EVALUATION OF SCTP FOR SPACE NETWORKS

SHAOJIAN FU, OPNET
MOHAMMED ATIQUZZAMAN, UNIVERSITY OF OKLAHOMA
WILLIAM IVANCIC, NASA GLENN RESEARCH CENTER

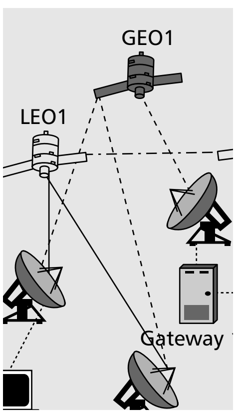

SCTP has recently been standardized as a new transport layer protocol in the IP protocol suite. SCTP is based on the TCP protocol, but incorporates a number of advanced and unique features which are not available in TCP.

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

ABSTRACT

The Stream Control Transmission Protocol 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 that 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 article 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.

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, geostationary Earth orbit (GEO) and low Earth orbit (LEO), 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 uplink are transparently piped to the corresponding downlink (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 onboard switching and routing units [1]. This is also consistent with the current efforts of the National Aeronautics and Space Administration (NASA) and the European Space Agency (ESA).

The long propagation delay of GEO satellites makes them less desirable for real-time applications, such as voice communications. The concept of a LEO satellite constellation was introduced in the 1990s to provide satellite services at a lower obit by utilizing a larger 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 nonstationary 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 challenges 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 [2, 3]. The recent increase in interest in transmitting voice over IP (VoIP) networks has led to the development by the Internet Engineering Task Force (IETF) of a new transport layer protocol, called Stream Control Transmission Protocol (SCTP) [4] for the IP protocol suite. Although, the initial objective of developing SCTP was to provide a robust protocol for the transport of VoIP signaling 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 that can coexist with TCP in the same network. The design of SCTP absorbed many strengths and features (window-based congestion control, error detection and retransmission, etc.) of TCP and its enhancements [2] 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 later), that are not available in TCP. Some of these unique fea-

 1536-1284/05/$20.00 © 2005 IEEE

tures may also help SCTP to achieve 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 [5]. For example, study of the coexistence of SCTP and TCP in the Internet has shown that SCTP traffic has the same impact on the congestion control decision of TCP connections as normal TCP traffic. Study of the effect of SCTP multihoming on the recovery of SS7 network linkset failures has shown that multihoming in SCTP can help endpoints to detect link failures earlier than traditional approaches and is transparent to upper-layer applications. Research on SCTP multistreaming in reducing the latency of streaming multimedia in high-loss environments shows that multistreaming results in 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 networks [6] and wireless multihop networks [7] has been studied. The performance of SCTP in Mobile IP (MIP) was investigated by Fu *et al.* [6], and it was shown that the support of a large number of SCTP gap acknowledgment (GapACK) blocks in its selective ACK (SACK) chunks can expedite error discovery and lost packet retransmission, and result in better performance than TCP-Reno and TCP-SACK. In [7] 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) [8], which utilizes the multihoming feature of SCTP to achieve lower handover latency, lower packet loss rate, and higher throughput than MIPv6 enhancements. In [9] SIGMA is also applied to space networks to support intersatellite handovers in LEO/MEO constellations. Despite 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 article is to evaluate and highlight those advanced features of SCTP that make it suitable for data communication over satellite networks. This article *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 article can be used to enhance the performance of SCTP for data communications over satellite networks.

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

- Evaluation of *standard* SCTP features. These are features that are either available in TCP or have been proposed as enhancements to improve TCP's performance in satellite networks.
- Evaluation of *unique* SCTP features. These are features that are not available in TCP, but might help SCTP achieve high throughput in satellite networks.

In addition to evaluating the suitability of SCTP for satellite networks, a *secondary objective* of this article 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 article 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 article is organized as follows. 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, respectively. Simulation results on the performance of the unique SCTP features over satellite constellations is presented. Recommendations on using SCTP over satellite networks and our conclusions from this research are presented.

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 [2, 10]. Because SCTP and TCP use similar congestion control, retransmission, and round-trip time (RTT) estimation algorithms, the characteristics have many similar effects on the two protocols. The following are the satellite link characteristics that are of interest in this article.

Long propagation delay: The propagation delay between an Earth station and a GEO satellite is around 120–140 ms, 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 96,500-byte size pipe. The fundamental performance problems with the current TCP over LFN links were discussed in [11].

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 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 RTT and link handovers: The ground

A number of satellite link characteristics may limit the performance of transport protocols over satellite networks. 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.

