Tom Normand, Director of International Sales and Marketing for Los Angeles-based Mentat Inc., investigates the challenge presented by data communications via satellite.

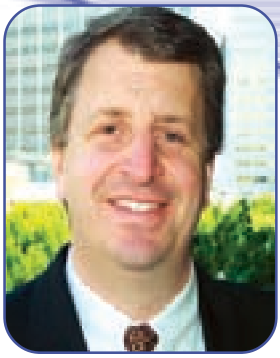

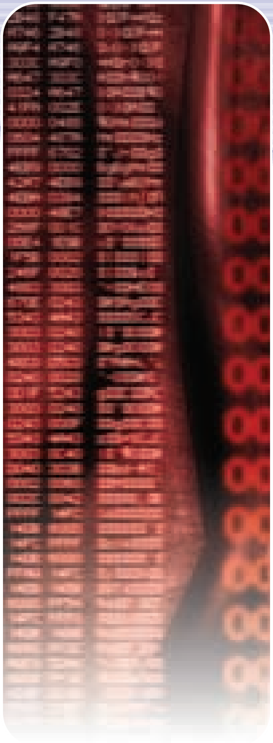

Delivering IP traffic

With the growing importance of Internet access for businesses, governments and consumers, satellites are playing an increasingly critical role in providing Internet Protocol (IP) connectivity across Asia and around the world.

However, the long transmission delay and other conditions unique to satellite communications reduces the efficiency of the IPs used to transmit data over computer networks, leading to sluggish performance and slow downloads. This often creates a perception that satellite links are slow and unsuitable for data networking.

Fortunately, these performance challenges can be overcome through the use of protocol gateways which take advantage of a variety of performance-enhancement techniques to make satellite networks seem as fast and responsive as terrestrial networks.

Satellites and TCP

Satellites are able to provide ubiquitous coverage across the globe and are often the only feasible solution for building voice and data links to the remote villages, islands, offshore drilling platforms and many other areas across Asia where fibre does not reach.

Using geosynchronous communications satellites orbiting at an altitude of 36,000km, the Radiofrequency (RF) signal takes more than half a second to bounce off the satellite, reach the destination, and a response to return. In contrast, round trip times on terrestrial Wide Area Networks (WANs) average about a tenth of second and local area networks usually have round trip times of less than a thousandth of a second.

The wireless signal can also become corrupted as it passes through the atmosphere, adding bit errors in the data stream. This makes it necessary to verify the validity of all data that arrives and retransmit any lost or corrupted packets. In addition, especially for consumer Internet connectivity, satellite networks often use asymmetric bandwidth, with a large pipe to the user and a thin return channel.

Transmission Control Protocol (TCP) is the protocol used for the reliable transmission of files, images, e-mail, Web content and other data over IP networks. The long delay, bit errors, and asymmetric bandwidth of satellite networks interact with the window sizing, congestion avoidance algorithms, and data acknowledgment mechanisms of TCP. Together, these can limit the performance and efficiency of TCP when run-

