TCP Over Satellite: Getting It Right

Introduction

Using satellite links for corporate and ISP network infrastructures has advantages that are familiar to many. However, as applications move from the traditional transaction exchanges toward Internet/Intranet access, TCP performance over satellite links presents new challenges.

Experience shows that over satellite links, TCP is limited in terms of data-transfer speeds. In addition, TCP over satellite can be relatively expensive in terms of resource consumption.

Since TCP is the de-facto standard protocol for reliable data transfer, it seems we must accept these limitations. Fortunately, that is not the case, and the situation can be improved.

This article will show how usage of TCP over satellite links can be made more efficient using compression and acceleration technologies.

TCP over satellite – what are the issues?

The root of the problem is that TCP was not designed for use over satellite links. Indeed, TCP was specifically designed not to be tied to any specific link technology. The downside to that flexibility is that TCP is not optimal on any of these links. Satellite links in particular prove to be a medium where TCP provides faces extreme challenges.

In this article, we will address only the most prominent causes for TCP's inefficiencies over satellite connections: [1]

Speed-limiting effects

- ?? "3-Way Handshake"
  Every TCP connection is established by a "3-Way Handshake" between the Receiver and the Sender. On satellite links, this fixed overhead means that even very short data exchanges take at least a few seconds to be completed.
- ?? "Slow Start"
  TCP is unaware of the link bandwidth. Therefore, whenever a connection is made, an "acceleration" process called "Slow Start" is performed. This means TCP starts sending at a slow rate and tries faster and faster rates until the link limit is reached. Since this process can take 7-15 Round Trip Times (RTTs), which for a satellite link can mean 3-7 seconds, the link is underutilized, and is perceived by the user as being slow (See figure 1). Moreover, the entire transaction may be over before TCP reaches "cruising speed."

[1] A full description can be found in the form of RCC 2488 from the Internet Engineering Task Force (IETF) widely available on the Internet, e.g. at www.ietf.org.

Flash Networks
Holmdel Corporate Plaza
2137 State Highway 35 N
Holmdel, NJ 07733
(P) 732.203.4060
(F) 732.203.4069
www.flashnetworks.com

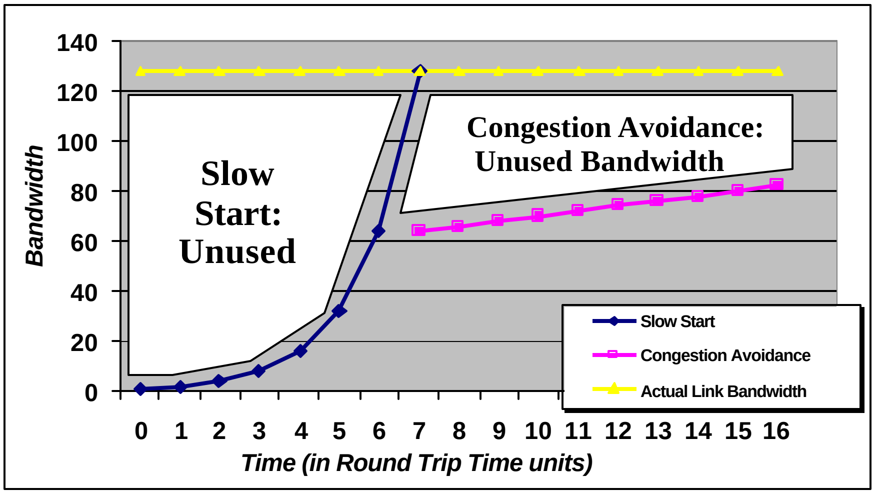

?? "Congestion Avoidance" response to packet loss

TCP assumes any packet loss is caused by congestion. To eliminate the perceived congestion whenever a packet is lost, TCP reduces the current transmission rate by 50 percent. In addition, the "acceleration" method is changed into a much slower system. If the packet loss occurred due to link noise, or other transient conditions, TCP will still react in this manner. If the packet loss occurs early on in the "slow start" process, ramping up to full speed will take significantly more than normal.

?? Reaction time

TCP's feedback loop and reaction time to any event, such as a packet loss, are measured in units of Round Trip Time. On satellite links, TCP's reaction is sluggish, which makes interactive applications problematic, and leads to sub-optimal packet-loss handling.