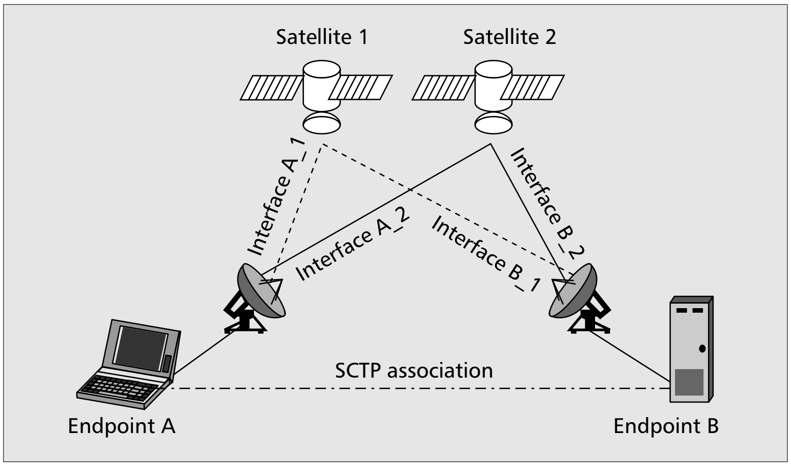

Figure 1. *An SCTP association with multihomed endpoints.*

stations in a 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.

We discuss the SCTP features that can be used to alleviate the effect of the above satellite link characteristics on the transport layer performance.

Standard Features of SCTP for the Space Environment

A number of TCP's built-in features and enhancements have been recommended for use in satellite networks [2]. These features, which are also available in SCTP (we call them *standard* SCTP features[1]), 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.

Support for Path MTU Discovery

Path minimum transmission unit (MTU) discovery employs a DF bit in IP header and a "*fragmentation needed and DF set*" ICMP message to discover the most appropriate path MTU to be used for data transmission. 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.

Congestion Control Mechanisms

Like TCP, SCTP also uses slow start and congestion avoidance algorithms [4] 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 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.

SCTP Selective Acknowledgment

Unlike TCP, the use of SACK is mandatory in SCTP. In SCTP, all data are carried in a structure called a *chunk* which is fully described by the *chunk type*, *chunk flags*, *chunk length*, and *chunk data* fields (see [4] for detailed SCTP SACK chunk format). 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, is 2^{16} bytes. Subtracting the space used by first 16 bytes of SACK chunk, the maximum space available for gap blocks is $2^{16} \times 16$, 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. 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. 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.

Large Initial Congestion Window

RFC 2960 [4] recommends that an SCTP receiver use delayed SACK in acknowledging user data. This requires an ACK to be generated for every second segment received, or within 200 ms of the arrival of any unacknowledged segment. If the ini-

[1] *The unique features of SCTP will be described later.*

tial congestion window (*cwnd*) is one segment, the receiver must wait for the 200 ms timer to expire before acknowledging the first received segment. Because SCTP requires the initial *cwnd* ≤ 2 segments [4], 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.

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 65,535 bytes. Because a DS1-speed GEO satellite channel has a 96,500-byte pipe, TCP cannot fully utilize the channel bandwidth. As a result, the window scaling option was proposed [11] to extend the TCP usable window size to $65{,}535 \times 2^{14}$ bytes, and has been recommended for use in satellite communication [2].

The *advertised receiver window credit* field in the SCTP SACK header has a length of 32 bits [4]. It enables a usable receiver window of up to 2^{32} bytes, compared to $65{,}535 \times 2^{14}$ 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 makes it suitable for satellite networks.

Unique Features of SCTP for Space Environment

In this section we describe the *unique* SCTP features (i.e., SCTP features 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 later.

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. 1, 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 case of failure of the primary or blackout periods when the primary satellite is cut out of communication due to shadowing, satellite handovers, and so on. 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.

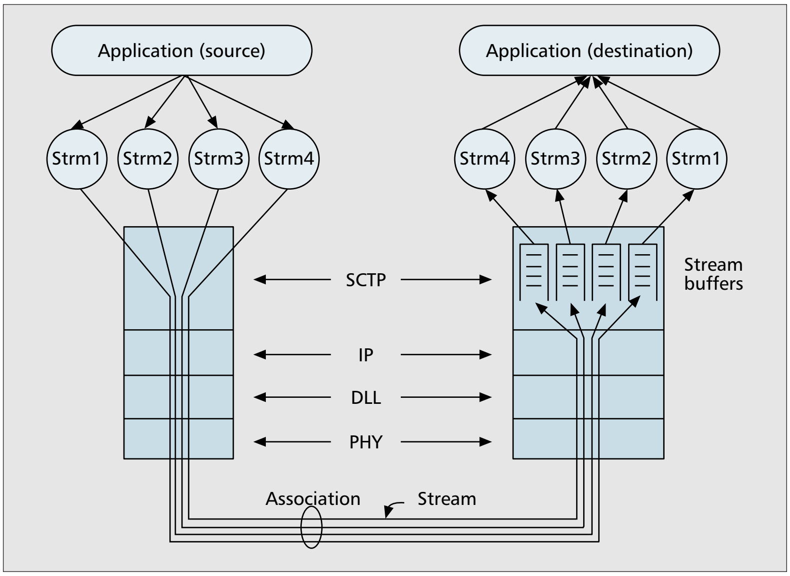

Figure 2. *An SCTP association consisting of four streams carrying data from one upper layer application.*

