. Fu and M. Atiquzzaman, *TraSH: A Transport Layer Handover Protocol for Mobile Terrestrial and Space Networks*. Invited book chapter, Kluwer Academic Publishers, 2005.
- [17] S. Fu and M. Atiquzzaman, "SCTP: State of the art in research, products, and technical challenges," *IEEE Communication Magazine*, vol. 42, no. 4, pp. 64–76, April 2004.
- [18] R. Stewart and C. Metz, "SCTP: New transport protocol for TCP/IP," *IEEE Internet Computing*, vol. 5, no. 6, pp. 64–69, November/December 2001.
- [19] A.L. Caro, J.R. Iyengar, and P.D. Amer et. al, "SCTP: a proposed standard for robust Internet data transport," *IEEE Computer*, vol. 36, no. 11, pp. 56–63, November 2003.
- [20] V. Jacobson, R. Braden, and D. Borman, "TCP extensions for high performance." IETF RFC 1323, May 1992.
- [21] J. Mogul and S. Deering, "Path MTU Discovery." IETF RFC 1191, November 1990.
- [22] 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.
- [23] K. Ramakrishnan and S. Floyd, "A Proposal to add Explicit Congestion Notification (ECN) to IP." IETF RFC 2481, January 1999.
- [24] A.C. Clarke, "Extra-terrestrial relays," *Wireless World*, vol. LI, no. 10, pp. 305–308, Oct 1945.
- [25] J.L. Grubb, "Iridium overview - the traveler's dream come true," *IEEE Communications Magazine*, vol. 29, no. 11, pp. 48–51, Nov 1991.
- [26] C.E. Fossa, R.A. Raines, G.H. Gunsch, and M.A. Temple, "An overview of the IRIDIUM low earth orbit (LEO) satellite system," *IEEE National Aerospace and Electronics Conference (NAECON)*, pp. 152 – 159, July 1998.
- [27] T. Pratt, C.W. Bostian, and J. Allnutt, *Satellite communications*, New York: Wiley, 2 ed., 2003.