

**A**DVANTAGES OF THE LATEST AUDIO FORMATS are apparent for silicon, software, and technology providers; system manufacturers; and media producers. But benefits for consumers are less obvious. Couple that with the formats' diversity and incompatibility, and they all may end up as niches. Listen up.

With high-resolution formats (those with sample sizes larger than 16 bits, sample rates higher than 48 kHz, or both), the audio industry hopes it's concocted the secret sauce that'll get consumers buying again. The pitches sure sound alluring: 24-bit samples, 192-kHz and 2.822-MHz sampling rates, and the like. The PC industry has historically used an analogous bigger-numbers-are-better pitch to sell (and upsell) consumers on its latest and greatest, most expensive microprocessors and the computers containing them. But, as the PC industry has lately discovered, there's an upper-end price threshold beyond which, especially in the absence of compelling applications, customers aren't willing to travel.

Will music aficionados refresh their collections with high-resolution versions and purchase the more expensive (and not coincidentally, copy-protected) high-resolution variants of new tunes they acquire, if their ears and brains can't discern a significant-enough difference from good, old Red Book audio CDs? (See sidebar "Sig-


nal to noise: calculating the high-resolution audio reality-to-hype ratio.") And will they rush out to buy new players, receivers, and speakers if surround sound, not ultrasonic sound, is the compelling new feature, and their existing gear will therefore suit them just fine? Ironically, these questions come just a few months past the 20th anniversary of the first CD-player sales in Japan. Recent trends don't give encouraging answers.