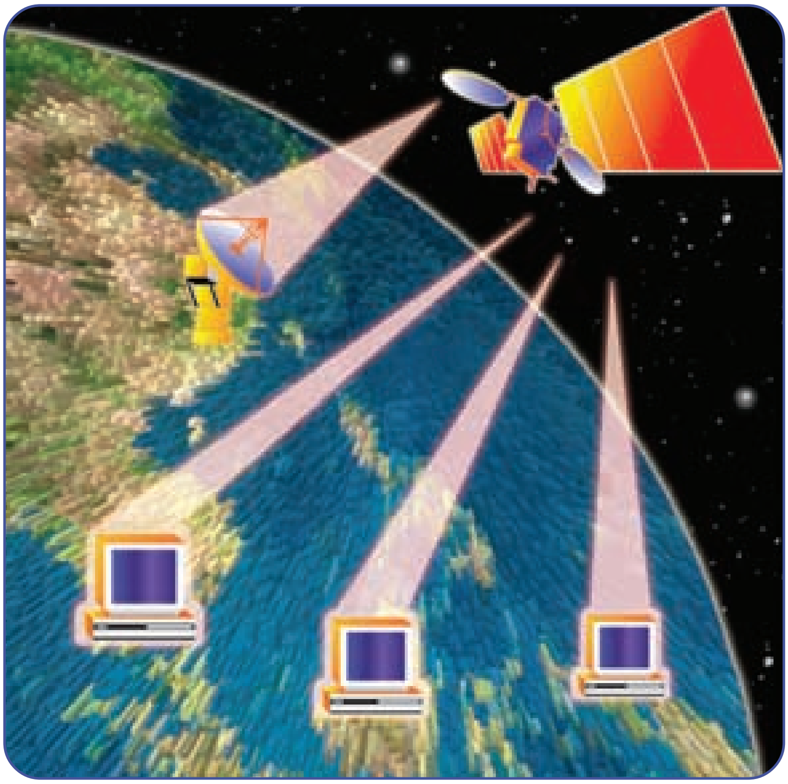

Internet via satellite. Image courtesy of Mentat Inc.

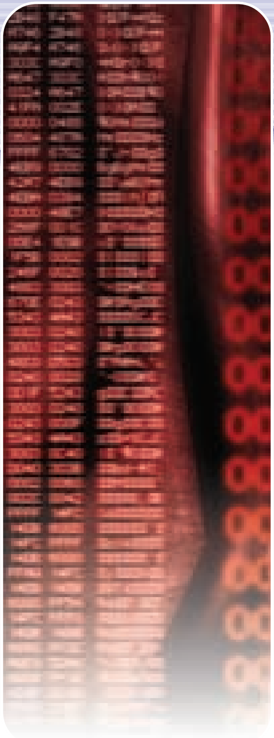

Figure.1. Protocol gateway network integration

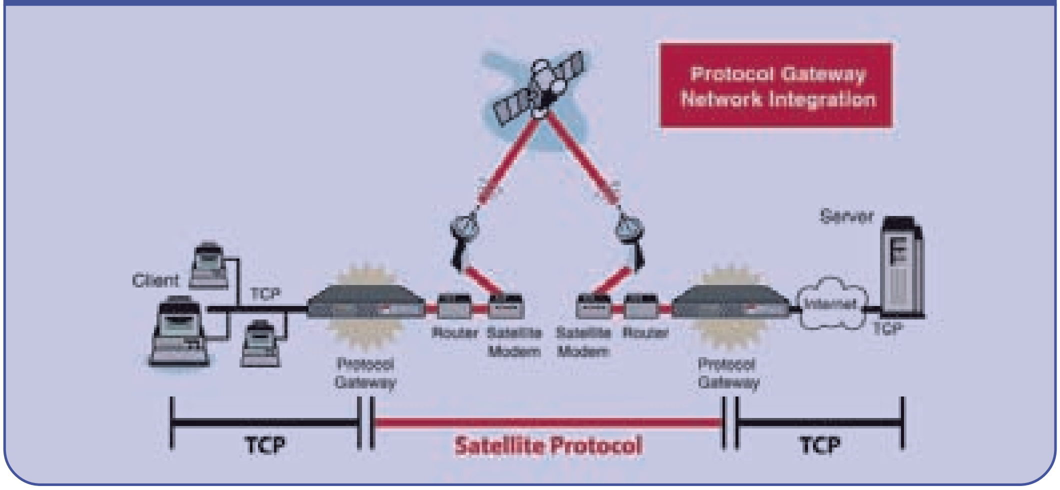

ning over satellite, which users then perceive as unresponsive networks and slow transfers.

The most obvious limitation to users is a result of the interaction of TCP's windowing mechanism with the satellite delay. In order to guarantee that all data eventually arrives without corruption, the sender retains a copy of the original data in a buffer until it gets an acknowledgement from the receiver that the data has arrived. If the sender does not receive an acknowledgement, it transmits the data again. Because computers had limited memory available at the time TCP was originally designed, TCP only buffers a small amount of data for each connection while it waits for acknowledgements. The size of the buffer depends on the operating system, but is usually pre-configured to a value of between 8kbyte and 64kbyte. Once the buffer becomes full, TCP stops transmitting until acknowledgements return and data can be removed from the buffer.

