Minimizing Latency in Low Symbol Rate DVB-S2 Applications

Wallace A. Ritchie WALLY.RITCHIE@GMAIL.COM
Deltona, FL 32738 USA

Abstract

Most of the Applications for the DVB-S2 (Digital Video Broadcasting) standard operate at relatively high symbol rates where the inherent latencies introduced by FEC decoding are of little concern. The normal operating range for DVB-S2 in broadcast applications is typically 10 to 30 MBaud. There are however some applications, in particular Radio Amateur Satellite Service applications, where symbol rates well below 1 MBaud are of interest. At these lower symbol rates the latency inherent in FEC decoding the relatively large DVB-S2 frames can be of significant concern, particularly when interactive two-way voice communication is involved. This paper discusses the various latencies associated with DVB-S2 applications and several methods that are available to reduce such latencies for applications operating in the 80 kBaud to 1 MBaud range.

1. Introduction

Amateur Radio Satellite Service Transponders will begin operating from GEO (Geosynchronous Earth Orbit) in 2018 with the upcoming launch of the Qatar Satellite Company's Es'hail-2 Satellite (ESHAILSAT, 2015). In addition to its Ku and Ka transponders providing commercial service, Es'hail-2 will carry a pair of transponders operating in the Amateur Radio Satellite Service bands. The pair includes a so-called wideband transponder with a bandwidth of 8 MHz and a narrowband transponder with a bandwidth of 250 kHz. Both transponders will have uplinks in S-Band and downlinks in X-Band. The narrowband linear transponder will support operation in narrowband analog and digital modes. The wideband linear transponder is intended to support up to 2 or 3 DVB carriers. Other AMSAT GEO satellites are in the planning stages. There is also increasing interest in utilizing DVB-S2 for Amateur Radio Service Satellites operating at microwave frequencies from HEO (High Earth Orbit) and LEO (Low Earth Orbit). DVB-S2 is also of interest for operation in the Amateur Radio Service microwave and UHF frequencies on the ground.

1.1. Commercial Video Broadcast Configuration

In commercial video broadcast applications the satellite system is generally designed to maximize the QEF (Quasi Error Free) data rate for economic Earth terminals in the 60 cm to 90 cm range operating in Ku band. Using DVB-S2 a single carrier from the satellite can provide multiple HD and SD television broadcast channels through a single transponder. With DVB-S2, a 30 MBaud sample rate fits nicely within a standard 36 MHz transponder bandwidth providing more than 60 Mbit/s of QEF transport for multiple MPEG streams. These signals are easily received on the ground using mass produced low cost LNBs mounted on small dishes combined with low cost ASIC (Application Specific Integrated Circuit) based set top boxes.

1.2. Amateur Radio Satellite System Configuration

In Amateur Radio Satellite Service applications the transponder bandwidths are significantly narrower, approaching 10 MHz at the high end and 100 kHz at the low end. The available DC power is often well below that available within a commercial satellite - typically less than 100 W and often 5 W or less. Optimizing the system around the available DC power typically requires operation of the RF Final Amplifier at or near saturation. This generally dictates single carrier operation on the downlink. To close the link with small stations, the system design may require a symbol rate well below that of normal DVB-S2 Broadcast.

1.3. Why DVB-S2?

The second generation Digital Television standard, DVB-S2 (ETSI, 2014), was designed to support both broadcast TV and other applications e.g. IP transport. DVB-S2 is both flexible and high performance. Its forward error correction, based on a combination of LDPC (Low Density Parity Coding) and BCH (Bose, Chaudhuri, and Hocquenghem) coding, provides performance very close to theoretical limits. With 100s of millions of terminals, there is wide availability for much of the equipment. This includes economical ASIC (Application Specific Integrated Circuit) devices that perform nearly all of the required functions. The actual standard documents are also widely available at no cost (DVB). All of these factors argue strongly for adapting DVB-S2 for Amateur Radio Service applications.

1.4. Satellite Latency

Geo-synchronous orbits are located directly above the equator and approximately 40 km above the Earth's surface. The transmission delay of the uplink and downlink signal to or from a particular location on the surface is on the order of 130 ms in each direction. This results in a minimum one-way station to station delay on the order of 260 ms. The inherent round trip delay, without introducing any added latency in the electronics, is on the order of 520 ms or roughly 1/2 second. For two-way interactive speech this means that when a speaker stops talking it takes at least 1/4 second before the other party hears the pause. Even if that party replies immediately it will be 1/2 second before the first party hears the start of that reply. These delays are known to be objectionable for telephone like service. In many cases, the two parties will end up talking and pausing at the same time causing an annoying interruption to the conversation.

