Audio Precision APx555 B Series Audio Analyzer

Expanded Capabilities and Software

Audio Precision recently loaned us one of its newest analyzers, the APx555 B Series. Stuart Yaniger had one month to test it, and in this article he provides readers with a hint of the new expanded capabilities—both of the APx555 B Series hardware and the new APx500 5.0.1 software.

With the announcement of its new APx555 B Series two-channel audio analyzer, Audio Precision also unveiled the updated version 5.X of its APx audio measurement software. Since December 2018, new APx analyzers, whether the new B Series or Legacy, ship with version 5.0 software and one year of software maintenance, effectively licensing that instrument for APx version 6.0 when it is released, as well as any minor releases that occur between versions 5.0 and 6.0.

By
Stuart Yaniger
(United States)

I was recently involved in an online discussion of the term "Gold Standard" in the context of test and measurement. Regardless of eristic quibbles, I think there would be near unanimity that in the realm of audio measurement, Audio Precision (AP) has long been the gold standard. Its earlier analyzers had spectacular performance and versatility, but were notoriously difficult to use. AP addressed this in its APx line of analyzers, which are far more intuitive and simple to use, and the company also expanded its instruments' capabilities from electronics to the realm of acoustic and transducer measurement.

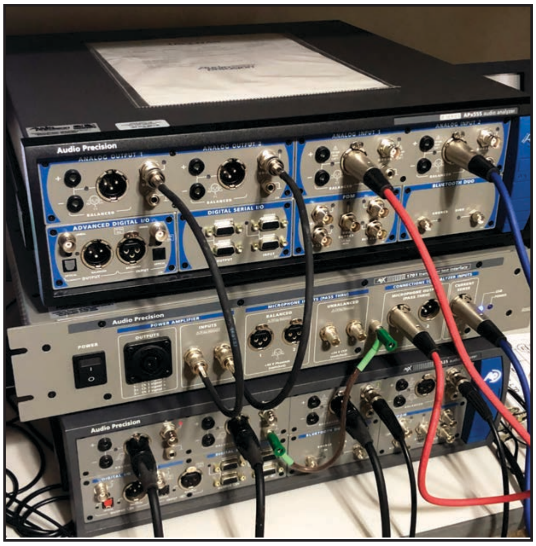

Photo 1: For four weeks, I was the coolest kid on the block. Top to bottom: APx555 B Series audio analyzer, the APx1701 transducer test interface, and the APx525 audio analyzer.

Part of being the gold standard is recognizing that markets and needs for test and measurement expand and shift, and effectively respond to those changes. The growth of smart speakers over the past few years has been remarkable—this product category has gone from nonexistent to nearly 100 million units in less than five years, and the growth and proliferation of brands do not seem to be slowing down.

Smart speakers present a significant measurement challenge compared with more traditional loudspeakers. The measurement paradigm of "feed it a test signal, then measure what comes out" will usually fail, either because of the need to address the speaker via cloud storage with its variable and uncertain timing or the built-in programming that detects when signals are something other than speech or music (e.g., sine waves) and mutes the speaker.

Another change in measurement needs is the new generation of electronics: amplification; analog-to-digital converters (ADCs); digital-to-analog converters (DACs); and signal processors, which can have distortion so low that existing instruments have difficulty resolving it.

Recognizing these and other new needs, AP has introduced an updated version of its flagship APx555 analyzer, the APx555 B series. At the same time, AP has made a major update to the APx 500 software, which incorporates enough new features and operational improvements to justify a new version number, 5.0.1.

After working with AP equipment and its applications people for the past several years, I was fortunate enough to be loaned one of the new APx555 B Series analyzers with a full software upgrade (see **Photo 1**). Sadly, I only was able to keep it for a month, but that allowed me time to experiment with some of the new features and see some of what the latest and greatest can do. I have only scratched the surface here, but I hope that this article will at least give a hint of the new expanded capabilities both of the APx555 B series hardware and the new APx500 5.0.1 software.

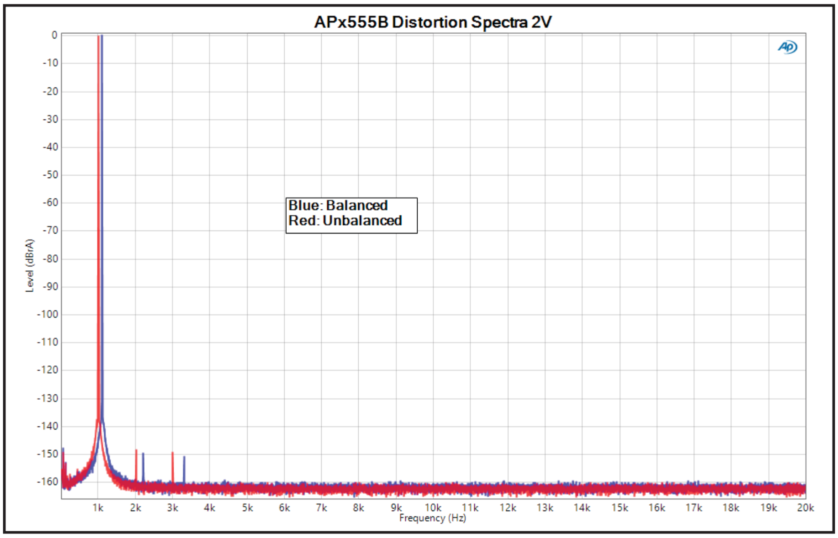

Figure 1: The distortion spectrum at 2 V output shows remarkably low harmonic content and noise floor.

