9 Digital

Digital signals are nonlinear. The transition area between the high and low states of a digital signal is undefined when determining the value of an individual bit. To maximize the chances of reliable detection, we want sufficient signal amplitude so noise does not cause errors, and small amounts of attenuation (a few hundred feet of cable loss) do not cause the highs and lows to fall into the undefined area.

The transitions, or "edges", between states are just as important. These transitions enable clock recovery from the bit stream in a self-clocking signal such as SMPTE 259M. Without a clock to use at the receiving end, there is no way to know when to check the status of an arriving bit. An algorithm is used to scramble data as it leaves the transmitter to create as many edges as possible. This assists the receiver's PLL circuitry in generating a local clock synchronized to the transmit clock.

The error correction and error masking in modern digital equipment insures that digital signals do not gradually degrade with increasing attenuation in the signal path as analog signals do. Instead, a digital transmission path continues to work perfectly, up to the point where it suddenly does not work at all. This is the well known "Cliff Effect", illustrated in figure 1.

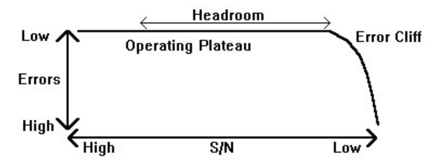

Figure 1

SDI (Serial Digital Interface) signals that are experiencing few or no errors are somewhere on the operational plateau shown in Figure 1. Operation stays uneventful until you reach the "Error Cliff". As the path traverses over the knee of this cliff, errors go rapidly from non-existent to enough to swamp recovery efforts, making the path unusable. As little as 3 extra feet of coax can be enough to send a signal over the cliff. Many things determine where you are on that operational plateau.

Bandwidth & Signal Requirements

Although the way the SDI signal is used is "digital" in nature, many "analog" attributes of the signal can be used to predict how close to the error cliff a particular digital path is. Figure 2 shows a typical SDI signal. The portion shown here is of three successive "1"s in a serial video digital bitstream.

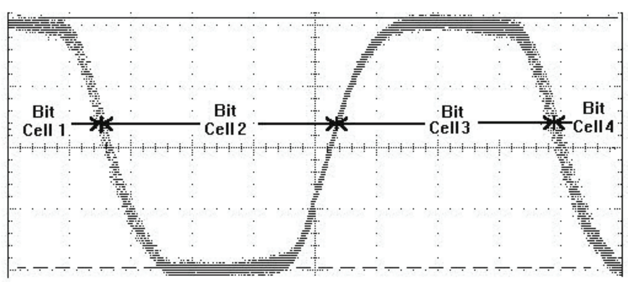

Figure 2

A digital video data stream changes state at the start of each bit cell if the bit cell has a data value of "1". The three transitions shown here indicate the transition changes at the start of each bit cell. This transmission scheme is known as NRZI (Non Return to Zero Inverted), which means that the receiver doesn't need to worry about the polarity (high or low) of the incoming bit stream. This approach also yields a constant high or low if a string of zeros is encountered, hence the use of a bit scrambling algorithm. The peak to peak value of this signal should be 0.8V, and the risetime, or transition time from low to high, or vice-versa should be between 20 and 40% of total cell time. If the transmission path of this signal had infinite bandwidth, and no phase delay, or group delay, it would be a perfect squarewave.

Figure 2 shows three successive bit cells, two of which, in succession, constitute a positive and negative half of a square wave with a frequency of ½ the bit rate. So for HD, that frequency is 1500 mbps/2 or 750 MHz. Therefore, that 750 MHz squarewave has a 750 MHz sinewave as its fundamental (Figure 3).

Square Wave Fundamentals

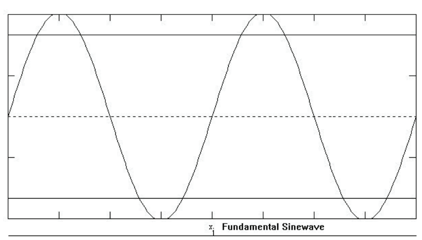

*Figure 3 If the third harmonic is added to our fundamental (3 * 750 MHz) the waveform in figure 4 is produced.*

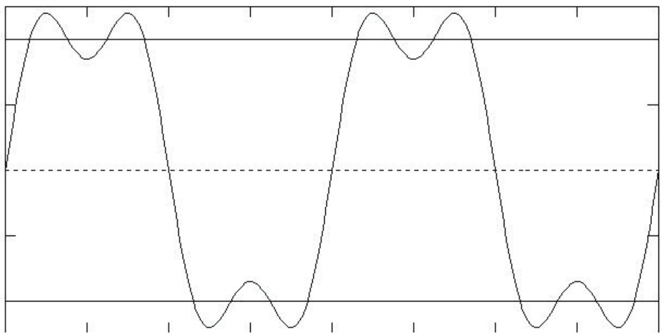

Figure 4 If some additional odd harmonics are added in the correct amplitude and phase the waveform in Figure 5 results.