For two-way, push-to-talk voice communication, half second round trip delays are generally considered acceptable - but just barely so. Added delays measured in seconds, however, are often considered unacceptable to most users.

For satellites in LEO, the inherent transmission delays are relatively short. However, adding seconds of delay can render even these links unacceptable.

The added latencies from FEC decoding, and methods for mitigating them, are the focus of this paper.

1.5. System Configuration

In commercial applications DVB-S2 is used in a so-called ”bent pipe” transponder configuration. The DVB-S2 signal originates on the ground, arrives on the uplink after 130 ms delay, is amplified and frequency translated to the downlink, and arrives on the ground after another 130 ms delay. The added delay at the satellite is minimal. The AMSAT-DL Phase 4A transponder on Es'Hail-2 also operates in this fashion. Future AMSAT microwave satellites may use an all-digital regenerative repeater in place of the traditional bent-pipe transponder. A regenerative repeater operates in a somewhat different fashion. The downlink is a single DVB-S2 carrier that originates not on the ground but at the satellite. This single downlink multiplexes the data received from multiple narrowband digital uplinks. Each of the uplinks is demodulated and error corrected by the satellite's onboard multi-channel receiver. The individual payload data from all channels is then multiplexed onto the single downlink. Uplink and downlink can use entirely different modulation and forward error correction schemes.

In this paper we are concerned only with downlinks operating with DVB-S2. While digital uplinks will introduce their own added latencies, these are relatively small as the frames are far shorter than those of DVB-S2.

2. DVB-S2 Signal Formats

DVB-S2 has a flexible signal format that allows a single carrier to be utilized with modulation and coding that can vary on a frame by frame basis. All DVB-S2 carriers transmit continuous back to back frames. With VCM (Variable Coding and Modulation) the modulation scheme and coding can be varied with every frame. The modulation scheme and coding formats for DVB-S2 range from QPSK with 1/4 rate FEC all the way to 32APSK with 9/10 FEC. The C/N (Carrier to Noise) required to receive QEF frames in these formats ranges from about -2 dB to +16 dB. This spans an 18 dB range in receive station performance. For reference, this is approximately the difference between an antenna aperture of 0.6 vs 4.5 meters. Using DVB-S2 with VCM, a single carrier can simultaneously serve a wide range of station capabilities ranging from small dishes or patch antennas capable of receiving 1/4 rate QPSK - to much larger dishes able to receive 32APSK with 9/10 FEC. The trade-off for operation at lower C/N (i.e. with smaller dishes) is a lower bit efficiency. In DVB-S2, these bit efficiencies range from 0.25 to 2.35 bits/symbol, a 9 to 1 range.

2.1. Error Correction for QEF

In video broadcasting using DVB-S2, a single carrier normally occupies a full 36 MHz transponder bandwidth operating at a symbol rate of about 30 MBaud. For broadcast applications, each frame will typically contain 64800 raw bits of data. This includes both the payload data bits and the additional bits required for error correction. DVB-S2 is designed to deliver QEF transport which is essential for highly compressed digital television where even a single bit error can corrupt the quality of the video. QEF is also necessary for IP where a single bit error will invalidate an entire IP packet causing both retransmission and added delay. QEF transport is realized in DVB-S2 by devoting a significant fraction of the transmitted bits to error correction. This ranges from 10% to 75% of the transmitted bits.

2.2. Latency

In DVB-S2 both the data bits and the error correction bits are spread over the frame. The entire frame must be available to the FEC decoder in order to begin the error correction process. To maintain the continuous data flow, the error correction process for one frame must normally be completed before the next frame has completely arrived. Demodulation and error correction together impose a typical delay through a DVB-S2 demodulator ASIC that is on the order of 72000 symbols. This typically fixed delay is the result of the implied requirement for handling the worst case sequence of frames over the full supported range of modulation schemes and frame sizes. Generally, two frames of the longest possible size must be buffered in order to maintain a continuous data flow under the worst case sequencing of variable sized frames.

For QPSK with the normal frame size of 64800 bits with pilots[1] the length of each frame is 32400 symbols, plus an additional 792 symbols of pilots, plus 90 symbols of PLHEADER[2] for a total of 33282 symbols. The typical 72000 symbol fixed delay allows for the storage of two complete QPSK long frames plus several thousand additional symbols in the demodulator in front of the FEC decoder. At the moment the ASIC begins transferring the error corrected output bits of a frame it will have the entire decoded frame buffered for output plus the complete subsequent frame that will just be beginning the process of forward error correction. A few thousand additional symbols will be buffered in the signal processing stages of the demodulator inf front of the FEC decoder. In other words, from the IQ sample of the device to the output of the corresponding data bits from the device there is a fixed 72000 symbol delay.