Distortion Measurements

The B series now features a new High Performance Sine Generator and a High Performance Sine Analyzer, which AP claims reduce THD+N to better than -120 dB, with a 1 MHz bandwidth and Fast Fourier Transforms (FFTs) of up to 1.2 million points. This is a significant improvement over its previous series of instruments, which were certainly not slouches in the distortion and noise department. One might reasonably ask about the practicality of this improvement, but there are certainly reasons why this could be useful. I'll provide a couple examples.

Last year, *audioXpress* Technical Editor Jan Didden was here for his annual visit and brought along a prototype of his sound card measurement interface (the Autoranger). We hooked it up to the APx525 in order to characterize its performance. We seemed to hit a limit at -115 dB third harmonic and couldn't figure out why, since simulations had predicted much lower distortion. After consideration of this observation over a glass of single cask Irish whiskey (Jan is the best house guest ever!), it occurred to us that we hadn't run a loopback measurement to determine the APx distortion floor. And... yes, it was -115 dB. To that point, the APx525 distortion was lower than any device under test I had thrown at it (and to be fair, -115 dB is pretty darn low, about 0.0002%), but we finally had something that the APx525 couldn't measure.

Likewise, I recently reviewed the superb Cambridge Audio Edge W power amplifier. I measured distortion at -115 dB, and remembering what Jan and I had seen, I wrote that the amp's distortion was likely better than that. But was it, and how much better?

Having learned my lesson, I started by performing loopback measurements between the APx's outputs and inputs. There are two ways to accomplish this, either by running a cable directly from out to in, or more conveniently, checking the "Loopback" box in the APx500 software, which automatically connects output to input. I tried it both ways and saw no differences.

I first checked distortion and noise in both balanced and unbalanced modes at 2 V and 1 kHz. **Figure 1** shows the distortion spectra at 2 V for both balanced and unbalanced; I displaced the test frequencies by 100 Hz for clarity. The largest component is the second harmonic at nearly -150 dB from the fundamental. This is astounding performance, only slightly more than 0.000003%—and no, the decimal is not misplaced, that's 3 ppm. The noise over a 24 kHz bandwidth is under 1 µV and dominates the THD+N measurement, which was slightly under 0.0001%. Clearly, AP's new methods of distortion reduction are effective! Of course, the High Performance Generator and Analyzer are slower than the standard performance option, but that's a small trade-off.

Then, I heaved the Edge W up onto the test bench, connected the balanced output of the APx555B to the Edge W input, attached an 8 Ω dummy load, and measured the distortion at 1 W (2.83 V) output. I chose that low power deliberately since crossover distortion would be greater than at full power. The distortion spectrum is shown in **Figure 2** and proved my suspicions correct. This amplifier has lower distortion than the APx525 I used to measure it! I'm not sure of what impresses me more, the amp's performance or

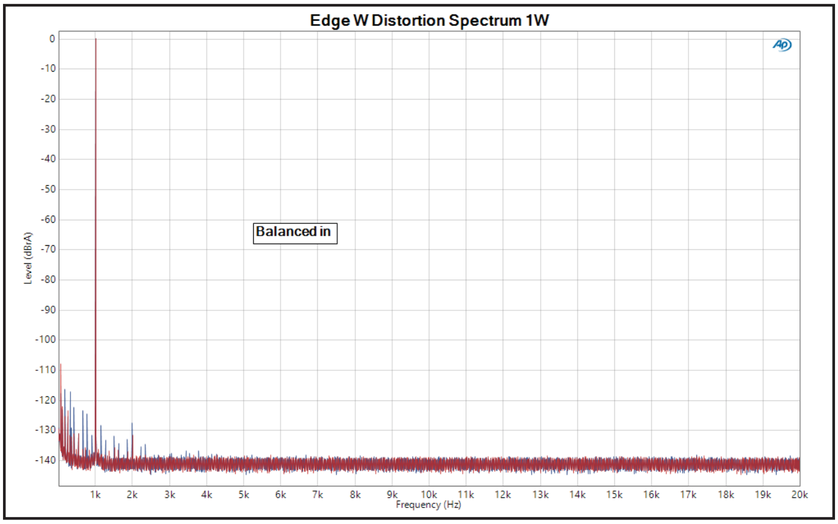

Figure 2: Re-measuring the Cambridge Edge W with the APx555 B Series showed that my previous measurements were, as I suspected, limited by the analyzer, not the amplifier.

