Characterizing PC audio devices -- some challenges and their solutions

By Wayne Jones, Director of Applications Engineering, Audio Precision
PlanetAnalog
Jan 04, 2005 (8:36 PM)
URL: http://www.planetanalog.com/showArticle?articleID=56900543

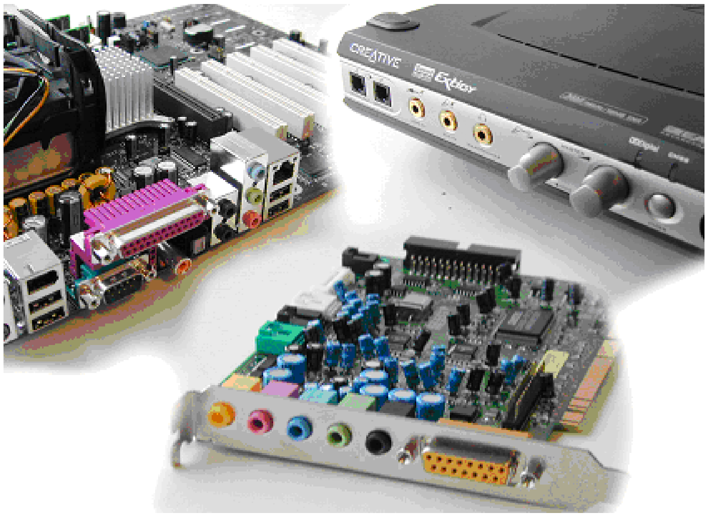

Click to Enlarge

The ubiquitous PC is no longer just a sophisticated word processor or email tool; it's migrating into a personal jukebox or even the entertainment center of the home. New media-centric software and operating systems, more sophisticated audio and video hardware, and attractive capabilities well beyond those available from conventional stereo devices are making the PC attractive as the new audio and video delivery system.

Home entertainment hardware long ago reached a level of technical maturity that provides high-quality, reliability, and excellent value. But the PC adds new levels of source material organization and accessibility that go well beyond CD or vinyl album collections. The PC lets you enjoy their media anywhere in your house, sharing the music collection with several family members independently and simultaneously. Selections can be called up quickly and play lists created for unprecedented convenience.

Of course, few of us are willing to trade audio quality for this added convenience. There is a high expectation from consumers based on what even a low-cost conventional stereo can deliver. Unfortunately, too often even a low-end stereo system can outperform the overall audio quality of a PC-based audio system, which is hard to believe when so many sound cards on the store shelves these days proclaim impressive numbers like "24-bit," "96 kHz" and "100 dB THD."

But what level of audio performance do today's PC audio systems actually deliver? Certainly the digital-to-analog converters are no longer the limiting factor they were years ago; 24-bit converters with excellent THD+N specs are now low in cost and easily available. But the noisy PC environment, the inadequate supporting circuits used with converters and even the interconnections between the PC and external audio equipment can degrade the overall system performance.

In this article, we'll look at how to characterize the audio performance of a PC audio system, some of the sources of error, and the difficulties in making some measurements in the PC environment. We'll go over common measurement practice, what should be measured, and how to report the results.

Deconstructing the PC audio device

First, let's be clear just what a PC audio device is. The most obvious form is the familiar "sound card", a plug-in card that adds PC audio capability to a PC. But it can also be a motherboard implementation, with essentially the same plug-in sound card components mounted directly on the motherboard. A PC audio device can also be an external box connected to the PC by a USB or IEEE-1394 interface. We'll treat all of these devices as the same since, in general, their functionality is similar. Figure 1 shows the block diagram of a typical PC audio device.

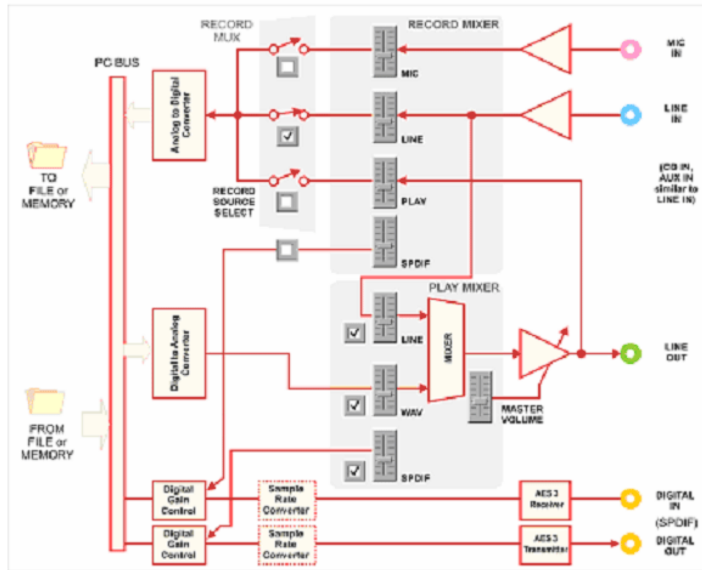

Figure 1. Block diagram of a typical PC audio device.