Multistreaming

Multistreaming allows data from an upper layer application to be split into multiple *streams* in an association as shown in Fig. 2. 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 still be delivered to upper layer applications when they arrive at the destination. This avoids the head of line (HOL) blocking [4, 5] 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.

Byte Counting in Acknowledgments

As discussed earlier, delayed SACK is recommended by RFC 2960. A *byte counting* algorithm increases *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 is illustrated by Fig. 3, where the PMTU is 1500 bytes. For the SCTP association in Fig. 3a, the SACK chunk acknowledges 1072 bytes (less than PMTU), so

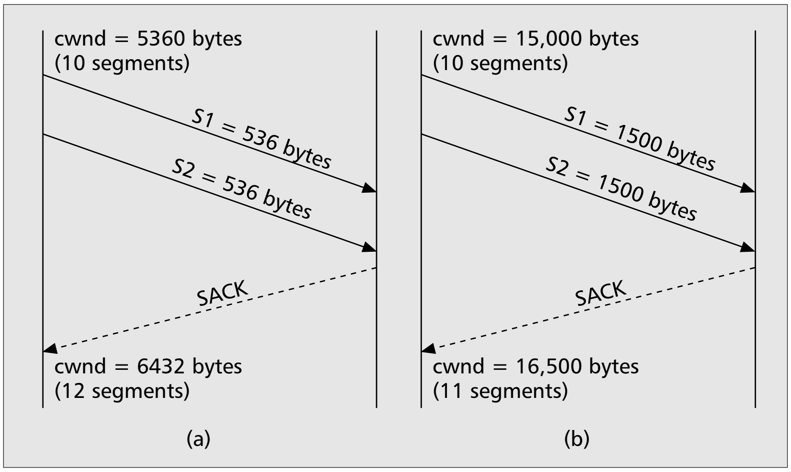

Figure 3. *The effect of BCL on* cwnd *with a) segment size = 536 bytes; b) segment size = 1500 bytes.*

