Internet over Satellite Optimization with XipLink White Paper

November 12th, 2004
Release 1

Author: Charlie Younghusband - XipLink Product Manager

Contributors	
Karim Fodil-Lemelin	Kernel Development Lead
Igor Mordkovitch	Proxy Development Lead
Joshua Lamorie	Algorithms Developer
Alex deVries	Algorithms Developer
Emily Birrell	Sales and Marketing Support
Rob Mason	Technical Support

www.xiplink.com

Abstract

This paper provides an overview of Internet over Satellite optimization with XipLink technology. XipLink utilizes three core elements to optimize Internet over satellite: transport layer TCP Acceleration and optimization, HTTP Acceleration with prefetching and high ratio data compression. Engineered together, these technologies deliver a superior DSL-like user experience while making more efficient and effective use of available bandwidth. The XipLink solution is scalable, customizable, TCP/IP compatible and transparent as well as being standards-oriented. It is available as hardware XipLink Gateways or integrated into terminals on a variety of platforms.

Overview

XipLink utilizes three core elements to optimize Internet over satellite: transport layer TCP Acceleration and optimization, HTTP Acceleration with prefetching as well as high ratio data compression. Engineered together, these technologies deliver a superior DSL-like user experience while making more efficient and effective use of available bandwidth.

Satellite communication is a medium that sanctions latency and loss properties unlike the terrestrial Internet. What's more, there is generally a much higher price premium on access to this bandwidth, given the expensive cost of satellites and ground equipment. This makes the business case for performance enhancement to the point where such technology is a virtual requirement for efficient Internet over satellite. The ROI on such technology is generally very short, down to a matter of months.

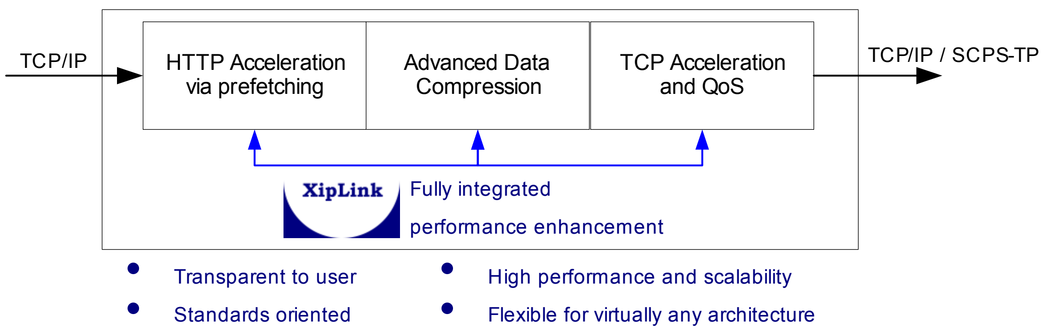

This document introduces the various link enhancement technologies. The full integration of these technologies working together permits exceptional performance. Such technology is not magic; the intent in part of this document is to better describe the key elements required of a high performance PEP and the benefits of such technology with XipLink.

For maximum performance and flexibility, XipLink technology is installed at both ends of the satellite link, bracketing the link. TCP/IP traffic is transparently intercepted from a network and XipLink performance enhancement is applied in concert with another XipLink Gateway at the other side of the link. The protocols used are generally open-standard and completely dynamic with no configuration required.

HTTP Acceleration

Even with a TCP-level fast connection start, browsing through satellites can result in poor performance despite the broad bandwidth offered by satellite links. The core reason for this is the interactivity required by the Hypertext Transfer Protocol (HTTP) to download pages with many embedded objects. These are most frequently images on media rich web sites, a very common occurrence among the top web sites where there can be twenty or more objects. Graphic artists frequently use images to replace text so they can fully control the end result despite differences in font utilization by various browsers and operating systems.

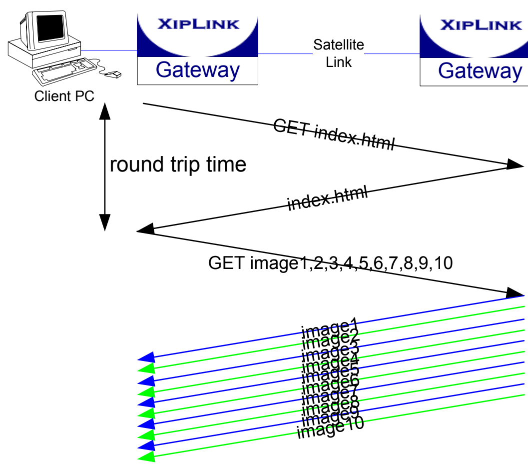

Figure 1: XipLink HTTP Acceleration

A web browser's page download process essentially consists of creating a TCP connection to the server and downloading a primary HTML page from the HTTP server. The web browser parses the HTML and the objects within the page are then requested from the server one by one.