This windowing algorithm works well on low delay links since the acknowledgements return quickly and the buffer never fills up. Over satellite links, where acknowledgements are slow to return, this buffer is usually much too small, especially for high speed links and broadband networks, setting a limit on the maximum throughput rate.

As an example, with a 16kbyte buffer, TCP will fill the buffer then wait half a second for the acknowledgements to return before sending another 16kbyte. This limits the throughput to 32kbyte each second, equivalent to 256kbit/s. No matter how much bandwidth is available, with a 16kbyte window, TCP is limited to a maximum throughput of 256kbit/s per TCP connection. The remainder of the bandwidth will remain unused and wasted.

Through the use of recent extensions to TCP that are available on new operating systems, the buffer size can be adjusted up to approximately 1Gbyte, but this requires modification of registry variables or operating system parameters on both client and server. Each application also has the option to select its own buffer size, overriding the pre-set parameters.

In addition to the window size limitation, TCP also assumes that all data loss is the result of a congested network. Even when data is lost instead due to bit errors, TCP reduces its transmission rate in an attempt to relieve congestion. For this reason, TCP performance is very sensitive to any bit errors on the satellite link. While this problem can be alleviated through the use of forward error correction to reduce the bit error rates, high levels of FEC reduce the amount of bandwidth available.

Further, TCP depends on frequent acknowledgements returning from the receiver to the sender to verify that data has arrived and adjust the transmission rate to avoid congestion. The small return channel common for consumer satellite networks is often not large enough to support this steady stream of acknowledgements, and this can become the bottleneck of the system.

On top of the TCP limitations, Web page design is also not conducive to quick display over long-delay networks. Most Web pages include large numbers of graphical elements and other embedded objects, each of which must be downloaded separately. Over a long delay link, the process of requesting, receiving, and acknowledging each small object, one at a time, can be very time consuming, rendering Web browsing particularly slow over satellite.

Fortunately, protocol gateways are able to take advantage of a variety of TCP, Web, and multicast optimisation techniques to overcome these performance limitations and allow users to take best advantage of the expensive satellite bandwidth.

TCP optimisations

In order to overcome TCP's window size limitation as well as other inefficiencies of TCP over satellite, protocol gateways intercept TCP traffic and transmit the data over the satellite using a protocol especially designed to operate efficiently under typical satellite con-

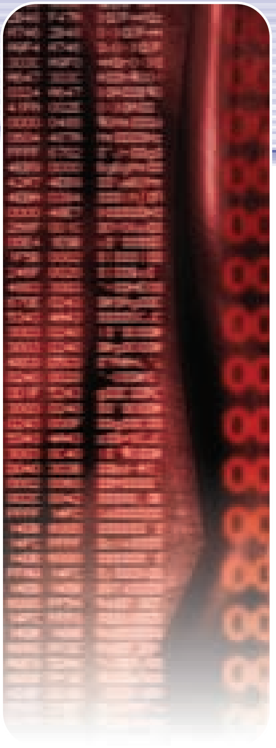

ditions. This process is known as protocol acceleration or TCP Performance Enhancing Proxy (PEP).

As shown in Figure 1, protocol acceleration works by intercepting the TCP connection from the server. The data is then transmitted over the satellite link using the satellite-optimised protocol. The protocol gateway on the opposite side of the satellite link places the data in a new TCP connection for delivery to the client.

In order to obtain ideal performance under satellite conditions, satellite protocols use much larger buffers than TCP to ensure that the buffers do not become full when transmitting at the maximum speed of the link. Other differences from TCP usually include a more efficient acknowledgement algorithm to quickly signal for retransmission of any corrupted data and a rate control mechanism which can find an appropriate transmission rate without depending on a steady stream of acknowledgements.

The effect of protocol acceleration is to allow the TCP connection to fully utilise the available bandwidth regardless of the TCP buffer size on the end nodes. For example, a file transfer over a 2Mbit/s link with a 16kbyte buffer would be limited to 256kbit/s with TCP. By placing a protocol gateway on either side of the link, the same connection would be able to fully utilise the bandwidth, providing eight times faster transfers.