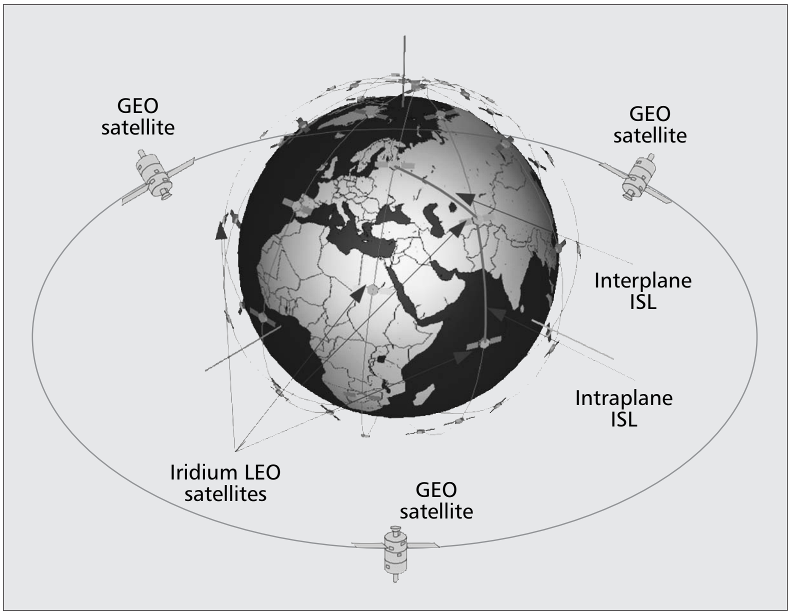

Figure 4. *Mixed constellation of Iridium and GEO.*

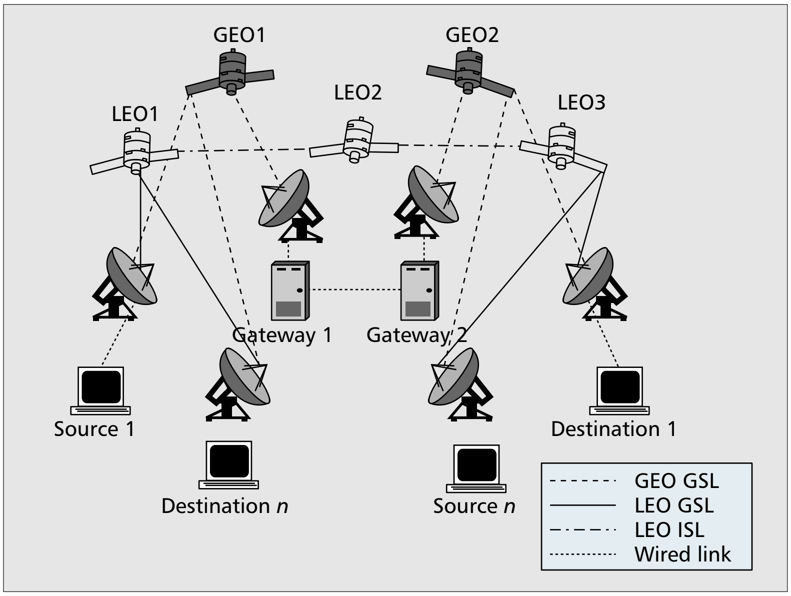

Figure 5. *Simulation topology with three LEO and two GEO satellites.*

cwnd increases by two segments. As a comparison, in Fig. 3b the SACK chunk acknowledges 3000 bytes, but *cwnd* can only be increased by one segment. We therefore recommend increasing BCL to two PMTUs in order to speed up the slow start phase when delayed SACK is used.

EXPLICIT CONGESTION NOTIFICATION

SCTP defines an explicit congestion notification (ECN)-capable type-length-value (TLV) (the optional parameters in an SCTP chunk use the TLV format [4]) in both INIT and INIT-ACK chunks 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 is enabled, detecting and responding to congestion in SCTP are almost similar to those defined in IETF RFC 2481. The difference is when the SCTP receiver detects the *congestion experienced* bit in the IP header of a received segment. It will use an ECN 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. 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 congestion window, which is an important advantage in long delay satellite networks.

PERFORMANCE EVALUATION OF SCTP FOR SATELLITE NETWORKS

Earlier, 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.

SATELLITE CONSTELLATIONS

We consider two types of satellite constellations in our simulation: a GEO constellation proposed by the Clarke model [12], and a LEO constellation called Iridium [13]. The GEO constellation resides at an altitude of 35786 km, and each satellite has onboard processing capability to route the packets. We choose Iridium as the LEO constellation in this article because it is the first operational LEO system that provides truly global coverage. Figure 4 shows the satellites and their orbits in both a GEO and the Iridium LEO constellation.

The Iridium constellation consists of 66 satellites, grouped into 6 *planes* with each plane having 11 satellites. Each satellite has four 25 Mb/s intersatellite links (ISLs), which operate in the frequency range of 22.55–23.55 GHz. Two of the ISLs (called *intraplane* ISLs) connect a satellite to its adjacent satellites in the same plane, and the other two ISLs (*interplane* ISLs) connect it to the satellites in the neighboring corotating planes. The interplane ISLs are temporarily deactivated near the poles because of antenna limitations in tracking these ISLs in polar areas [14]. Each Earth endpoint can be connected to a GEO and/or LEO satellite through a ground-to-satellite link (GSL). In connection to LEO, GSL links experience periodical handovers to accommodate the relative movement of the LEO satellites and the Earth.

Simulation Setup and Parameters

The orbit and link characteristics of the GEO and LEO satellite constellations used in our simulation are shown in Table 1, respectively, and SCTP protocol parameter values are summarized in Table 2. The simulations were repeated a number of times by varying a number of the parameters in Tables 1 and 2, while choosing new sets of positions for the Earth endpoints for every simulation run. For ease of understanding the interconnections between the senders, receivers, and 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. 5. 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-receiver pairs are set up between randomly chosen Earth endpoints within a 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 [14].

A GEO GSL link is set up 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).

Figure 6 shows a complete snapshot of the ISLs and GSLs in an Iridium-GEO constellation through which 30 SCTP associations between 30 pairs of endpoints are set up. 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°, 30°, and 150°. Since each of the GEO satellites has, on average, 20 GSLs set up to support 30 SCTP associations requiring 60 GSLs, the GSLs connected to the three GEO satellites appear to be denser than the LEO GSLs.