The 72000 symbols delay from demodulation and FEC decoding at 30 MBaud amounts to a latency of only 2.4 ms. This is only a small fraction (1%) of the GEO transmission delay (240 ms) and is of little consequence. At 3 MBaud the added latency rises to 24 ms which is still a modest 10% of the transmission delay. However, at 80 KBaud the added latency amounts to a whopping 900 ms! This brings the round trip delay to 2.3 seconds for GEO. For LEO the delay is still a significant 1.82 seconds. If we wish to realize the 11 dB or so C/N improvements that can result from operation at lower symbol rates, we need to find an approach that can reduce this added latency to acceptable levels.

3. Latency Reduction Approach

Commercial ASIC devices are generally focused on broadcast applications. While some newer devices support many of the non-broadcast features of DVB-S2, including lower symbol rates, the basic architectures of the devices must generally remain compatible with all common broadcast formats. The 72000 symbol fixed delay described in the previous section is a consequence of the length (in symbols) of the normal size (64800 bit) frame coded in QPSK (2 bits per symbol). By restricting frames to the short (16200 bit) frame size we can reduce the longest frame to 8100 symbols. Adding pilots and PLHEADER, the maximum length short frame will then be 8370 symbols. At 80 kBaud the delay for a single frame is now 105 ms.

It is always necessary to buffer a complete demodulated frame as the FEC decoder generally requires all of the symbols of the frame to begin decoding. This will always add a delay equal to the 105 ms transmission duration of the frame at 80 kBaud. If it were still necessary to buffer an additional frame, the added delay would be about 210ms. However, it is possible to eliminate the need for this second full frame buffer as we will discuss below.

3.1. LLR Values

DVB-S2 demodulation does not directly produce digital ”bits.” Rather, the demodulator proper produces a soft decision LLR (Log Likelihood Ratio) for each individual bit. An LLR is a typically 6-bit value that represents the probability that a particular bit is a zero divided by the probably that the bit is a one. So the 8100 QPSK data symbols of a short frame are decoded by the demodulator into 16200 LLR values, one for each bit. The FEC decoder then processes these LLR values to correct all of the payload data bits of the frame. For short frames, the number of payload bits can range from 3072 bits for Rate 1/4 FEC to 14232 bits for Rate 8/9 FEC[3].

3.2. FEC Decoder

The FEC decode process in a typical ASIC device must obviously be able to operate at the maximum rates supported by the device. A device that supports 45 MBaud and 8PSK (3 bits per symbol) must be capable of FEC decoding at 135 million LLRs per second. A device that supports two such channels from a single decoder must support decoding at a rate of 270 million LLRs per second. A decoder with a performance of 270 million LLRs per second can process a single frame in a very short time, far less the transmission time for a short frame at 80 kBaud. When only FEC is considered, the 16200 LLRs in a QPSK short frame can be fully processed in as little as 60 μsec consuming only a tiny fraction of the capacity of the decoder. Given this high level of FEC decoding performance, the added latency for FEC decoding can be potentially reduced to a few milliseconds. Harnessing a high performance ASIC decoder for this task can bring the total added latency down to the 110 ms range for an 80 kBuad symbol rate.

[1] Pilot blocks are optionally added to frames to aid in synchronization.

[2] The PLHEADER flags the frame boundary and indicates the modulation and coding for the frame.

[3] The 9/10 FEC rate is not supported for short frames.

With this approach the one-way delay (up and down) for GEO can be reduced from more than a second to about 370 ms. The corresponding round trip delay can be reduced from 2.3 seconds to less than 3/4 second. The one-way delay for LEO can be reduced from about 908 ms to 130 ms with its round trip delay reduced from 1.8 seconds to 230 ms. Even allowing for additional delays for speech coding and processing delays in the transmitter, the total delays can be brought down to levels that should be acceptable for conversational speech, especially in a push to talk environment.

