INNOVATIONS IN AUDIO • AUDIO ELECTRONICS • THE BEST IN DIY AUDIO

www.audioxpress.com

audioxpress

ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY

Standards Review

Audio Network Development: An Introduction

By João Martins

Immersive Sound

A Look at Auro-3D

By Ron Tipton

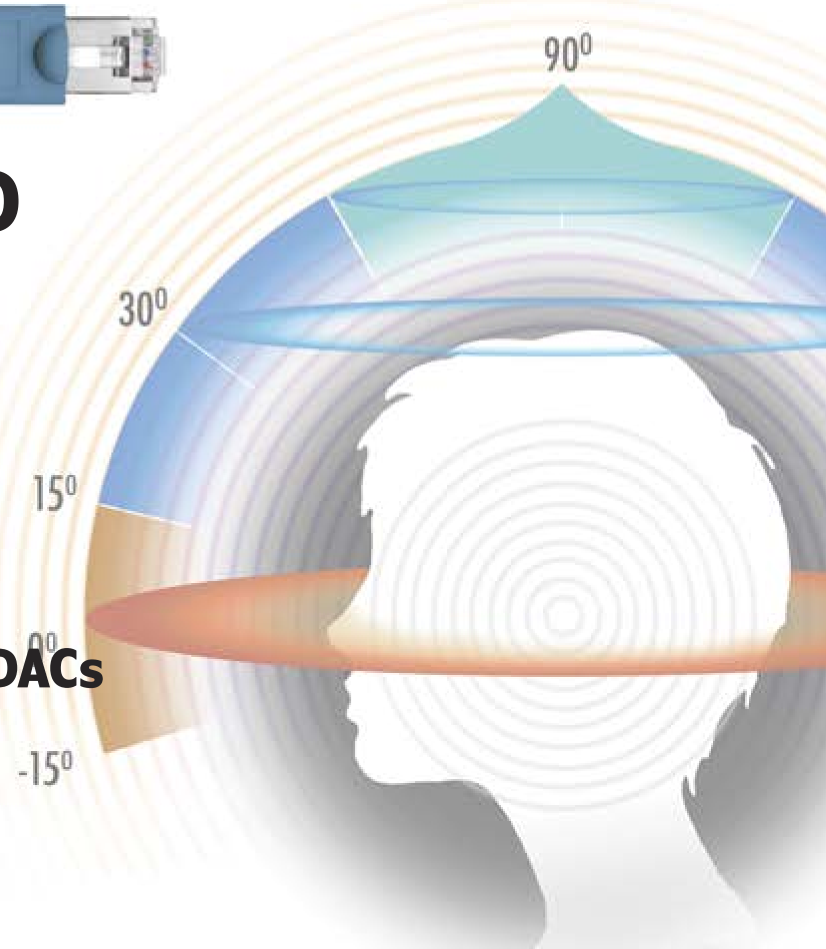

Sound Control

Audio Tools from Studio Six Digital

By Richard Honeycutt

Practical Test & Measurement

Testing Audio ADCs and DACs

By David Mathew

Pitting the CLIO Pocket Against a Sound Card

By Stuart Yaniger

You Can DIY!

A Low-Cost High-Performance Intergrated Circuit Audio Preamplifier

By Bruce Heran

Hollow-State Electronics

Creating Low-Noise Amplifiers

By Richard Honeycutt

Audio Electronics Tubes

Single-Ended Type-A Monobloc with 2 KT88 Tubes

By Gerhard Haas

$9.00US $10.00CAN

DECEMBER 2015

Testing Audio ADCs and DACs

Learn more about ADCs and DACs, and the methods engineers use to procure high-quality results by correctly evaluating those two critical stages, with conventional and more evolved measurements used in audio tests.

Audio precision®

By
David Mathew
(United States)

In 2015, there's not much question about audio storage, transmission, or streaming: it's digital. Apart from rare sightings of vinyl or open-reel tapes in boutique sales or creative enclaves, audio is digital. Done right, digital audio is flexible, robust, and of very high quality. Pulse code modulation (PCM) recording, lossless surround formats, and even lossy compression (at least at high data rates) provide the soundtracks for our lives.

But, of course, sound in air is not digital. The pressure waves created by a human voice or a musical instrument are recorded after exciting a transducer of some sort, and the transducer responds with an electrical voltage that is an analog of the pressure wave. Likewise, at the end of the chain the digitized audio signal must eventually move air, using a voltage that is the analog of the original sound wave to drive a transducer that creates a pressure wave.

Near the beginning of a digital chain, we must use an ADC (analog-to-digital converter) to transform the analog electrical signal to a digital representation of that signal. Near the end of the chain, we must use a digital-to-analog converter (DAC) to transform the digital audio signal back into an analog electrical signal. Along with the transducers, these two links in the chain (the ADC and the DAC) are key in determining the overall quality of the sound presented to the listener.