?? "Advertised Receive Window"

An application using TCP to send data is only allowed to send as many bytes as are explicitly allowed in the last packet seen from the intended destination. This value is known as the "Advertised Receive Window." Typically, this is set to about 8 Kbytes, which severely limits the speeds that a TCP connection can achieve over a satellite. For example, consider a 4 Mbps satellite link with 500ms delay. The sender is allowed to send 8 Kbytes, and then no more may be sent until the next packet from the destination is seen, carrying the next 8 Kbytes allowance. On average, only 128 Kbps will be sent on a 4 Mbps link. Even setting the "Advertised Receive Window" to the theoretical maximum of 64 Kbytes will still limit TCP's speed for satellite links to under 1 Mbps.

?? Asymmetry and TCP transfer speed

TCP can essentially send data only if a constant stream of acknowledgements flow back from the receiver to the sender. When links are highly asymmetric (say more than about 48:1) there is simply not enough "return" bandwidth to keep TCP sending at full speed.

Flash Networks
Holmdel Corporate Plaza
2137 State Highway 35 N
Holmdel, NJ 07733

(P) 732.203.4060
(F) 732.203.4069
www.flashnetworks.com

Resource consumption effects

?? Variable response time

TCP performance is very dependent on measuring the Round Trip Time between Sender and Receiver. On multiple-access satellite networks, such as TDMA or DAMA-based access methods, the time needed to set up the link or solve channel contention can cause TCP to experience a variable Round Trip Time. This causes TCP to assume that packets were lost when they were in fact correctly received. TCP then needlessly retransmits these packets. In one trial on a TDMA network, as much as 13 percent of the traffic sent was actually redundant and therefore unnecessary.

?? TCP acknowledgment content and frequency

TCP's acknowledgement (ACK) mechanism is very simple, and carries little information in each acknowledgement packet. In fact, most TCP implementations generate an ACK packet for every two data packets seen. On a satellite network, especially a TDMA-based one, this is a significant amount of overhead traffic and consumes a significant amount of resources.

Table 1 shows a sample network sizing calculation for a TDMA satellite network used for Internet/Intranet browsing. Almost 70 percent of the required bandwidth is actually used only to carry TCP's acknowledgements – essentially it is overhead.

Description	**Unit**	**Inbound**	**Outbound**	**System**
# of VSAT terminals				400
Transactions every S seconds per terminal	Seconds			60
Packets/Transaction		50	69	
Bytes/Transaction	Bytes	87	1212	
Net Traffic	BPS	231,680	4,460,160	
Nominal Carrier bit rate	BPS	56000	2000000	
B/W per carrier	kHz	100	400	
# of Carriers Needed		27	3	
Total BW Needed	**kHz**	**2,700**	**1,200**	**3,900**

Table 1 - Network Sizing for Web Browsing

The model assumes that at peak time, each user downloads 2 Web pages per minute, composed of 20 objects each, with an average object size of 5 Kbytes. The number of packets and their average size are what would be generated by a typical Browser/TCP combination.

Is TCP suitable for use over satellite links? The answer is a resounding, "It depends." While most TCP-based applications will work to some degree, the performance and efficiency achieved may or may not be considered acceptable and cost effective.

What can be done?

Tuning TCP parameters for satellite links

TCP's performance can be tuned, within limits, to make better use of the available link conditions.[2]

On a satellite communications link, the following parameters should be considered carefully:

- ?? Advertised Receive Window Size (RWIN)
  This value determines how many bytes the receiver allows the server to transmit. When TCP is used over satellite links, this value should generally be set as large as possible. It should be noted that under packet loss conditions, increasing RWIN will have a reduced impact, due to the "Congestion Avoidance" mechanism.
- ?? T2 – "Acknowledgement" or "piggyback" timer
  As previously mentioned, TCP needs a constant stream of acknowledgements to be sent from receiver to transmitter. If traffic exchange is bi-directional, these acknowledgements can be added to the data being sent in the correct direction with no additional overhead. (This is known as "piggybacking," hence the timer's name).

