WHAT EXACTLY IS 8-VSB ANYWAY?

By David Sparano

This is the third edition of an article that originally appeared in 1997. Previous editions have appeared on the Harris Broadcast website and the Miller Freeman Guide to Digital Television.

The Grand Alliance proposed it, the FCC accepted it, stations around the United States are transmitting it, but what is 8-VSB anyway? Simply put, 8-VSB is the RF modulation format utilized by the DTV (ATSC) digital television standard to transmit digital bits over the airwaves to the home consumer. Since any terrestrial TV system must overcome numerous channel impairments such as ghosts, noise bursts, signal fades, and interference to reach the home viewer, the selection of the right RF modulation format is critical. The 8-VSB format is the cornerstone upon which the DTV standard is based; developing a basic understanding of 8-VSB is imperative for those who will be working around DTV in the future.

In the alphabet soup world of digital communications, there are two big names to remember when thinking about the complete DTV system: 8-VSB and MPEG-II. 8-VSB is the RF modulation format and MPEG-II is the video compression/packetization format used in DTV. That is, there are two distinct stages of processing needed to convert high-definition video into a form suitable for over-the-air broadcast: *MPEG-II encoding* and *8-VSB modulation*. Accordingly, two major pieces of equipment form the heart of a DTV transmission system: an MPEG-II encoder and an 8-VSB exciter.

The MPEG-II encoder takes baseband digital video and performs bit rate compression using the techniques of discrete cosine transform, run length coding, and bi-directional motion prediction -- all of which are beyond the scope of this article. The MPEG-II encoder then multiplexes this compressed video information together with pre-coded Dolby AC-3 audio and any ancillary data to be transmitted. The result is a stream of highly compressed MPEG-II data packets with a data frequency of only 19.39 Mbit/Sec. This is by no means a trivial task since the high-resolution digital video (or multiple programs of standard resolution video) input to the MPEG-II encoder could easily have a data rate of 1 Gbit/sec or more. This 19.39 Mbit/sec data stream exiting the MPEG-II encoder is known as the *DTV Transport Layer*. It is transmitted from the encoder to the 8-VSB exciter in serial form via a 75-ohm coaxial cable, according to the SMPTE-310 interface standard.

Although MPEG-II compression techniques can achieve stunning bit-rate reduction results, still more tricks must be employed to squeeze the 19.39 Mbit/sec DTV Transport Layer signal into a slender six MHz RF channel for over-the-air transmission. This is the job of the 8-VSB exciter.

Figure 1 is a block diagram of a typical 8-VSB exciter. In this article, we will walk through the major processes that occur in the 8-VSB exciter, identifying the major components of the 8-VSB signal and explaining how the 8-VSB signal is generated.

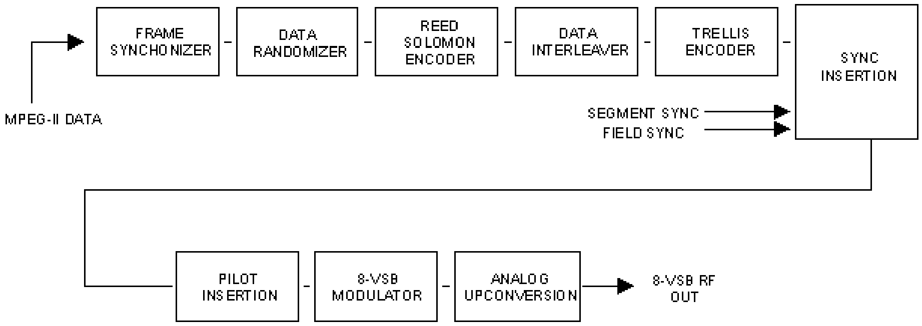

FIGURE 1: BLOCK DIAGRAM, 8-VSB EXCITER

DATA SYNCHRONIZATION

The first thing that the 8-VSB exciter does upon receiving the MPEG-II data packets is to synchronize its own internal circuits to the incoming signal. Before any signal processing can occur, the 8-VSB exciter must correctly identify the starting and ending points of each MPEG-II data packet. This is accomplished using the MPEG-II sync byte. MPEG-II packets are 188 bytes in length with the first byte in each packet always being the sync byte. The MPEG-II sync byte is then discarded; it will ultimately be replaced by the ATSC segment sync in a later stage of processing.

DATA RANDOMIZER

With the exception of the segment and field syncs (to be discussed later), the 8-VSB bit stream must have a completely random, noise-like nature. This is because the transmitted signal frequency response must have a flat noise-like spectrum in order to use the allotted RF channel space with maximum efficiency. If the data contained repetitious patterns, the recurring rhythm of these patterns would cause the RF energy content of the transmitted signal to “lump” together at certain discrete points in the frequency spectrum, thereby leaving holes at other frequencies. This implies that certain parts of the six MHz channel would be overused, while other parts would be underused. Moreover, the large concentrations of RF energy at certain modulating frequencies would be more likely to create discernible beat patterns in an NTSC television set, if DTV-to-NTSC interference were experienced.