Testing Audio Converters

The conventional measurements used in audio tests can also be used to evaluate ADCs and DACs. These measurements include frequency response, signal-to-noise ratio (SNR), interchannel phase, crosstalk, distortion, group delay, polarity, and others. But conversion between the continuous and sampled domains brings several new mechanisms for nonlinearity, particularly for low-level signals. This article looks at problems seen in audio analog-

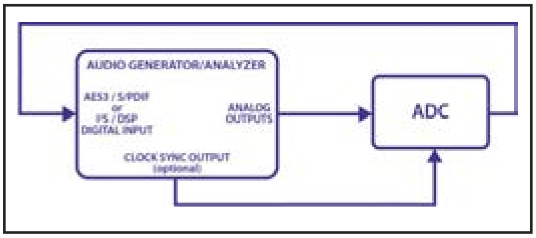

Figure 1: This is a typical ADC test block diagram.

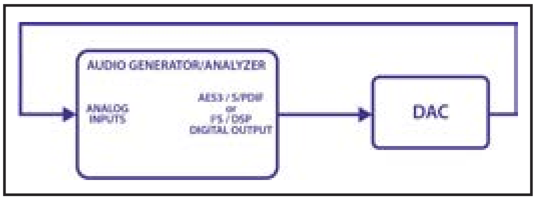

Figure 2: Here is a typical DAC test block diagram.

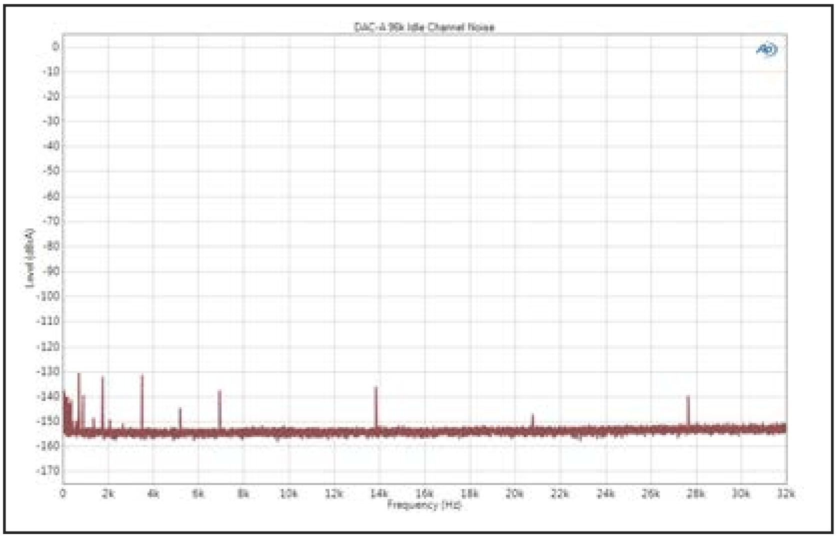

Figure 3: Here is the Fast Fourier Transform (FFT) idle channel noise, DAC "A."

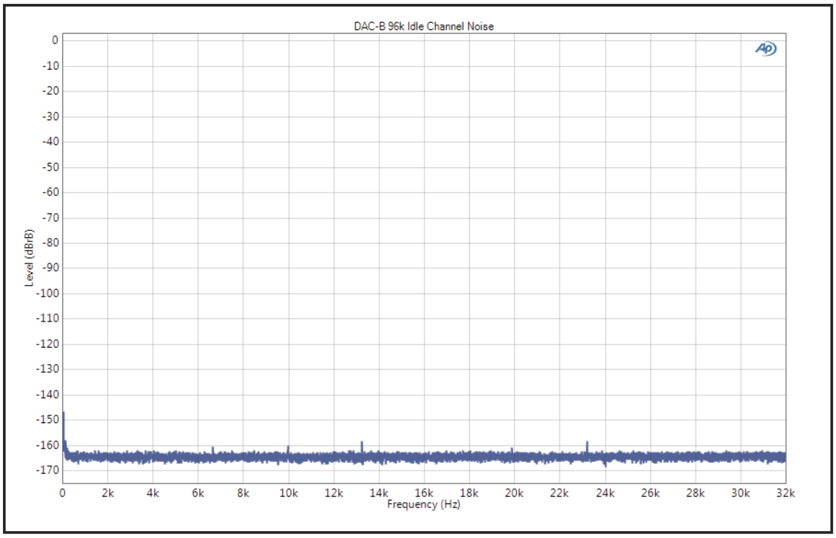

Figure 4: This is the Fast Fourier Transform (FFT) idle channel noise, DAC "B."