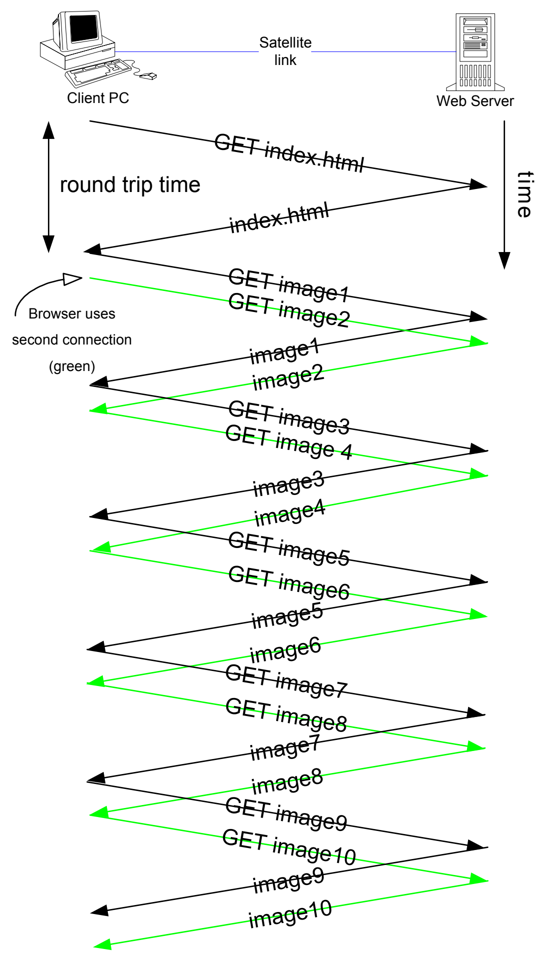

Figure 2: Normal HTTP Page Download

Later versions of HTTP allows the use of 'pipelining' which allows existing connections to be re-used and objects to be sent in over the same TCP link. This represents a substantial improvement over HTTP 1.0 where separate connections are used for objects, each being subjected to TCP's 3-way handshake (although this is corrected with XipLink's Connection Fast Start).

XipLink's HTTP Acceleration can offer dramatically improved page download time. XipLink is able to prefetch images from multiple servers, which is common for load-distributed web servers where some images or other content may come from different web servers to make up a complete page. XipLink's prefetching algorithm is also attractive as it hands images to the client browser as they are received; a user will immediately see the HTML content and images as they are downloaded, as opposed to a packing strategy where the page is only displayed when everything has been downloaded.

Data Compression

XipLink offers streaming data compression on the TCP payload. This compression is separate from the TCP Header Compression, which is a transport-layer SCPS-TP optimization. To achieve a high compression ratio XipLink performs compression on large chunks of data at one time. XipLink uses a tuned algorithm to ensure that additional latency is not introduced through the buffering required for data compression. For example, if only a few hundred bytes of data were received during communications, it would be briefly buffered. The compression algorithm would then be performed on the available data in a matter of milliseconds and the data output towards the satellite link.

This compression is unlike any other payload compression strategies that only compress data by individual IP packets yielding lower compression ratios. Compression is most effective when performed on larger data sets as more patterns inside the data can be found and removed.

Compressing the data prior to TCP/SCPS-TP output also proves advantageous in that output algorithms are applied to the compressed data, so the output segments to the satellite link are still consistent regardless of the compression ratio. More directly, XipLink will continue to output at the prescribed rate (either as prescribed by rate control or estimated by Vegas) regardless of whether the compression ratio varies from packet to packet.

There are other compression strategies that can yield higher compression ratios, such as by putting all traffic into a common compressing tunnel.

However, such approaches are a poor choice for satellite networks given the loss coupled with the latency of retransmissions. XipLink is engineered to be loss and latency tolerant.

