JUPITER™ System Bandwidth Efficiency

It is well understood by satellite network operators that "bandwidth efficiency" for a VSAT system is a critical element in achieving profitability, as higher efficiency will increase the amount of traffic supported over a given amount of satellite capacity. Not always understood, on the other hand, is that spectral efficiency alone, measured in bits per hertz, is only one factor in determining overall IP network efficiency or net data throughput for a given satellite bandwidth. As illustrated in Figure 1, each of the seven layers of the entire OSI model for networking offers the potential to implement efficiency gains, beyond just the spectral efficiency techniques applied at the physical layer. With this in mind, Hughes has implemented efficiencies across all layers of the OSI model to yield the highest possible overall IP network efficiency.

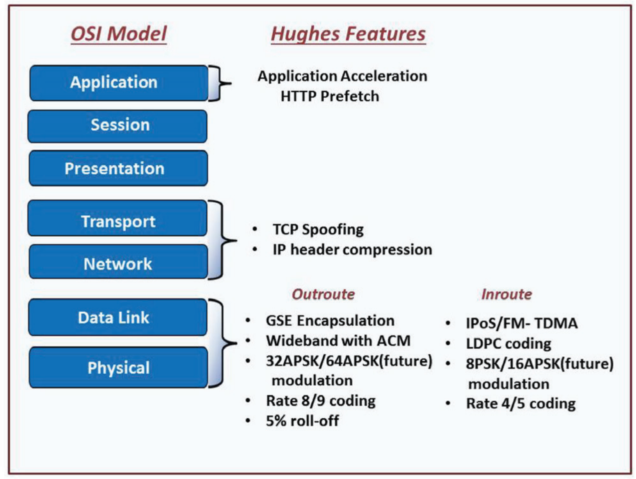

Figure 1. OSI Seven-Layer Network Model

In addition, the JUPITER™ System implements various compression schemes to reduce the amount of data that is transmitted across the space segment. Redundant data streams always exist and are intelligently removed before transmission over the satellite. If a stream of data has been transmitted in the past, a much smaller representative set of data is sent along with a lookup location of past packet streams. Leveraging these techniques results in a reduction factor from a given number of "router bits" (those packets delivered to router) to the least "satellite bits" (those transmitted through satellite).

Physical Layer Optimization

The Hughes JUPITER System utilizes the DVB-S2X wideband forward channel standard with Adaptive Coding Modulation (ACM), supporting rates from 1 Msps to 225 Msps. DVB-S2X is the latest industry standard from the DVB standards committee and is generally recognized to provide higher efficiencies over the DVB-S standard. As illustrated in Figure 2, these efficiencies generally come from the introduction of more granular modcods (modulation and coding combinations), which enable JUPITER to achieve performance very close to the theoretical Shannon curve. The DVB-S2X forward channel supports QPSK, 8PSK, 16APSK, and 32APSK, as well as 64APSK (future capability) modulation schemes, and can operate with a 5% rolloff, thereby enabling very high efficiency.

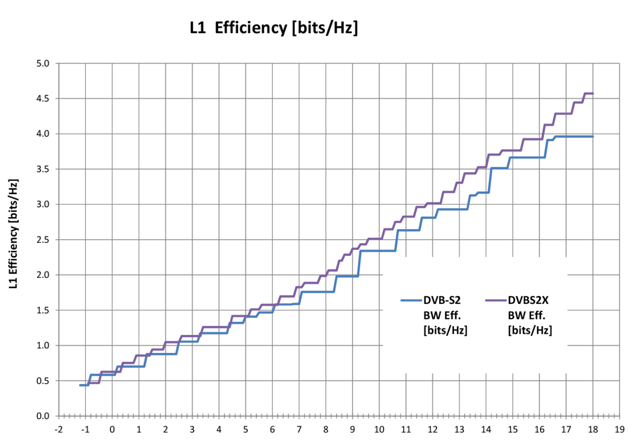

Figure 2. DVB-S2X vs. DVB-S2 Efficiency

Another key efficiency of the JUPITER forward channel is the ability to support wideband channels of more than 200 Msps. As compared to a legacy VSAT system using DVB-S2 with 20% rolloff, the JUPITER forward channel with a single 235 Msps carrier on a 250 MHz channel with 5% rolloff will provide more than 10 percent efficiency. The 10% efficiency gain comes from a combination of lower output backoff, lower channel spacing and lower channel overhead (one channel versus five channels). Figure 3 illustrates the DVB-S2 solution using multiple carriers versus a single Hughes DVB-S2X with one carrier.