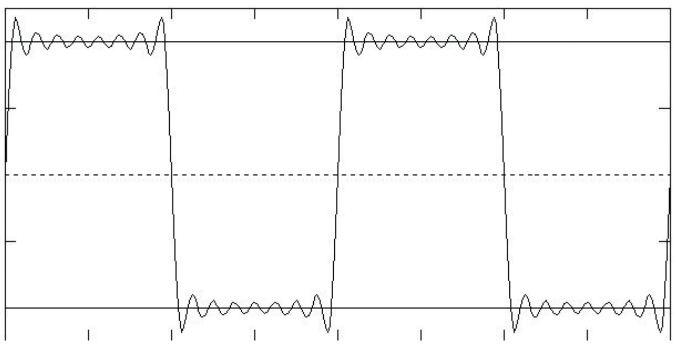

Figure 5 Notice that we only worried about the odd harmonics in building the above square wave. A square wave with 50% duty cycle (that is half the time high and half the time low) has the fundamental frequency, plus the 3rd, 5th, etc. or the odd harmonics. The 2nd, 4th, etc. are missing.

Notice that only the odd harmonics are needed in building the above square wave. Figure 6 shows why.

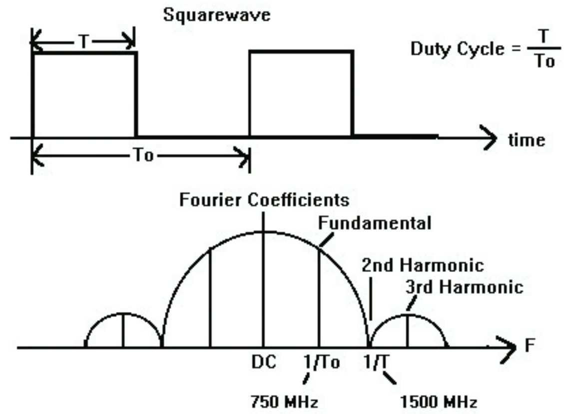

Figure 6 A square wave with 50% duty cycle has the fundamental frequency, plus the 3rd, 5th, etc. or the odd harmonics. But notice, the 2nd, 4th, etc. are missing. A pulse with its duty cycle less than 50% will have the response demonstrated in Figure 7.

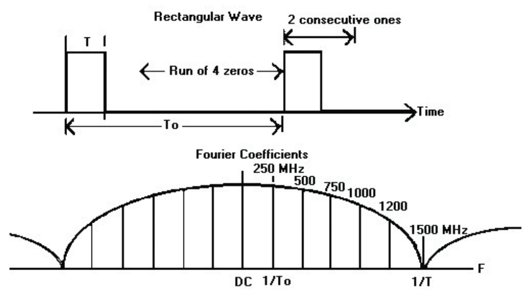

Figure 7 It takes many more harmonics than are needed for a simple square wave to describe the function at the top of Figure 7.

Thus, if only a continuous string of ones are sent, the spectrum shown in Figure 6 would occur. The fact of the matter is that various runs of "ones" and "zeros" are sent, which results in a spectrum more resembling Figure 7 being sent. It takes a lot more low frequency energy to create a signal with the duty cycle shown in Figure 7, versus the one in Figure 6.

Many things still happen at traditional television rates. Pairs of SAV, and EAV (Start of Active Video, End of Active Video timing reference signals) occur at the horizontal line rate. The patterns encountered during the vertical interval still occur at the field rate. This ensures that considerable energy will occur at fairly low frequencies.