In the data randomizer, each byte value is changed according to known pattern of pseudo-random number generation. This process is reversed in the DTV receiver to recover the proper data values.

REED-SOLOMON ENCODING

Reed Solomon encoding is a Forward Error Correction (FEC) scheme applied to the incoming data stream. Forward error correction is a general term used to describe a variety of techniques that can be used to correct bit errors that occur during transmission. Atmospheric noise, multipath propagation, signal fades, and transmitter non-linearities may all create received bit errors. Forward error correction can detect and correct these errors, up to a reasonable limit.

The Reed-Solomon encoder takes all 187 bytes of an incoming MPEG-II data packet (the packet sync byte has been removed) and mathematically manipulates them as a block to create a sort of “digital ID tag” of the block contents. This “ID tag” occupies 20 additional bytes, which are then tacked onto the tail end of the original 187-byte packet by the encoder. These 20 bytes are known as Reed-Solomon parity bytes.

The DTV receiver compares the received 187-byte block to the 20 parity bytes to determine the validity of the recovered data. If errors are detected, the receiver determines that the “ID tag” no longer corresponds to the packet contents and searches for a similar packet (i.e. with only a few bit positions changed) that matches the received tag.

This is somewhat like the license plate on automobiles: An eyewitness describes the car driven by bank robber as being a red 1999 Toyota, license plate ABC123. The police database has no record of a 1999 Toyota with a plate matching (or even similar to) ABC123. There is, however, a record of a red 1999 Nissan with license plate ABC123. The logical conclusion is that the eyewitness mistook one type of car for another, very similar type of car. The search is on for a red 1999 Nissan. The Reed-Solomon decoder in the DTV receiver performs a similar operation by comparing parity bytes to determine the most likely transmitted packet.

Unfortunately, this type of error correction has its limits. The greater the discrepancy between the Reed-Solomon bytes and the packet at the receiving end, the greater the chance of error in matching the correct ID tag to the correct packet. Continuing with the car analogy: What if license plate ABC123 belonged to a red pick-up truck? …or a green SAAB? …or something completely different, such as a silver SUV? These vehicles are not as easily mistaken for a red Toyota. Maybe the eyewitness is mistaken about the license plate number, instead (?). At some point, the ambiguity becomes too great to draw a reasonable conclusion as to the correct vehicle.

The Reed-Solomon coding scheme used in DTV can correct up to 10 byte errors per packet. If too many byte errors are present in a given packet, the receiver can no longer match the parity tag to any packet with a sufficient level of certainty. The validity of the data can no longer be confirmed, and the entire MPEG-II packet must be discarded.

DATA INTERLEAVER

The data interleaver scrambles the sequential order of the data stream and disperses the MPEG-II packet data throughout time (over a range of about 4.5 msec through the use of memory buffers) in order to minimize the transmitted signal’s sensitivity to burst type interference. The data interleaver then assembles new data packets incorporating tiny fragments from many different MPEG-II (pre-interleaved) packets. These reconstituted data packets are the same length as the original MPEG-II packets: 207 bytes (after Reed-Solomon coding).

This is the equivalent of spreading all of your eggs (bytes) over many different (time) baskets. If a noise burst punches a hole in the signal during propagation and “one basket” is lost (i.e. several milliseconds), many different MPEG-II packets lose one egg instead of one MPEG-II packet losing all of its eggs. This is known as *time diversity*. If each packet only loses a tiny number of bytes, the Reed-Solomon decoder in the DTV receiver is able to correct the errors and recover the data, as described in the preceding section.

Data interleaving is done according to a known pattern; the process is reversed in the receiver to recover the proper data order.

TRELLIS ENCODER

Trellis coding is yet another form of forward error correction. Unlike Reed-Solomon coding, which treated the entire MPEG-II packet simultaneously as a block, trellis coding is an evolving code that tracks the progressing stream of bits as it develops through time. Accordingly, Reed-Solomon coding is known as a *block* code, while trellis coding is a *convolutional* code.

For trellis coding, each 8-bit byte is split up into a stream of four, 2-bit words. In the trellis coder, each 2-bit word that arrives is compared to the history of previous 2-bit words. A 3-bit binary code is mathematically generated to describe the transition from the previous 2-bit word to the current one[1]. These 3-bit codes are substituted for the original 2-bit words and transmitted over-the-air as the eight level symbols of 8-VSB (3 bits = 2^3 = 8 combinations or levels). For every two bits that go into the trellis coder, three bits come out. For this reason, the trellis coder in the 8-VSB system is said to be a 2/3-rate coder.

The trellis decoder in the receiver uses the received 3-bit transition codes to reconstruct the evolution of the data stream from one 2-bit word to the next. In this way, the trellis coder follows a "trail" as the signal moves from word to word. The power of trellis coding lies in its ability to track a signal's history through time and discard potentially faulty information (errors) based on a signal's past and future behavior.