Click to Enlarge

Notice that there are several signal paths. We can organize these into two groups: the record or analog-to-digital path and the play or digital-to-analog path. Associated with both of these path groups are two mixer panels that are created by the device driver and can be made visible in Microsoft Windows. A typical mixer panel, as shown in figure 2, provides a fader or level control for each source. It also typically includes a mute or selection check box to enable or disable the particular source. For most devices, the Play mixer is a true mixer that allows complete freedom to mix any variation of sources. However, the Record mixer is usually not a mixer but a multiplexer that allows only one of the available sources to be selected at a time, although there are separate faders for each source.

Every audio device has a specific set of capabilities, including the types and number of inputs and outputs. Just about every device has a line output; some also include surround outputs: the rear, front center and sub-woofer channels. Most devices have a line input and many also have a microphone input. Notebook computers may only have a mic input but no line input. Many sound cards have additional line-level inputs such as Auxiliary In, CD In, etc.

Figure 2. Typical PC Mixer panel. This example is the play control.

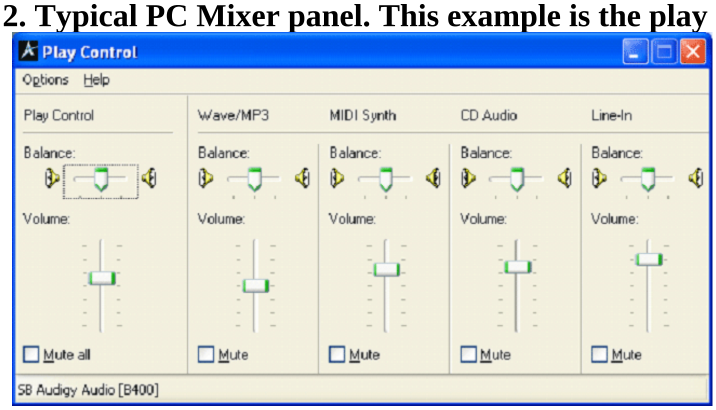

Click to Enlarge

In addition to these analog paths, newer devices may also include digital paths. More common is a digital output, although some may also have a digital input. These are usually labeled as "SPDIF" for Sony Philips Digital Interface. This digital format follows the definition described in the standard described in IEC60958, AES3 and AES3id. The digital output is typically included to provide Dolby Digital or DTS encoded surround sound 5.1 or 7.1 data to an external decoder that may be present in a receiver.

These digital inputs and outputs are also shown on the block diagram in figure 1. Many devices may include a gain control in the digital path that is controlled by a fader on the Play or Record mixer panel. Of course, this fader is not mixing this source in the same way as the other analog faders but does provide level adjustment.

All digital signals have a characteristic sample rate, and all converters operate at a specific sample rate. Sometimes it is necessary to exchange data between devices that are not operating at the same sample rate. In this case, as sample rate converter is used. Usually, a PC audio device will run its internal converters at a fixed sample rate, typically 44.1 kHz or 48 kHz. But the device may be able to accept external signals at a wide variety of sample rates (perhaps ranging from 8 kHz to 192 kHz) and provide output signals across a similar range of rates. This is generally accomplished by leaving the converters at their native rate but using a sample rate converter to provide these additional rates to and from the external interface.

Sample rate converters can be a source of signal degradation, suggesting that testing be done at various sample rates.

Potential areas of concern

With performance expectations high, there is a strong need to test PC audio devices and see how they match up. Although a PC audio device is in many ways an ordinary audio device, much like a tape recorder in fact, testing can present some new and unusual challenges. There are unique conditions in the PC environment that can make testing difficult: an audio-unfriendly and noisy environment, difficulty in setting levels and routing signals, shared resources, and hard-to-access software interface. And, considering the types of connectors involved and the system grounding, the actual test setup is even more important than with conventional audio devices.