Spectral Analysis

The weak link in most serial digital systems is the path from the transmitter in one box to the receiver in the next box. The physical layer used to transport the data from one "box" to the next is comprised of mostly coax. Some connectors and perhaps a jackfield might be included in an average path. But coax provides the greatest exposure to problems for a video data stream. Coax can be thought of as an infinite network composed of inductive and resistive components in series, with distributed shunt capacitance to ground. This works out to be a low pass filter whose poles increase in number, and move closer to zero with length. This means that the longer the cable, the greater the attenuation of all frequencies, with the roll off increasing as a function of frequency. As an illustrative example, if a given coax had the response of Figure 8:

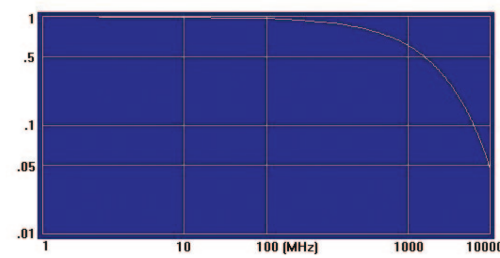

Figure 8 It would show the response of Figure 9 if its length were increased by a factor of ten.

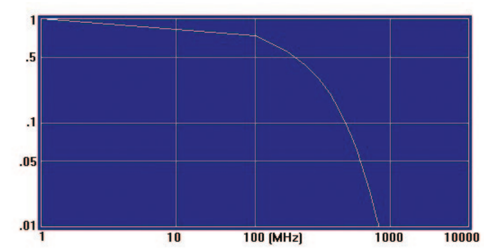

Figure 9

It is obvious that another decade increase would not be good. Since the signal is attenuated as the frequency increases, the upper harmonics of our signal disappear, and our squarewave data signal starts to look more like a sinewave. Additionally, with all the low frequency (DC) energy still available our squarewave starts to look like figure 10.

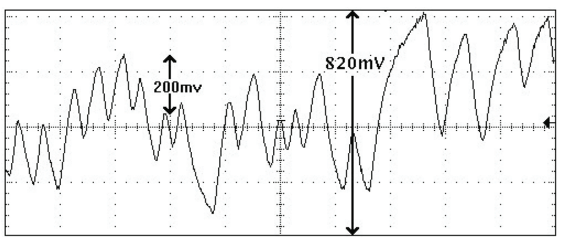

Figure 10

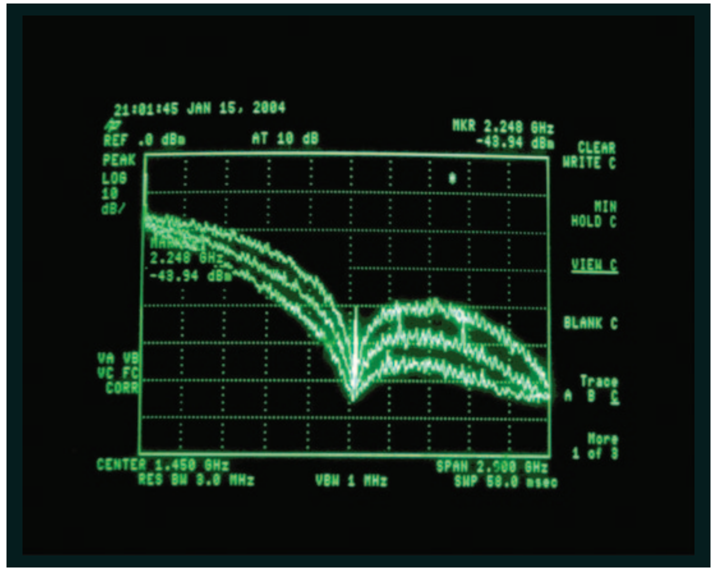

Figure 11 Figure 11 shows super-imposed spectrum waveforms at distances of 50, 100, and 150 feet over 1855A coax cable. The left half of the displays shows the fundamental energy band, and the right half shows the third harmonic energy band.

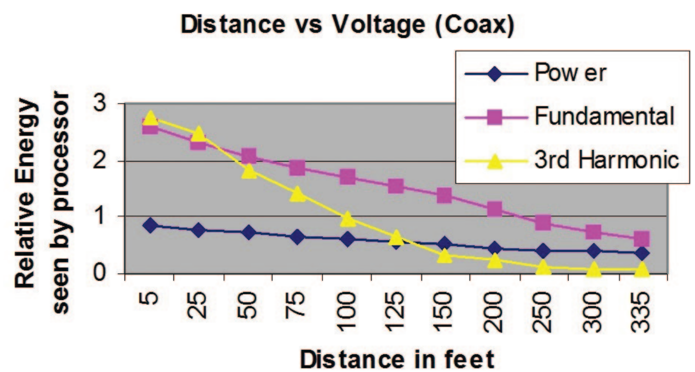

Figure 12 Figure 12 is a chart that shows how quickly the third harmonic rolls off compared to the fundamental. When the third harmonic is gone, the digital squarewave is now a sinewave and no useful information will be gleaned by the receiver.