About the Author

Stuart Yaniger has been designing and building audio equipment for nearly half a century, and currently works as a technical director for a large industrial company. His professional research interests have spanned theoretical physics, electronics, chemistry, spectroscopy, aerospace, biology, and sensory science. One day, he will figure out what he would like to be when he grows up.

the B Series APx555's ability to measure it, but either way, I'm impressed.

Of course, I had to now retest the Autoranger. And as predicted, the distortion was indeed fabulously low, about -145 dB second and -135 dB third harmonic. No wonder we couldn't reach its limits with the APx525.

Transfer Function

Traditionally, electroacoustic equipment was tested in some sort of synchronous mode: a test signal was applied and the output was measured. In the case of conventional amplification, signal processors, or transducers, there might be a constant delay, but the relation between input and output in terms of timing is pretty constant.

Devices such as Bluetooth headphones and smart speakers present a new challenge. For example, if I want to measure the frequency response of a conventional loudspeaker, I can apply a test signal like a log-frequency sweep (also called a "chirp"), derive the impulse response, and after performing a Fourier transformation, simultaneously acquire frequency response, distortion, and an energy-time curve from that single data set (the Farina method, see Resources).

Now, how do I apply the signal to a smart speaker? The only practical way is to put the desired signal in a file in the cloud, which the smart speaker can retrieve and play back. Unfortunately, this isn't in any way synchronous, and the time delay is quite variable.

Last year in version 4.6 of the APx software, AP introduced the Open Loop function, which allowed a pre-recorded chirp signal to be played back and captured by the APx analyzers for analysis. This took the variable (and long!) time delay involved in recording the chirp, uploading it to the cloud, then having the smart speaker retrieve and replay it out of the equation.

Problem solved? Well... not quite. This approach helps a great deal, but there are limitations on the sorts of test signals that can be used and how they can be accessed. For example, chirp signals can often be sensed as sine waves and suppressed by algorithms in the "smart" part of the signal processing. Likewise, multitone signals can be mistaken for noise and also suppressed. Clearly, this is also true of actual noise signals (e.g., pink or white noise). And since real-world signals such as speech and music can have high crest factors, there's always a question about correlating the measurements with test signals to actual performance in the end use.

The transfer function method allows the use of any arbitrary file (e.g., recorded speech or music) as the test signal (termed the "reference file"), with the proviso that the test signal be broadband to cover the frequency range of interest. The acquired asynchronous device under test (DUT) output, typically

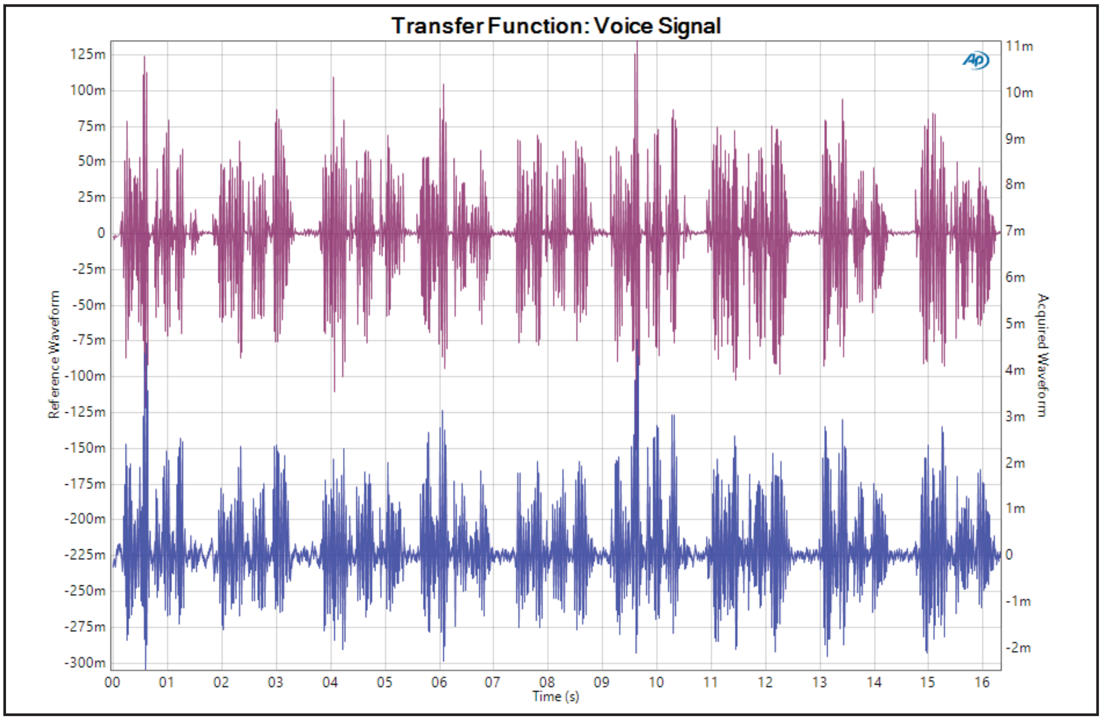

Figure 3: The transfer function measurement module automatically lines up the acquired signal (lower trace, blue) and the reference file (upper trace, red) in time.