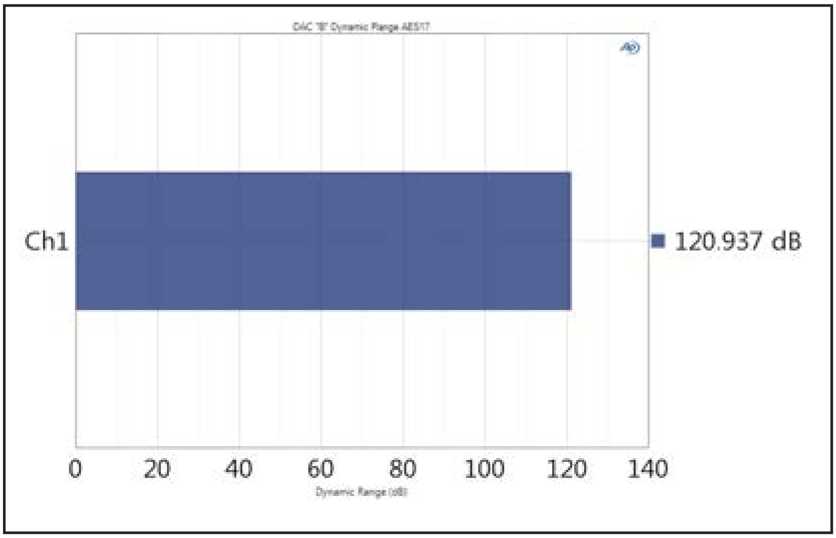

Figure 5: The signal-noise ratio (SNR) for DAC "B" is shown here.

to-digital and digital-to-analog conversion and some of the methods that have evolved to address these issues.

Of course, ADCs and DACs are used in a great number of non-audio applications, often operating at much higher sampling rates than audio converters. Very good oscilloscopes might have bandwidths of 33 GHz and sampling rates up to 100 GS/s, with prices comparable to a Lamborghini. Although audio converters don't sample at anywhere near that rate, they are required to cover a much larger dynamic range, with high-performance ADCs digitizing at 24 bits and having SNRs over 120 dB.

Even a high-end oscilloscope typically uses only an 8-bit digitizer. A 24-bit conversion pushes the measurement of noise and other small-signal performance characteristics to the bleeding edge. Consequently, measurements of such converters require an analyzer of extraordinary analog performance.

Test Setups

The typical test setups are straightforward. For ADC testing, the analyzer must provide extremely pure stimulus signals at the drive levels appropriate for the converter input. For converter ICs, the analyzer must have a digital input in a format and protocol to match the IC output, such as I^2S, DSP, or a custom format.

For a commercial converter device, the digital format is typically an AES3-S/PDIF-compatible stream. For devices that can sync to an external clock, the analyzer should provide a clock sync output.

For DAC testing, the analyzer must have a digital output in the appropriate format, and analog inputs of very high performance. The graphs in this article were created by testing commercial converters, using the AES3 digital transport. The analyzer is the Audio Precision APx555.

As previously mentioned, ADCs and DACs exhibit behaviors unique to converters. The Audio Engineering Society (AES) has recommended methods to measure many converter behaviors in the AES17 standard. The following examples examine and compare several characteristic converter issues.

Idle Tones

Common audio converter architectures, such as delta-sigma devices, are prone to have an idling behavior that produces low-level tones. These "idle tones" can be modulated in frequency by the applied signal and by DC offset, which means they are difficult to identify if a signal is present. A Fast

Interpretation of Noise in the Fast Fourier Transform Power Spectra

By David Mathew

In audio systems, noise figures are generally measured and reported as:

–103 dB rms noise, 20 Hz to 20 kHz BW, A-weighted

The noise signal is measured with an RMS detector across a specified bandwidth. The 0 dB reference for the noise measurement might be a device's nominal operating level, but more typically it is the maximum operating level, which produces a more impressive number. Similarly, weighting filters usually produce better noise figures and are often used in marketing specifications. A weighting filter attempts to approximate the response that we humans perceive.

So let's measure the noise of DAC B using these methods. We'll make an RMS measurement first across the full signal bandwidth, then across a limited bandwidth, and with an added weighting filter.

–120.3 dB RMS noise, DC to 48 kHz.
–123.9 dB RMS noise, 20 Hz to 20 kHz.
–126.2 dB RMS noise, 20 Hz to 20 kHz, A-weighted.

But if you look at the Fast Fourier Transform (FFT) spectrum (see **Figure 4** in the main article), the average noise floor appears to be about –163 dB or so. That's a big difference.

Conventionally, an audio FFT amplitude spectrum is displayed by scaling the vertical axis so that a bin peak indicates a value that corresponds to the amplitude of any discrete frequency components within the bin. This calibration is not appropriate for measuring broadband signals (e.g., such as noise power) without applying a conversion factor that depends on the bin width and on the FFT window used.

In this case, each bin is 0.375 Hz wide (the sample rate of 96 kS/s divided by an FFT length of 256K points). The window spreads the energy from the signal component at any discrete frequency, and the Y-axis calibration takes this windowing into account.

For the Audio Precision Equiripple window used here, the calibration compensates for the power being spread over a bandwidth of 2.63 bins.