XipLink's Data Compression and RFC 2393 – IP Payload Compression Protocol (IPComp) compression strategies can be used in combination. Using IPComp, one can output XipLink TCP compressed packets into a compression tunnel as well. IPComp is nice in that compression does not need to be performed if the data is not compressible. IPComp can compress UDP packets, although bandwidth intensive UDP traffic such as RTP usually enjoys only a minimal or zero compression ratio due to the high degree of pre-existing compression of its payload. Besides a lower compression ratio, IPComp however introduces unnecessary overhead due to encapsulation overhead and additional management issues, whereas XipLink's compression is negotiated dynamically.

XipLink's compression algorithm utilizes minimal memory and can be tuned to match the available computing resources. Decompression utilizes little or no additional memory utilization and considerably less computing power than compression so it is ideal in resource constrained remote terminals that typically receive more data then they send to the satellite link.

The resulting compression is completely dependent on the nature of the data stream. There is a great deal of excessive marketing with respect to compression technologies – the fact is that it depends completely on the data. Random data is not compressible. Text is highly compressible at ratios well over 10 – 1, often 40 times. Most XipLink users see between a 40% and 300% compression ratio depending on the usage patterns. This allows more users to be fit into the same satellite bandwidth and data is transferred more quickly.

TCP Acceleration Techniques

Introduction

A detailed description of the transport layer enhancements (also known as "TCP Acceleration"), SCPS-TP and the resulting is provided by the sister document, SCPS-TP Transport Layer Optimization with XipLink. However this section will provide an overview.

XipLink optimizes the transport layer for satellite communications in several ways. The conceptual algorithms and air interface are based largely on the SCPS-TP standard. XipLink's implementation of TCP acceleration uses advanced algorithms that are the result of years of engineering and advanced simulation to ensure optimum data transport.

XipLink addresses the loss and latency of wireless communications by applying algorithms and other optimization techniques to optimize data transfer.

Here are some problems with TCP when used over satellite:

- Slow connection setup
- Slow to discover available bandwidth
- Erratic approach to available bandwidth
- Slow to respond to loss
- Over reaction to loss at higher bandwidths by dramatically reducing output rate
- Hard coded limitations such as window sizes
- Difficulties in responding to burst loss of many data stream 'holes'
- Protocol header overhead
- Considerable bandwidth required for acknowledgements
- No build in quality of service, and normal QoS schemes may not normally integrate with TCP Acceleration

XipLink mitigates these issues with a combination of approaches that work together and offers several congestion controls strategies that are more appropriate to satellite communications. XipLink's loss recovery scheme is highly efficient for lossy links and works with the congestion control algorithms to adapt appropriately to loss. Header compression and acknowledgement frequency reduction offer low protocol overhead. Large windows that scale with the network load ensure there are no other limits on communications and latency is not further introduced by using buffers that are too large. TCP/IP and network transparency is maintained through the use of SCPS-TP.

⇒ Bandwidth Rate Control

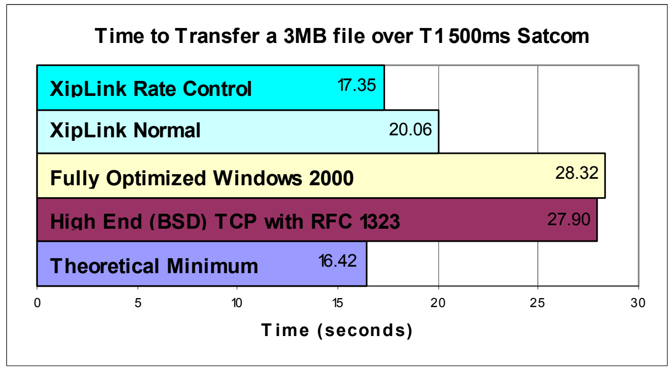

When the bandwidth is known, rate control simply makes full use of the link capacity. Rate control allows you to configure the outbound traffic rate to a fixed rate, whether there is only one connection or hundreds of users. Rate control is fundamental for combating poor performance due to latency of satellite communication, so that performance does not decrease as the round trip time climbs even into the seconds range. XipLink Normal, also shown, is used when the bandwidth is unknown or varies. It uses the TCP Vegas algorithm to dynamically discover the available bandwidth. Note: we do not show the benefit of stream compression in this graph since compression is dependant on the randomness of the data.