This is somewhat like charting someone's travels by checking their travel itinerary. Imagine that a travel agent issues an itinerary for a trip to India. The itinerary calls for the following flights:

UNITED	FLT 100	Los Angeles – New York
UNITED	FLT 010	New York – Paris
AMERICAN	FLT 111	Paris – Bombay

However, a check of the master route schedules for the various airlines shows that there is no flight UNITED 100 from Los Angeles to New York, nor is there a flight UNITED 010 from New York to Paris. There are, however, flights on Delta Airlines matching these numbers from Los Angeles to Newark to Paris. The conclusion: Someone mistook New York for Newark and United for Delta. The trellis decoder in DTV receiver performs a similar function by using the 3-bit transition codes (transmitted as 8-level symbols) to reconstruct the signal trajectory from each 2-bit word to the next. When some of the 3-bit transition codes are corrupted during transmission, resulting in impossible "flight-destination" combinations, the trellis decoder will consider several alternative signal trails to find the most likely candidate.

SYNC & PILOT INSERTION

The next step in the signal processing chain is the insertion of the various "helper" signals that aid the DTV receiver to accurately locate and demodulate the transmitted RF signal. These are the ATSC pilot, segment sync, and field sync. The pilot and sync signals are inserted after the randomization and error coding stages so as not to destroy the fixed time and amplitude relationships that these signals must possess to be effective.

Recovering a clock signal in order to decode a received waveform has always been a tricky proposition in digital RF communications. If we derive the receiver clock from the recovered data, we have a sort of "chicken and egg" dilemma. The data must be sampled by the receiver clock in order to be accurately recovered. The receiver clock itself must be generated from accurately recovered data. The resulting clocking system quickly "crashes" when the noise or interference level rises to a point that significant data errors are received.

---

[1] The trellis coder in the DTV system actually leapfrogs ahead twelve symbols at a time to determine the next symbol transition. There are then twelve different trellis codes operating in parallel (e.g. symbol 0 links with 12, 24, 36…. symbol 1 links with 13, 25, 37… symbol 2 links with 14, 26, 38… etc.). This is yet another form of interleaving and offers some additional protection against burst-type noise. This scheme was designed to work well in conjunction with an NTSC interference rejection filter in the receiver that makes use of a twelve-symbol tapped delay line.

When NTSC was invented, the need was recognized to have a powerful sync pulse that rose above the rest of the RF modulation envelope. In this way, the receiver synchronization circuits could still “home in” on the sync pulses and maintain the correct picture framing - even if the contents of the picture were a bit snowy. (Everyone saw the need for this except the French; sync in France is the weakest part of the signal – *être comme le reste du monde, ça serait trop facile, quoi*). NTSC also benefited from a large residual visual carrier (caused by the DC component of the modulating video) that helped TV receiver tuners zero in on the transmitted carrier center frequency.

8-VSB employs a similar strategy of sync pulses and residual carriers that allows the DTV receiver to “lock” onto the incoming signal and begin decoding, even in the presence of heavy ghosting and high noise levels.

The first “helper” signal is the ATSC pilot. Just before modulation, a small DC shift is applied to the 8-level baseband signal (which was previously centered about zero volts with no DC component). This causes a small residual carrier to appear at the zero frequency (unmodulated carrier) point of the resulting modulated spectrum. This is the ATSC pilot. This gives the RF PLL circuits in the DTV receiver something to lock onto that is independent of the transmitted data. Although similar in nature, the ATSC pilot is much smaller than the NTSC visual carrier, consuming only 0.3 dB or 7 percent of the transmitted power.

The other “helper” signals are the ATSC segment and field syncs. An ATSC data segment is comprised of the 207 bytes of an interleaved data packet. After trellis coding, the 207-byte segment has been stretched out into a stream of 828 eight-level symbols. The ATSC segment sync is a four-symbol pulse that is added to the front of each data segment and replaces the missing first byte (packet sync byte) of the original MPEG-II data packet. The segment sync appears once every 832 symbols and always takes the form of a positive-negative-positive pulse swinging between the +5 and -5 signal levels (see Figure 2)[2]. Correlation circuits in the 8-VSB receiver home in on the repetitive nature of the segment sync, which is contrasted against the background of pseudo-random data (Remember the data randomizer processing stage). The recovered segment sync is used by the receiver to regenerate the system clock and sample the received signal. Because of their high frequency of repetition, large signal level swing, and extended duration, the segment syncs are easy for the receiver to spot. Consequently, accurate clock recovery can be had at noise and interference levels well above those where data recovery is impossible (up to 0 dB SNR - data recovery requires at least 15 dB SNR). This robust synchronization system, along with ATSC pilot, allows the DTV receiver to recover lock-up quickly during channel changes and other transient conditions. Figure 2 shows the make-up of the ATSC data segment and the position of the ATSC segment sync.