an acoustic output measured with a microphone, is compared with the reference file. That seems simple, but the asynchronous nature requires a trick in order to match up the reference and acquired signal, namely the use of a cross-correlation function (see the sidebar article). The APx software allows the matching criteria to be set at High, Medium, Low, or None, depending on the desired degree of precision balanced against criticality of measurement conditions. Likewise, you can set the time allowed for the matching before the software "gives up" and returns an error. This can happen if there's too much noise or the reference (and/or acquired signals) have insufficient bandwidth.

Once the matching is complete and the reference and acquired waveforms are aligned in time, the acquired signal is automatically truncated to remove the portions outside of the desired output (e.g., the delay between starting the acquisition and the beginning of the actual signal output), and the two files can then be compared using a complex Fourier transform to extract DUT frequency response amplitude and phase information, as well as coherence and both amplitude and power spectral densities of the reference and acquired data.

To demonstrate the power of this method, I created a speech signal consisting of alternating male and female voices reading standard phrases (specifically, stitched together phrases from the ABC-MRT16 speech intelligibility test- see Resources). This was saved as a .wav file and used as a reference signal. The reference signal was loaded into the B Series APx555's generator module, and the balanced output was connected to a Kali LP-6 loudspeaker that I have been using in my lab for mastering. I then connected a calibrated PCB Piezotronics 376A31 microphone to the APx1701 transducer interface input and the microphone was positioned about 1 m away from the loudspeaker.

I started the acquisition, then played the reference file at about 70 dBSPL at the microphone position. The transfer function measurement module automatically lined up the acquisition and the reference signals in time as shown in **Figure 3**. Now, here's the magic of the transfer function measurement: this is a voice signal, yet the measurement yields the driver frequency response, which matches very closely with the frequency response measured by more conventional synchronous methods, in this case, a log-sweep (see **Figure 4**).

The transfer function module also allows end-to-end testing of smart speakers. This requires a calibrated acoustic source to output the "wake word" (e.g., "Hello, Siri!") and command to stimulate the DUT to play the desired file, generally a recorded speech, music, or test signal file. If you need to capture the

