

Reprint of the article

“Measuring Digital Audio”

from the November 2002 issue
of Broadcast Engineering Magazine





Pat Shearer, chief engineer at KWBP in Beaverton, OR, monitors the quality of the station's audio signals from the studio control room.

Digital audio measurements for broadcast

BY WAYNE JONES

Transmitting and storing audio signals in the digital domain is well-established in the broadcast industry. Analog audio has given way to the AES3 and Sony/Phillips Digital Interface Format (SPDIF). AES3 datastreams are also embedded in SDI television signals.

Handling audio in the digital domain offers many advantages over analog methods. An analog signal

studios had a large investment in signal-transmission infrastructure, specifically in two-conductor, shielded cables interconnecting systems, equipment and studios. So the standard uses a self-clocking, polarity-insensitive technique to allow studios to transmit digital audio over these existing cables.

Conventional digital signals at the circuit level commonly define a logic “high” or “one” as a particular absolute voltage

The notion that a signal comprised only of ones and zeros is immune to degradation is seductive, but misleading.

incurs progressive degradation as it passes through a chain of circuits. Converting the analog signal into digital and converting back to analog as late in the chain as possible overcomes this degradation.

But keeping audio in a digital format does not automatically guarantee perfection. The notion that a signal comprised only of ones and zeros is immune to degradation is seductive, but misleading. Digital signals are affected by crosstalk, noise, cable length, poor circuit design and other factors that can ultimately translate into audible problems.

Many of these mechanisms of degradation are not obvious. This article describes what can go wrong, what effects to look for and what you can do about it.

First, let's take a look at the digital audio signal itself so we can better understand how problems can develop. The AES3 digital interface format was initially designed to simplify the transition from analog to digital in a studio environment. Its designers recognized that

level, perhaps +5V, and a “low” or “zero” as 0V. Such an interface format is polarity-sensitive and would be difficult to use in a studio environment that might have occasional inadvertent polarity inversions in infrastructure wiring.

To circumvent this problem, the AES3 interface relies on level transitions rather than absolute levels to define logic states. It defines a “unit interval” or UI as the smallest quantized time interval in the format. At the common 48kHz sample rate used in professional digital audio, this interval is 163ns long. The standard says that if the signal level remains unchanged for two UIs, the logic level is a zero; if it changes state from one UI to the next, the logic level is a one.

AES3 defines the actual protocol of the serial datastream representing the digitized audio signal. Figures 1 and 2 illustrate the sequence of the serial data bitstream. It starts with a preamble — a series of bits that identify the start of a frame, a subframe or a block. The next 24 bits are reserved to represent the audio signal. Finally, at