Parameter	GEO	Iridium
Number of planes	1	6
Number of satellites/plane	3	11
Altitude	35786.1 km	780.0 km
Period time	24 h	100.4 min
Longitude separation	120°	31.6°
Minimum elevation angle	N/A	8.2°
ISLs per satellite	0	4
GSL link bandwidth	2 Mb/s	1.5 Mb/s
Bandwidth between gateways	100 Mb/s	N/A
ISL link bandwidth	N/A	25 Mb/s
Link queue size	50	50
Path BER	10^{-4} to 10^{-9}	10^{-4} to 10^{-9}

Table 1. *Orbit and link characteristics of satellite constellations.*

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 *cwnd*	2 segments
Initial *ssthresh*	20 segments

Table 2. *Simulation parameters for the topology of Fig. 5.*

The Effect of Multihoming

In this section we study the effectiveness of SCTP multihoming in improving the end-to-end throughput in satellite networks. 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:

- Standard Iridium (where only the Iridium constellation is used) with single-homed SCTP associations
- 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
- 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}. BERs of LEO GSLs and LEO ISLs varied between 10^{-4} and 10^{-9}. We kept the BER of the GEO GSL fixed since we are mainly interested in the effect of SCTP multihoming under different BERs in the 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.

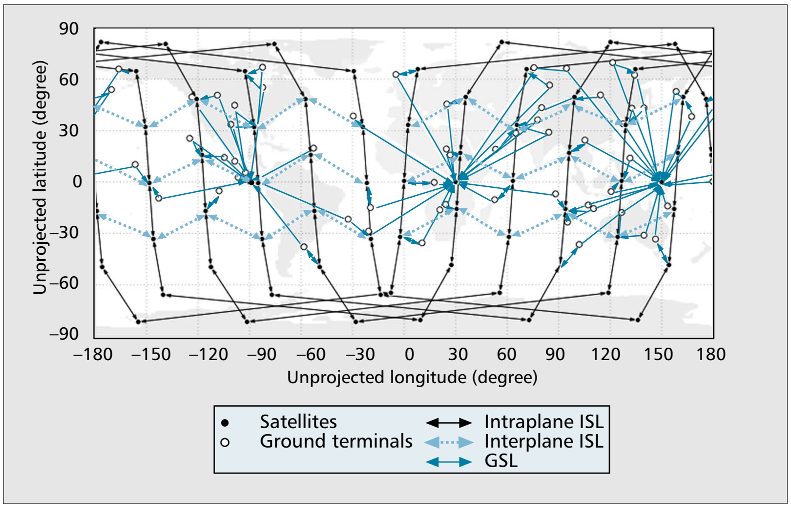

Figure 6. *Unprojected map of mixed Iridium/GEO constellation. Background map 1 (outline) provided by Xerox Parc Map Viewer; background maps 2 and 3 (grayscale photo montages) provided as samples by Living earth; background map 4 (land masses) courtesy of the footprint generator from L. Wood.*

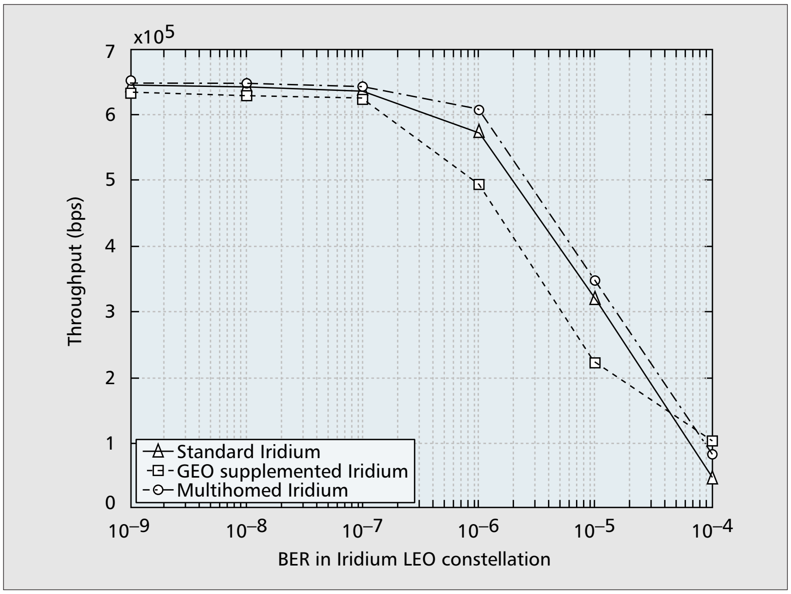

Figure 7. *End-to-end throughput comparison of standard Iridium and GEO supplemented Iridium for GEO GSL BER of 10^{-6} and 30 SCTP associations.*

Figure 7 shows the end-to-end throughput per association as a function of the BER in the Iridium LEO constellation. For low BER, standard Iridium and GEO supplemented Iridium have similar performance. This is because of very few packets 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 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. 8 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 was changed by varying the number of associations between 10 and 40. As expected, higher LEO link BER results in larger numbers of retransmitted packets due to larger numbers of corrupted packets. The larger number of retransmitted packets results in large bandwidth requirementd of GEO GSLs. Considering the highest value of BER (10^{-4}) and dividing the total bandwidth by the number of associations, the maximum bandwidth requirement of a GEO GSL is found to be about 10 kb/s 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.