⇒ Operation on Shared or Dynamic Bandwidth Networks, TCP Vegas

TCP Vegas, like TCP, uses an adaptive congestion control algorithm but one that is ideally suited for satellite networks. XipLink's implementation of TCP Vegas is the result of years of development. It uses a superior algorithm that is well suited to satellite networks to determine the available bandwidth and respond dynamically. Vegas measures buffering occurring within the network and controls the flow rate accordingly. It is usually able to anticipate and reduce its rate before congestive loss occurs. Vegas is an excellent alternative to TCP for TDMA and other variable bandwidth links.

⇒ Optimization for Dynamic Bandwidth Systems: Dynamic Rate Control

The rate control setting can be varied multiple times per second. When integrated with a satellite modem or other external mechanism, this variable rate control permits coordination with the available bandwidth to the terminal through varying levels of network load. For example, the rate control setting can be varied using the size of the outbound buffer as a feedback system to ensure full use of available bandwidth without driving the link into loss due to buffer overflow and network congestion. QoS mechanisms can also be fully integrated. Full performance on variable bandwidth links can still be achieved.

⇒ Fast Connection Start

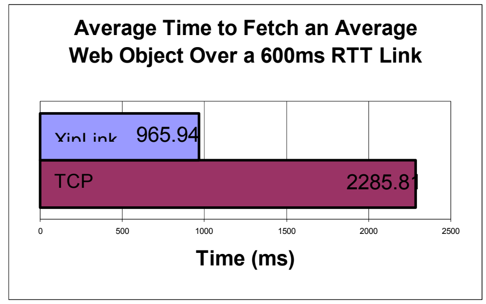

Using XipLink fast start technology, data requests can be sent in the initial connection opening, allowing for single round trip time web object fetches and server network application while maintaining data integrity. XipLink's fast start is based on the efficient integration of RFC 1644. Standard TCP meanwhile uses a 3-way handshake before data communications can begin. The benefits of fast start are most apparent for communications where multiple connections are used together in sequence, such as HTTP and various client-server applications.

⇒ Selective Negative Acknowledgements

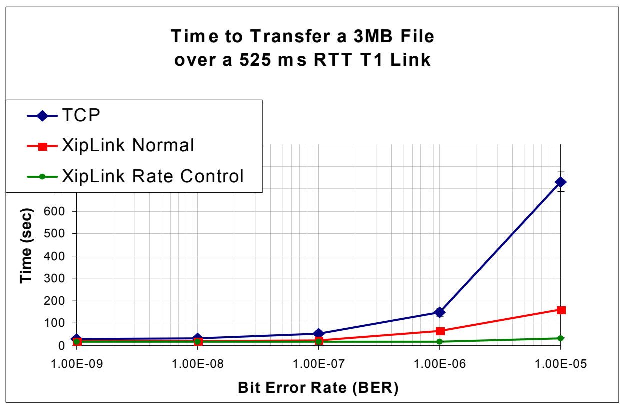

Selective Negative Acknowledgements (SNACK) allow for swift response to packet loss. Communications can continue robustly even at very high bit error rates. This makes XipLink Accelerators very resilient to rain fade and other weather conditions. SNACK is more bit-efficient than SACK and is more effective against multiple losses in a bandwidth-delay product, which is common with wireless communication interfere loss.

⇒ Quality of Service Support

XipLink offers several different Quality of Service (QoS) capabilities for maximum flexibility. XipLink's TCP QoS support connections meeting a matching criterion are offered preferred access to the available bandwidth. This is an ideal solution for ensuring mission critical or performance sensitive TCP applications receive the bandwidth they need. Performance sensitive traffic, such as RTP traffic, also receives priority access to the link ensuring performance is retained. Data passed through the gateway without acceleration, like UDP traffic, can be rate limited. XipLink can also be integrated with other QoS systems to coordinate differentiated users classes and or to limit bandwidth abusers. XipLink also makes a variety of other conventional QoS disciplines available.

⇒ TCP Header Compression

TCP Header Compression as specified by SCPS-TP works by removing redundant TCP header data. This results in approximately 40% savings on your TCP headers, removing communications overhead. While the protocol overhead savings for a TCP sender is only fractional, this offers considerable benefit to reducing the overhead introduced by a TCP receiver sending only acknowledgements.