For weight issues, mini-coax such as Belden 1855A cables are often used on trucks. The physics of smaller diameter cable results in increased losses as a function of frequency. Whereas Belden 1694A has 5 dB loss per 100 feet at 750 MHz, 1855A has 9.59 dB of loss. At 2250 MHz (the center of the third harmonic band) 1694A has 9.14 dB per hundred feet, while 1855A has 16.0 dB over the same distance.

Figure 11 shows that while the fundamental energy drops as distance increases, the third harmonic energy drops at a much greater rate. Figure 12 shows this reduction in energy graphically. While overall power will drop by only half as cable is increased from 0 to over 300 feet, the fundamental will drop to a third of its original value, and the upper harmonic, the energy that allows the HD serial bit stream to be recovered, will fall to the noise floor over that distance.

Coax that is robust at high frequencies is called for in a digital plant. The center conductor should be solid copper (better skin effect then stranded), and the shield should be braided (with a coverage figure near 100%), and also have a layer of foil, which has better skin effect at higher frequencies. The dielectric should produce as low a shunt capacitance value as possible, decreasing the high end roll off, and increasing the velocity of propagation. A trade off here is that some dielectrics do this by using air pockets to lower the dielectric constant, which can lead to center conductor migration, affecting the impedance along the length of the cable, especially at sharp bends, leading to reflections.

Another issue involving paths and frequency is impedance. The impedance is determined by the series inductance and shunt capacitance of cable, connectors, and jackfields. The problem is that the capacitance and inductance can change with frequency. Usually that change is not much. But occasionally spot damage to cable, or a connector not installed correctly, or a jackfield that did not normal up as well as it should have, can cause impedance to literally spiral way out of tolerance causing what will be a "flakey" path.

As can be seen in Figure 13, the resistance, and reactance is changing very little as a function of frequency. This is a polar plot of a healthy jackfield. Coax cables have plots that are generally tighter around the origin than this. When the trace is not on the X axis the total impedance will not be equal to the resistance. When impedance in one segment of a digital path is not equal to impedance of the next segment of that path, reflections result. Anytime the plot on a polar chart is not tightly centered around the origin, reflections will result. When the impedance moves farther away from the origin, reflections result in reflected power back to the source, which lowers Power Transmitted, which leads to transmission loss.

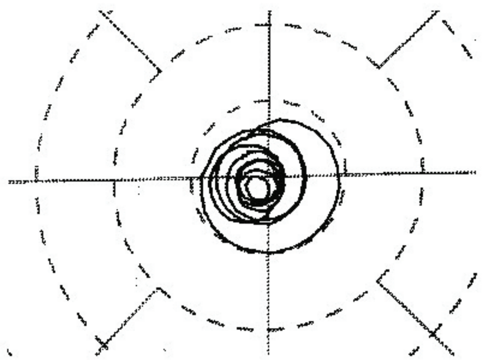

Figure 13 The plot in Figure 14 shows a jackfield where the reactance function is much worse than the one shown in Figure 13. As can be seen here at certain points on the plot the reactance starts to become a significant enough value that the total impedance verses the resistance will be such that significant reflections can occur along that path.

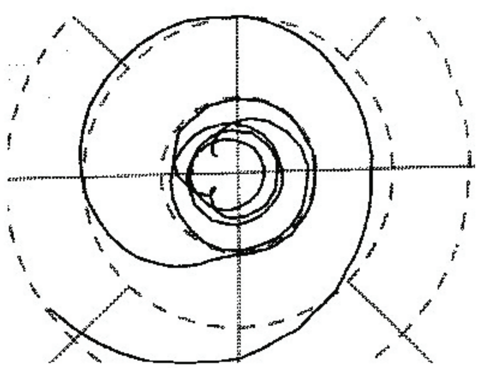

Figure 14

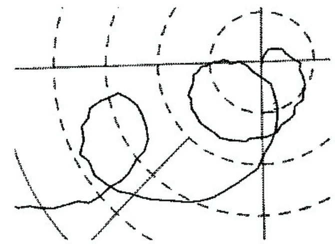

Figure 15 Figure 15 represents a horror story when it comes to reflections. This has significant reflections by the 270 MHz point (near the bottom Y intercept). This plot is of a jackfield that would greatly reduce the distance a digital signal could travel without errors.