This can be converted to the power in a 1 Hz bandwidth, or the power density, by adding a scaling factor in decibels that can be calculated as follows:

$$\begin{aligned}\text{scaling factor} &= 10\log(1 / \text{Window Scaling} \cdot \text{Bin Width})\\ &= 10\log(1/2.63 \cdot 0.375)\\ &= 4.2\text{ dB}\end{aligned}$$

To estimate the noise from a device based on an FFT spectrum, you can integrate the power density over the frequency range of interest. For an approximately flat total noise (where the noise power density is roughly constant) it is possible to estimate the sum of the power in each bin within reasonable accuracy, by estimating the average noise power density and multiplying by the bandwidth.

Figure 4, for example has a noise floor that is approximately in line with about –163 dB on the Y-axis. The noise power density is the apparent noise floor minus the window conversion factor or:

–163 dB – 4.2 dB = –167.2 dB per Hz

The integration to figure the total noise over a given bandwidth is simple, if the noise is spectrally flat. Multiply the noise power density by the bandwidth, which in this case is 20 kHz. For decibel power (dB = 10logX), this is the same as adding 43 dB =10log(20000), as follows:

–167.2 + 43 dB = –124.2 dB

This calculation provides a result only 0.5 dB different from the 20 Hz to 20 kHz unweighted measurement cited earlier.

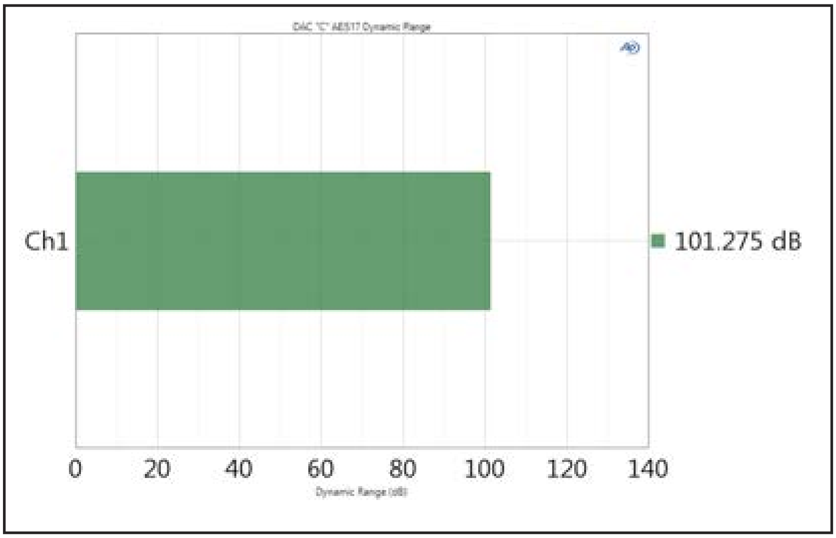

Figure 6: This is the signal-noise ratio (SNR) for DAC "C."

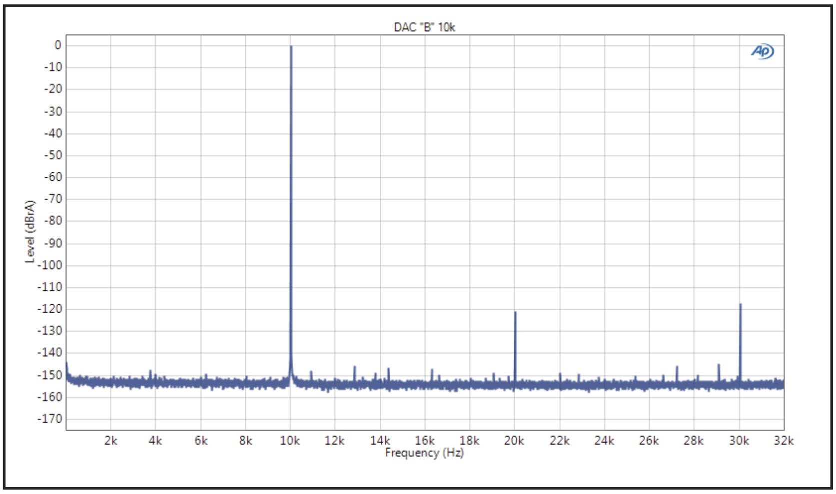

Figure 7: DAC "B" shows no jitter sidebands at 10 kHz.