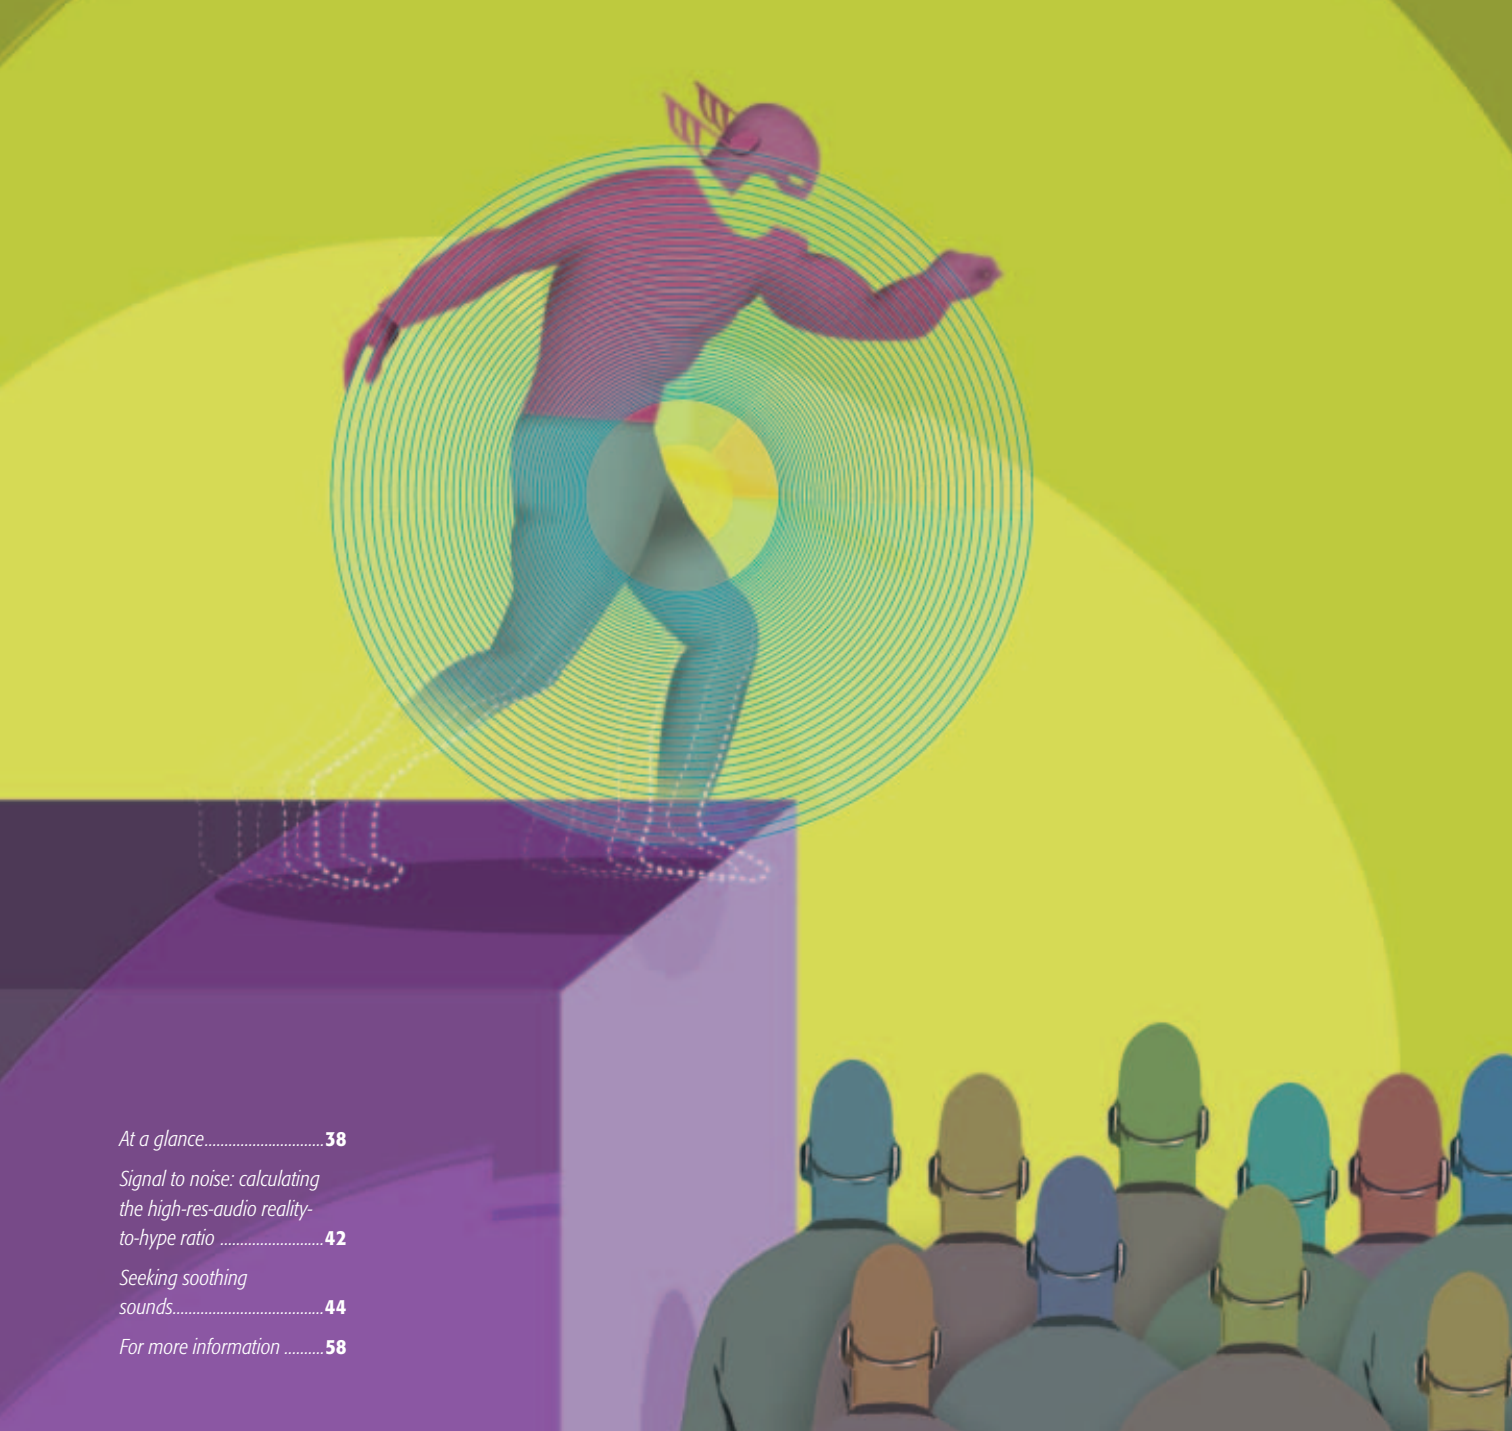
Warner Brothers recently announced plans to lower its DVD-Audio disc prices to CD-like levels. Do you think the company planned to make this drastic a price cut this soon, when it first decided to enter the DVD-Audio market years ago? DVD-Audio player prices, near \$1000 just one year ago, are now less than \$150 at retail. Granted, for suppliers, this situation is an improvement over the \$50 DVD-Video players crowding store shelves. But who thought DVD-Video players would be at \$50 this quickly? DVD-Audio isn't the only victim; SACDs (Super Audio CDs) and their players are