---

[2] The numerals {-7, -5, -3, -1, 1, 3, 5, 7} are used to represent the eight symbol levels. These are the eight smallest integer values that are both equally spaced and centered about zero. When modulation takes places, these numbers are proportional to eight levels of signal voltage. (That is, they represent voltage, as opposed to power)

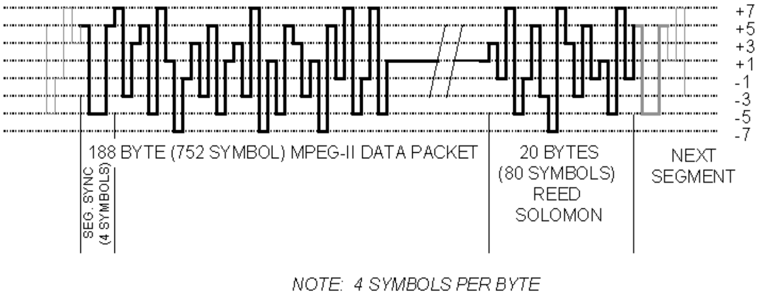

FIGURE 2: ATSC BASEBAND DATA SEGMENT

An ATSC data segment is roughly analogous to an NTSC line; ATSC segment sync is somewhat like NTSC horizontal sync. Their duration and frequencies of repetition are, of course, completely different. Each ATSC segment sync lasts 0.37 µsec; NTSC sync lasts 4.7 µsec. An ATSC data segment lasts 77.3 µsec; an NTSC line 63.6 µsec. A careful inspection of the numbers involved reveals that the ATSC segment sync is somewhat more "slender" when compared to its NTSC counterpart. This is done to maximize the active data payload and minimize the time committed to sync "overhead."

Three hundred and thirteen consecutive data segments are combined to make a data field. Figure 3 shows the make-up of an ATSC data field. The ATSC field sync is an entire data segment that is repeated once per field (24.2 msec) and is roughly analogous to the NTSC vertical interval.[3] The ATSC field sync has a known data symbol pattern of positive-negative pulses and is used by the receiver to eliminate signal ghosts caused by poor reception. This is done by comparing the received field sync with errors against the known field sync sequence before transmission. The resulting error vectors are used to adjust the taps of the receiver ghost-canceling equalizer. Like segment syncs, the large signal level swing and repetitive nature of field syncs allow them to be successfully recovered at very high noise and interference levels (up to 0 dB SNR).

---

[3] Note, however, that unlike NTSC, the ATSC syncs do not play any role in the framing of the displayed image on the receiver picture tube. This information is encoded digitally as part of the MPEG packet address information.

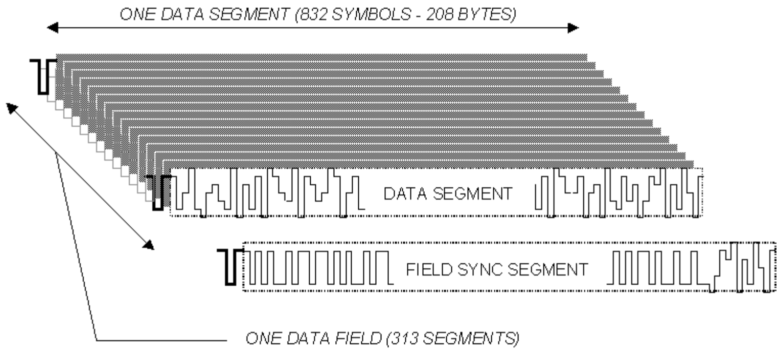

FIGURE 3: ATSC BASEBAND DATA FIELD

At the end of each field sync segment, the last twelve symbols from the last data segment are repeated to restart the trellis coder in the receiver.

The robustness of the segment and field syncs permits accurate clock recovery and ghost-canceling operation in the 8-VSB receiver, even when the active data payload is completely corrupted by poor reception conditions. This allows the adaptive ghost-canceling equalizer to "hunt around in the mud" and recover a useable signal before the data payload has been successfully decoded - thus eliminating the "chicken and egg dilemma" described earlier.

AM MODULATION

The eight-level baseband DTV signal, with syncs and DC pilot shift added, is then amplitude modulated onto an intermediate frequency (IF) carrier. This creates a large, double sideband IF spectrum about the carrier frequency, as is shown in Figure 5. The occupied bandwidth of this IF signal is far too wide to be transmitted in the assigned six MHz RF channel. Fortunately, there are tricks that can be employed to filter away a large part of this spectrum without destroying any of the vital digital information.

A simple inspection of Figure 5 reveals the high degree of redundancy in the double sideband IF spectrum. The various sidelobes are simply scaled copies of the center spectrum, and the entire lower sideband is a mirror image of the upper sideband. This makes it possible to discard almost the entire lower sideband and all of the sidelobes in the upper sideband without any loss of information. The remaining signal (upper half of the center spectrum) can be further cut in half by virtue of the Nyquist Theorem, which states that only a ½ frequency bandwidth is required to transmit a digital signal at a given sampling rate.[4]

The job of trimming the double sideband IF spectrum down to size falls to the next processing stage, the Nyquist filter[5].