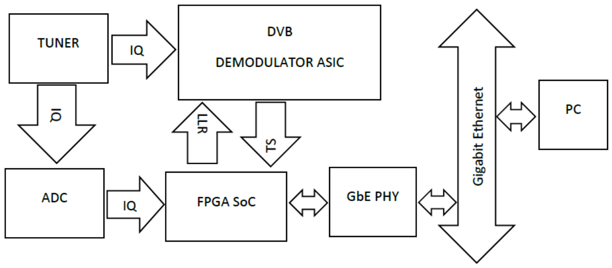

Figure 1. Block Diagram

4. Implementation

Fortunately, some ASIC DVB-S2 demodulator ASICs have a feature known as external demodulation or demodulator bypass. This feature essentially breaks the connection between the internal DVB-S2 demodulator and the internal FEC decoder such that the LLR values can be sourced from an external demodulator. The LLR values from the external demodulator can be loaded directly to the FEC decoder for subsequent downstream processing. This processing typically includes both the FEC decoding and the de-framing and formatting of the error corrected data bits into a normal DVB-S2 transport byte stream. This transport byte stream can consist of either MPEG Transport Streams as used in broadcast applications or raw BBFRAMES (Baseband Frames). This support for BBFRAMES allows access to all of the non-broadcast functionality of DVB-S2 including Generic Stream Encapsulation and Multiple Transport Streams.

For low symbol rate applications, an external demodulator can be operated with a symbol rate of 80 kBaud or lower. Unlike the ASIC's internal demodulator, which has to be optimized for broadcast rates and consumer grade LNBs, the external demodulator can be optimized instead for low symbol rates in a specific application. Such an external demodulator can handle unusual requirements e.g. the high Doppler shifts present in LEO.

Using an external demodulator, a full frame of LLRs must still be buffered before FEC decoding can commence and this requires about 105 ms at the 80 kBaud rate. But once the full frame of buffered LLRs is available to the decoder, the FEC decoding can proceed at a very rapid rate. This allows the latency to be reduced from 2 full QPSK long frames to a little over 1 QPSK short frame. This can improve added latency by a factor of 7 or more.

5. Low Latency DVB-S2 Receiver Hardware

A complete flexible DVB-S2 receiver can be implemented by combining an ASIC demodulator chip set with a relatively low cost dual ADC and an FPGA SoC (System on a Chip). A simplified block diagram of the hardware is illustrated in Figure 1. The hardware, FPGA HDL, and software for this receiver are currently being developed as part of the AMSAT-NA Phase4Ground effort (Phase4Ground). The implementation illustrated uses a GbE Ethernet link for connection to a PC. This link can be shared or dedicated. Alternative designs could use USB for the link.

The DVB-S2 receive function is split between the hardware and the SDR signal processing executing in the PC. The PC signal processing will be based on the open source GNU Radio (GNURadio).

Modern multi-core PCs with advanced SIMD (Single Instruction Multiple Data) instructions sets are particularly good at performing the types of signal processing required. GNU radio takes advantages of these instructions through the volk library (VOLK). Since a PC is typically required for other function, this configuration is both flexible and economical. As DVB-S2 based amateur radio satellites become operational, integrated configurations can be evolved to perform everything within a single hardware box.

5.1. Normal Symbol Rate Operation

For symbol rates in the 1 to 45 MBaud range, all of the DVB-S2 processing can be handled entirely by the ASIC. The decoded transport byte streams are simply packaged onto Ethernet frames and delivered to the PC. Further processing of the byte steam will be handled by a driver or user-space software on the PC. Since all of the standard broadcast MPEG transport streams for both DVB-S2 and DVB-S are already widely supported, these MPEG streams can be passed to existing Windows or Linux video driver frameworks. The BBFRAMES of non-MPEG applications do required different processing. This will be handled by new drivers and new GNURadio blocks that decode the BBFRAMES, perform de-encapsulation of Generic Streams, and make the data available to specialized applications.

It should be noted that operational DVB-S2 transmit blocks already exist within GNURadio. These currently handle all of the transmit functions including frame formatting, FEC coding, modulation, and filtering. All of the new GNURadio processing blocks for the receiver will initially be available from the Phase4Ground(Phase4Ground) site as out of tree modules and will eventually be submitted for inclusion in a future release of GNURadio.

5.2. Low Symbol Rate Operation

For symbol rates in the 80 kBaud to 1.5 MBaud range (or even lower), the DVB-S2 processing is split between the FPGA, the GNU Radio based signal processing software at the PC, and the ASIC. This split processing requires that data first be moved from the hardware to the PC and then back from the PC to the hardware. This is the primary function of the FPGA SoC which provides the necessary interface and control signals to the ASIC.

5.3. Tuner

The Tuner IC handles the received signal from the dish mounted LNB. The Tuner translates this L-Band signal from the down-converter to IQ baseband. This device also includes an integrated low pass filter programmable from 5 to 36 MHz. The bandwidth is large enough to encompass both the DVB-S2 carrier's occupied bandwidth and the worst case expected frequency offset.