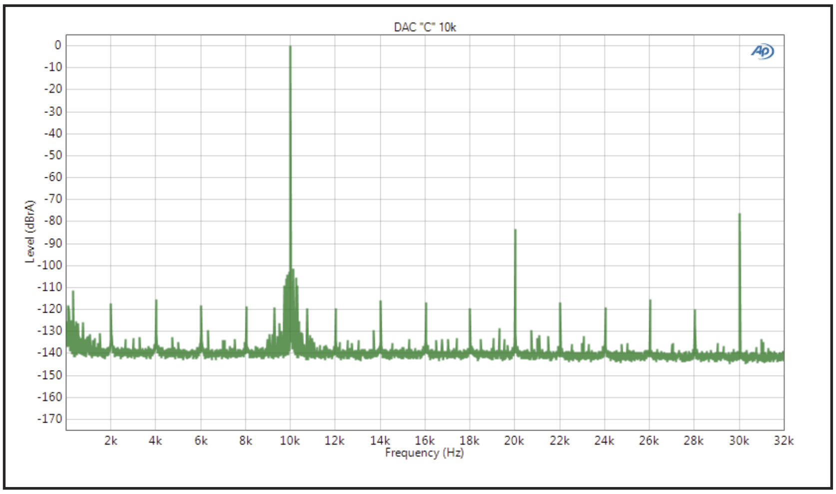

Figure 8: DAC "C" shows strong jitter sidebands at 10 kHz.

Fourier Transform (FFT) of the idle channel test output can be used to identify these tones.

The DAC shown in **Figure 3** has several idle tones, some with levels as high as –130 dB. The idle tones (and the noise floor) of the DAC shown in **Figure 4** are much lower.

Signal-to-Noise Ratio (Dynamic Range)

For analog audio devices, a SNR measurement involves finding the device maximum output and the bandwidth-limited RMS noise floor, and reporting the difference between the two in decibels.

With audio converters, the maximum level is usually defined as that level where the peaks of a sine wave just touch the maximum and the minimum sample values. This is called "full scale" (1 FS), which can be expressed logarithmically as 0 dBFS. The RMS noise floor is a little tricky to measure because of low-level idle tones and, in some converters, muting that is applied when the signal input is zero. AES17 recommends that a –60 dB tone be applied to defeat any muting and to allow the converter to linearly operate. The distortion products of this tone are so low they fall below the noise floor, and the tone itself is notched out during the noise measurement. The IEC 61606 standard recommends a similar method, but calls the measurement dynamic range.

Figure 5 and **Figure 6** show a comparison of the signal-to-noise measurements of two 24-bit DACs operating at 96 kS/s, using this method. As can be seen, some converter designs are much more effective than others.

Jitter

For ADCs, clock jitter can occur within the converter, and synchronization jitter can be contributed through an external clock sync input. For DACs receiving a signal with an embedded clock (e.g., AES3 or S/PDIF), interface jitter on the incoming signal must be attenuated.

Sinusoidal jitter primarily affects the audio signal by creating modulation sidebands—frequencies above and below the original audio signal. More complex or broadband jitter will raise the converter noise floor. A common measurement that reveals jitter susceptibility is to use a high-frequency sinusoidal stimulus and examine an FFT of the converter output for jitter sidebands, which are symmetrical around the stimulus tone. DAC "B" shows no sidebands (see **Figure 7**). However, DAC "C" (see **Figure 8**) shows strong sidebands. Note that the strong tones at 20 kHz and 30 kHz are products of harmonic distortion and are not jitter sidebands.

Jitter Tolerance Template

AES3 describes a jitter tolerance test, where the capability of a receiver to tolerate defined levels of interface jitter on its input is examined. A digital audio signal is applied to the input. The signal is jittered with sinusoidal jitter, swept from 100 Hz to 100 kHz. As the jitter is swept, its level is varied according to the AES3 jitter tolerance template. Jitter is set at a high level up to 200 Hz, then reduced to a lower level by 8 kHz, where it is maintained until the end of the sweep.

An interface data receiver should correctly decode an incoming data stream with any sinusoidal jitter defined by the jitter tolerance template shown in **Figure 9**. The template requires a jitter tolerance of 0.25 UI peak-to-peak at high frequencies, increasing with the inverse of frequency below 8 kHz to level off at 10 UI peak-to-peak below 200 Hz.

In this case, jitter is set to about 9.775 Unit Intervals (UI) at the lower jitter frequencies, and drops to about 0.25 UI at the higher frequencies. The blue trace is the THD+N ratio (distortion products of the 3 kHz audio tone), which remains constant across the jitter sweep, indicating good jitter tolerance in this DUT. As jitter level rises, poor tolerance will cause a receiver to decode the signal incorrectly, and then fail to decode the signal, occasionally muting or sometimes losing lock altogether.

Filter Effects

Figure 10 shows the response of the anti-alias filter in ADC "C." A tone at the input of the ADC is swept across the out-of-band (OOB) range of interest (in this case, from 40 kHz to 200 kHz) and the level of the signal reflected into the passband is plotted against the stimulus frequency. A second trace shows the converter noise floor as a reference.

For **Figure 11**, spectrally flat random noise is presented to the DAC input. The analog output is plotted (with many averages) to show the response of the DAC's anti-image filter. In this case, a second trace showing a 1 kHz tone and the DAC noise floor is plotted, scaled so that the sine peak corresponds to the noise peak.