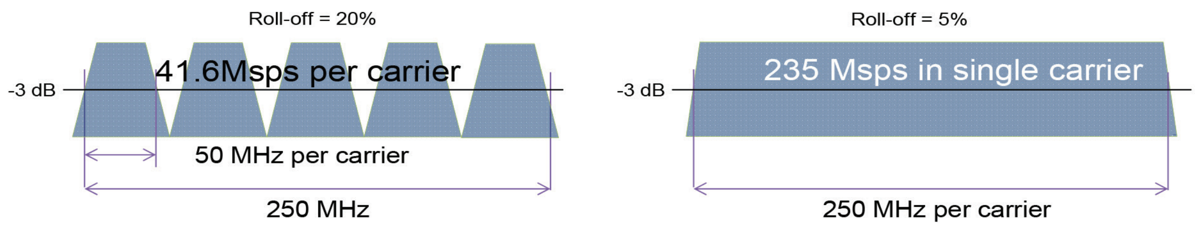

Figure 3. DVB-S2 Multiple Carries vs. Hughes DVB-S2X Single Carrier

Another significant benefit of a wideband forward channel is the ability to deliver better statistical multiplexing. As illustrated in Figure 4, depending on the distribution of required service plans, the statistical multiplexing utilization improves as the data rate of the channel increases. Assuming a link budget closure at 8PSK with FEC Rate 2/3, a standard DVB-S2 50 MHz carrier would achieve a data rate of 83.2 Mbps, while a JUPITER 250 MHz carrier would achieve a data rate of 450 Mbps. In the illustration, we can plot that a 15 Mbps service plan at 70% maximum rate during busy hours shows utilization of 80% for standard versus 95% for JUPITER, thereby demonstrating the greater statistical multiplexing efficiency of the JUPITER wideband carrier.

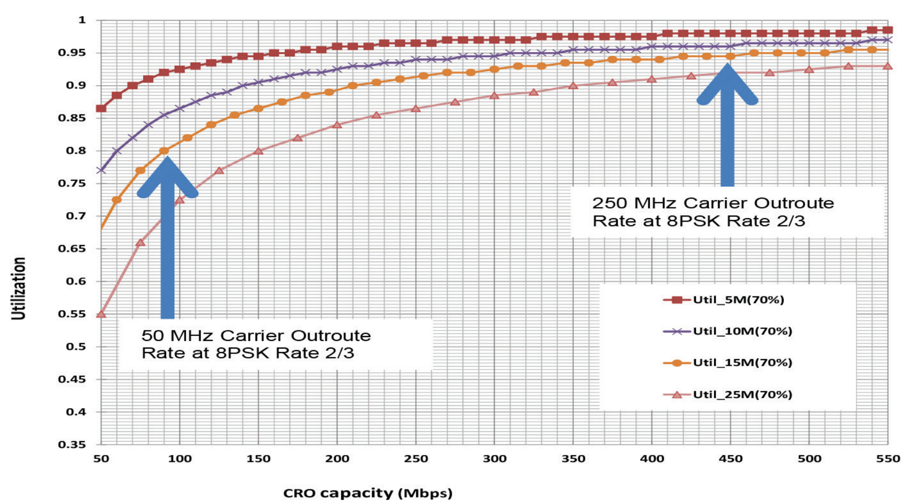

Figure 4. Utilization Efficiency Standard vs. Wideband Carrier

Networks are typically configured with multiple "gateways," with each gateway transmitting to remote terminals over the entire bandwidth allocated to forward channels or outroutes; and remote terminals transmitting back to the gateway over return channels, or inroutes, which can be assigned by the gateway.

The return channel is based on the well-established IPoS standard. The inroute transmission scheme uses Multiple Frequency (MF)/Time Division Multiple Access (TDMA) with Adaptive Low-Density Parity Check (LDPC) coding and Closed-Loop Power Control (CLPC). LDPC is widely recognized as one of the most efficient coding technologies for use over satellite, as evidenced by its adaption into the DVB-S2X standards. The Hughes JUPITER System utilizes the same LDPC coding technology for the return channel and consequently has the ability to apply code rates as high as 9/10 (the highest in the industry for return channel TDMA systems). As illustrated in Figure 5, compared to TurboCodes, the LDPC enables Hughes to achieve a 33% bandwidth efficiency increase over TDMA return channel systems.