The Effect of Multistreaming

To study the effect of SCTP multistreaming (described earlier) on buffer size requirements at the receiver endpoint, and end-to-end throughput over GEO and LEO satellite networks (having different characteristics as shown in Table 1), the number of streams per association was varied between one and 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 HOL blocking resulting from TCP's strict byte-order delivery. Each stream is a kind of *subflow* within the overall data flow, where the delivery of packets in a subflow is independent of other subflows. Under error-prone satellite link conditions and limited receiver buffer size, multistreaming can significantly reduce the receiver buffer size requirements and increase end to end throughput. This is illustrated in Fig. 9, which shows 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 LEO 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, HOL blocking will result in buffer overflow at the receiver and reduction in throughput.

Figure 10 shows the end-to-end throughput as a function of BER for Iridium LEO satellites. We also varied the number of streams and buffer 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 (eight 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 of HOL blocking, and the number of streams has less effect on the end-to-end throughput.

The results for the effect of SCTP multistreaming on buffer size requirement and end-to-end throughput under a GEO satellite environment are very similar to those under an Iridium LEO environment. Due to space limitations, the GEO results are not shown in this article. Interested readers can refer to our full-sized report [15].

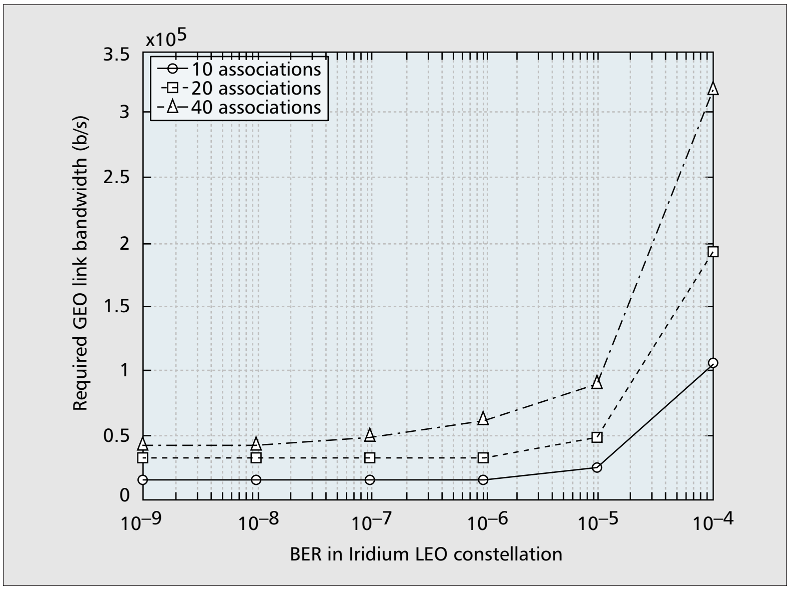

Figure 8. *Requirement on GEO link bandwidth when used as backup path.*

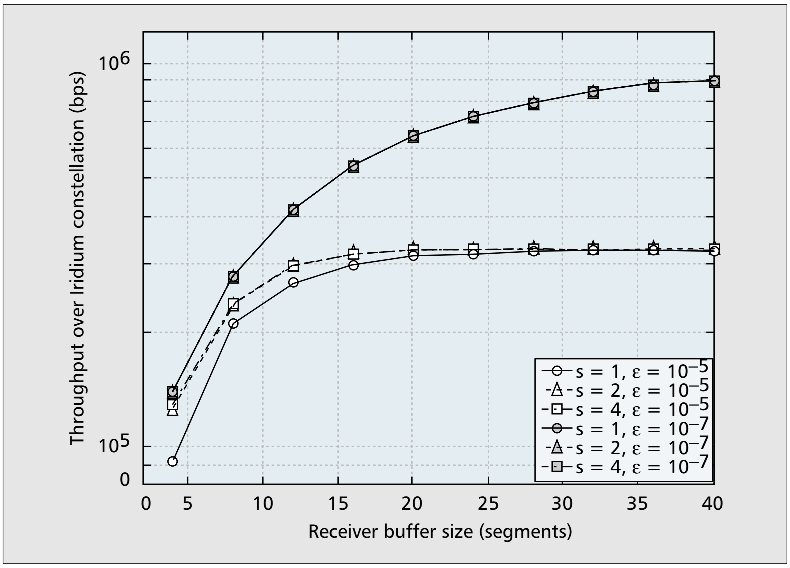

Figure 9. *The effect of multistreaming on Iridium throughput as a function of receiver buffer size.*

Recommendations and Conclusion

SCTP is a new transport layer protocol being developed by the IETF. In this article we study the suitability of its various features in the space environment. In addition to a number of features in common with TCP known to help transport layer performance in a space environment, we have also investigated the suitability of some of its unique features for that environment.