Polarity

Audio circuits (including converters) often use differential (balanced) architectures. This opens the door for polarity faults. An impulse response (IR) stimulus provides a clear observation of normal or reversed polarity (see **Figure 12**).

Summary

Tests for the high-level nonlinear behavior of an ADC are similar to those for nonlinearities in analog electronics, using standardized tests for harmonic distortion and intermodulation distortion. But audio converters bring new mechanisms for non-linearity, particularly for low-level signals. AES17 and Audio Precision's Technote 124 describe effective testing methods for audio converter measurements.

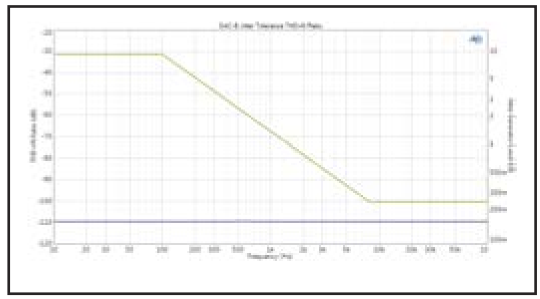

Figure 9: This is the total harmonic distortion plus noise (THD+N) for DAC "B" during the jitter tolerance sweep.

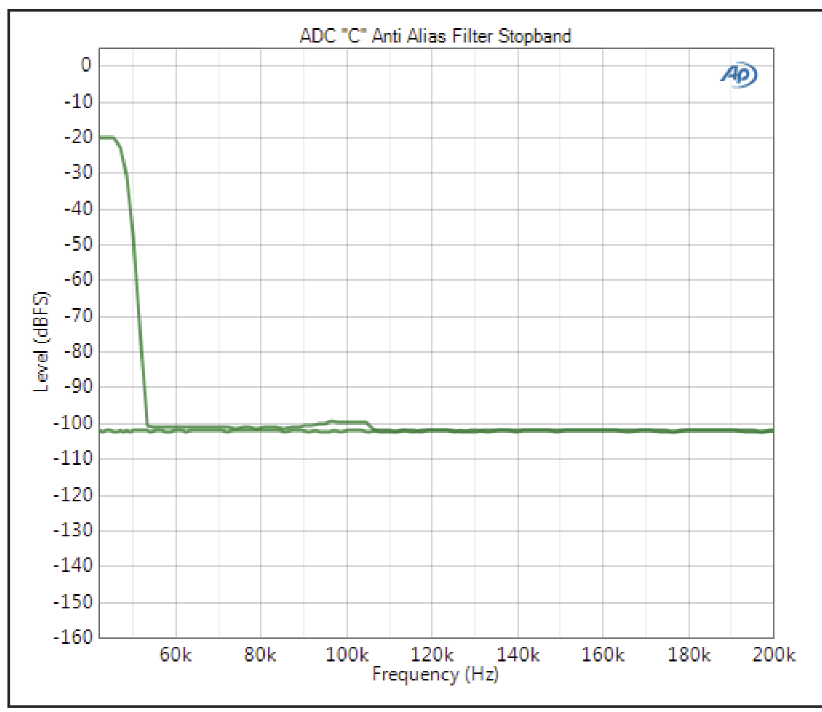

Figure 10: Here is the response of the anti-alias filter in ADC "C."

Resources

AES3, AES standard for digital audio engineering—Serial transmission format for two-channel linearly represented digital audio data, 2009 (rev. 2014), www.aes.org.

AES17, AES standard method for digital audio engineering—Measurement of digital audio equipment, 1998, rev. 2009, (Revision of AES17-1991), www.aes.org.

J. Dunn, "Application Note #5: Measurement Techniques for Digital Audio," Audio Precision, 2004, www.ap.com.

S. Peterson, "Technote 124, Measuring A-to-D and D-to-A Converters with APx555," Audio Precision, 2015, www.ap.com.

Glossary

AES3, S/PDIF—In the consumer and professional audio field, digital audio is typically carried from point to point as a bi-phase coded signal, commonly referred to as AES3, AES/EBU or S/PDIF. There are electrical and bitstream protocol differences among the variations of bi-phase coded digital audio, but the various signals are largely compatible. Variations are defined in the standards AES3, IEC 60958, and SMPTE 276M.

anti-alias filter—In sampled systems, the bandwidth of the input has to be limited to the folding frequency to avoid aliasing. Modern audio ADCs normally have this anti-alias filter implemented with a combination of a sharp-cutoff finite impulse response (FIR) digital filter and a simple low order analog filter. The digital filter operates on a version of the signal after conversion at an oversampled rate, and the analog filter is required to attenuate signals that are close to the oversampling frequency. This analog filter can have a relaxed response, since the oversampling frequency is often many octaves above the passband.

anti-image filter (reconstruction filter)—Digital audio signals can only represent a selected bandwidth. When constructing an analog signal from a digital audio data stream, a direct conversion of sample data values to analog voltages will produce images of the audio band spectrum at multiples of the sampling frequency. Normally, these images are removed by an anti-imaging filter. This filter has a stopband that starts at half of the sampling frequency—the folding frequency.