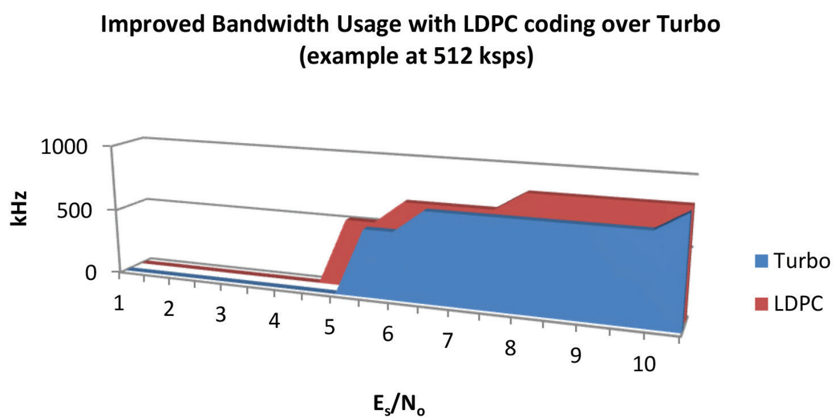

Figure 5. LDPC vs. Turbo Return Link Efficiency

To improve the spectral efficiency further for the return channel, Hughes has implemented "adaptive coding" on the return channel. Similar to ACM on the forward channel, adaptive coding on the return channel enables a remote terminal to dynamically adjust its transmissions to handle fade conditions, in parallel with the hub using ACM to handle forward channel fade.

Adaptive coding, also known as "Code Rate Change on the Fly," is where the return channel demodulator is able to demodulate all bursts on the same channel, no matter what coding rate is used. As shown in Figure 6, this feature enables the JUPITER Gateway to demodulate, decode, and process bursts of varying coding rates within the same TDMA frame. The remote terminal, using feedback from the hub, including the received Es/No levels, selects the most efficient coding rate that enables the transmission to be received by the demodulator without error. For some remote terminals this may be Rate 2/3, for other remote terminals this may be Rate 4/5, and for others this may be Rate 9/10. Using adaptive coding gives a satellite operator at least 20% increased throughput over the satellite because the return channel coding rate does not have to be configured with extra rain fade margin.

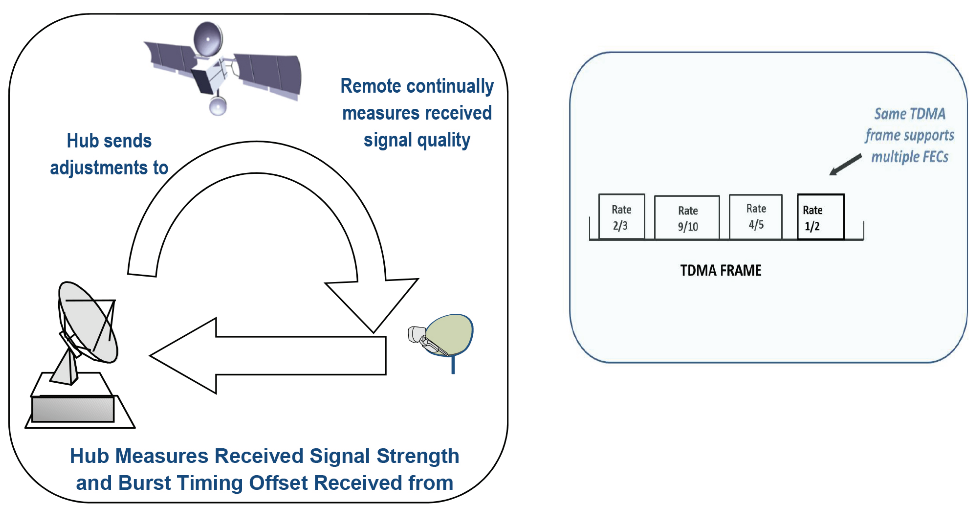

Figure 6. Multiple FECs within One TDMA Frame

As outlined in Figure 6, Hughes has implemented the most powerful modem coding schemes available in order to achieve the highest possible spectral efficiency. But, Hughes also applies a number of powerful techniques at higher layers of the OSI model. The application of these techniques is not reflected in a higher spectral efficiency, but rather in a higher IP throughput figure.