---

[4] Reverse the logical order of this statement and you have the principal behind the minimum "2x frequency response" sampling rate for CD players and other digital audio devices.

[5] There are several different ways to implement the AM modulation, VSB filtering, and pilot insertion stages of the 8-VSB exciter, some of which are completely digital and involve direct digital synthesis of the required waveforms. All methods aim to achieve the same results at the exciter output. This particular arrangement was chosen in the interest of providing a clear, easily understandable signal flow diagram. Note: The Harris CD series of 8-VSB exciters employs all-digital signal synthesis.

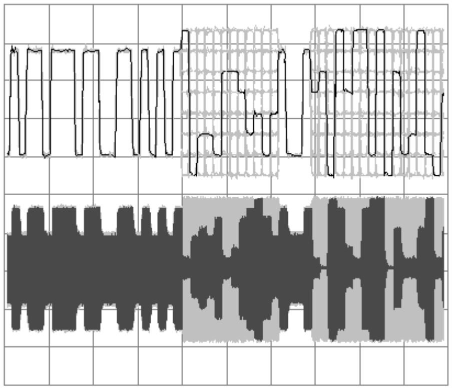

FIGURE 4: BASEBAND EIGHT LEVEL SIGNAL IS AM MODULATED ONTO IF CARRIER

Top: The eight level baseband signal at the end of a field sync segment. The light gray traces in the background show the history of the signal.

Bottom: A (different) field sync segment after AM modulation onto the IF carrier. Note how the modulated envelope is not symmetrical about the zero carrier. For example, the alternating +5 and –5 symbols to the left do not have same absolute RF amplitude after modulation. A small DC shift is added to the baseband signal before modulation. This creates a small residual pilot carrier at the unmodulated carrier frequency.

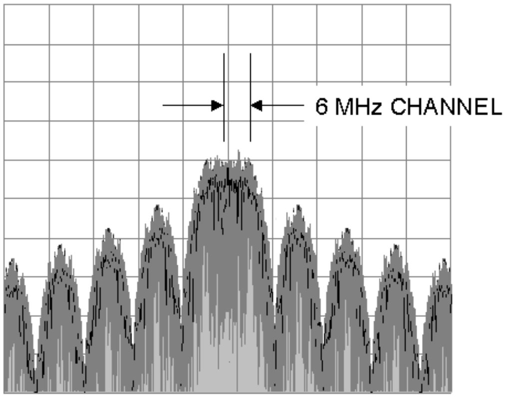

FIGURE 5: DOUBLE SIDEBAND SPECTRUM CREATED BY AM MODULATION

The double sideband RF spectrum created by the AM modulation of the baseband signal onto the IF carrier is far too wide to fit into a six MHz RF channel.

NYQUIST FILTER

As a result of the data overhead added to the signal stream in the form of forward error correction coding and sync insertion, the total data rate of the signal stream balloons from 19.39 Mbit/sec at the exciter input to 32.28 Mbit/sec at the output of the trellis coder. Since 3-bits are transmitted in each 8-level symbol, the resulting symbol rate is 32Mb / 3 = 10.76 Million symbols/sec. By virtue of the Nyquist Theorem, it is possible to transmit 10.76 Million symbols/sec in a vestigial sideband signal (VSB)[6] with a minimum frequency bandwidth of ½ * 10.76 MHz = 5.38 MHz. Since the allotted channel bandwidth is 6 MHz, it is possible to relax the steepness of the VSB filter skirts slightly and still fall within the 6 MHz channel. This permissible excess bandwidth (represented by α, the Greek letter alpha) is 11.5% for the ATSC 8-VSB system. That is, 5.38 MHz (minimum bandwidth per Nyquist) + 620 kHz (11.5% excess bandwidth) = 6.00 MHz (channel bandwidth used). The higher the alpha factor used, the easier the hardware implementation is, both in terms of filter requirements and clock precision for sampling.

The resulting frequency response after the Nyquist VSB filter is shown in Figure 6. Note how the 8-VSB format, like traditional NTSC, uses a vestigial sideband approach to conserve spectrum space. Unlike NTSC, however, 8-VSB takes this concept to greater extremes: the lower RF sideband is almost completely removed.

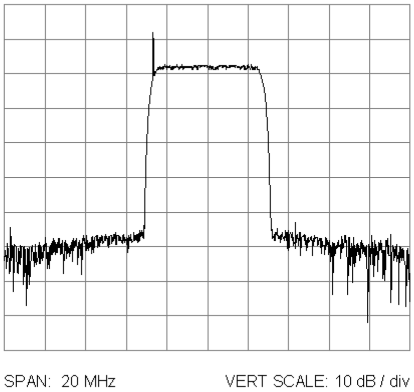

FIGURE 6: 8-VSB RF FREQUENCY SPECTRUM

Note presence of ATSC pilot at lower edge of channel. The lower sideband (below pilot frequency) is almost completely removed.