Modern audio DACs usually have this anti-imaging filter implemented with a combination of two filters: a sharp cut-off digital finite impulse response (FIR) filter, followed by a relatively simple low-order analog filter. The digital filter is operating on an oversampled version of the input signal, and the analog filter is required to attenuate signals that are close to the oversampling frequency.

unit interval (UI)—The UI is a measure of time that scales with the interface data rate, and is often a convenient term for interface jitter discussions. The UI is defined as the shortest nominal time interval in the coding scheme. For an AES3 signal, there are 32 bits per subframe and 64 bits per frame, giving a nominal 128 pulses per frame in the channel after bi-phase mark encoding is applied. So, in our case of a sampling rate of 96 kHz:

1 UI = 1 / (128 x 96000) = 81.4 ns

The UI is used for several of the jitter specifications in AES3.

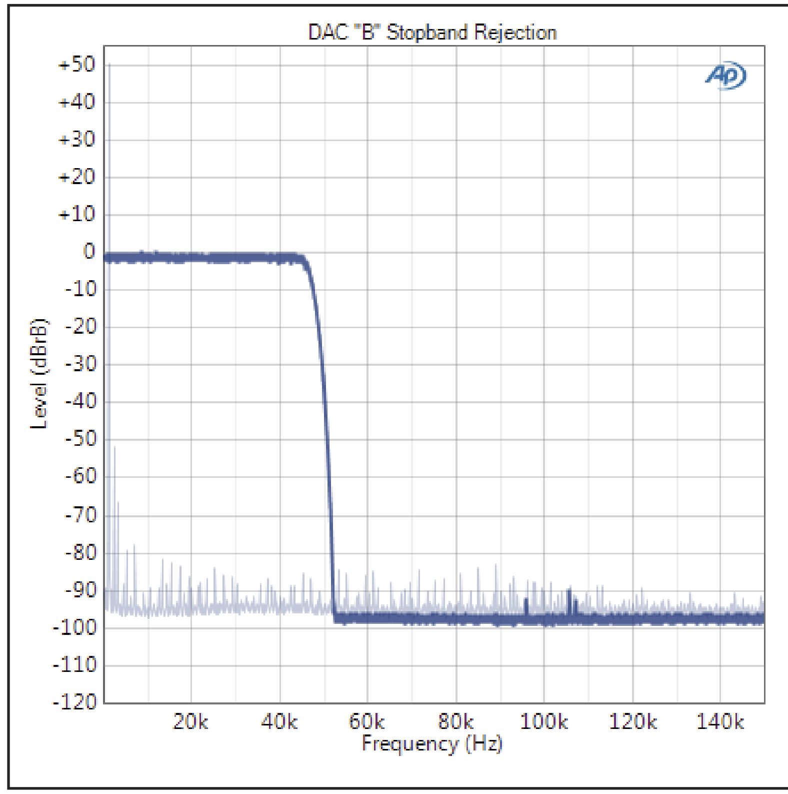

Figure 11: Spectrally flat random noise is presented to the DAC input. This figure shows the DAC "B" anti-image filter out-of-band rejection.

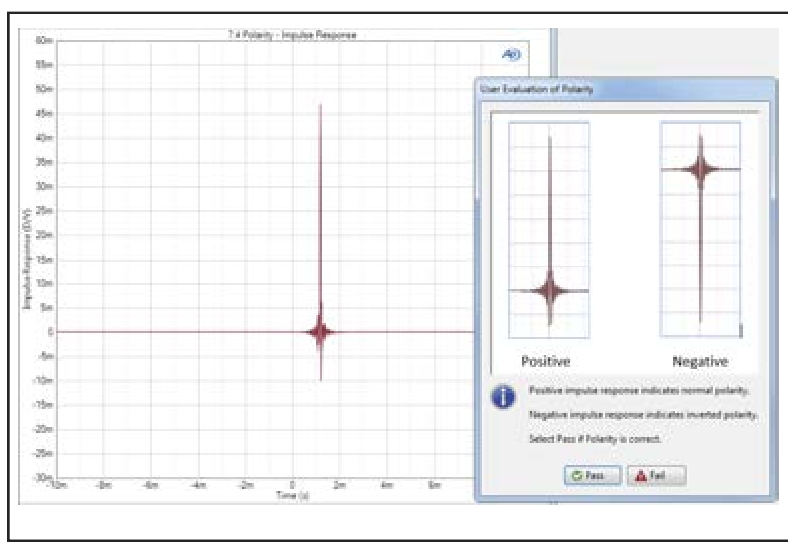

Figure 12: An impulse response (IR) is used to check polarity.

About the Author

David Mathew is Technical Publications Manager and a senior technical writer at Audio Precision in Beaverton, OR. He has worked as both a mixing engineer and as a technical engineer in the recording and filmmaking industries, and was awarded an Emmy for his sound work in 1988.