**High-resolution audio  
strides toward an  
unclear future**

# DESTINATION DISTORTION

Illustration by Mike O'Leary





*At a glance*..... 38

*Signal to noise: calculating the high-res-audio reality-to-hype ratio* ..... 42

*Seeking soothing sounds*..... 44

*For more information* ..... 58

close behind in price, and hybrid SACD-plus-DVD-Audio players have already dipped below \$500 (see sidebar “Seeking soothing sounds”).

Anticipated customer acceptance and the price at which you’ll achieve that acceptance are important issues to consider when designing your next audio-augmented system. Extra features invariably add cost, but if you can’t recover that cost (and, ideally, earn a profit) with a higher price tag, you’re in trouble. Which of the format contenders have the greatest chance for commercial success and why? When will that success happen and in which applications will it occur—pro-

fessional audio, high-end or mainstream home audio, car audio, or portable audio? The auditions are open, and the contestants have arrived. Let’s listen in and rate their performances.

Pacific Microsonics was perhaps the first company to develop a high-resolution audio technology that achieved broad industry adoption. HDCD (High Density Compatible Digital) audio builds on industry-standard audio-CD and -DVD foundations (Figure 1). By embedding control information within the altered LSBs (least significant bits) of, on average, less than 5% of the formats’ stored audio samples, the developers

claim that HDCD can dramatically expand the audio’s dynamic range if the playback device contains an HDCD-aware decoder. Otherwise, the altered LSBs take the form of uncorrelated noise, analogous to dither and are effectively inaudible.

The substance behind HDCD’s “20-bit” marketing sizzle takes three primary forms. First, HDCD encoding dynamically selects, as music characteristics vary over time among four anti-alias FIR (finite-impulse-response) filters with different characteristics between 16 and 22 kHz. The intent, in balancing frequency rejection and transient response, is to

make the filter act as sonically neutral as possible over a wide range of possible audio conditions. Some of the LSB control bits drive a corresponding configurable filter array at the HDCD decoder.

Other HDCD control bits extend peak levels as much as 6 dB beyond the normal 16-bit audio limits, via a compress-on-encode-and-expand-on-decode mechanism that conceptually works similarly to Dolby and other noise-reduction systems. Likewise, the remaining HDCD control bits extend low levels as much as 7.5 dB below an audio CD's dynamic-range bottom end. After reading the company's documentation and applying the 6.02-dB-per-bit rule of thumb, it's still unclear how HDCD marketers can claim 20-bit results, versus the slightly more than 18-bit dynamic range that peak- and low-level extension would imply.

When considering the DVD-Video format as a high-resolution audio-storage platform, some folks forget that they don't necessarily need to resort to lossy-compression schemes, such as Dolby Digital or DTS (Digital Theater Systems). It's possible to shoehorn at least two, and as many as six, high-resolution audio channels within DVD-Video's 6.144-Mbps maximum-allowable audio bit rate (Table 1). Given, though, that most DVD-Video discs (as the name implies) hold image-plus-sound presentations, that high video bit rates are critical to acceptable image quality, and that both filmmakers and viewers value the flexibility of multiple audio-track options within the stream (representing different languages, audio formats, and other qualities), lossy audio compression dominates the DVD-Video scene.

