

# HN/HX 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. But not always understood 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. This white paper summarizes the techniques that Hughes has implemented in its industry-leading HN and HX System product lines across all layers of the OSI model to yield the highest possible overall IP network efficiency.

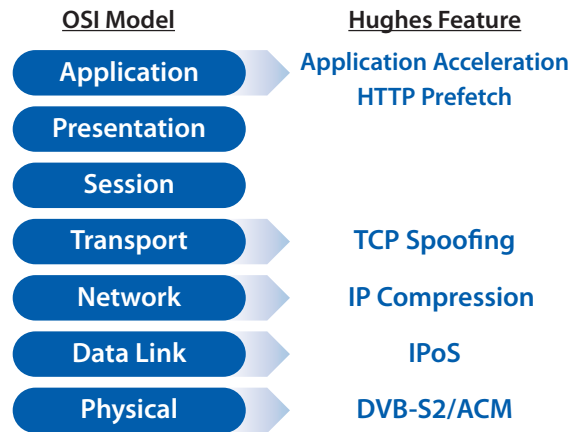


Figure 1. OSI Seven Layer Network Model and Corresponding Hughes Features

## Physical Layer Optimization

Hughes HN and HX Systems are fully compliant with IPoS/ DVB-S2, including adaptive coding and modulation (ACM), the world’s most successful satellite air interface, approved by TIA, ETSI, and ITU. Networks are typically configured with a central “hub” that communicates to remote terminals over the entire bandwidth allocated to forward channels or outroutes, and with remote terminals that communicate back to the hub over return channels, or inroutes, which can be assigned by the hub individually or be shared by terminals on an assigned or contention basis. IPoS outroutes employ a statistical multiplexing scheme compliant with DVB-S2 for sharing among multiple remotes, while inroutes use a demand-assigned, multifrequency Time Division Multiple Access (TDMA) approach to allow remotes to transmit to the hub.

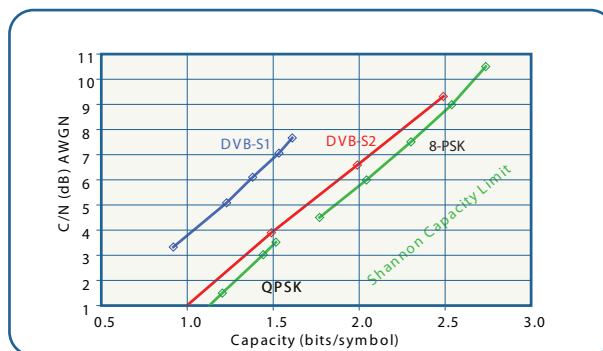


Figure 2. DVB-S2 Vs DVB-S Link Efficiency

As illustrated in Figure 2, the DVB-S2 (ETSI EN 302 307) standard is an improved, second generation of DVB-S. With enhancements such as more sophisticated modulation techniques and low-density parity-checking correction codes (LDPCs), the improved DVB-S2 coding achieves a 30 percent bandwidth efficiency increase over DVB-S.

In addition to the improved efficiency achieved by using LDPC, DVB-S2 enables the use of “Adaptive Coding and Modulation” or ACM. ACM allows the system to dynamically vary the modulation and coding of the forward channel for each transmission. Using information gained via a closed loop feedback mechanism from the remote to the hub enables the system to transmit to each remote using the optimal combination of coding and modulation.