⇒ Acknowledgment Frequency Reduction

On highly asymmetric links enabling this feature will reduce the number of acknowledgments the receiver will send back as data is received. When it is activated, it will use the delayed acknowledgement time instead of the TCP-described every second packet. Acknowledgements can be reduced to as much as is practical.

Other Internal Technology Benefits

⇒ Dynamic Buffering

XipLink's dynamic buffer allocation allows for full bandwidth scalability while minimizing delay at all load profiles. Statically allocated buffers introduce buffering delay when the connection count increases or cannot scale to make full use of the bandwidth when there are few connections. XipLink's Accelerators use advanced memory management technology that can scale efficiently from a 486 processor to tens of thousands of simultaneous connections at over 45Mbps for hub servers.

⇒ Burst Connection Handling

XipLink's burst connection handling accommodates periods of high loads through highly efficient I/O techniques. For example, even on an already loaded system, high end XipLink Gateways can handle well over 500 new connections per second per CPU.

⇒ Operational Modes

XipLink Accelerators can operate in two different modes: as a gateway or router or as "bump-in-the-wire" (bridge). These two different modes allow for even greater flexibility when deploying in existent networks or when building new ones. XipLink can also offer a variety of automatic routing algorithms such as OSPF and BGP.

⇒ Gateway and Network Appliance General Features

XipLink can be remotely configured, managed and monitored using SNMP, web interfaces and SSH. XipLink Accelerators have multiple ways of handling failure, such as a redundant configuration and fail-to-wire.

Standards Based: SCPS-TP

XipLink provides the highest performance Space Communications Protocol Specification – Transport Protocol (SCPS-TP) gateway solution. It has been demonstrated to adhere to the specification by 3rd party inter-operability tests. SCPS-TP is further recognized as an ISO and CCSDS standard and most recently by the SatLabs group (www.satlabs.org) as the basis for the I-PEP standard.

Because SCPS-TP is essentially an 'upgraded' TCP, there are many other benefits to its use. For example, its transparency inside the network allows it to be analyzed using conventional techniques during the course of network operations, unlike tunneled protocols like XTP. It can also be distributed for servers located at Internet gateways as well as corporate gateways in conjunction with VPN technology. Of course, XipLink is also fully transparent to end-users. Using an upgraded TCP also allows it to be deployed in mixed environments where there are not necessarily XipLink or other SCPS-TP at all nodes, while realizing partial benefits can be realized by the one-sided PEPs and full benefits by cooperative PEPs on each node. It also allows acceleration between all nodes on meshed networks.

SCPS-TP is internationally recognized as ISO recommendation 15893:2000 and Consultative Committee on Space Data Systems CCSDS 714.0-B-1 and MIL-STD-2045-44000 among others. It has been recommended as the standard technique for performance enhancement by the U.S. Department of Defense for MILSATCOM IP communications.

Combined Results

The integration of TCP Acceleration, HTTP optimization and data compression that XipLink offers is state of the art and second to none. The following diagram represents a simple case of downloading a basic web page with XipLink.