For normal symbols rates, the IQ output is digitized by the 10 bit ADC converters on the demodulator ASIC, typically at 135 Msamples/s. This allows support for symbol rates up to 45 MBaud with 3 samples per symbol.

For low symbol rates, the ASIC demodulator is bypassed and the baseband IQ output of the tuner is routed to a pair of external Analog to Digital Converters. The tuner IC's internal low pass filters have a minimum programmable cutoff frequency of 5 MHz. Their minimum attenuation is typically specified at 2.4 times the cut-off frequency or 12.5 MHz minimum. To avoid aliasing, the 12 bit ADC channels sample the IQ baseband at 40Msps and transfer the digital values to the FPGA. The baseband IQ sample stream is then decimated by 8 in the FPGA to produce a low pass filtered 5 Msps IQ stream. This sample rate is sufficient to handle symbol rates of 1 MBaud or lower together with frequency offsets up to about 250kHz, depending on the sample rate. The frequency offset is corrected in subsequent signal processing as is traditional in DVB-S2 demodulation.

5.4. Moving IQ Samples to the PC

The 5 Msps IQ stream contains 20MB/sec of data. Using gigabit Ethernet with jumbo packets, received sample blocks can be transferred from the FPGA 2000 samples at a time at a nominal rate of about 2500 packets/second. The delay through the decimating filter is minimal as is the per packet transmission time over the wire (64 μsec).

The 5 Msps IQ stream is received at the PC where it is then made available for processing in GNU Radio signal processing blocks. These blocks will perform a course frequency correction and then further decimate the input stream to a lower sample rate. Following any decimation, the GNU Radio blocks will perform the required demodulator functions of fine frequency correction, symbol timing recovery, framing, demodulation, descrambling, deinterleaving, and de-mapping. Ultimately, the demodulation will produce soft-decision LLR values for each bit.

When only short frames are used, each demodulated frame will contain exactly 16200 LLR values together with a single MODCOD that indicates the modulation scheme and FEC rate for the frame. The duration of the frame will vary by the modulation order (2, 3, 4, or 5). At the 80 kBaud symbol rate and QPSK the duration is about 105 ms. This delay represents the largest component of the added latency of the receiver.

5.5. Moving LLR Values to the ASIC

LLR values for an entire frame can be transported from the PC back to the hardware receiver over GbE in 2 jumbo packets. The transit time on the wire for these two packets is about 128 μsec. There are additional delays in the PC network stack. There many also be short delays in any intervening switch hardware which will typically buffer one or more frames at each switching node.

After these short delays, the receiver hardware assembles the Layer 2 packet from the wire and transfers the LLR data directly to the ASIC. These transfers occurs at a nominal 100 MHz. Since the frame will always contain 16200 LLRs, this transfer requires only about 162 μsec and does not contribute significantly to the total delay. 162 μsec later the payload data frame will be fully error corrected and a short time later payload bytes will appear on the transport interface, again at 100 MHz.

For the final step, the FPGA SoC delivers the complete demodulated and error corrected BBFRAME over the GbE to the PC. The subsequent PC processing is then handled in the same manner as for the normal symbol rates where processing is performed entirely by the ASIC. In other words, the downstream processing of the transport stream or BBFRAME is the same for high sample rates using the internal demodulator and low sample rates using the external demodulation chain.

6. Summary

DVB-S2 latency is important in low symbol rate applications e.g. Amateur Radio Satellites. Economic receivers that minimize the low-symbol rate latency are possible by combining relatively low cost hardware with the significant signal processing capabilities of a modern PC using GNU Radio. With significant improvements in DVB-S2 receiver latency, system designs are possible where the overall round-trip latencies are low enough to support applications e.g. two way conversation speech.

References

DVB. https://www.dvb.org/.

ESHAILSAT. ESHAIL-2 Satellite AMSAT Payload, 2015. URL http://www.itu.int/en/ITU-R/space/workshops/2015-prague-small-sat/Presentations/Eshail-2.pdf. ITU-R Conference and Workshop on the Small Satellite Regulation and Communication Systems, Prague, Czech Republic, 2-4 March 2015.

ETSI. EN 302 307 1 v1.4.1 Digital Video Broadcasting (DVB); Second generation framing structure, channel coding and modulation systems for Broadcasting, Interactive Serivces, News Gathering and other broadband satellite applications; Part 1: DVB-S2, 2014.

GNURadio. https://www.gnuradio.org/.

Phase4Ground. https://phase4ground.github.io/.

VOLK. http://libvolk.org,.