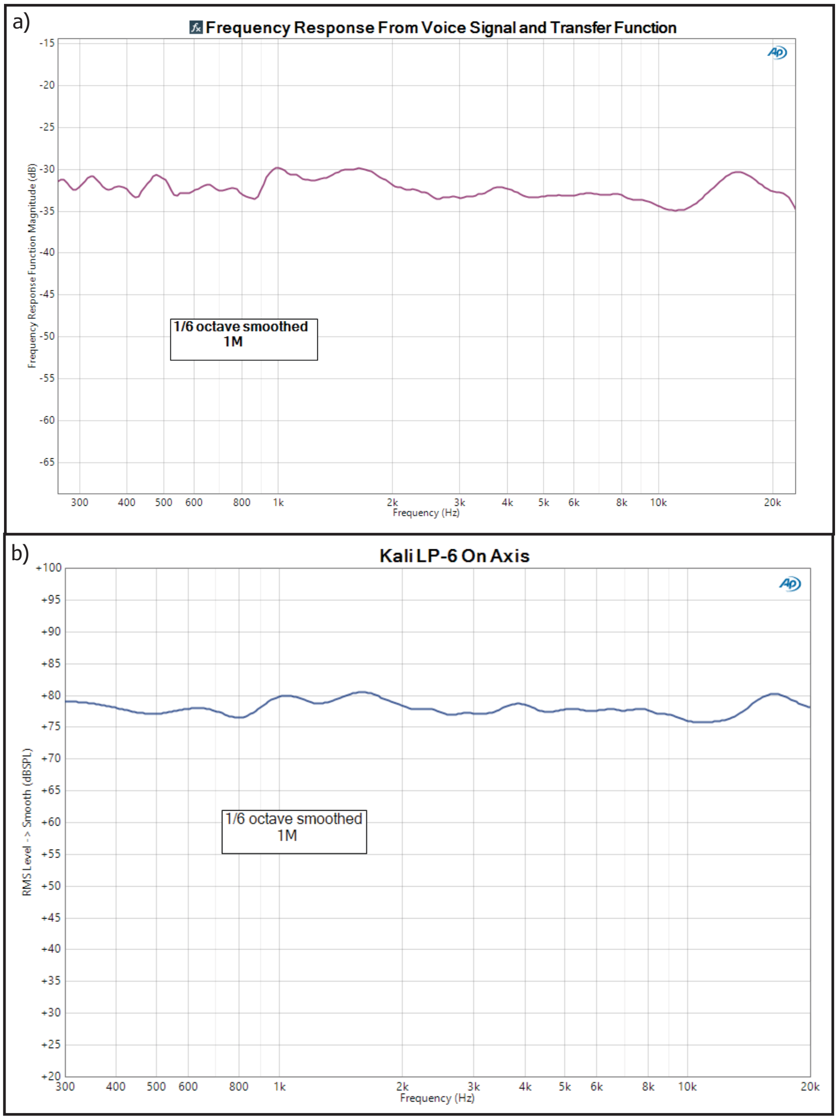

Figure 4: The loudspeaker frequency response measured using a voice signal and the transfer function module (a) compares closely with a measurement taken using a conventional log-frequency sweep (b).

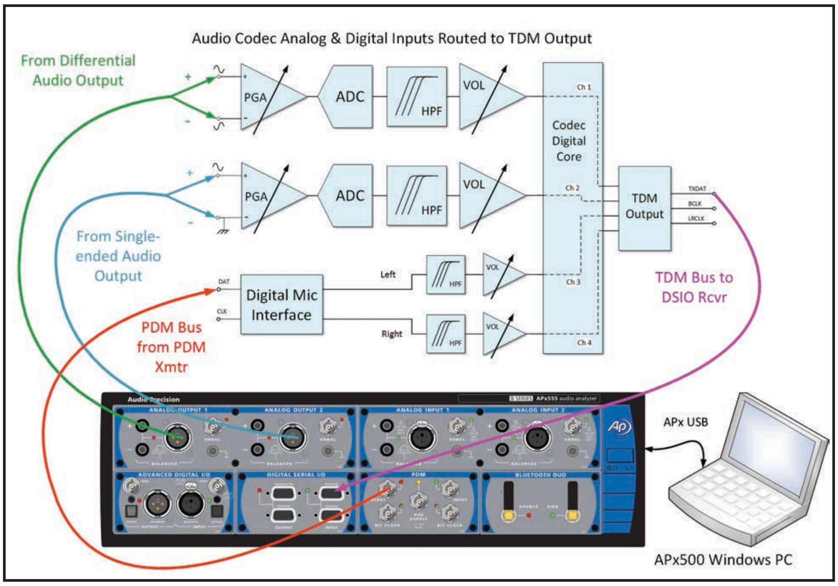

Figure 5: The APx555 B Series can be configured to do simultaneous measurements of codecs with single-ended, differential, and digital inputs. (Image courtesy of Audio Precision)

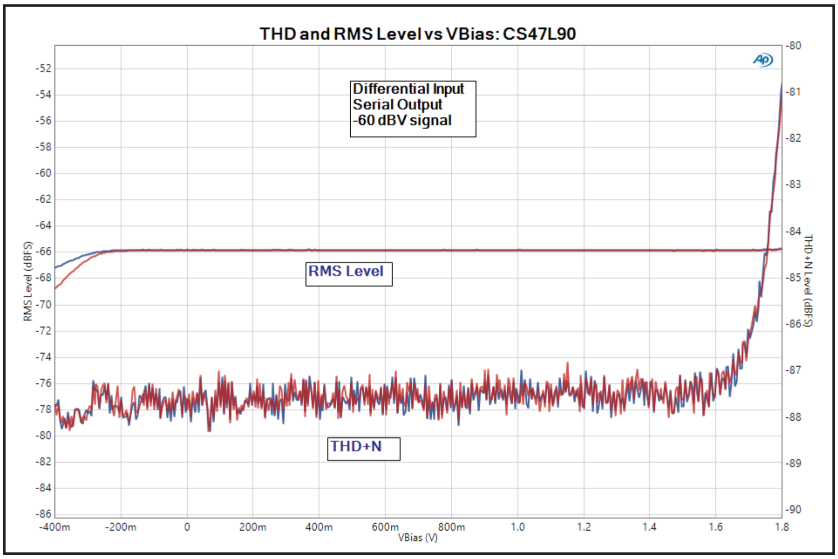

Figure 6: Using the swept DC offset function of the APx555 B Series in differential mode, the sensitivity of a codec to variations in DC offset voltage can be quantified.

speech and replay it (e.g., to characterize the beam-forming function of a microphone array), this sort of test will generally require an anechoic chamber to perform. But the issues of asynchronous capture and analysis can proceed exactly analogously to my rather simple demonstration.

ADC Test

Conventional ADC testing involves injecting signals, then examining the digital files following conversion. A good review of the traditional measurement methods appeared in *audioXpress* a few years ago (see Resources). The continuing proliferation of low power ADCs for applications such as smart phones has created a need for tests that go beyond the previous methods; for example, most ADCs in portable devices work off a single supply and need to be biased at about one-half the supply voltage (Vdd) to operate optimally.