Dolby Digital compression delivers a 48-kHz sample rate, a "claimed" 20-bit dynamic range, and a 5.1-channel (five full-range and one for low-frequency effects) audio presentation via a 384- to 448-kbps bit rate (Reference 1). Because many of the formats that this article discusses from this point on are "lossy" (perceptual) in nature, perhaps it's time to elaborate on the "claimed" qualifier in the preceding sen-

**AT A GLANCE**

- ▶ Advanced audio formats' success won't necessarily come from their high-resolution features.
- ▶ Format-backward compatibility enables consumers to gradually upgrade their gear.
- ▶ Conversions to, from, and within the digital domain are a key piece of the audio puzzle.
- ▶ Processing options balance performance with price to deliver a design-specific optimum approach.

tence, which applies not only to Dolby Digital. A perceptual encoder might accept, for example, a 24-bit, 192-kHz sampled input, and its corresponding decoder might spit back out a 24-bit, 192-kHz stream. But who's to say what audio information the intermediary compressed file has preserved?

Perceptual encoders will, after all, discard audio information they judge to be inaudible to listeners (references 2 to 4). An encoder could decide, to achieve a certain overall level of quality at a target

bit rate, to throw away the dynamic range potential beyond 16 bits (96 dB, or 93 dB if LSB-dithered) or to filter out the two octaves' worth of information beyond 48 kHz, because "nobody can hear it anyway." This decision is pragmatic and by no means an inappropriate trade-off, especially at aggressive bit-rate targets.

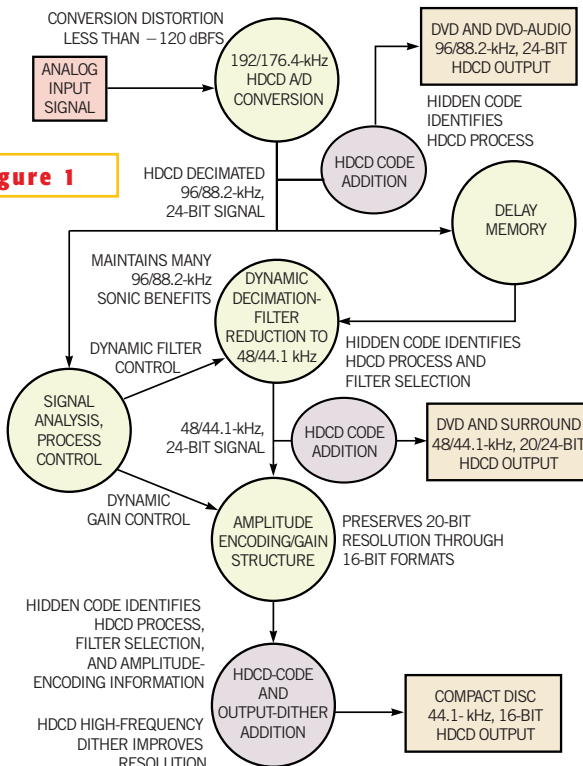
DTS's Coherent Acoustics lossy-compression scheme, Dolby's primary surround-sound competitor on DVD-Video, creates a 1.5-Mbps or 768-kbps compressed bit stream with a claimed 24-bit dynamic range at a 48-kHz sampling rate. DTS compression also appears on audio CDs at 1.234 Mbps, with 20- or 24-bit dynamic range and a 44.1-kHz sampling rate. Unlike HDCD, DTS audio CDs aren't backward-compatible with the Red Book audio-CD format, so decoders that aren't DTS-aware will output audio that sounds like the hiss of random noise.

**SUPERSONIC TONIC**

The DTS bit-stream format was, from the very beginning, designed for future extensibility, with corresponding backward compatibility to prior-generation decoders. The company, for example, uses this extension mechanism to add another channel's worth of data in the DTS-ES Discrete 6.1 format. It also finds use in storing the extra octave's worth of frequency information in the DTS 96/24 format. The first number in the name suggests a 96-kHz sampling rate. DTS 96/24 takes the difference between the 48- and 96-kHz signals, compresses it, and sends it in an optional field, which a first-generation decoder ignores (Figure 2). DTS 96/24-aware decoders sum the fields to recreate a semblance of the original audio.