Consumer PC audio devices usually provide 3.5 mm tip-ring-sleeve connectors for all analog inputs and outputs and RCA or phono coaxial connectors for digital interface. These analog connectors do not provide a robust ground connection, often leading to ground loops, noise and interference pickup. To minimize this problem, supplement the ground connection between the device under test and the test instrument with a heavy gauge (#12 or larger) cable securely bonding the PC chassis to the instrument ground. A large alligator clip at the PC end usually helps. Our experience has shown as much as a 6 to 12 dB reduction in measured noise with the addition of this supplementary ground. See figure 3 for an example of how a ground cable reduces the mains hum component in an actual system.

Figure 3. Low-frequency spectrum analysis graph of noise with and without a supplemental ground connection between the PC and the measurement instrument.

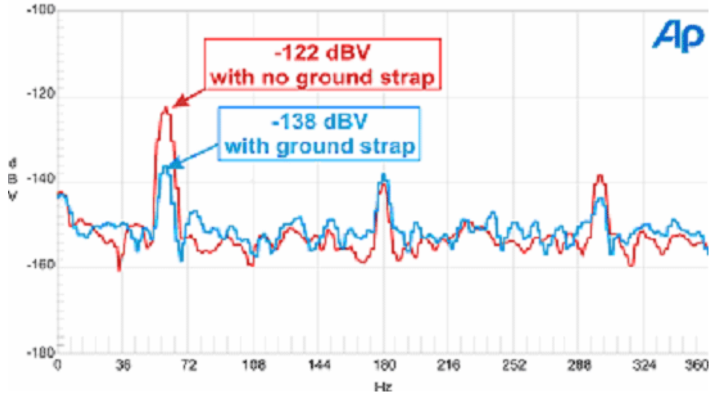

Click to Enlarge

Another area of concern is noise interference and out-of-band noise. These can come from a variety of sources. Of course, the PC chassis is full of possible noise sources: clocks, switching power supplies, hard drive motors, monitor horizontal and vertical scans, and over-sampling converters. This last source, in our experience, is one of the most problematical. Most present-day converters use a delta-sigma design that pushes the noise up out of the audio band, providing a quieter audio band in trade for significant but inaudible noise beyond 20 kHz. Although this technique provides a significant audio performance advantage, it can create problems for measurement. While human hearing may ignore energy above 20 kHz, most audio analyzers have flat response well beyond—often out to 100 kHz or even 500 kHz. These analyzers will react to the out-of-band energy and can occasionally produce incorrect readings. The Audio Engineering Society in their standard AES17 recommends the use of a sharp 20 kHz low-pass filter for measurements of devices with digital-to-analog converters. Figure 4 shows an example of the noise spectrum of a device with and without this filter.

Figure 4. Spectrum analysis graph of the output of an oversampling D-to-A converter showing rising out-of-band noise beyond 20 kHz (red curve) and roll off of this noise with the addition of an AES17 low-pass filter.

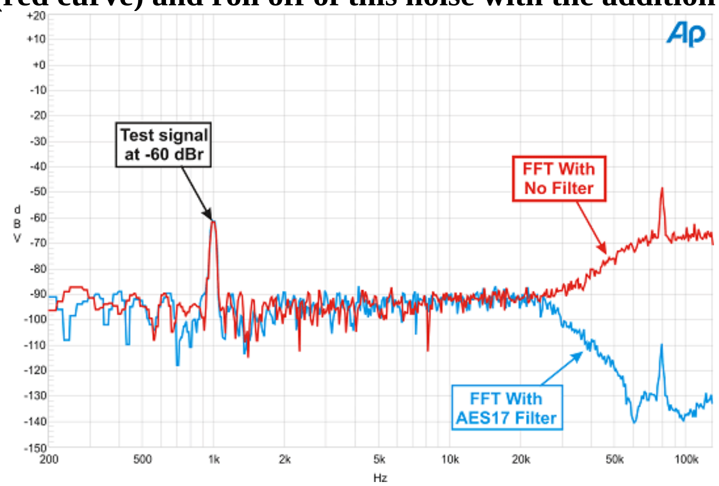

Click to Enlarge

Characterizing a PC audio device

To measure PC audio device performance, it may be tempting to simply send an analog audio test signal to the line input, record this on the hard drive, then play this signal back through the line output to an analog audio analyzer. This will characterize both the A-to-D and the D-to-A converter paths, but is not able to derive individual measurements of either converter path, or to characterize the digital path or any sample rate converters. To measure these paths individually, the signal must be routed directly to or from the PC digital bus.