We first outlined satellite link characteristics that may limit the performance of transport protocols, followed by unique SCTP features that may help better utilize 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 Table 3, including recommendations for both standard (i.e., those also available in TCP) and unique (i.e. those not available in TCP) SCTP features. In this table 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.

There are a number of unresolved issues for both TCP and SCTP in a space environment, such as SCTP/IP header compression in a high BER environment, bias against long RTT associations during congestion avoidance, and the interaction between SCTP retransmissions and link layer automatic repeat request (ARQ). Such issues require further research to improve the performance of SCTP over satellite links. Moreover, TCP enhancements, such as protecting against wrapped sequence (PAWS) numbers and RTT measurement (RTTM) [11], require the timestamp option [11], which is not available in SCTP. In order to use these features in SCTP, a new timestamp chunk type should be considered in future developments of SCTP.

References

[1] L. Wood *et al.*, "IP Routing Issues in Satellite Constellation Networks," *Int'l. J. Sat. Commun.*, vol. 19, no. 1, Jan./ Feb. 2001, pp. 69–92.

[2] M. Allman, D. Glover, and L. Sanchez, "Enhancing TCP over Satellite Channels Using Standard Mechanisms." IETF RFC 2488, Jan. 1999.

[3] N. Ghani and S. Dixit, "TCP/IP Enhancements for Satellite Networks," *IEEE Commun. Mag.*, vol. 37, no. 7, July 1999, pp. 64–72.

[4] R. Stewart *et al.*, "Stream Control Transmission Protocol," IETF RFC 2960, Oct. 2000.

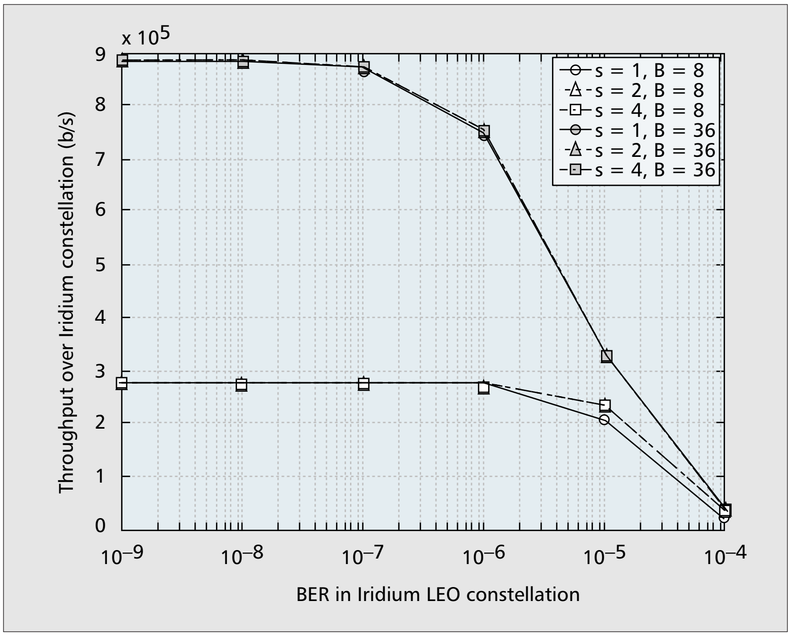

Figure 10. *The effect of multistreaming on Iridium throughput as a function of BER.*

Type	Feature name	Use	Where
Standard features	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 cwnd	Recommended	S
	Large receiver window	Implicitly used	S, R
Unique features	SCTP multihoming	Recommended	S, R
	SCTP multistreaming	Recommended	S, R
	Byte counting	Implicitly used	S, R
	Larger BCL	Recommended	S
	ECN	Recommended	S, R

Table 3. *Summary of recommendations for SCTP features.*

[5] S. Fu and M. Atiquzzaman, "SCTP: State of the Art in Research, Products, and Technical Challenges," *IEEE Commun. Mag.*, vol. 42, no. 4, Apr. 2004, pp. 64–76.

[6] S. Fu and M. Atiquzzaman, "Improving End-to-end Throughput of Mobile IP using SCTP," *Wksp. High Perf. Switching and Routing*, Torino, Italy, June 2003, pp. 171–76.