This virtual elimination of the lower sideband, along with the narrowband filtering of the upper sideband, creates very significant changes in the RF waveform that is ultimately transmitted. The 8-VSB IF envelope undergoes a transformation and loses the neat, “8-level stairstep” appearance it had before filtering. The train of “squared-off” symbol pulses that made up the double-sideband IF signal is modified by the impulse response of the narrowband Nyquist filter. This is shown in Figure 7.

---

[6] Note: 8-VSB = **8** level - **V**estigial **S**ide **B**and

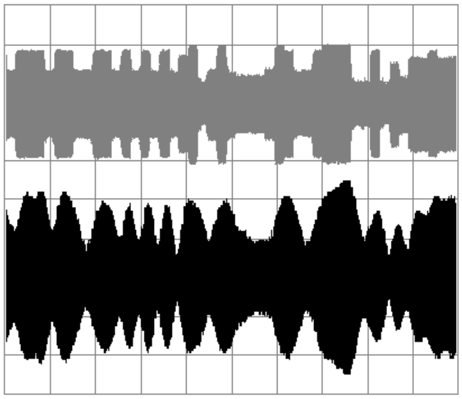

FIGURE 7: EFFECT OF NYQUIST FILTER ON 8-VSB IF ENVELOPE

Top: Double sideband IF envelope before Nyquist filtering.

Bottom: The same IF signal after Nyquist filtering. The squared-off transitions are lost and the envelope acquires a noise-like appearance.

After considering Figure 7, the natural reaction is to ask, “What happened to the eight levels?” and to wonder if all of the information contained in those eight levels has been lost forever. The answer is no; the following paragraphs explain why:

As any video or transmitter engineer knows, when a square signal pulse is frequency band-limited, it will lose its square edges and “ring” (oscillate) in time before and after the initial pulse. For the 8-level 8-VSB signal, this would spell disaster as the pre- and post-ringing from one symbol pulse would interfere with the preceding and following pulses, thereby distorting their levels and disrupting their information content.

Fortunately, there is still a way to transmit the 8-VSB symbol pulses, if we observe that the eight-level information is only recognized during the precise instant of sampling in the receiver. At all other times, the symbol pulse amplitude is unimportant and can be modified in any way we please -- so long as the amplitude at the precise instant of sampling still assumes one of the required eight amplitude levels.

If the narrowband frequency filtering is done correctly according to the Nyquist Theorem, the resulting train of symbol pulses will be *orthogonal*. This means that at each precise instant of sampling, only one symbol pulse will contribute to the final RF envelope waveform; all preceding and following symbol pulses will be experiencing a zero crossing in their amplitude at that point in time. This is shown in Figure 8. In this way, when the receiver clock samples the RF waveform, the recovered voltage will represent only the current symbol’s amplitude (one of the eight possible levels).[7]

---

[7] A subtle clarification: The 8-VSB system actually employs a matched pair of Nyquist filters – one in the exciter (to reduce transmitted bandwidth), another in the receiver (to eliminate adjacent channel interference). Each Nyquist filter provides one-half of the orthogonal impulse response described above. I.e. The rolloff in each filter is “half-strength.” The effect shown in Figure 8, therefore, does not fully exist when the RF signal is transmitted, but rather only after it is filtered for a second time in the receiver.