With DVD-Audio, the entire 9.6-Mbps bit stream is available, if desired, for audio information. This higher bit rate enables the media to store two channels of 24-bit, 192-kHz uncompressed audio. The bit rate is not, however, high enough to enable the uncompressed transfer of six channels' worth of 24-bit, 192-kHz or, for that matter, even 96-kHz surround audio. Numerous ways to get around this limita-

**Figure 1**



**HDCD processing is applicable to both CDs and DVDs (courtesy Microsoft).**

tion exist. You could reduce the sampling rate or sample size of all channels. Thanks to the flexibility built into the DVD-Audio specification, you can also selectively reduce the sampling rate or sample size of only *some* of the channels; they need not all have the same characteristics. Consider, for example, a channel carrying primarily vocals (center), reverberation effects (surrounds), or low-frequency information (subwoofer), and the bit-slimming techniques possible in such cases.

The third option involves the use of MLP (Meridian Lossless Packing), which as its name implies was developed by Meridian Audio. Unlike most other high-resolution audio-compression algorithms this article discusses, MLP is lossless (Reference 5). In most cases, MLP in conjunction with buffering will enable

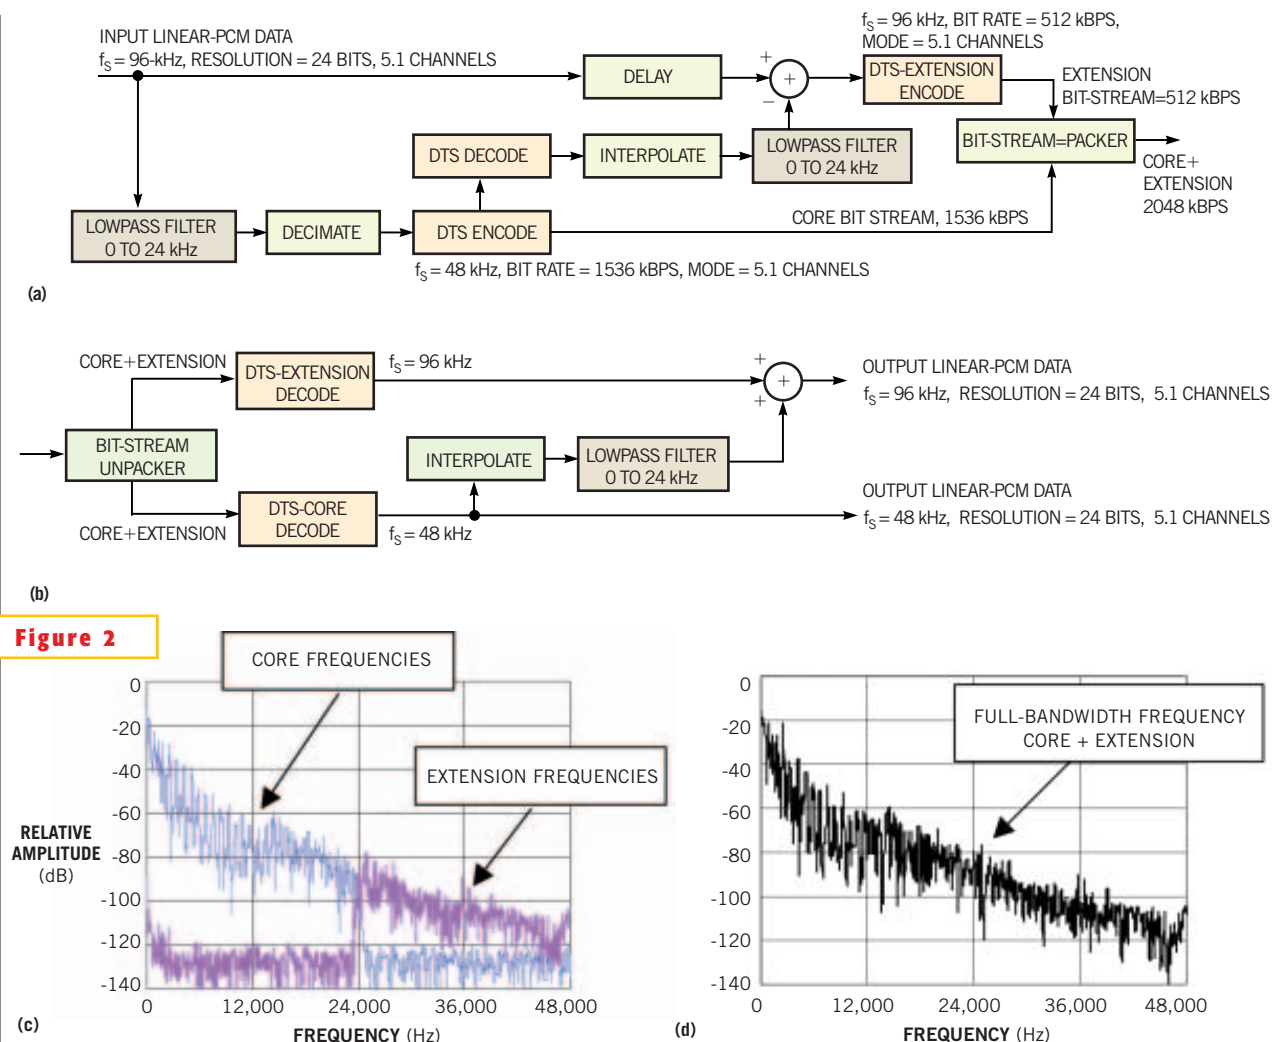
six-channel, 24-bit, 96-kHz audio storage (normally requiring a 13.8-Mbit bit rate) to fit within DVD-Audio's peak transfer-rate envelope (Table 2). Lossless-compression ratios depend on the source-media entropy characteristics, though, so there are no guarantees—you'll need to test the process yourself.