An additional problem with the simple analog-to-analog composite path is that there is no absolute way to assess signal levels at the digital bus. Signal operating levels can have a significant impact on performance, as we will see later. The Record and Play mixer faders are generally not calibrated, preventing any way of setting specific gains, a necessary part of characterization.

There are a number of ways to gain access to the signal at the PC digital bus. One method is to use software tools to generate tests signals in the digital domain, and to analyze signals in the digital domain. When this method is used in conjunction with an external analog test generator and analog analyzer, the D-to-A Play path and the A-to-D Record path can be independently characterized. Some audio measurement systems can produce generator test signal files that can be "played" on a PC through the D-to-A path and can accept recorded signal files for analysis. Some systems provide direct connectivity with the PC bus for real-time signal generation and measurement. The Audio Precision PC Audio Test application provides a direct path to and from the PC digital bus allowing measurements to be made in real time without the need to record and subsequently playback test signals.

What we are really doing here are cross-domain measurements. We would like to generate a digital test signal with defined characteristics and specific test levels, stream (or "play") this test signal through the D-to-A converter, and measure the analog result with the analog analyzer. Then we need to generate a defined analog test signal, send it to the line input and record the signal, and analyze this digitally recorded signal in the digital domain. This allows independent analysis of the two primary paths and, with the proper software, accurate control of test levels. The Audio Precision 2722 Series audio measurement instruments are true dual-domain architecture allowing simultaneous generation and measurement of both analog and digital signals.

Reference levels

As anyone who has tested analog tape or cassette recorders knows, establishing reference test levels is very important when characterizing such devices. Every audio component has an optimal operating level. Too high, and you approach or exceed clipping or overload; too low, and you get into the noise region. Operating comfortably below the maximum level before clipping or overload will achieved the best performance. Measurement standards are very clear on setting reference levels before making any measurements.

A reference level is typically established by feeding a test signal to the input of the device and gradually increasing the level of this signal while monitoring the level and distortion at the output of the device. Various standards and common practices set definitions for maximum signal level. This may be just the onset of clipping, usually defined as just below the 1percent (-40 dB) total harmonic distortion point, or 1 or 3 dB below that point.

Here is where one of the biggest difficulties with PC audio sound card testing lies. With a simple analog amplifier, it is usually very easy to establish the onset of clipping. A graph of THD versus level will show a very pronounced knee in the curve at this point and a steep rise in distortion beyond this point. In a digital system, the maximum level is not called clipping but is actually the full scale value of the digital word. In an analog system, the clipping point can be somewhat soft; in a digital system, hitting full scale is very abrupt. While occasional minor overloads in an analog system may be audibly offensive, overloads to full scale in a digital system are completely intolerable. Recognizing this, many A-to-D systems include some form of compression or limiting that gently constrain the maximum signal to some point below full scale. This can make the determination of maximum level difficult, since it is not possible to look for digital full scale when it can never be attained. The graph in figure 5 illustrates what can happen in the line input circuit of a PC audio device that incorporates limiting in the A-to-D converter.

Figure 5. Typical transfer curves for an A-to-D converter path in a PC audio device. Note non-linear compression as input signal gets close to the full scale value but never reaches this value. The red curve is a plot of THD+N showing a sharp rise as the signal level approaches full scale.

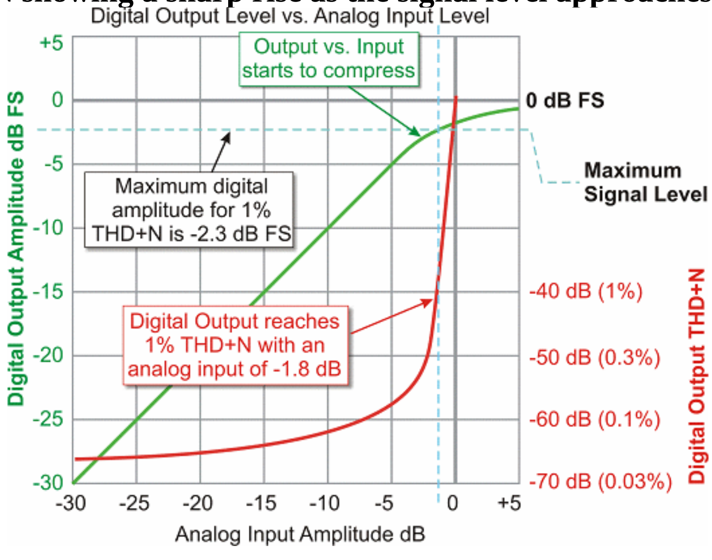

Click to Enlarge

The Audio Engineering Society has addressed this problem and recommended the correct practice to follow in such systems to establish a reference level. Standards AES17 and AES6id recommend a method of determining the maximum level, first determining the signal level that just produces 1 percent THD (-40 dB), and then reducing the level 0.5 dB from that point. This level is defined as the maximum level that will now be comfortably below digital full scale, but still high enough to provide the best noise performance. Specific measurements are then made relative to this reference point. For example, distortion is to be measured with a signal at 1 dB or 3 dB below this reference level.