[7] G. Ye, T. Saadawi, and M. Lee, "SCTP Congestion Control Performance in Wireless Multi-hop Networks," *MILCOM2002*, Anaheim, California, Oct. 2002, pp. 934–39.

[8] S. Fu *et al.*, "Architecture and Performance of SIGMA: A Seamless Mobility Architecture for Data Networks," *40th IEEE ICC*, Seoul, Korea, May 2005.

[9] S. Fu and M. Atiquzzaman, *SIGMA: A Transport Layer Handover Protocol for Mobile Terrestrial and Space Networks*, Invited book chap., Kluwer, 2005.

[10] C. Partridge and T. J. Shepard, "TCP/IP Performance over Satellite Links," *IEEE Network*, vol. 11, no. 5, Sept./Oct. 1997, pp. 44–49.

[11] V. Jacobson, R. Braden, and D. Borman, "TCP Extensions for High Performance," IETF RFC 1323, May 1992.

[12] A. C. Clarke, "Extraterrestrial Relays," *Wireless World*, vol. LI, no. 10, Oct 1945, pp. 305–08.

[13] C. E. Fossa *et al.*, "An Overview of the IRIDIUM Low Earth Orbit (LEO) Satellite System," *IEEE Nat'l. Aerospace and Elect. Conf.*, July 1998, pp. 152–59.

[14] T. Pratt, C. W. Bostian, and J. Allnutt, *Satellite Communications*, 2nd ed., Wiley, 2003.

[15] S. Fu, M. Atiquzzaman, and W. Ivancic, "Evaluation of SCTP for Space Networks," Tech. rep. TR-OU-TNRL-05-101, Univ. of OK, Telecommun. & Network Res. Lab., http://www.cs.ou.edu/~netlab, Jan. 2005.

Biographies

SHAOJIAN FU received a B.E. degree in transportation engineering in 1997 and an M.E. degree in systems engineering in 2000, both from Northern Jiaotong University, Beijing, P. R. China. During 2000–2001 he worked with Bell Labs China, Lucent Technologies in the area of network surveillance and performance management. He earned his Ph.D. degree at the School of Computer Science, University of Oklahoma, and currently works at OPNET. His research interests are in the areas of wireless communications, IP mobility, transport protocols, network simulation, and embedded reconfigurable computing.

MOHAMMED ATIQUZZAMAN (atiq@ou.edu) received M.Sc. and Ph.D. degrees in electrical engineering from the University of Manchester, England. Currently he is a professor in the School of Computer Science at the University of Oklahoma. He is Co-Editor-in-Chief of *Computer Communications Journal*, and serves on the editorial boards of *IEEE Communications Magazine*, *Telecommunications Systems Journal*, *Wireless and Optical Networks Journal*, and *Real Time Imaging Journal*. He has guest edited many special issues in various journals, and organized special sessions in conferences. He was technical co-chair of HPSR 2003 and the SPIE Quality of Service over Next-Generation Data Networks Conference (2001, 2002, and 2003). He also serves on the technical program committee of many national and international conferences including IEEE INFOCOM, IEEE GLOBECOM, and IEEE International Conference on Computers and Communication Networks. His current research interests are in wireless, satellite, and mobile networks, QoS for next-generation Internet, broadband networks, multimedia over high-speed networks, TCP/IP over ATM, multiprocessor systems, and image processing. He is a coauthor of the book *TCP/IP over ATM Networks*. He has taught many short courses to industry in the area of computer and telecommunication networking. His research has been supported by state and federal agencies like NSF, NASA, U.S. Air Force, Ohio Board of Regents, and DITARD (Australia). He has over 130 refereed publications in the above areas, most of which can be accessed at http://www.cs.ou.edu/~atiq.

WILLIAM IVANCIC is a senior research engineer at NASA's Glenn Research Center. He has over 23 years' experience in network and system engineering for communication applications, communication networking research, state-of-the-art digital, analog, and RF hardware design and testing, and firmware design. He currently directs hybrid satellite/terrestrial networking, space-based Internet, and aeronautical Internet research. He is leading a research effort to deploy commercial-off-the-shelf technology in NASA missions including the Space Exploration Initiative, the Airspace Systems Program, and the Aviation Safety Program. In addition, he is performing joint research with Cisco Systems on advanced routing research for space-based and aeronautic-based networks. Of particular interest is large-scale secure deployment of mobile networks including MIP, mobile router technology, and ad hoc networking using IPv6 technology. He is also a principal of Syzygy Engineering, a consulting company specializing in communications systems and networking as well as advanced technology risk assessment.