Another complication is that, given the variety of smart devices and IoT appliances that surround us,

Some Basic Signal Processing Theory Behind the Transfer Function

Although an understanding of the basic math behind the transfer function module in the APx software is not necessary for making good use of it, it's probably worthwhile to at least peek behind the curtain, if you're the curious sort (as I am). This is all elementary to engineers, and I'm sure I'll get some significant wincing at my explanations, but for non-specialists like me, a review of the meaning of the terms used and the graphical data that the transfer function provides may be worthwhile.

The cross-correlation between two time dependent functions A and B is defined as:

$$R_{AB}(\tau) = \int_{-\infty}^{+\infty} A(t)B(t+\tau)\,dt$$

where A(t) is the reference signal and B(t) is the acquired signal. It is evident that as the time delay between the signals A and B approaches zero, the cross-correlation will reach a maximum. And this is basically how the transfer measurement aligns the reference and acquired signals; the maximum of the cross correlation function determines the time offset needed to align the reference and acquired waveforms. In the special case where A(t) and B(t) are the same signal, the function is referred to as the auto-correlation, and its maximum will of course occur at τ = 0.

For the loudspeaker measurement discussed in the article, the cross-correlation function is shown in **Figure 1**, where the APx transfer function module has lined up the signals.

Coherence is a measure of causality between reference and acquired waveforms, that is, what proportion of the acquired waveform is due to the stimulus signal and what part is due to noise or other non-causal factors. It will take a couple steps to get to the actual definition, so please bear with me.

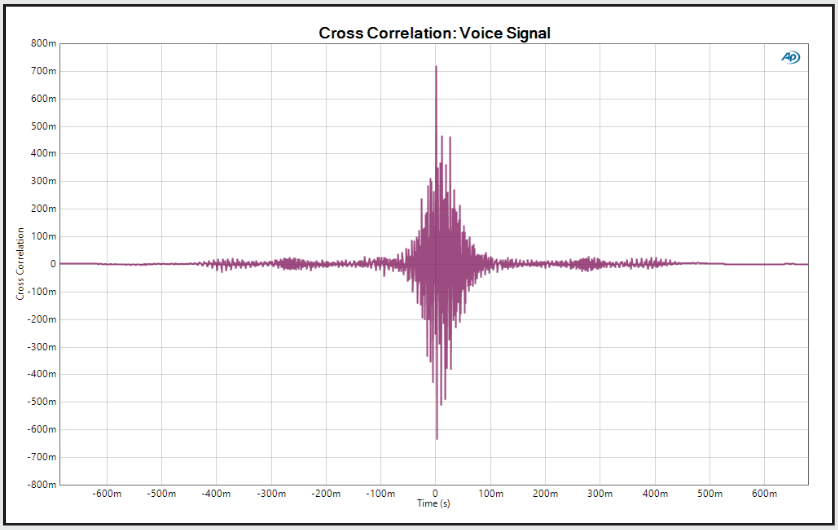

Figure 1: The cross-correlation function calculated by the APx500 software between the reference file and the acquired acoustic signal allows alignment of the two signals in time so that frequency and phase response can be determined.

ADCs are often used in either differential or single-ended mode, and often both at once in the case of multichannel chips. The testing in all of these modes can be rather lengthy and tedious, especially when throwing in the DC offset variable. Here's where the multichannel modes and flexible I/O capabilities of the APx555B come into play. **Figure 5** shows a typical (though complicated!) test setup using the APx555B to characterize a Crystal Semiconductor CS47L90 codec using the Crystal demo board. In this example, all three types of inputs are exercised, differential, single-ended, and digital microphone inputs, with the serial output of the demo board connected to the AP analyzer's digital serial input port.

AP provides several project files as part of its TN136 "Measuring Low Power ADCs" Applications Note (see Resources), so making the connections and the measurements is nearly foolproof. The particular part that is new is the ability to have the analog outputs of the APx555B generate a DC offset voltage riding on the