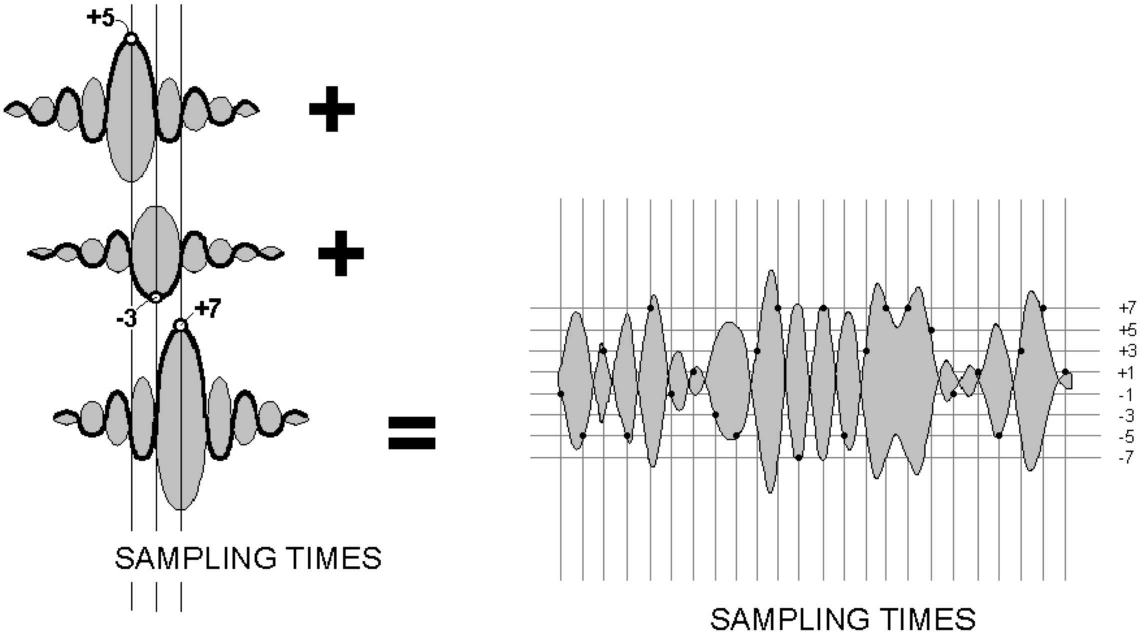

FIGURE 8: ADDITION OF NARROWBAND, ORTHOGONAL SYMBOL PULSES

At any given sampling time (vertical line), only one symbol pulse contributes to total signal amplitude, all other pulses experience a zero crossing. The resulting RF envelope corresponds to the eight digital levels only during the precise instant of sampling. Note: The symbol pulses are mirrored (double-sided) because we are now dealing with a modulated RF envelope.

At all times in-between the sampling times, the total RF envelope waveform is the addition of the “ringing” of dozens of previous and future symbols (since all symbols have non-zero amplitudes between sampling times). Note that, in the interest of simplicity, Figure 8 shows the narrowband symbol pulses as ringing for only 10 sampling periods; in reality they ring for a much longer time. These non-zero values (between sampling times) from dozens of symbols can add up to very large signal voltages. The result is a very “peaky” signal that most closely resembles white noise. This is shown in Figure 9. The peak to average ratio of this signal can be as high as 8 - 10 dB, although RF peak clipping in the transmitter can limit this value to 6 - 7 dB with minimal consequences.

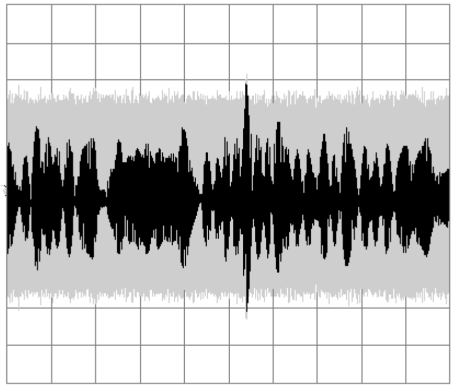

FIGURE 9: THE 8-VSB RF WAVEFORM AT THE EXCITER OUTPUT

The black regions represent the current oscilloscope trace. The gray regions show the stored value of all past traces.

8-VSB EYE DIAGRAM

One popular representation of the 8-VSB signal that emphasizes the principals just discussed is the 8-VSB eye diagram. This diagram turns up in many articles written about 8-VSB and on the screens of many pieces of 8-VSB test equipment. The eye diagram is the overlay of many traces of the received RF signal voltage at the instant of sampling. Since the RF signal must attain one of eight possible levels whenever a sampling time occurs (somewhat like finding one of eight chairs whenever the music stops in a game of musical chairs), the convergence of the many signal traces forms seven "eyes" that coincide with the occurrence of clock pulses in the receiver. This is shown in Figure 10.

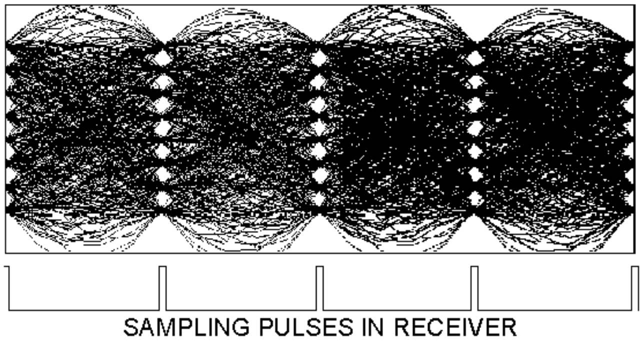

FIGURE 10: THE 8-VSB EYE DIAGRAM

At each sampling time, the demodulated RF amplitude assumes one of eight possible levels. The resulting display creates seven vertical "eyes." If the 8-VSB signal is corrupted during transmission, these "eyes" will close up and disappear, as the RF signal will no longer posses the correct amplitude at the right instant.

8-VSB SIGNAL CONSTELLATION

Another popular representation of the 8-VSB signal that is common to many pieces of test equipment is the 8-VSB signal constellation. This is a two-dimensional graphical representation of the 8-VSB RF carrier amplitude and phase at each sampling time.

In 8-VSB, the digital information is transmitted exclusively in the amplitude of the RF envelope and not in the phase. This is unlike other digital modulation formats, such as QAM, where each point in the signal constellation is a certain vector combination of carrier amplitude and phase. A QAM-like arrangement is not possible in a vestigial sideband system like 8-VSB, because the carrier phase is no longer an independent variable under our control, but is rather "consumed" in suppressing the lower sideband.