Frequency response is measured at 20 dB below this reference point. These measurement levels have been established to provide the optimal operating point for the particular measurement.

Controlling the card

During the process of making measurements on a device, certain parameters of the device must be controlled. Signal routing must be established for the particular measurement; that is, the appropriate mixer check boxes must be checked or unchecked. Path gains must be set and optimized for best performance for each particular measurement. The difficulty here is that the gain setting elements (the faders on the mixer panels) are not calibrated. To complicate matters further, their control laws are usually inconsistent between manufacturers, and may change at different positions of the fader. The actual gain element has discrete steps that could be as coarse as 1.5 dB or more in some devices. This can make precision settings difficult.

Some measurements may require iterative adjustment of more than one gain element to achieve optimal performance. For example, the best noise and distortion readings may require careful balancing of the generator test level, line in fader level, and master fader level. Each of these will typically affect the signal level in only one portion of a path, and to balance signal levels throughout the system for best overall system performance, each must be optimally set.

Making the measurements

The actual audio performance measurements on a PC audio device are quite similar to those on conventional audio devices. The primary measurements are frequency response, noise, and distortion. Additional measurements that are particularly relevant in a stereo or surround sound system are interchannel phase and interchannel crosstalk.

Frequency response is most common and perhaps best understood measurement. It is an expression of the flatness of the system and how it may color the sound. It is usually presented as a graph with the horizontal axis being log frequency 20 Hz to 20 kHz. Frequency response can also be reported numerically in one of at least two ways: as amplitude deviation over a specific bandwidth, or as the bandwidth that provides a specific amplitude variance. Here are two examples of this expression:

Frequency Response (-20 dBr): ±1.5 dB relative to 1 kHz, 20 Hz to 20 kHz

Frequency Response (-20 dBr): +1, -3 dB relative to 1 kHz, 18 Hz to 19.5 kHz

Both measurements indicate that they were performed at the recommended 20 dB below reference level. The first one indicates that the response over the 20 Hz to 20 kHz audio band deviates no more than 1.5 dB from the normalized level at 1 kHz. The second expression indicates that the "3 dB down bandwidth" of the device is 18 Hz to 19.5 kHz, meaning that the response relative to a normalized 0 dB level at 1 kHz is within +1 dB and -3 dB within that band.

Some recommendations also call for measurement of passband ripple. This is a specialized frequency response measurement that focuses on the response errors that converter reconstruction and anti-aliasing filters may contribute. In the early days of converters, this ripple could be a significant contributor to response flatness, but is much less of an issue with today's converters. Improved filter design and over-sampling converters have reduced passband ripple to insignificant levels.

The measurement of noise is a little more involved. Here is where the strict establishment of a reference level is important, since this will directly impact the result. Noise is always expressed relative to a reference. In the expression "Signal-to-Noise ratio: -75 dB," the signal part of this ratio is the reference level.

For analog circuits, noise is measured by establishing a reference level, terminating the input with a low impedance (that is, no signal), measuring the noise and reporting the difference between the two. That method cannot be used in the digital paths in a PC audio device. Simply terminating the analog input will cause the converter to shut down the circuit, producing an unrealistic and incorrect result. It is necessary to keep the circuit open while measuring noise. We must, in an apparent contradiction, measure noise in the presence of signal.

There is a simple trick to accomplish this: we measure THD+N using a very low stimulus signal. This technique will carefully remove the stimulus signal from the measurement using a narrow notch filter set to the same frequency. If the level of the test signal is low enough, the resulting distortion components will fall below the noise level, and the analyzer will only see the noise. That is, the N of THD+N will be the dominant part, while the low-level test signal will still keep the path open. The standards recommend using a test signal of -60 dBr. Because of the different methodology used, this noise measurement is called Dynamic Range rather than signal-to-noise ratio. It is a measurement of the difference between the reference level and the noise level, expressed in decibels.