Data Link Layer Optimization

While the physical layer optimization increases spectrum efficiency, the JUPITER System further employs a number of optimization techniques to maximize the use of available bandwidth for user IP traffic. As noted above, bandwidth efficiency is more than the theoretical calculation of bits per hertz. Bandwidth efficiency can be considered as the percentage of available bandwidth used for user IP traffic. Bandwidth efficiency is diminished if bandwidth is allocated and unused due to control and signaling overhead.

The JUPITER System has the capability to support large return carriers of up to 8 Msps as well as the ability to support both OQPSK and 8PSK modulation (8PSK modulation requires use of a linear radio at the remote location).

Hughes typically achieves 85–90% bandwidth packing efficiency on the return channel. The efficiency of the return channel is achieved by using a variable burst length architecture and use of a contention channel (Aloha) via which a satellite terminal requests bandwidth and becomes active in the network. The use of Aloha allows Hughes to operate without dedicated bandwidth for idle terminals, providing significant bandwidth savings. The JUPITER System uses real-time bandwidth allocation algorithms to determine how much bandwidth to allocate to the terminals for return channel burst transmissions. A variable burst length architecture allows bandwidth allocation to be tailored to the exact size required by a terminal to send user traffic.

The return channel allocation granularity supported in the JUPITER System is one TDMA time slot, and the minimum allowed burst size consists of five such time slots. This type of allocation with smaller granularity provides statistically a better packing efficiency for user data when compared with the 53 bytes, ATM cell granularity for DVB-RCS, or some higher values of allocation granularity, such as 125 bytes for other vendors' proprietary implementations. Figure 7 compares return channel packing efficiencies between Hughes and others mentioned above for various sizes of user packets, and representing a typical satellite system used mainly for Internet access.

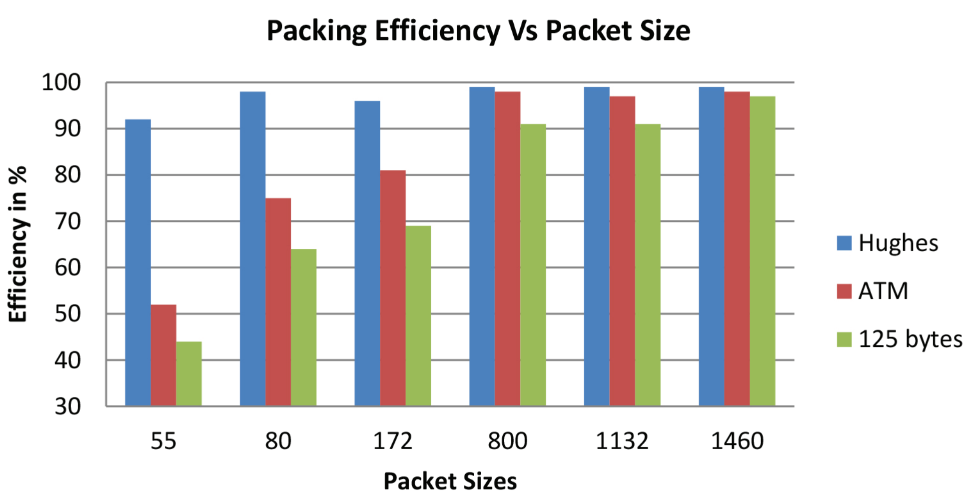

Figure 7. Return Channel Packing Efficiencies

As can be seen in Figure 7, compared to other systems in the market Hughes TDMA access achieves much higher packing efficiencies as packet sizes become smaller. Internet access and VoIP support are two key applications driving the growth of satellite networks around the world. In the case of Internet access, most of the return channel traffic consists of small packets containing "gets" while VoIP packets are even smaller. Thus, the Hughes TDMA access scheme generates higher efficiencies for these applications.

Network Layer Optimization

The JUPITER System delivers more user data over a given available satellite link capacity by incorporating various bandwidth optimization techniques at the network layer, both in the forward and return directions.

Embedded within each Hughes terminal is a set of applications designed to provide superior user experience using Hughes' proprietary acceleration and compression techniques. These techniques compensate the effects of satellite latency (delay) and improve transmission efficiency, providing Hughes users with terrestrial equivalent—or better—performance on their IP applications. Furthermore, these applications do not require reconfiguration of connected devices, which means users can enjoy high performance right out of the box.