The APx555 and Benchmark's AHB2
One Perspective

By Eric Hodges

(United States)

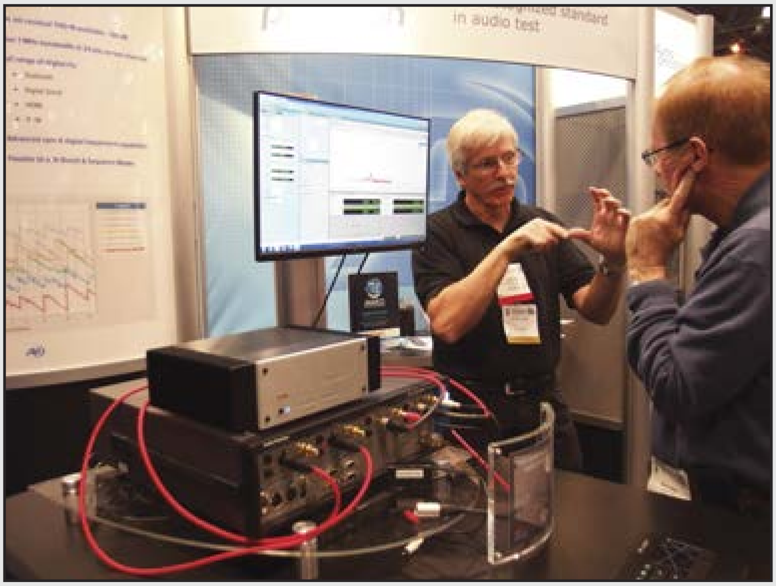

Audio Precision was one of the exhibitors at the 139th International Audio Engineering Society (AES) Convention in New York. The APx555 analyzer was demonstrated with the Benchmark AHB2 amplifier. Here Benchmark's John Siau takes the opportunity to show the measurements to Bob Cordell.

It was July 30, 2014, and I'd been with Audio Precision for a whopping nine days. A co-worker had just forwarded me an article on Benchmark's AHB2 power amplifier, highlighting a section in the piece that described the challenges of measuring the amplifier's performance with commercially available instrumentation.

A little more than a year later, this still makes me smile. The entire Audio Precision team was focused on final preparations for the introduction of our APx555, slated for early September 2014. Here was a well-respected audio organization validating one of the key elements of the APx 555's design—the need for analog test performance that surpassed that of the long-standing leader, our own 2700. The only frustrating part was the need to remain silent on the point until September. I made a mental note for a follow up and returned to the task of preparing for our product launch.

Fast forward a few months to a Thursday evening in October, and I'm in Los Angeles, CA, at the 137th International Audio Engineering Society (AES) Convention supervising the finishing touches in the booth set-up process. Imagine my surprise—seriously, the mental note I'd made in July was a bit dusty—when a trio from Benchmark and THX (a patent holder for amplifier and error-correction technologies used in the AHB2) arrived at our booth wondering if we'd like to introduce the APx555 and the AHB2. In the ensuing flurry of activity to set up a test of the AHB2 with the APx555—and effectively a test of the APx555's analog performance as indicated by its total harmonic distortion plus noise (THD+N) measurement capability—I'm not sure which engineers were more excited by the trial unfolding in booth #1212. And, as a singular group, this was a noteworthy collection of engineering talent, comprised as it was of Laurie Fincham and Jayant Datta (THX), John Siau (Benchmark), and Tom Kite (Audio Precision).

During this time, the corner of the booth occupied by the AHB2 and the APx555 radiated a palpable sense of anticipation from all involved. While I'm not sure how long the experiment set-up process actually took, focused as I was on my own show-prep duties, it seemed the test was ready in no time at all. With all eyes focused on the THD+N meter in the APx555's Bench Mode display, Tom enabled the instrument's generator and stepped back. As the generator kicked in, the bright green digits of the meter quickly stabilized at −118 dB and the resultant grins spreading across each participant's face appeared to include equal measures of pride and joy with a bit of relief thrown in as well. Just as they'd said, the AHB2 had market-leading performance. Likewise, the APx555 delivered, validating its own capabilities as it did the same for the AHB2.

The Benchmark AHB2 amplifier at the 137th AES convention in Los Angeles, CA.

For a spontaneous, completely unscripted event, it worked out quite well. Almost as though it had been engineered that way. For more information on the Benchmark AHB2 amplifier, the complete *audioXpress* review is available at http://audioxpress.com/article/Fresh-from-the-Bench-Benchmark-AHB2-Stereo-Power-Amplifier.html.

More information about Audio Precision's APx555 audio analyzer is available at www.ap.com/products/apx555.

About the Author

Eric Hodges is the Marketing Manager for Audio Precision in Beaverton, OR. His 18 years as a marketing professional have been exclusively in the high-tech industry, across applications as varied as test and measurement, microelectronics, broadcast, post-production, and audio.