Measurement bandwidth and frequency weighting are important when measuring noise or dynamic range. The measurement bandwidth should be limited to the audio band, 20 Hz to 20 kHz in order to exclude inaudible out-of-band noise above the audio band, and any so-called flicker noise below 20 Hz. Frequency weighting may also be applied, to produce a result that is more closely aligned to the human perception of noise. For example, at low sound pressure levels the human ear is less sensitive to high and low frequencies. Since noise is typically at a low level, it makes sense to filter the noise measurement with a response similar to that of the human ear at low levels.

Two frequency weighting curves are in common use: "A-Weighting" and "CCIR-468" or "ITU 468". Both of these significantly roll off the low and high frequencies, as shown in their response curves in figure 6. The A-Weighting curve is commonly used in North America, while the CCIR-468 curve is popular elsewhere.

Figure 6. Noise weighting curves in common usage. Both curves approximate the response of the human ear at low sound levels.

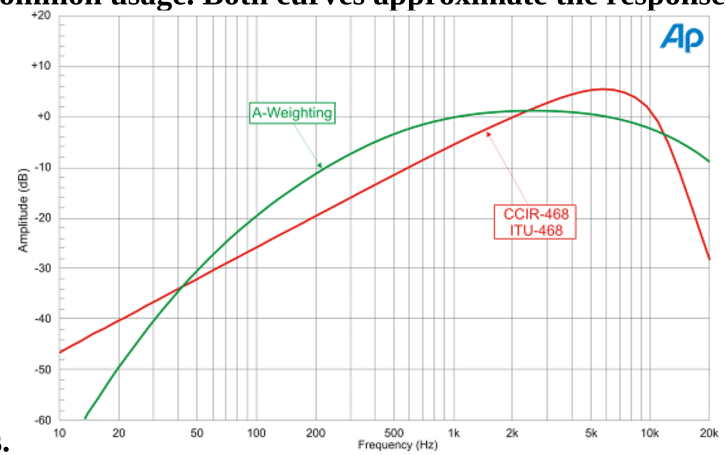

Click to Enlarge

When reporting noise or dynamic range, it is important to include the measurement bandwidth and what weighting filters, if any, were used. Here is an example of a correct way to report dynamic range:

Dynamic Range: 85 dB, A-Weighted

This indicates that the difference between the established reference level and the noise in an A-Weighted band is 85 dB.

Measuring distortion

The most common distortion measurement method is THD+N (total harmonic distortion plus noise). This uses a single, pure sine wave as the stimulus. The analyzer removes this fundamental with an analog or digital band-reject or "notch" filter and measures everything that remains, which includes the harmonics produced by the device and noise within the measurement bandwidth. This type of measurement is a good indicator of quality, but there are issues that should be considered.

Distortion should be measured at more than one frequency. Different elements in a device can contribute to distortion. Each of these elements may exhibit different non-linear behavior at different parts of the spectrum; some, for example, may show increased distortion only at lower frequencies or only at higher frequencies. For completeness, it is useful to measure THD over a wide frequency range.

This brings up an inherent problem with THD+N measurements in band-limited devices. Virtually all devices using converters are band-limited since they must include anti-aliasing and reconstruction low-pass filters. For a THD measurement to be valid, at least the second harmonic and preferably the third harmonic as well must be measured. If the bandwidth of the device is 20 kHz, this would indicate that the highest valid test frequency that can be measured is approximately 6 or 10 kHz. Measurements at higher frequencies will not be meaningful, because their harmonics fall outside of the limited passband. Of course, this doesn't mean that there is no perceived distortion at higher frequencies; in fact, many devices would have increasing distortion at higher frequencies. So how can this be measured?

Intermodulation distortion or IMD methods measure the distortion products produced by the interaction of two or more signals. There are several techniques for measuring IMD, but the most useful is to use two high frequencies of equal amplitude. The distortion components will be the sum and difference frequencies of the two test signals and the sum and difference frequencies of the harmonics of the test signals. The key point here is that the difference components will fall within the audio band. For example, using a test signal of 18 kHz and 20 kHz, the second order difference product will fall at 2 kHz. This is the technique that can be used to characterize high frequency distortion performance of a band limited device. The test signal simultaneous can be near the upper band edge of the device and produce in-band distortion components.

Distortion measurements, either THD+N or IMD, should never use frequency weighting. Weighting the distortion components would result in unrealistic readings. However, band limiting is useful when measuring THD+N to reduce the contribution of out-of-band noise.