The Performance Enhancing Proxy (PEP) conforms to recommendation (RFC 3135) mechanisms for TCP ACK reduction and three-way handshaking, but has expansions beyond the standard to maximize "filling of the pipe" without congestion. The system also supports industry standard RObust Header Compression (ROHC) (RFC 3095) header compression tuned to the satellite link for IP, TCP, UDP, PEP, and RTP protocol suites. In addition, the ITU-T V.44 algorithm used for payload compression provides another 40% efficiency gain on top of the gain provided by the header compression for compressible payload data, particularly with Web traffic.

Other techniques to generate efficiency include DNS caching and preload, where the DNS information is stored in the remote terminal and thereby eliminates the need for a transmission across the satellite.

To stay competitive, Hughes implemented a standard feature in the JUPITER System called TCP Performance Enhancement Proxy (TPEP). The JUPITER PEP establishes a PEP Backbone Protocol (PBP) between the gateway and remote to accelerate and compress TCP traffic. TPEP is implemented at either side of the link and comprises several separate features that are designed to improve the performance and response time of various widely used protocols while minimizing the required bandwidth.

- TCP acceleration: TPEP transparently converts TCP into a satellite-friendly backbone upon entering the satellite network and restores it to TCP prior to leaving the network. The following are some techniques used by TPEP to accelerate TCP:
    - Bidirectional TCP spoofing: Spoofing the TCP handshaking eliminates the effect of latency on bulk transfer throughput due to the TCPs slow start and window sizing mechanisms. TCP data segments are locally acknowledged by the gateway when received. Because this occurs at local LAN speeds (without any satellite delay involved), the sending IP host can very quickly grow its TCP window to its maximum value. PEP also handles local retransmission of lost packets, providing much faster recovery from such losses.
    - Acknowledgment reduction: TCP acknowledgments are cumulative in nature; hence, significant bandwidth savings can be achieved by not transmitting individual acknowledgements. PBP sends PEP acknowledgements every N millisecond-- where N is a definable parameter. With this feature, the acknowledgment bandwidth required for a 2.5 Mbps outbound file transfer is reduced from approximately 100 kbps to around 8 kbps. Figure 8 illustrates how the Hughes TPEP feature works to accelerate IP traffic.

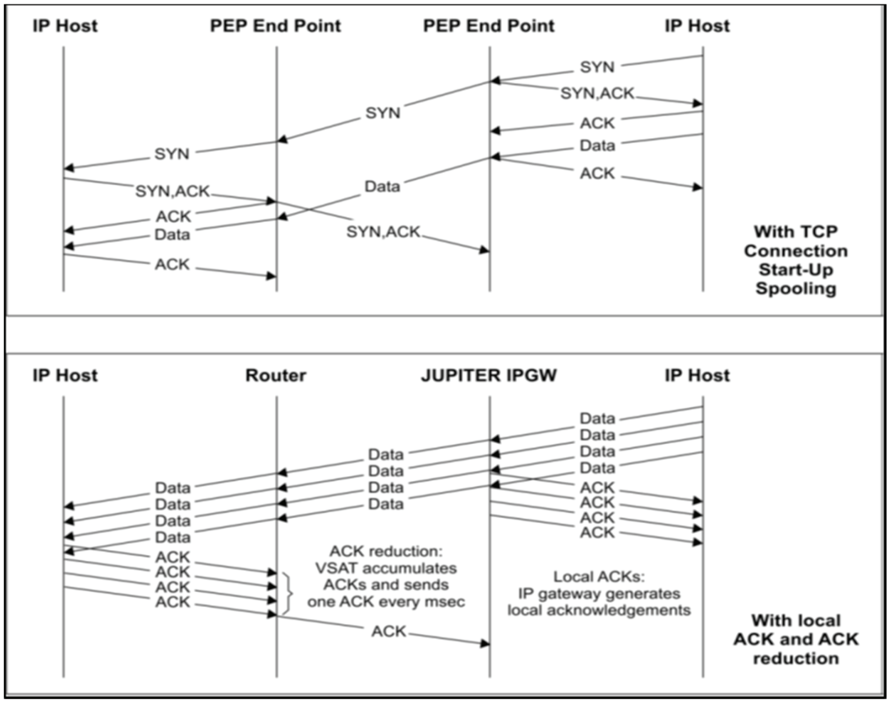

Figure 8. Hughes TPEP

- DNS caching and preload: DNS caching and preload is incorporated into every terminal to eliminate satellite latency introduced by DNS lookup queries. The DNS cache can be preloaded by the gateway and is also populated via local DNS lookups initiated from the local LAN. When WAS is used, the DNS cache is also loaded with the DNS responses associated with preloaded objects.