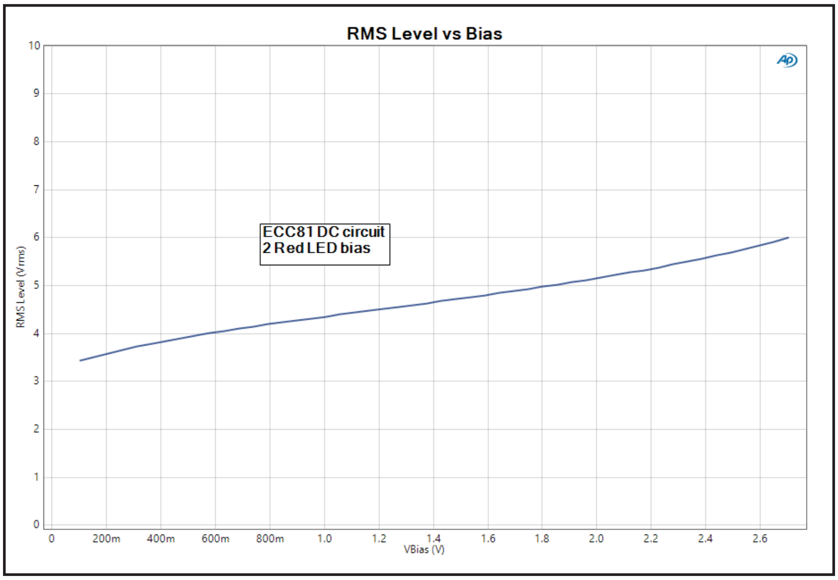

Figure 7: The DC offset sweep function also works to determine the effect of biasing on the gain of an analog circuit block (in this case, a vacuum tube circuit).

First, we need to define the cross power spectral density, S_{AB}. It is basically the Fourier transform of the cross-correlation function:

$$S_{AB}(\omega) = \int_{-\infty}^{+\infty} R_{AB}(\tau) e^{-j\omega t} dt$$

where, as usual, ω = 2 πf. The coherence is then defined as:

$$C_{AB}(f) = \frac{|S_{AB}|^2}{S_{AA}S_{BB}}$$

Clearly, if A and B are identical (perfect causality), the coherence function is 1 at all frequencies. Where A and B are completely unrelated, the numerator vanishes and the coherence is zero at all frequencies. In places where some of the acquired signal, B, differs from A because of noise or droops in frequency response, the coherence will be less than 1 but greater than or equal to zero.

For the loudspeaker measurement discussed in the article, the coherence function is shown in **Figure 2**; you can see that it's close to unity through much of the audible frequency range, but drops at low frequencies where the driver rolls off and noise dominates the measurements. It likewise has dips at the 60 Hz line frequency and its harmonics. The coherence function is very useful as a basic measurement quality check—if it's significantly less than 1 over most of the range of interest, the acquired data is suspect because of lack of causality between applied and acquired signals.

An additional note: the APx software offers three different calculation mode options, which it terms H1, H2, and Magnitude Only. H1 is the most common choice, chosen when the reference signal has a higher signal-to-noise ratio (SNR) than the acquired signal, which is usually the case, and is S_{AA}/S_{AB}. H2 is more appropriate if the reference signal is noisy (e.g., a sound capture from the microphone of a smart speaker in a room) and is S_{BB}/S_{AB}. If the reference and acquired signals have no common clock reference, you can select the Magnitude Only option, but of course, all phase information is lost.

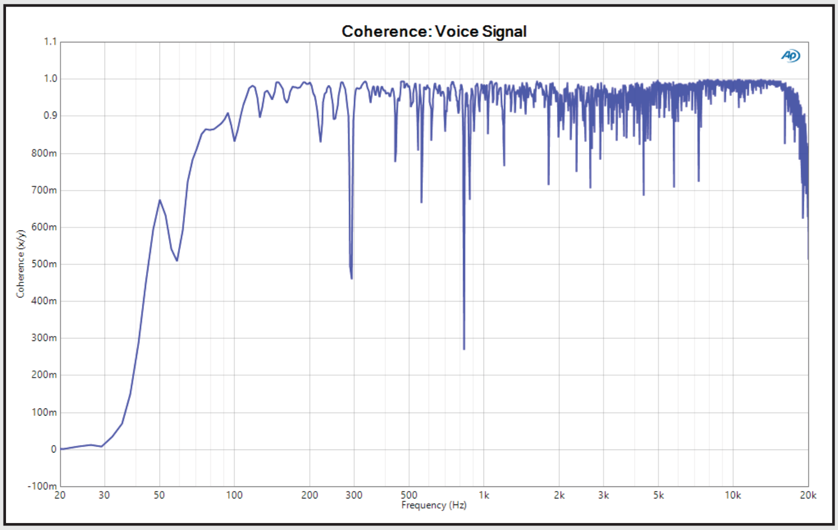

Figure 2: The coherence function calculated by the APx500 software between the reference file and the acquired acoustic signal is a measure of causality of the applied signal to the measured acoustic output.

test signal that can be swept during a measurement. This can be particularly useful for systems using an electret or MEMS microphone that will drop a small voltage. Although 0.5 Vdd is generally the optimum offset, the transducer voltage drop will affect this, and it's important to know how much offset will cause an increase in distortion. The APx software also allows safe limits ("pin protection") for the total offset plus peak signal voltage to be set in order to avoid damaging the DUT.