Stereo-related measurements

The three primary measurements described above will give a good picture of the performance of a device. Two additional measurements that, while less common, are particularly useful in stereo and surround sound systems are interchannel crosstalk and interchannel phase difference. Poor performance in either of these categories can impair spatial imaging, an important attribute of both stereo and multi-channel surround sound experiences. Analog component matching, digital sample sequencing and digital clocking can affect these parameters.

Interchannel phase is measured by stimulating all channels with a pure sine wave and measuring the phase difference between all channels with respect to a reference channel. It is important to make this measurement at several frequencies across the audio band, as most phase degradation will occur at frequency extremes.

Interchannel crosstalk is measured by stimulating one or more channels and measuring the leakage into undriven channels. For accurate results, the measurements should be made with a tracking bandpass filter tuned to the stimulus signal. This method will exclude noise and other signals, providing a reading of crosstalk only.

Special measurements

The five types of measurements discussed to this point all relate to audio performance. In addition to these, there are some specialized tests that are particular to the PC environment.

The PC's primary function is not audio-centric so it is generally not optimized for audio performance. It is an electrically noisy environment and it shares resources for efficiency. Most non-audio applications are tolerant of the consequences of this reality. If a file transfer is interrupted for half a second, it isn't even noticed. But interrupt an audio program delivery for a few milliseconds and most listeners will consider this unacceptable. In general, the electrical noise created by switching supplies and hard drive motors does not affect program operation but it can add unacceptable noise to an audio program.

Many of these conditions are new to the audio world, so tests are not in common use to quantify them. Nevertheless, there are some tests that can be performed to assess overall performance.

Dynamic range (the signal-to-noise test optimized for digital audio circuits) can be performed while exercising various capabilities of the PC to provoke interference. THD+N measured over a long duration can expose occasional "glitches" caused by resource overload conditions. The notch filter used on virtually all THD+N analyzers is an excellent signal magnifier that is very sensitive to amplitude or phase changes over time. Signal dropouts, repeats, transient spikes or other disturbances will show up as an abrupt increase in distortion. Monitoring the distortion residual over time while provoking the system to stimulate such disturbances can be an effective test.

Tying it all together

With an enhanced expectation for audio quality in a multimedia PC, the need to test increases. The use of testing techniques based on traditional methods and adapted to the special requirements of the PC can be a sharp tool to characterize quality.

For effective PC audio device testing, pay close attention to the test setup and perform a comprehensive series of tests. With careful interpretation of the results, the PC audio device designer or manufacturer can satisfy demanding customers by providing audio quality matching that of a high-quality home stereo system, but with the added features and convenience that the PC brings.

The Audio Precision PC Audio Test application in conjunction with a 2722 Series measurement instrument is able to perform all of the measurements described here as specified in the various international measurement standards. It provides direct signal connectivity at both the analog and PC bus digital points to permit independent and comprehensive characterization of each path. It controls the device under test and can automatically establish reference levels as discussed above using a sophisticated multi-stage iterative algorithm. Figure 7 is an example of a test report produced by this application.

Figure 7. Sample Test Report for D to A Path produced by Audio Precision's PC Audio Test application.