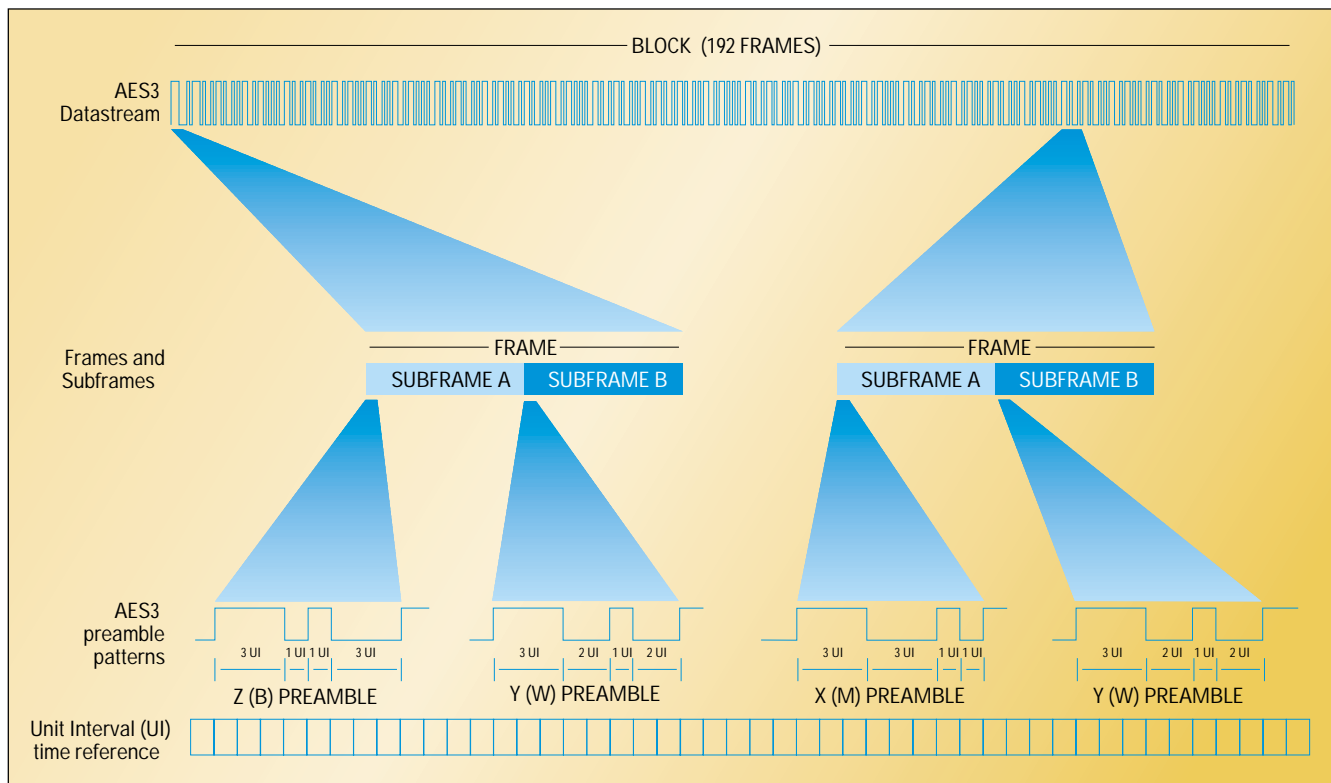


Figure 1. This breakdown of the AES serial data bitstream reveals the frames and subframes and their components.

the end of the sequence are administrative bits that identify characteristics about the data itself, such as sample rate, consumer or professional format, linear or compressed audio, and other information.

What can go wrong?

Some common errors include incorrect sample rate, excessive jitter, insufficient amplitude, poor data integrity caused by incorrect termination, or the use of poor cable. There can be problems in the analog-to-digital conversion process itself, although the quality of converters has progressed substantially. The most common design implementation used for converters these days uses an over-sampling technology. This method can substantially reduce residual noise and distortion by using a noise-shaping technique that shifts the noise upwards in the spectrum beyond the audio band. But, in return, it does produce substantial “out-of-band noise” that can cause problems in some areas. Good practice suggests that these out-of-band components be filtered out at

the source before they can pollute subsequent devices. Being aware of possible artifacts can help when troubleshooting obscure problems.

Sample rate shouldn't be a problem — provided you use the correct rate. Two sample rates are in common use: 48kHz in the professional environment and 44.1kHz in the consumer world. One scenario that can cause problems is if a faulty transmitting device has a sample rate sufficiently different from what it should be. This could prevent a subsequent device from locking onto the signal. Or perhaps someone is using the incorrect sample rate, perhaps 44.1kHz instead of the required 48kHz. Sample-rate converters can correct the latter problem.

Jitter is perhaps the biggest problem in digital audio transmission. As mentioned earlier, the AES3 bitstream is self-clocking: The

AES3 receiver derives its clock from the transitions of the datastream itself. If the interface pulses received were perfect rectangles, the time of the fast-rising vertical transitions would be clearly defined. But because the cable has capacitance and the AES3 transmitter has finite source impedance, the level transitions have a finite rise time. Since modulation of the datastream typically produces asymmetrical level states, the DC level of the interface waveform shifts with the data content.

Figure 3 shows how these facts affect the zero-crossing detection time. As the DC value moves up and down with time, the transition time, and thus the embedded clock edges, vary. This variation from cycle to cycle results in




Figure 2. Each block of audio data contains a preamble, the audio information and administrative data.

a constantly varying phase shift called cable-induced clock jitter. This problem can be reduced by using a cable specifically designed for AES3 transmission, rather than standard microphone or analog audio cable. Most cable manufacturers make such a cable with the correct 110Ω impedance. It's important to have the cable properly terminated at the far end. Data-pulse integrity, unlike analog-signal transmission, is affected by proper source, cable and termination impedances. Figure 4 shows an AES3 datastream with and without proper termination.