- Payload compression: Advanced data compression algorithms incorporated into the PEP functionality significantly increase compression ratios resulting in improved throughput across the JUPITER System.

- IP header compression: The JUPITER System takes full advantage of the ROHC technique to provide high performance compression. The ROHC IP header compression saves inroute bandwidth by compressing headers of common protocols, such as IP, UDP, TCP, RTP, and the Hughes proprietary PEP backbone protocol.

Web Acceleration

The JUPITER System provides integral HTTP Web page acceleration that improves page-load response time performance for both static page content and dynamic (i.e., JavaScript) content. This Web acceleration prefetches and caches Web page objects within the remote terminal memory, and delivers those objects to the user Web browser when the browser asks for them

The architecture of Web acceleration is to place two HTTP proxies, the Web Acceleration Client (WAC) within the VSAT and the Web Acceleration Server (WAS) within the gateway, in the path between the Web browser and the Web server located on either side of the space link.

WAS works by reducing the chattiness involved in fetching objects that are part of a Web page. The WAC intercepts Web requests on the remote router and talks to a WAS at the gateway. The normal process would involve waiting for a remote PC to parse the initial HTML page, sending a DNS (domain name server) request for each server that has an object, such as an image or flash file, and then initiating multiple requests to each of those other servers to retrieve each required object. Instead, the WAS prefetches the objects and caches them temporarily at the terminal, providing a local delivery of the requested objects rather than requiring an end-to-end request and response. WAS therefore assures the freshest content from the Web server, while delivering lightning-fast performance.

In addition, Hughes implemented ActiveCompression™, a powerful two-stage compression feature available in the JUPITER WAS and WAC to achieve significant compression on HTTP traffic. To achieve the high compression savings ActiveCompression incorporates a long-range and short-range compression scheme with dynamic real-time selection of the optimal algorithm. Compression functionality of the two-stage compression is implemented on the WAS Server at the hub station and the decompression functionality is implemented by the WAC on the VSAT.

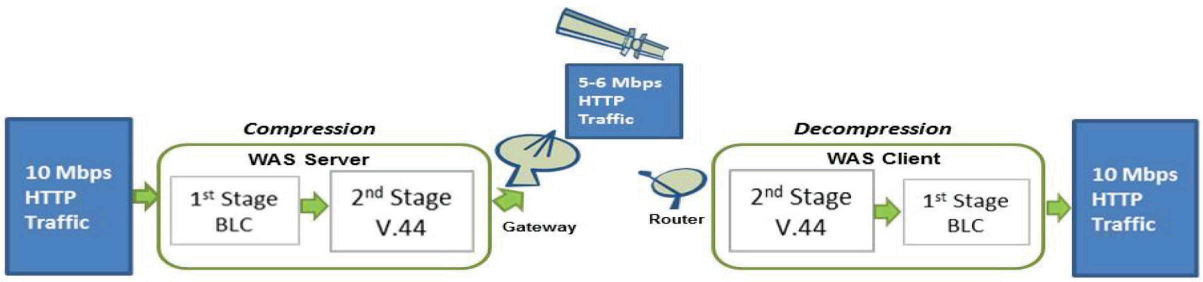

Figure 9. Two-Stage Compression

As illustrated in Figure 9, the first stage is a Byte-Level Caching (BLC) algorithm, which is a lossless compression scheme that exploits duplication of byte sequences in a data stream. A large cache is used by the BLC compression scheme that allows for detecting redundancy in data streams transmitted several megabytes in the past. BLC therefore can provide compression gains on both compressible (text/html) and noncompressible (images) data alike. The second stage is implemented using V.44 compression scheme. The dynamic use of single-stage compression with BLC only or two-stage compression with BLC+V.44 on a block-by-block basis improves Web page response by selecting the compression scheme that optimizes compression gain and decompression time.

Results

The overall efficiency of a VSAT network is the product of much more than just spectral efficiency at the physical layer, and operators should fully understand all the capabilities and efficiencies in assessing a vendor's products. The Hughes JUPITER System employs an extensive set of efficiency techniques across all communication layers to yield the highest possible overall IP network efficiency. Leveraging various intelligent compression techniques, including header compression, result in significant bandwidth saving over the space link. Taken together, these features and technologies enable operators to achieve more throughput than any other system on the market.

Proprietary Statement