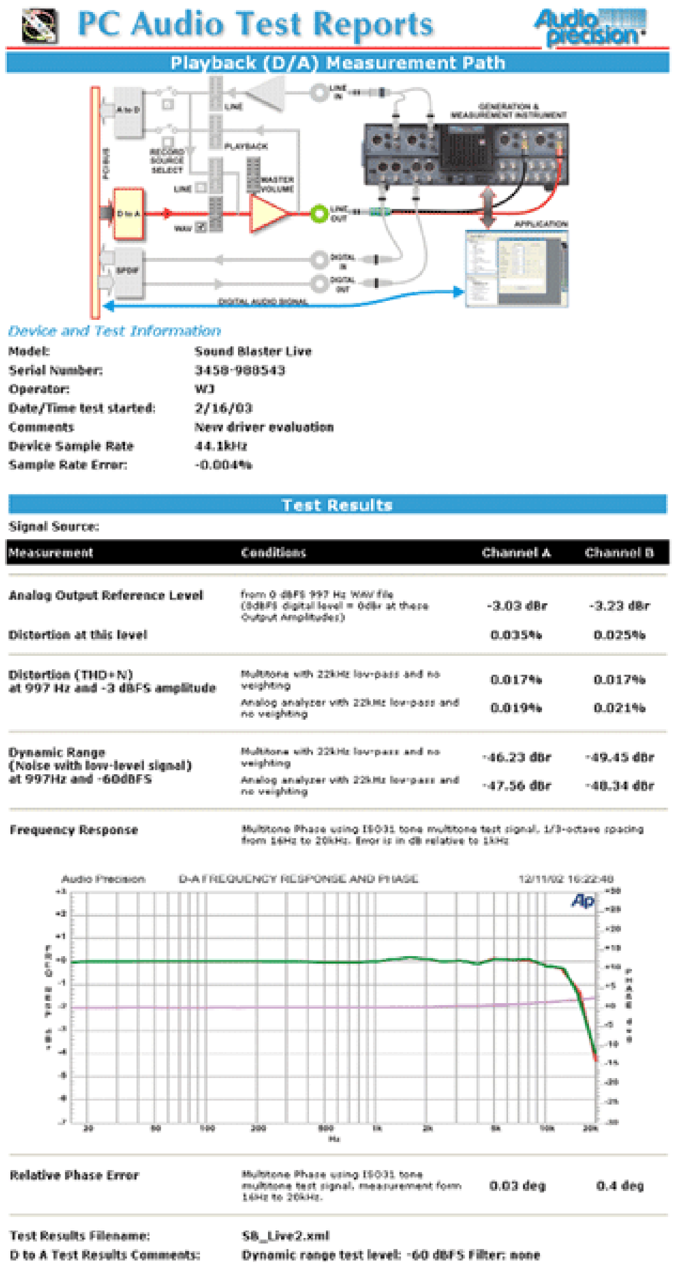

PC Audio Test Reports

Playback (D/A) Measurement Path

Device and Test Information

Model:	Sound Blaster Live
Serial Number:	3458-988543
Operator:	WJ
Date/Time test started:	2/16/03
Comments	New driver evaluation
Device Sample Rate	44.1kHz
Sample Rate Error:	-0.004%

Test Results

Signal Source:

Measurement	Conditions	Channel A	Channel B
Analog Output Reference Level	from 0 dBFS 997 Hz WAV file (0dBFS digital level = 0dBr at these Output Amplitudes)	-3.03 dBr	-3.23 dBr
Distortion at this level		0.035%	0.025%
Distortion (THD+N) at 997 Hz and -3 dBFS amplitude	Multitone with 22kHz low-pass and no weighting	0.017%	0.017%
	Analog analyzer with 22kHz low-pass and no weighting	0.019%	0.021%
Dynamic Range (Noise with low-level signal) at 997Hz and -60dBFS	Multitone with 22kHz low-pass and no weighting	-46.23 dBr	-49.45 dBr
	Analog analyzer with 22kHz low-pass and no weighting	-47.56 dBr	-48.34 dBr
Frequency Response	Multitone Phase using IS031 tone multitone test signal, 1/3-octave spacing from 16Hz to 20kHz. Error is in dB relative to 1kHz		

Measurement	Conditions	Channel A	Channel B
Relative Phase Error	Multitone Phase using IS031 tone multitone test signal, measurement form 16Hz to 20kHz.	0.03 deg	0.4 deg
Test Results Filename:	SB_Live2.xml		
D to A Test Results Comments:	Dynamic range test level: -60 dBFS Filter: none		

Click to Enlarge

Figure 8. PC Audio Test application produced by Audio Precision's PC Audio Test application.

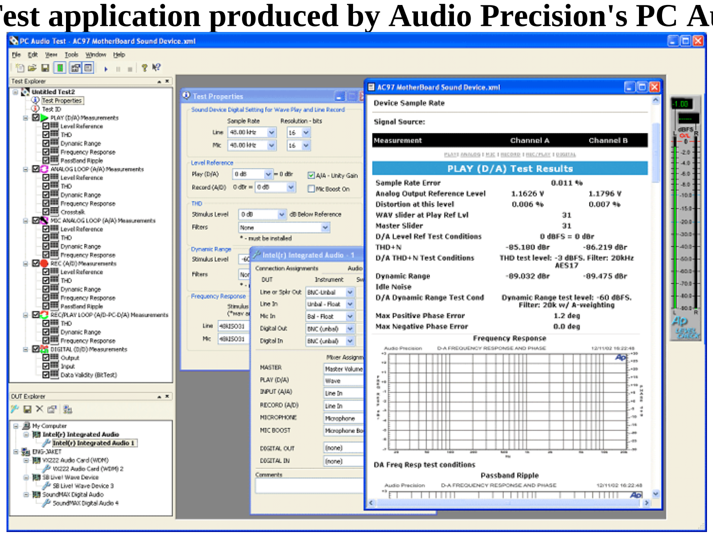

Click to Enlarge