  For request/response situations (e.g., file download, web browsing, etc.), T2 should be set as small as possible to accelerate acknowledgements.
- ?? Maximum Transmission Unit (MTU)
  This value controls the amount of data to be sent in a single data-link frame. In general, this value should be set as large as possible, but no larger than the capability of the data-link technology. The correct MTU depends on the data-link technology. For example on Ethernet LANs the MTU is 1518 bytes.
- ?? Maximum Segment Size (MSS)
  This value is the measure of how much TCP will count as one "unit" in its transmission window calculations (e.g., when performing "Slow Start," traffic is ramped up in MSS units). This value should be set as large as possible, but no larger than MTU minus the bytes consumed by protocol headers.

Giving TCP less traffic to transport over the satellite link

A major source of improvement is to be gained by simply making sure that TCP has a smaller amount of data to transport over satellite links. This is actually easy to do. Here are a few ways to accomplish this:

- ?? Compression

  A compression system directly reduces the amount of traffic that needs to be transported, saving bandwidth, speeding up response time, and shortening transfer times. Compression gains depend on the content transferred. Typical savings for general Internet/Intranet browsing are around 30 percent; for text-based data (e.g., e-mail messages) 60 percent gains are common.

[2] Tuning methods vary according to the operating system.

Flash Networks (P) 732.203.4060
Holmdel Corporate Plaza (F) 732.203.4069
2137 State Highway 35 N www.flashnetworks.com
Holmdel, NJ 07733

?? Caching

Consider a corporate Intranet over a satellite link. It is very likely when a query is carried out, it is a repeat of a previous query. With a caching system, the results of the previous query are saved and are still up to date, the locally installed cache server will respond. This avoids the need to carry the data over the satellite link again and again. Savings of 30 to 50 percent are commonly achieved using cache servers.

TCP "Spoofing" and protocol conversion systems

While compression and caching minimize the amount of data to be transferred, the remaining data is subject to all of the inefficiencies and limitations listed above and must be optimized for satellite transport.

?? Spoofing

A popular approach to this problem is TCP "spoofing", which mitigates some of the problems. A "spoofing" system involves a device at the "near end" of the satellite link pretending to be the intended destination for the TCP connection. Data sent to that device is transferred over the satellite link to the "far end" where a device at the "far end" playing the role of the original source builds a separate TCP connection to the intended destination. "Spoofing" systems are typically available as an integral part (or at least as an option) of satellite equipment. The performance and savings provided vary greatly, and typically users have limited control over them, sometimes limited to a simple enable/disable choice. It is necessary to carefully test these systems for suitability in conditions as close to the intended environment as possible.

?? Protocol Converters

Protocol conversion systems are a form of "super spoofing" that provide all the benefits of spoofing but are more flexible, controllable and tunable, and offer additional features and benefits. One benefit is that they are independent of the satellite equipment vendor, providing full cross-vendor and cross-product compatibility.

These systems work by converting the TCP protocol over the satellite link to another protocol, specifically designed and optimized for satellite data transfer in terms of speed and resource usage. As with spoofing, traffic at the "far end" of the satellite link is translated back into TCP so both the source and destination of the connection can keep using their existing TCP-based software with no changes.

Flash Networks
Holmdel Corporate Plaza
2137 State Highway 35 N
Holmdel, NJ 07733

(P) 732.203.4060
(F) 732.203.4069
www.flashnetworks.com

Protocol conversion systems can operate as a proxy or be embedded into the satellite modem. These systems allow a TCP connection to fully utilize link capacity, thus speeding up data transfer rates by 200 to 500 percent, possibly even more, depending on available bandwidth.

Figure 1 below shows a throughput comparison test for traffic transferred over a satellite link between TCP and a link enhanced with a protocol conversion system.