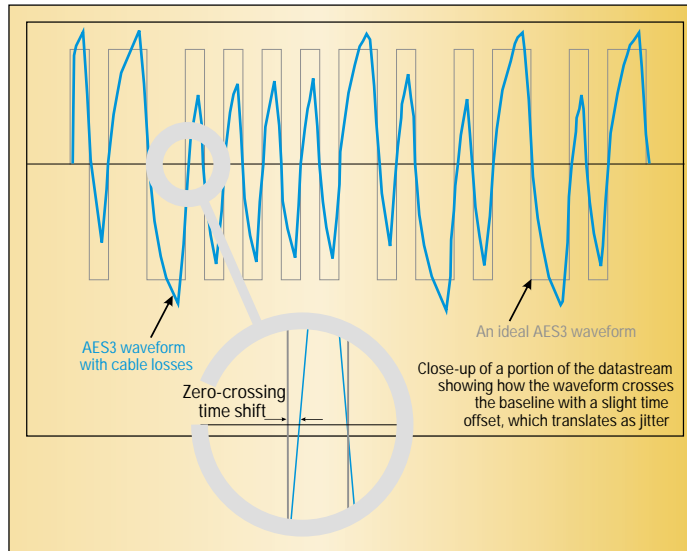


Figure 3. This diagram shows how the cable's capacitance and the transmitter's finite source impedance affect the signal's waveform, which in turn affects the zero-crossing detection time.

Jitter can also be caused by noise or crosstalk added to the datastream. The noise or crosstalk will cause ambiguity in the zero-cross transition of the data pulses, again causing jitter. In a broadcast facility, program synchronization is very important. To ensure a common time reference, "house sync" is typically distributed and used by all systems. If there is jitter on this clock signal, it can be transferred to any audio device that uses it in the process of trying to maintain audio/video synchronization. Another phenomenon, "jitter accumulation," occurs when several digital devices are cascaded — a common situation in a broadcast facility. Each device can pass on the jitter it receives while

signal would reveal any crosstalk and noise components due to jitter accompanying the signal. The digital-to-analog converter usually uses the clock signal extracted from the digital datastream as its sampling clock. In such cases, the jitter will modulate the conversion process. This can raise the noise floor or add unwanted frequencies to the audio signal.

More seriously, if jitter reaches too high a level, some data receivers will begin to malfunction, eventually losing lock. Often, this situation may occur in a large facility due to a particular interconnection of several subsystems with the resulting jitter-accumulation effect. This makes after-the-fact maintenance diffi-

are agile enough to follow the jitter in the signal and display a stationary waveform. Fortunately, there are several diagnostic and analytical tools available. Stand-alone AES3 data analyzers are available, as well as complete audio-measurement systems that include sophisticated data-integrity-measurement functions. Such instruments measure the level of jitter directly, usually expressing the result in UI or time. They also usually measure the more obvious, but

nonetheless, important, parameters like data signal level and sample rate.

A particularly valuable test is an "eye pattern" display. This function averages multiple pulses and displays a statistical average. The measurement system extracts the clock from the datastream but regenerates it using a phase-locked loop to produce a "perfect" clock reference — free of jitter but of the same frequency and phase as the embedded clock. This reference synchronizes the display, but the actual datastream is displayed without correction and thereby shows its actual jitter. If the signal were perfect, the display would be a rectangle with thin traces. A real-world signal will show a rise time and the jitter, all of which result in a display with fat traces as several successive pulses are overlaid. The rise time gives a triangular appearance to the display, which is where the term "eye pattern" comes from. The size of the opening in the center of the eye directly measures the integrity of the signal. Slow rise times and high jitter make the opening smaller. The AES3 standard defines a minimal eye size for reliable performance. Figure 5 shows a typical eye pattern with an AES3 limit shown in red.

In addition to characterizing an AES3 datastream with such an analyzer, data impairment simulation allows you to

Jitter is perhaps the biggest problem in digital audio transmission.

adding its own accumulation of jitter.

So why do we care?