Figure 6 shows the results of sweeping the offset from -0.400 V to +1.8 V with a 1 kHz sine wave test signal at 1 mV (-60 dBV) applied to the Crystal codec. As is clear, over most of the DC offset range, the total harmonic distortion (THD) and gain remain constant and low. Above about +1.6 VDC, the THD starts to rise, and below -0.250 VDC, the gain starts to drop. This sets the small signal limits for this particular ADC, although if the signal voltage were to increase, the offset limits would likewise decrease because of the sum of the peak signal voltage and offset impinging on the ADC limits.

Although this measurement is differential, the offset sweep can also be applied to single-ended signals. The Crystal codec in single-ended mode shows similar results compared to differential. But that got me thinking…

For those of us dinosaurs who still do discrete circuit design, one of the interesting design and verification challenges is optimizing bias. It occurred to me that this feature of the APx555 B Series could be adapted to determine the bias level for minimum distortion in an analog amplifier. To try this out, I slapped together a simple, non-optimized gain stage consisting of a 12AT7/ECC81 tube as the gain element (I use whatever is close to hand—this could equally well have been done with FETs or degenerated bipolar transistors) and cathode biased it at +3.4 V with respect to the grid using two red LEDs in series. I then applied a single-ended signal (0.141 VRMS at 1 kHz) to the grid and swept the DC offset from +0.1 V to +2.7 V, which combined with the fixed bias voltage at the cathode effectively swept the tube's overall grid-to-cathode bias from -3.3 V to -0.7 V, respectively. A graph of the output voltage versus applied bias with the drive voltage fixed is shown in **Figure 7** and one can see nearly a 5 dB gain shift over this bias range. Likewise, the THD versus bias is shown in **Figure 8** and has a broad minimum at +1.4 V, corresponding to a grid to cathode voltage of -2.0 V, which I've previously observed is a "sweet spot" for this tube type. If you're designing and optimizing any sort of analog gain stage, this is a very handy way to optimize distortion and gain.

It did not escape my attention that this feature could also be used to see the effect of offset voltage on the common mode rejection of a differential input, but unfortunately, I ran out of time before gathering this data. When I win the lottery and buy one of these for myself, I promise to show some examples!

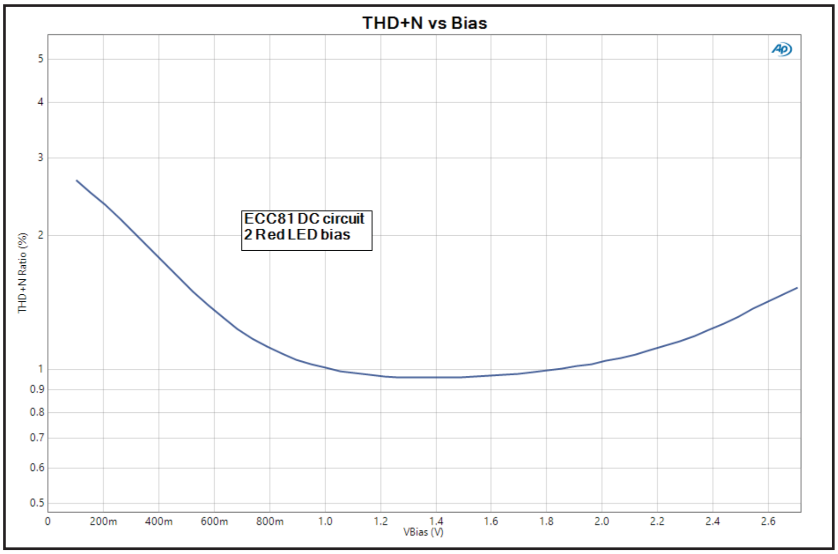

Figure 8: The DC offset sweep function can also be used to determine optimum bias point for an analog gain block to achieve minimum distortion.

Final Thoughts

The various new features of the APx555 B Series and the latest version of software were an absolute delight to use. Not only do they do exactly what they are intended for, their capabilities can also be extended into other uses. The upgrade from older APx555 versions is not cheap at about $6,000, but in my view, it is very much worth it for companies pushing the state-of-the-art in distortion or developing smart devices. It's unfortunate that I only had a short time to explore the new capabilities, but they have been sorely missed since returning the unit to AP. Maybe I'll get a second lottery ticket tonight…

Resources

"Application Note: Smart Speaker Acoustic Measurements," Audio Precision, 2019, www.ap.com

"CS47L90 Low Power, Highly Integrated Smart Audio Codec," Cirrus Logic, https://statics.cirrus.com/pubs/proBulletin/CS47L90_PB_20160919.pdf

D. Mathew, "Testing Audio ADCs and DACs," *audioXpress*, December 2015.

G. Munster and W. Thompson, "Smart Speaker Macro Model Update," Loup Ventures, LLC, June 2013, https://loupventures.com/smart-speaker-market-share-update

"Transfer Function Measurements with APx500 Audio Analyzers," Audio Precision Tech Note TN138, www.ap.com

S. Yaniger, Speech test file, SYclotron Audio, https://syclotron.com/sound-clip-supplement-to-audioxpress-review-of-APx555B