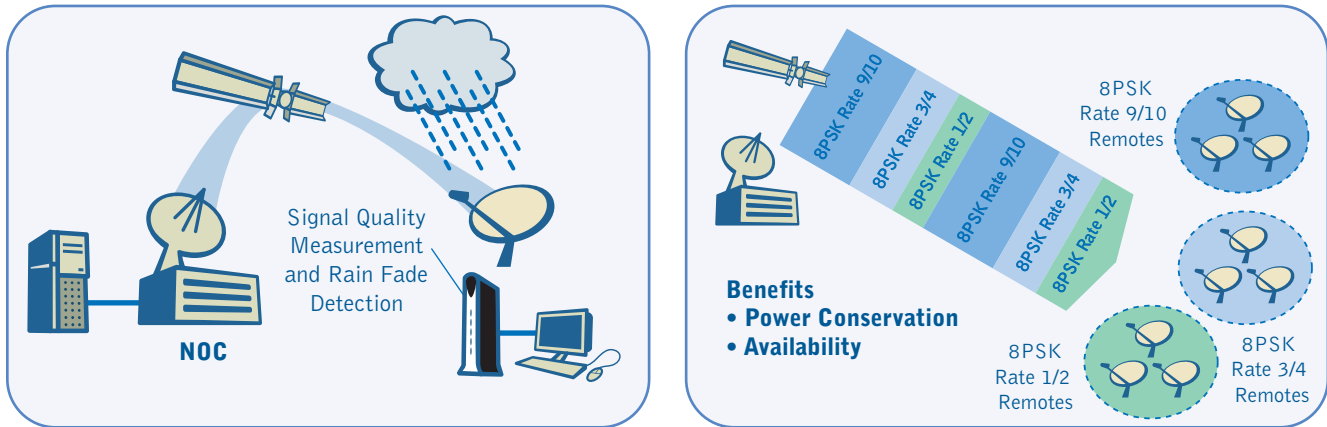


Figure 3. ACM Feature of the IPoS/DVB-S2 Standard

As shown in Figure 3, the remote terminal continuously evaluates signal quality to determine the optimal modulation and coding (also known as “modcod”) for the hub station to send data to that remote terminal. When conditions change or if the hub is not sending data at the optimal modcod, the remote will append a flag on return channel messages indicating that a different modcod should be used. In a rain condition, the hub can automatically adjust the modcod for the affected remote sites, allowing data transmission to continue, albeit at a lower rate for a short period of time. This is a very powerful feature because it allows satellite operators to size their networks to optimize individual site performance rather than be degraded to support the lowest performing site in the network. Additionally, higher availability across the network can be achieved with smaller antennas. Use of ACM provides up to 30 percent bandwidth gain on top of the 30 percent gain provided by LDPC.

For the return channels, which as noted use a multifrequency, TDMA Access (FD/TDMA) scheme, Hughes also offers short-block, low-latency, Low-Density Parity Check (LDPC) codes. As illustrated in Figure 4, using LDPC on these return channels enables Hughes to achieve an 8–12 percent bandwidth efficiency increase over existing Turbo Coded return channel systems.

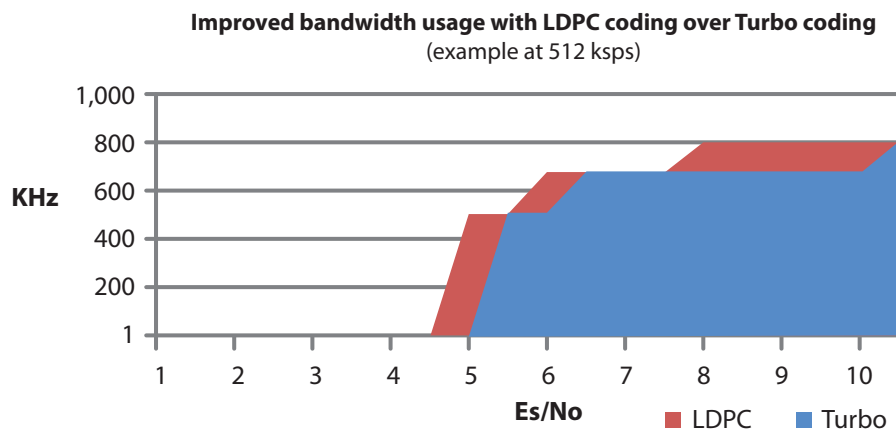


Figure 4. LDPC vs. Turbo Return Link Efficiency

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. Remote terminals also can be dynamically assigned to a return channel using a different symbol rate for better performance.