How can jitter affect audio performance? Jitter components, be they broadband noise, specific frequencies caused by crosstalk, or any of several other sources, can show up in the recovered analog audio signal. An FFT spectrum analysis of a recovered audio

cult. The trouble report may describe an intermittent signal loss, but, when the troubleshooter assesses the final device in the chain, he can find no fault because the complete path is no longer intact.

What can be done?

It is not possible to see jitter by looking at a time-domain waveform; the sync circuits of a conventional oscilloscope

evaluate how robust a device is when presented with a data signal that has noise, jitter and other problems. A sophisticated digital-domain audio analyzer usually includes the ability to add calibrated amounts of data-signal impairments to its AES3 output signal. It provides controls for rise and fall time, signal level, sample frequency, jitter level, and perhaps cable simulation. By adding known amounts of degradation, you can determine when a device will begin to have difficulty decoding the datastream and where it will eventually lose lock.

Additional measurements

We have talked about the integrity of the datastream itself — but what about traditional audio measurements? What errors do converters introduce? How can you characterize a cross-domain device? Again, a sophisticated dual-domain audio-measurement system can give you insight into the real audio performance of such devices.

To characterize a digital-to-analog device, such as you might find in a system that plays back a signal stored on a hard drive, you have to generate an AES3 digital test signal with an embedded audio signal. Generating such a signal completely in the digital domain can give it a dynamic range and residual distortion better than 140dB, far exceeding the performance of any converter and the best analog circuits. Measuring the resulting analog audio signal with a high-performance analog analyzer allows accurate characterization of




Figure 4. This diagram shows an AES3 datastream with proper source, cable and termination impedances (black) and without (blue).

the DAC.

For a more comprehensive analysis, add calibrated amounts of data impairments on the digital side and see the effect on the analog audio. Of course, the reverse is also true: characterize a recording device and its ADC by generating high-performance analog test signals and analyze them in the digital domain with an instrument capable of making such measurements. Some converter-specific measurements would include

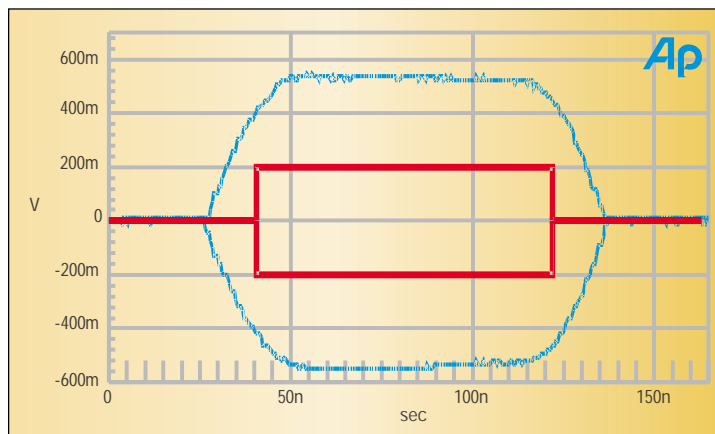


Figure 5. The blue pattern in this diagram is a typical eye pattern. The red box shows an AES3 jitter limit.

frequency response, with particular attention to the upper-band edge to characterize the anti-aliasing and reconstruction filters. Low-level linearity

measurements, noise modulation, quantization distortion and truncation artifacts can help you evaluate the analog-to-digital conversion circuitry of a device. All of the graphs and test results shown in this article were produced with an Audio Precision ATS-2 audio-measurement system.

Summary

Digital audio signal transmission and storage offers the advantage of higher initial quality, more robust delivery, virtually no progressive degradation with successive generations of storage or transmission, and more predictable quality at the far end. But these advantages can only be realized if the AES3 digital data transmitters and receivers in the individual devices in the chain are well-designed and if the transmission techniques follow good digital-data practices. Knowledge of the mechanisms of degradations and how they play out with the equipment in your facility will prevent unpleasant surprises and ensure high-quality audio delivery.

For more information on the subject of digital audio, including in-depth discussions on real-world problems, see the book *Measurement Techniques for Digital Audio* by Julian Dunn. This book is available from the Audio Precision Web site at:

audioprecision.com/publications/apnotes/index.htm. **BE**

Wayne Jones is vice president of applications support at Audio Precision.