MLP also provides a means by which you can extend the per-disc playback time (Table 3). Its use on DVD-Audio discs is optional, but its presence on DVD-Audio bit-stream decoders is required. DVD-Audio format standardization is by no means complete; the DVD Forum is, for example, finalizing specifications for DVD-AR (DVD-Audio Recordable). DVD-AR reportedly comprehends not only linear and packed (MLP-encoded) PCM (pulse-code-modulated) audio, but also six lossy formats:

Dolby Digital, DTS, MPEG-1 (and MPEG-2) Layer II, ATRAC-3, MP3PRO (a backward-compatible superset of MP3), and MPEG-2 AAC (Advanced Audio Coding).

Because DVD-Audio discs can contain DVD-Video partitions, many record labels also put Dolby Digital- or DTS-encoded versions of the audio in these partitions for backward compatibility with legacy players. Many record labels also embed Verance-developed watermarking information within the otherwise-pristine DVD-Audio data, using a perceptual technique analogous to the one that lossy compression employs. Although the watermarking is not so invasive as to create objectionable audio, ABX comparative testing reveals that it is audible under some circumstances (Reference 6).

Watermarking differences are among



**Figure 2**

The DTS extension mechanism enables the encoding (a) of a bit stream that's backward-compatible with first-generation decoders (b). Advanced decoders combine the core and extension data (c) to create an enhanced audio presentation (d) (courtesy DTS).

the *least* significant of the variations between DVD-Audio and SACD. The brainchild of CD pioneers Philips and Sony, SACD employs a physical, rather than perceptual, watermark that therefore doesn't alter the audio characteristics. Although many first-generation SACDs delivered only two-channel audio, an increasing number of discs contain surround-sound mixes, and many SACDs are multilayer hybrids that also work on CD players—albeit in a 16-bit, 44.1-kHz two-channel fashion (Figure 3).

Whereas DVD-Audio builds on the multibit PCM data-storage approach first employed on CDs, SACD switches to a single-bit 2.822-MHz-sampled PDM (pulse-density-modulation) scheme called DSD (direct stream digital), which works in conjunction with MLP-like lossless compression. Philips and Sony claim that, by bypassing the decimate-while-recording and oversample-during-playback stages in PCM-based audio, SACD delivers a higher quality result. Ironically, though, SACD still employs multibit PCM-like techniques during the mixing and mastering stages of most discs' creation process. This decision was probably driven by a desire for compatibility with today's computing hardware and software that, after all, best handles information when it's clustered in multibit groups.