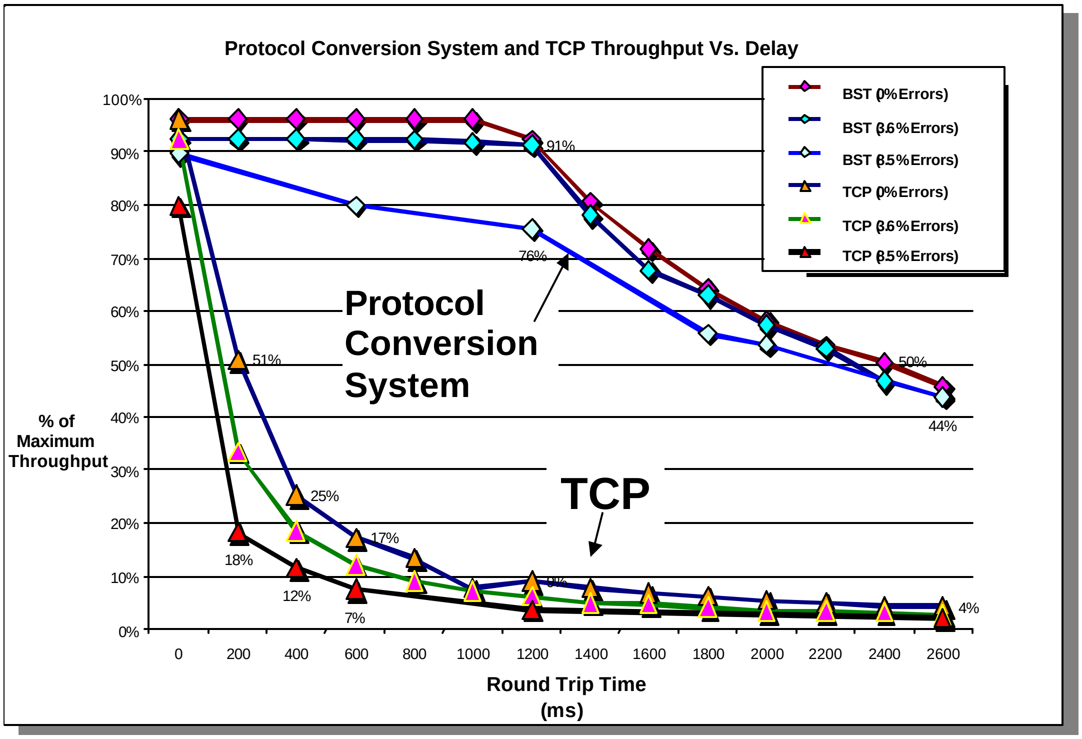

Speed is not enough

Getting a transmission speed to match the nominal link capacity is important and desirable, but by no means the full story. To provide full value, a link enhancement system, such as the protocol conversion system, should provide additional benefits and fulfill several conditions to be of real value. Features to look for should include:

?? Fault tolerance, support for redundant links and load balancing
?? Rate control
?? User mandated bandwidth allocation (QoS)
?? Compression
?? Caching

Cost Justification

While such a system can be easily justified by quantifying the time savings and improved productivity resulting from higher transmission speeds, ideally it could be justified on direct cost savings, especially for TDMA- or DAMA-based networks.

Flash Networks
Holmdel Corporate Plaza
2137 State Highway 35 N
Holmdel, NJ 07733

(P) 732.203.4060
(F) 732.203.4069
www.flashnetworks.com

Consider for example the effects of using a protocol conversion system in the same scenario described in Table 1. This system typically only needs to send about 10 percent of the acknowledgements, when compared with TCP (while keeping full reliability, of course), which has an immediate effect on the bandwidth requirements and consumption, hence, on the monthly communications bill.

This is clearly seen in Table 2, below, where the same traffic model is used as in Table 1, except that the packet counts and average sizes are what typically would be seen using a protocol conversion system.

Description	**Unit**	**Inbound**	**Outbound**	**System**
# of VSAT terminals Transactions every S seconds per terminal	Seconds			400 60
Packets/Transaction		9	69	
Bytes/Transaction	Bytes	191	1212	
Net Traffic	BPS	91,733	4,460,160	
Nominal Carrier bit rate B/W per carrier	BPS kHz	56000 100	2000000 400	
# of Carriers Needed		11	3	14
Total BW Needed	**kHz**	**1,100**	**1,200**	**2,300**

Table 2 - Adding a Protocol Conversion System to Browsing Model

There are significant direct savings in bandwidth consumption, which will translate into a direct cost saving in terms of bandwidth usage.

Conclusion

As applications and traffic over satellite networks move more and more towards an Internet/Intranet model, TCP and its behavior in a satellite-based network is becoming crucial. As has been demonstrated, using various techniques – ranging from tune-ups to protocol conversion – TCP can be made effective and efficient over satellite networks.

Flash Networks
Holmdel Corporate Plaza
2137 State Highway 35 N
Holmdel, NJ 07733

(P) 732.203.4060
(F) 732.203.4069
www.flashnetworks.com