As can now be seen, reflections can totally shut down a digital path, even when all the other attributes of the path are normal, and healthy. In fact, a casual observer of the eye pattern generated by the data through the jackfield in Figure 15 would probably notice nothing wrong. If you are having problems in a digital path, and don't have the luxury of owning a network analyzer, or even a Time Domain Reflectometer, you should remain critical of the analog attributes of the eye pattern. Just because we are now "digital", group delay, ringing, and pre/post overshoots are still important and should be monitored.

Triax is coax with an extra shield. The extra outside shield helps protect from outside RF interference or ingress into the cable. Many camera systems also use this outside shield as part of the camera power up safety check.

Triax has all the same attributes as coax. As such, a kink in the cable, blunt impacts by vehicles and carts running over it can affect its reliability, and usefulness. Changes in impedance can result in reflections, and while a single reflection isn't usually a problem, a number of them in the same cable or path can cause many strange and seemingly inexplicable phenomena.

This is because two or more impedance anomalies will cause some energy to be reflected back in the opposite direction. So if a signal encounters one of these, some energy is reflected back while the majority continues on. Now if a section impedance change is encountered, some energy is reflected back again. This energy reflected back at the second defect will hit the first where some percentage will be reflected back towards the second defect. As you can see, we will rapidly encounter a ringing effect having a fundamental wavelength based on how far the two defects are apart. Imagine if there are even more defects along the path. While the effects were easy to see with analog, remember ghosting? With digital, the way they will show up is if the wavelength of these signals affects some aspect of the digital data traversing the path.

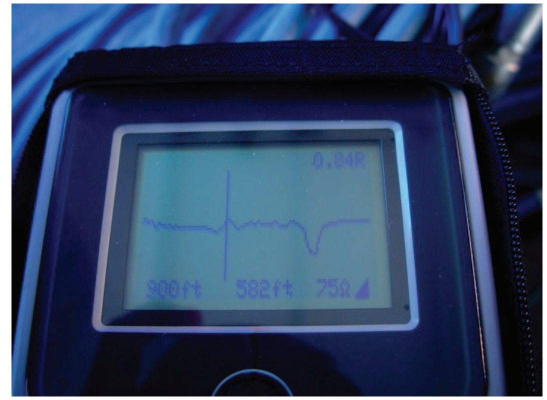

Figure 16 Figure 16 shows a Time Domain Reflectometer (TDR) looking down a 1000 feet of triax. The display shows that it sees an impedance anomaly 582 feet down that cable. A TDR works by sending a pulse down a cable and then times how long until any reflected energy arrives back at the TDR.

Organizing the Data

We looked at what the path can do to data. Now we'll look at how that data is used to represent digital video. Like their analog ancestors, digital video signals have many of the same components. The main component is the video payload, timing signals that let receiving equipment know when a new line or field is to start, with enhanced ways of putting additional data, metadata, in the horizontal and vertical blanking intervals.

We've already covered chroma sub-sampling but we will mention it here to re-enforce that while various compression and file formats use a number of formats such as 4:1:1 and 4:2:0, baseband digital video, which is all this section is about only uses 4:2:2. That means that active video samples are comprised of ½ luminance samples, and ½ chroma samples. Of the 3840 video samples per line of 1080i, or 2880 for 720p, half (1920 or 1440) are luminance and half are chroma of which ½ the chroma are for CB (B-Y) and the other half CR (R-Y).

The CB and CR samples are co-sited with the odd (1st, 3rd, 5th,etc.) Y samples. In the case of 1080, this creates 1920 multiplexed Y and CB \CR active video samples. Thus sample output data sequence for each horizontal line is CB,Y,CR,CB,Y,CR,Y…… The last sample in the line is a Y. Increasing data values for each luminance data means a brighter sample. Black is 40h and 3ADh is peak white. This allows for 877 discrete levels for luminance. Chroma is centered at 200h because originally analog chroma signals had positive and negative values. The highest chroma value is 3C0h and the lowest value is 40h. Chroma has 896 discrete levels.