The last several AES (Audio Engineering Society) conventions have been jam-packed with papers debating both the absolute and the relative merits of DVD-Audio and SACD. Predictably, many of the presentations come from company representatives with vested interests in the formats, such as those of

**TABLE 1—DVD-VIDEO'S HIGH-RESOLUTION-AUDIO OPTIONS**

Per-channel audio sample size and sampling rate	Max number of audio channels	Audio bit rate corresponding to maximum number of channels (Mbps)	Remaining bandwidth for audio (subject to 6.144-Mbps maximum restriction), video, and other data (Mbps)
16 bits, 96 kHz	four	6.144	3.456
20 bits, 48 kHz	six	5.76	3.84
20 bits, 96 kHz	three	5.76	3.84
24 bits, 48 kHz	five	5.76	3.84
24 bits, 96 kHz	two	4.608	4.992

Derk Reefman, who works for Philips Research Laboratories. Some of them, though, derive from independent, usually academic, sources, such as Professor Malcolm Hawksford from the University of Essex, UK, and Professors Stanley Lipshitz and John Vanderkooy from the University of Waterloo, Canada.

These AES presentations are overwhelmingly critical of SACD, noting shortcomings such as its inability to dither the single-bit data stream, the nonlinearities that DSD subsequently creates, and the incapacity of noise shaping to fully move these artifacts out of the 20- to 20,000-Hz frequency range. Although the presenters acknowledge that the levels of these distortions are likely so low as to make them inaudible, they point out that similar distortions do not exist with the alternative PCM approach. The counterarguments offered by SACD proponents, such as Professor James Angus of the University of Salford, UK, are, in my estimation, insubstantial in comparison.

**SHOW SONY THE MONEY**

If SACD is an inferior high-resolution format, as these and other audio pundits claim, then why did the Sony/Philips alliance introduce it instead of joining the rest of the industry in supporting DVD-Audio? The answer, in a word, is *money*. Philips and Sony have greatly benefited from the success of the CD format and the licensing revenues

that this success has generated. In DVD, they saw the seeds of CDs' likely eventual demise, and they therefore created SACD as a proprietary competitor, fueled by both companies' considerable consumer electronics muscle and by Sony's synergistic status

as a music producer. Ironically, at least for the moment, SACD is the healthier of the two formats, thanks to the aggressive media rollout strategy that Sony Music and its affiliate labels employ. When any of the format alternatives is "good enough," consumers will buy whichever one delivers the artists they prefer.

Turning your attention to streaming media, you should first be aware of the 24-bit MP3 decoders L3Dec and MAD (MPEG Audio Decoder). They claim to reduce the distortion caused by rounding approximations in traditional 16-bit MP3 decoders, and you can either play the results unaltered on a 24-bit-capable sound system or dither them down to 16-bit versions. The MAD plug-in for Winamp is perhaps the easiest way to explore the concept for yourself; a HEX-to-WAV transcoder for L3Dec is also available. Mindful of the diverse sample sizes and sampling rates that the heir-apparent AAC specification encompasses, the Fraunhofer Institute presented a paper on 24-bit, 96-kHz AAC compression at the December 2001 AES Convention in New York City.

Microsoft's WMA (Windows Media Audio) Professional is the new kid on the block in high-resolution audio. The company adapted its base WMA codec to handle as many as eight audio channels along with larger-than-16-bit sample sizes and greater-than-48 kHz sampling rates. WMA has undergone numerous

**SIGNAL TO NOISE: CALCULATING THE HIGH-RES-AUDIO REALITY-TO-HYPE RATIO**

Although this article covers the high-resolution audio-format alternatives and the silicon that supports them, it doesn't address the overriding questions that are probably on some of your minds, such as "Why does high-resolution audio theoretically sound better?" and "Does the theory hold water in the real world?"