With adaptive coding, also known as “Code Rate Change on the Fly,” the return channel demodulator is able to demodulate all bursts on the same channel, no matter what coding rate is used. 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, Rate 4/5, and for others, Rate 1/2. Using Adaptive Coding typically gives a satellite operator 20 percent increased throughput over the satellite because the return channel coding rate does not have to be configured with extra rain fade margin.

As outlined above, 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 Optimization

While the physical layer optimization increases spectrum efficiency, HN and HX Systems further employ 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.

Hughes typically achieves 85 percent 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 HN and HX Systems use real-time bandwidth allocation algorithms to determine how much bandwidth to allocate to the terminals for return channel burst transmissions. A variable length burst 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 HN and HX Systems is one TDMA time slot, and the minimum allowed burst size consists of three such time slots. One TDMA slot typically carries 15 bytes of user data. This type of allocation with smaller granularity provides a statistically 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 5 compares return channel packing efficiencies between Hughes and others mentioned above for various sizes of user packets, representing a typical satellite system used mainly for Internet access.

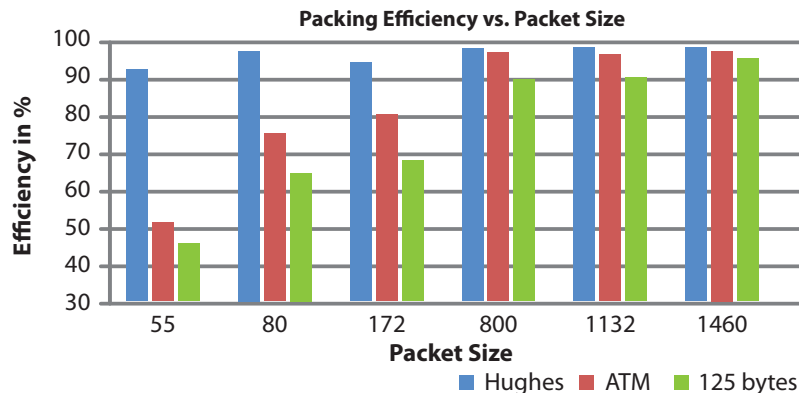


Figure 5. Return Channel Packing Efficiencies

As can be seen from the chart, Hughes 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 very 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

HN and HX Systems deliver 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.

The Performance Enhancing Proxy (PEP) uses standard (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 (RFC 3095) header compression tuned to the satellite link for IP, TCP, UDP, and RTP protocol suites. In addition, the ITU V.44 algorithm (a Hughes patent) used for payload compression provides another 40 percent efficiency gain on top of the gain provided by the header compression for compressible payload data, particularly with Web traffic. It should be noted that the V.44 scheme enables higher efficiency than other schemes such as “GZIP”.

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.

### TurboPage®

HTTP Prefetch, which is marketed by Hughes as TurboPage, is an application acceleration feature which enhances the performance of the HTTP (Web browser) protocol by pre-fetching the HTML (Hypertext Markup Language) objects embedded in a Web page into the remote terminal before the application requests them. This feature uses a hub traffic server to request HTTP objects in advance and on behalf of browsers at the remote terminal. The proxy server pushes these objects to the remote VSAT where it is cached and, therefore, immediately available when the client browser requests the object. This concept is illustrated in Figure 6 below. This feature provides exceptionally better user experience as Web pages are “painted” more quickly.




Figure 6. HTTP Prefetch

HTTP Prefetch is primarily aimed at improving web page response time rather than at improving bandwidth efficiency. However, it does contribute to the latter in two ways. First, because objects are pre-fetched, GET requests for the objects do not need to be sent across the return channel thus conserving bandwidth. Second, HTTP Prefetch applies V.44 compression to HTTP objects as they are forwarded, reducing forward channel bandwidth requirements.