Sources, such as cameras or graphics devices, use the following formulas for luminance and chroma signals:

Y = 0.21R + 0.72G + 0.07B

Cr = 0.63 (R-Y) = 0.63((1R + 0G + 0B) – (0.21R + 0.72G + 0.07B)

= 0.63((1R – 0.21R) + (0G – 0.72G) + (0B – 0.07B))

= 0.63 (0.79R – 0.72G – 0.07B)

= 0.50R – 0.45G – 0.04B

Cb = 0.54 (B-Y) = 0.50B – 0.39G – 0.11R

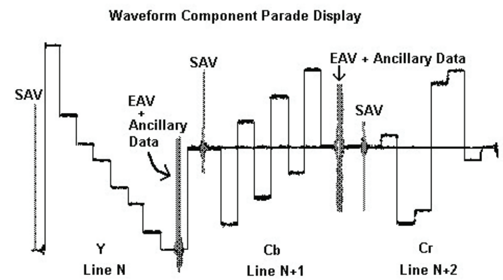

Figure 17 Analog representation of serial component digital video.

In place of sync pulses, groups of words called Timing Reference Signals (TRS) are inserted. There are 2 types of TRS strings, SAV (Start of Active Video) and EAV (End of Active Video). As their names imply they signify where the active video begins and ends. Active video and any ancillary data are not allowed to have values 000h – 003h, or 3FCh – 3FFh. Only TRS and Ancillary data headers can have values in those ranges. The first three TRS words are 3FFh, 000h, 000h. The 4th TRS word is a set of 3 flags. This word, referred to in the standard as the XYZ word uses 3 bits to indicate whether this TRS signal is:

1) An SAV or an EAV

2) In the vertical interval or not

3) Field 1 or Field 2

Most of the rest of the bits in the XYZ word are protection bits to prevent it from having illegal values below 003h and above 3FCh.

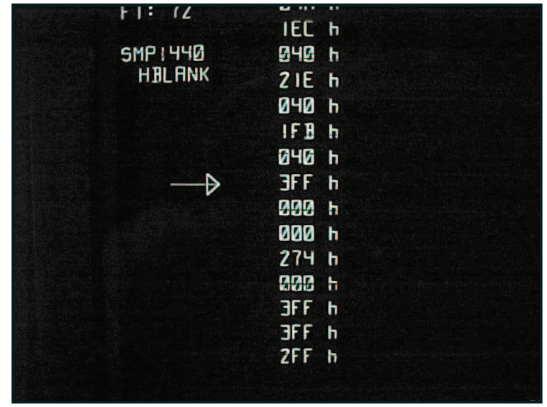

Figure 18 Hex decimal representation of EAV sequence in a SD signal

Figure 19 Same EAV sequence in Binary

The space between the EAV and SAV represents traditional horizontal blanking. 268 words are available during this time for use by ancillary data. The header for this data is 000h, 3FFh, 3FFh. Like the video Horizontal ANCillary (HANC) data consists of 10 bit words. The horizontal lines during the vertical interval can be used for ancillary data (VANC). Over the years, there has been some confusion about the use of VANC because early D1 VTRs, which where only 8 bit machines, used this data space. Originally to stay consistent digital standards specified 8 bit VANC. Therefore VANC was not used much.

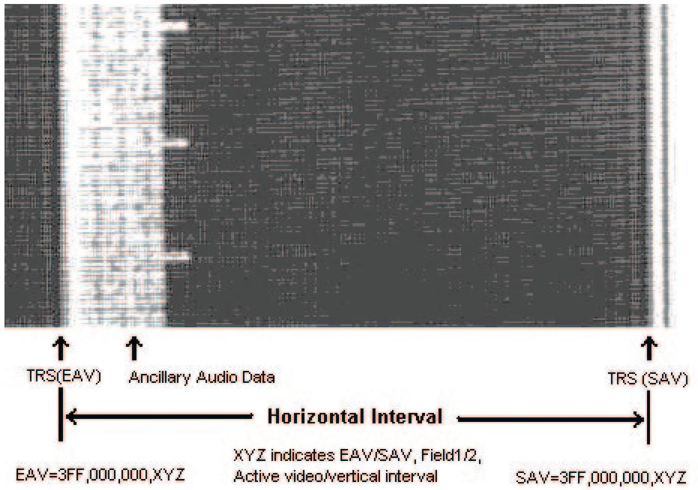

Figure 20 Video monitor in "Pulse-Cross" mode displaying embedded audio in the horizontal blanking interval. A common practice today is to embed the audio within the video—Figure 20 shows where the audio resides.