Protocol acceleration is classified into two different categories: proxy-based systems and transparent systems. Proxy-based protocol gateways enhance only particular applications, usually Web, E-mail, and FTP. Transparent solutions usually enhance all TCP traffic, which in addition to Web, E-mail, and file transfers also includes many corporate applications such as remote access, database connectivity, group collaboration software, Windows file sharing, CRM, and ERP systems. Protocol acceleration delivers the largest performance gains when transferring large blocks of data over high bandwidth or high error links.

Another TCP optimisation frequently available as part of protocol gateways is data compression. Because satellite bandwidth is usually the most expensive component of the network, compression can be useful for saving bandwidth as well as improving performance on congested links.

Compression functionality intercepts the data within the protocol gateway and compacts it prior to transmission over the satellite link. The data is then decompressed in the protocol gateway on the opposite side of the link and returned to its original format. This process is entirely transparent to the end users and unlike lossy compression techniques, does not remove any of the data. The benefits provided by this functionality depend on the compressibility of the data, with compression ratios of 2:1 typical for a mixture of Internet traffic.

Compression can be especially effective at reducing bandwidth costs by providing throughput rates greater than the link speed, allowing the operator to reduce bandwidth consumption or support larger numbers of users. Compression can enhance all TCP applications sending compressible data. Pre-compressed data including zip files, executables, and most graphics formats can not be further compressed.

Web Optimisations

To enhance the Internet browsing experience, protocol gateways usually include HTTP-specific optimisations such as Web prefetch functionality. Most Web pages consist of an HTML file, which includes the text and formatting instructions, along with links to the graphic elements and other embedded objects on the page. Web pages usually include many embedded object, each stored as a separate file and downloaded using a separate TCP connection. Although most objects are small and do not consume much bandwidth, the satellite delay slows the process of setting up each new connection, requesting the object, and downloading the object before moving to the next object. With 30 to 80 embedded objects per page downloaded one after another, satellite Web browsing can feel painfully slow.

Web prefetch enhances Web browsing performance by having the protocol gateway on the same side of the satellite link as the Web server pro-actively retrieves all of the embedded objects on the page. The protocol gateway then pushes all of the objects over the satellite link to the protocol gateway on the remote side of the network. The browser-side protocol gateway can then serve the objects when requested by the browser, avoiding the satellite delay. The use of prefetch functionality combined with protocol acceleration and other HTTP optimisations can reduce the time to display a typical Web page by more than three times.

Multicast optimisations

Due to the broadcast nature of satellite communications, the ability of satellites to simultaneously multicast to large number of receivers is a strategic