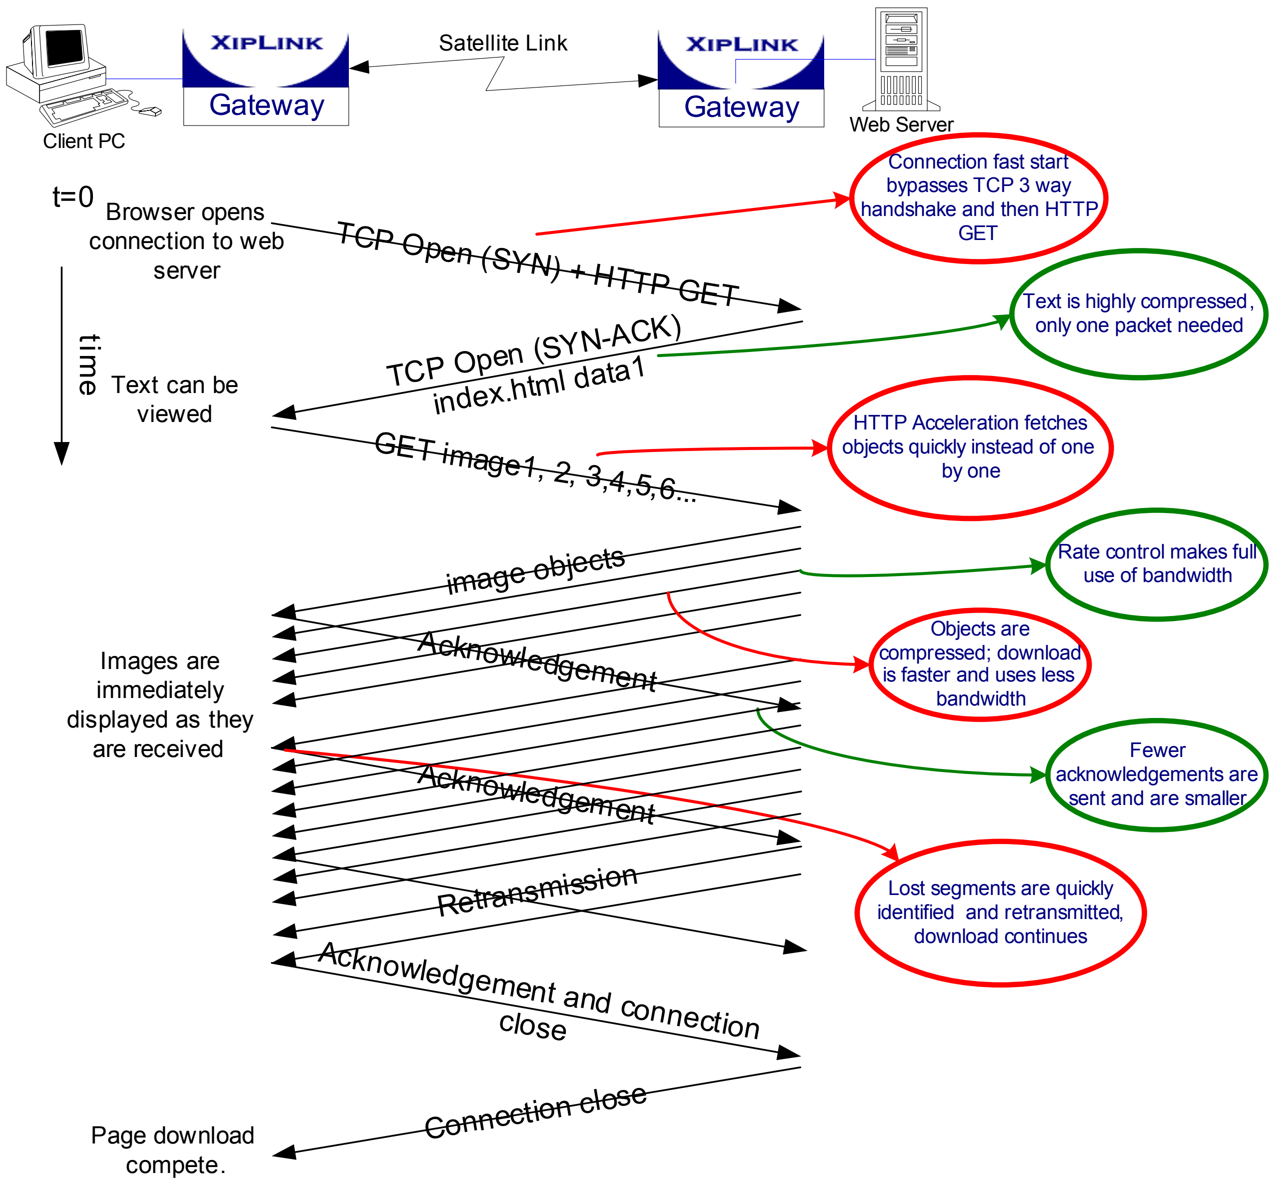

Figure 3: Example Web Page Download with Impact of XipLink Technology

Glossary

CCSDS	Consultative Committee on Space Data Systems
DSL	Digital Subscriber Line
HTTP	Hyper Text Transfer Protocol
IETF	Internet Engineering Task Force
ISO	International Standards Organization
IP	Internet Protocol
MIB	Management Information Base
NAT	Network Address Translation
PEP	Performance Enhancing Proxy
QoS	Quality of Service
RTP	Real Time Transport Protocol
SACK	Selective Acknowledgements
SCPS-TP	Space Communications Protocol Specification – Transport Protocol
SNACK	Selective Negative Acknowledgements
TCP	Transmission Control Protocol
UDP	Unicast Datagram Protocol