The 8-VSB signal constellation, as compared to 64-QAM, is shown in Figure 11. The eight symbol levels are recovered by sampling an in-phase synchronous detector (I channel axis).[8] The 8-VSB signal constellation diagram is therefore a series eight vertical lines that correspond to the eight transmitted amplitude levels.

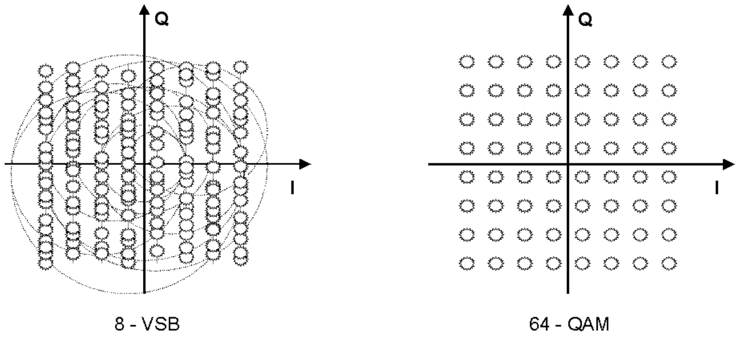

FIGURE 11: 8-VSB SIGNAL CONSTELLATION vs. 64-QAM

8-VSB signal constellation is a series of eight vertical lines on the I (in-phase) axis. The Q (quadrature) axis is not used to convey information. When the 8-VSB RF signal is corrupted, the eight vertical lines become blurred and errors are received. The light gray circular traces added to 8-VSB constellation show the instantaneous RF carrier amplitude and phase in a state of constant change; sampling in the receiver is like a strobe that "catches" the signal as it passes one of the eight amplitude levels.

ANALOG UPCONVERSION AND THE REST OF THE 8-VSB CHAIN

After the Nyquist VSB filter, the 8-VSB intermediate frequency (IF) signal is upconverted by traditional oscillator-mixer-filter circuits to the assigned channel frequency in the UHF or VHF band. The on-channel RF output of the 8-VSB exciter is then supplied to the DTV transmitter. The transmitter is essentially a traditional RF power amplifier, be it solid state or tube-type. A

---

[8] The synchronous detector is locked in phase with the ATSC pilot phase. Because the pilot is a tiny residual potion of the original unmodulated RF carrier, it retains the information as the reference (unmodulated) carrier phase. It also allows the receiver to determine the polarity of the recovered symbols from their instantaneous carrier phase. For example, in-phase with pilot equals positive symbol (e.g. +7), in-antiphase with pilot equals negative symbol (e.g. –7).

high-power RF output system filters the transmitter output signal and suppresses any spurious out-of-band signals caused by transmitter non-linearities. The last link in the transmitting chain is the antenna, which broadcasts the full-power, on-channel 8-VSB DTV signal.

In the home receiver, the over-the-air signal is demodulated by essentially applying in reverse the same principals that have already been discussed. The incoming RF signal is received, downconverted, filtered, then detected. The segment and field syncs are recovered. Segment syncs aid in receiver clock recovery and field syncs are used to train the adaptive ghost-canceling equalizer. Once the proper data stream has been recovered, it is trellis decoded, de-interleaved, Reed-Solomon decoded, and de-randomized. The result is the recovery of the original MPEG-II data packets. MPEG-II decoding circuits reconstruct the video image for display on the TV screen. Dolby AC-3 circuits decode the sound information and drive the receiver loudspeakers. The home viewer “receives his DTV” and the signal chain is complete.

CONCLUSION

The goal of this article has been to provide some insight into the inner workings of the 8-VSB transmission system. Like many things in life, 8-VSB can appear formidable at first, but is really quite simple “once you get to know it.” Hopefully, the knowledge conveyed in this article will dispel some of the fear factor that many NTSC engineers experience when faced with the unknown world of digital TV.

ACKNOWLEDGEMENTS

Special thanks go to Bob Plonka, Joe Seccia, and Ed Twitchell of Harris Corporation Broadcast Division for their contributions of time, material, and assistance in writing this article. Special thanks also go to Karl Black, Robert Uhl, and Eli Garcia for their proofreading assistance.

REFERENCES

[1] “The Harris VSB Exciter for Digital ATV” by Robert Davis and Edwin Twitchell. NAB 1996 Engineering Conference. April 15 - 18, 1996

[2] “ATSC Transmission System: 8-VSB Tutorial” by Richard Citta and Gary Sgrignoli. ITVS 1997 Montreux Symposium. June 12 - 17, 1997

[3] “The Fundamentals of Digital ATV Transmission” by Ron Totty, Robert Davis, and Robert Weirather. *ATV Seminar in Print*. Harris Corporation Broadcast Division, 1995.

AUTHOR

David Sparano is currently Principal Engineer with Harris Corporation Broadcast Division in Quincy, Illinois.