Figure.2. Time saving using Multicast Fan-Out

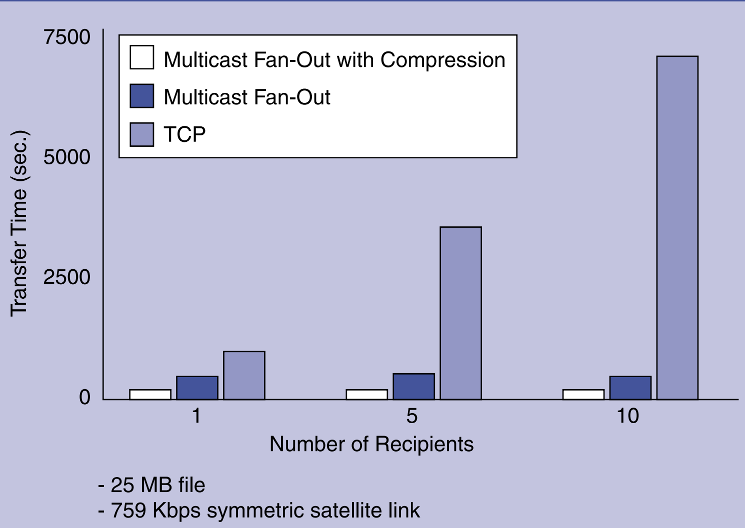

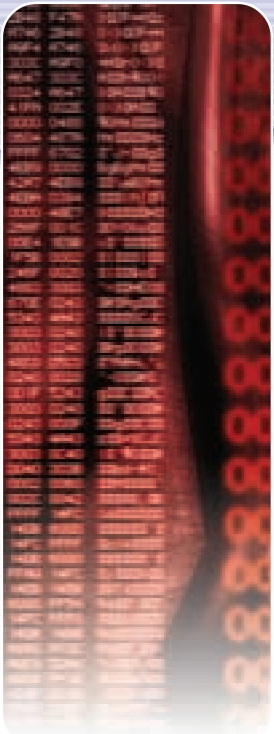

A heavy traffic of data. Image courtesy of PhotoDisc

> **"Protocol gateways take advantage of a variety of TCP, Web, and multicast optimisation techniques to overcome performance limitations."**

advantage that terrestrial networks have difficulty matching. However, multicast over satellite is hindered by two factors: the lack of a multicast-enabled terrestrial network and the absence of reliable multicast functionality within IP.

Since the end nodes of the transmission are usually located on a terrestrial network, the client and server applications are usually limited to unicast transfers, making it difficult for users to take advantage of the ability to multicast over the satellite portion of the network. Further, IP Multicast is based on UDP and only provides a non-reliable, best-effort service, which is not suitable for file transfers and other data applications.

One solution which overcomes both of these limitations is called Multicast Fan-Out. This offers the functionality of reliable multicast in a manner that is compatible with existing unicast-based TCP applications and easy to integrate with the terrestrial network.

Multicast Fan-Out works by transparently converting TCP unicast connections into reliable multicast transfers over the satellite network. The conversion from unicast TCP to reliable multicast protocol is performed by a protocol gateway at the hub. The actual content can be located on any content server, and a multicast network is not needed between the content server and the protocol gateway.

The protocol gateway then transmits the data over the satellite using a specialised reliable multicast protocol to any number of receiving protocol gateways. These receiving protocol gateways on the opposite side of the satellite link then repackage the data in a new unicast TCP connection and relay it to the actual destination device on that network.

Because the Multicast Fan-Out process is transparent to the end devices, any machine can originate or receive the multicast transfer without requiring any specialised multicast software, making Multicast Fan-Out compatible with FTP and other TCP applications which push data to a receiver.

The chart illustrates that Multicast Fan-Out can reduce the time required to deliver a file to multiple clients while slashing bandwidth costs. Over a 750kbit/s satellite link, sequential TCP transfers require 7300 seconds to deliver a 25Mbit/s file to ten remote sites. Using Multicast Fan-Out, the same file can be transferred to all sites in only 280 seconds. For compressible text files, the transfer time is only 66 seconds to transfer the file to all sites, less than one per cent of the time and bandwidth required by TCP.

Multicast Fan-out is ideal for file transfers, cache replication, video file distribution, content delivery networks, database replication, and other situation where it is desirable to distribute the same file or data to multiple sites.

Protocol gateways in Asia

Lawrence Tan, Managing Director at Singapore-based Suntze Communications finds that: "End users are becoming more knowledgeable when it comes to reliable data transfer over satellite. They realise that solutions like protocol gateways are available which overcome most of the problems associated with satellite communications. Governments, enterprises and consumers alike expect much better performance over their satellite links."

As Asia recovers from its recent economic malaise and computer networking becomes even more critical to their operations, organisations are again making the investments necessary to upgrade their network and improve the performance of their links.

David Atkinson, Sales Manager of STEP Electronics in Melbourne, Australia explains, "As the economic conditions in Asia improve, we expect demand for TCP/IP acceleration solutions to expand at a rapid clip. More end users understand that a network's performance over satellite links can, in many cases, be dramatically improved."

Michael Hou, Vice-President of Glocom, Inc. based in Beijing concurs. "There continues to be a gradual increase in demand for sophisticated, enhanced methods of delivering data reliably in mainland China, particularly in the commercial sector."

So while data networking over satellite faces significant performance challenges, protocol gateways help overcome these obstacles to provide a level playing field with terrestrial solutions. ■

Tom Normand can be reached at: tom@mentat.com.