### Application Layer Optimization

In addition to the payload optimization techniques already mentioned, Hughes has also implemented “Web Optimizers,” which work at the HTTP layer and are able to compress HTTP content. A Web Optimizer is typically implemented as a server at the HN or HX hub station and applies a number of data-specific compression techniques including image compression for JPG and GIF images. Through the application of the Web Optimizer, Hughes can reduce HTTP traffic volume by up to 30 percent. It should be noted that the application of compression on images results in loss of image quality. The higher the compression savings the greater the impact to image quality. Figure 7 illustrates that the Web Optimizer is placed between the Internet and the satellite network hub.




Figure 7. Web Optimizer Implementation

### Conclusion

Together, the combination of all the techniques described above results in significant efficiencies for overall IP throughput. Perhaps the best way to illustrate these savings is to look at a specific customer example. Hughes implemented a customer network with a DVB-S2 forward channel on 45 MHz of satellite capacity. The average satellite channel throughput over the forward channel achieved is about 90 Mbps. In this customer’s network there are 74 return channels, with each return channel operating at 256 kbps. Approximately half of the remotes in the network are able to transmit over the return channel using Rate 2/3, and the other half are able to transmit using Rate 1/2 coding. Thus, the overall return channel capacity is about 22 Mbps over 26 MHz. From the forward channel and return channel figures noted above, we can quickly calculate spectral efficiency at 1.9 bits per hertz for the forward channel and about .8 bits per hertz for the return channel.

But, as can be seen in Figure 8 below, the actual network throughput is much higher than what the spectral efficiency figure would otherwise indicate. The service operator of this network is actually driving 149 Mbps of IP content into the system in the forward channel direction. Applying 149 Mbps against 45 MHz gives us an IP efficiency of 3.2 bits per hertz, or over 50 percent higher throughput than indicated by spectral efficiency alone. In the return channel, the system carries 30 Mbps of IP content over the 26 MHz of return channel capacity, resulting in an IP efficiency of 1.13 bits per hertz, which is about 30 percent higher throughput than indicated by spectral efficiency alone.

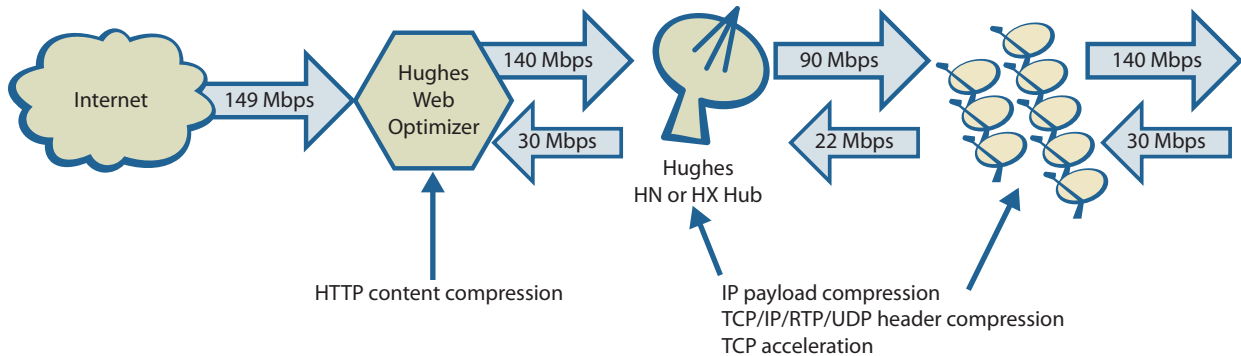


Figure 8. Compression Efficiency

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. Hughes HN and HX Systems employ an extensive set of efficiency techniques across all communication layers to yield the highest possible overall IP network efficiency.

## **Proprietary Statement**

---

All rights reserved. This publication and its contents are proprietary to Hughes Network Systems, LLC. No part of this publication may be reproduced in any form or by any means without the written permission of Hughes Network Systems, LLC, 11717 Exploration Lane, Germantown, Maryland 20876.