Since horizontal ancillary data in digital video is 10 bit words, and an AES frame consists of 64 bits containing up to two 24-bit audio words, some re-formatting must occur. This reformatted data will be inserted immediately after the EAV (End of Active Video) data words. They consist of three 10 bit words with the values 000, 3FF, 3FF (as opposed to 3FF,000,000 for EAV/SAV timing reference signals). Again only the EAV/SAV and ancillary data headers are allowed to have these values.

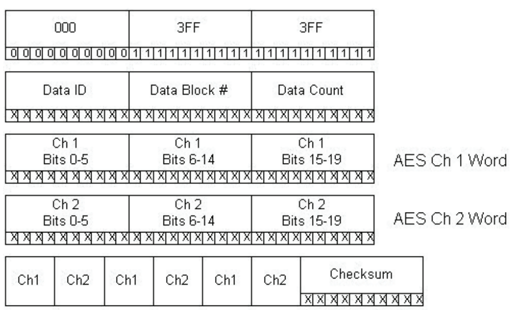

Figure 21 This chart shows how four samples of 16 – 24 bit AES audio can be mapped into 10-bit HANC space.

1440-1443	3FF	000	000	26C
1444-1447	000	3FF	3FF	2FF
1448-1451	200	212	1A0	2D3
1452-1455	21E	1A2	2D3	11E
1456-1459	170	2E8	21E	172
1460-1463	2D3	11E	198	2CC
1464-1467	21E	19A	2CC	11E
1468-1471	129	040	200	040
1472-1475	200	040	200	040
1476-1479	200	040	200	040

TRS (EAV)
Ancillary Data Preamble
Data ID
Data Block #
Data Count

Figure 22 This chart shows actual samples of AES audio being mapped into SD HANC space.

In Figure 22 you can see that the Ancillary Data Preamble immediately follows EAV. Next is the Data ID word. Its value indicates whether the audio data packets that follow are audio group 1 (channels 1-4), group 2 (channels 5-8), group 3 (channels 9-12), or group 4 (channels 13-16), or one of 4 extended audio data packets.

Next is a word that is the data block number. This number was meant to be incremented with each packet. Only 8 of the 10 bits are used to count to prevent illegal values, such as values reserved for TRS, and ancillary data preambles. The count value rolls back to 0 after reaching FF hex.

Some equipment uses a non-sequential value here from the previous packet to indicate that a switch has occurred in the audio data to a new source. A problem here early on, was that some products, such as some test equipment, did not increment this number with each packet. This caused intermittent muting in some receive equipment. If this value is not incremented, it is now kept at a value of 0h to eliminate any chance for confusion.

The next word indicates the number of data words to follow. A maximum of 255 bytes are possible because as with the data block number, only 8 of the 10 bits are used. The number of data words will depend on what horizontal line is being embedded with this data. Here is where it gets slightly confusing.

AES 20 bit samples are mapped into 3 successive 10 bit SMPTE259 words. The extra 10 bits left over ((3 * 10 bits) - (20 bit audio sample) = 10 bits) have AES audio channel status, and user bits mapped into four of them. Three bits indicate whether this is channel 1, 2, or a Z frame, which indicates the start of 192 sample AES block, and the final 3 are used to protect against illegal values, which again are reserved for EAV/SAV, or ancillary data preambles. This means that a single sample of channel 1, and 2 (1 AES frame) requires 6 SDI words. One way to implement this is to have most lines carry 6 AES samples. That's 36 SDI data words.

But the most common sampling frequency used in professional digital audio for television is 48 KHz (and the only one discussed here). This is not an exact multiple of the 29.94/59.94 field rates used commonly. This means that 8008 samples are required in the SDI stream every 10 fields. This works out to the necessary 48,000 necessary each second.

To spread these through the horizontal blanking, or horizontal ancillary data space, if most lines have 6 samples, every 12th line will have 8.

What about 24 bit audio? The first 20 bits are sent as just described, but the additional 4 bits of both channels in one AES frame are combined into 1 word and assembled into what is called an extended data packet. This extended data packet is sent immediately following the normal audio data packet. The extended data packet starts with the same preamble:

> (0000,3FFF,3FFF),
> data ID,
> Data Block Number,
> and Data Count words as the normal audio data packet.

Analog video is rapidly disappearing but many of the analog attributes still found at the digital physical layer lives on. Even the way digital video is organized is heavily reliant on methods pioneered during the analog era. Next we will look how all this data is organized into a system.