Disparate opinions on these topics are easy to obtain, but definitive answers are much harder to uncover. The answers also depend on the end applications and the thriftiness, or lack thereof, of potential customers in these applications. I'll explore the theory and the reality of high-resolution audio in an upcoming

how-it-works article in *EDN*. Until then, for more information go to:

- *Absolute Sound, Electronic Musician, Home Theater, Mix, Sound and Vision, Stereophile, Stereophile's Guide to Home Theater, and Surround Professional* magazines;
- trade shows, such as the Audio Engineering Society

- Conferences, the Consumer Electronics Show, the Home Entertainment show, and the NAMM (the International Music Products Association) Show;
- the publications of the Audio Engineering Society, available for purchase through the AES Web site; and
- rec.audio newsgroups.

quality-improving revisions during its brief life and is currently at version 9. Two-channel WMA (WMA Consumer) encoders create bit streams compatible with decoders stretching all the way back to version 2. Similarly, Microsoft hopes to “freeze” the WMA Professional bit stream at this initial version, so consumer-electronics customers can embed the necessary decoding hardware and software without fear of future obsolescence.

Just how large a bit stream WMA Professional requires is a matter of some ambiguity. The documentation Microsoft initially sent me suggested an encode bit rate of 384 to 700 kbps. When, not realizing that the literature already answered my bit-rate question, I again asked Microsoft for the information, the company provided the same 128-kbps-minimum-bit-rate response that it recommends for 16-bit, 44.1-kHz surround material. Asking the question a third time, for clarification, I learned that whereas 128 kbps

**TABLE 2—REASONS TO USE MLP**

Channels	96-kHz/24-bit	96-kHz/20-bit	88.2-kHz/24-bit	88.2-kHz/20-bit
6.0	Necessary to not exceed DVD-Audio player's 9.6-Mbit/sec read transfer capability and to ensure 74-minute maximum playing time.			
5.1				
5.0				
4.1	Necessary to ensure 74-minute maximum playing time.			
4.0				

would deliver six channels of audio, the quality for all but the least demanding material would be subpar compared with DVD-Audio or SACD.

Amir Majidimehr, Microsoft's General Manager for Digital Media comments, “It is generally true that as you increase the sample rate, the data rate should also go up to maintain the same level of distortion. However, this trend is usually for sample rates [of] less than or equal to CD's 44.1 kHz. It is unclear that distortion above 22.05 kHz is very audible. That is, as [you] go up to 96 kHz at the same data rate, there is bound to be more distortion at ultrasonic frequencies. But it

appears, in [Microsoft's] testing, anyway, that such distortion is not as audible as the distortion created, say, when going from 22.05-kHz sampling to 44.1 kHz.”

“The same is also true when you go from 16-bit to 24-bit [audio],” he continued. “Few systems, if any, can reproduce signals more accurately than 20 bits. So again, distortions in the least significant 4 bits may not be audible, and in fact they might even help to dither the DACs. The other key thing to consider is that one of the main benefits of using higher sampling rates is to encode the original high-quality digital master without re-sampling, not necessarily to preserve

## SEEKING SOOTHING SOUNDS

Those of you who've read some of my prior articles probably already know that I don't like to simply take the vendors' word on new technologies; I prefer to test-drive them for myself. High-resolution audio is no exception. Continuing along the path that I began in January 2001 (refer-ences 4 and 5), I've created a variety of large-sample-size and high-sample-rate test tones, supplemented them with some high-resolution music, and will run the suite through various formats'

encoders and decoders to see how they perform.

My study will be partially objective, or said another way, *quantitative*. As in my prior project, and again using the frequency- and time-based-analysis tools in Sonic Foundry's Sound Forge and Syntrillium Software's CoolEdit Pro, I'll compare the original clips with audio that's gone through lossy compression, watermark alteration, and other transformations. Format candidates include Dolby Digital (cour-

tesy of Sonic Foundry's Soft Encode), DTS (courtesy of Minnetonka Software's SurCode CD-DTS), and Microsoft's WMA Professional. Revisiting my past MP3 work, I also plan to compare conventional 16-bit decoders with the L3Dec and MAD (MPEG Audio Decoder) 24-bit versions.

DTS has offered to encode my high-resolution clips to its 96/24 format. I hope to work with Verance to watermark-encode the clips, which I'll then run through Minnetonka's MLP (Meridian Lossless Packing) Encoder and discWelder Chrome to create DVD-Audio discs. I also have the necessary software to create audio CDs and DVD-Video discs, and I have a Pioneer DVR-A04 DVD-RW drive. A Windows XP Professional-based PC with a HyperThreading-enabled 3.04-GHz Pentium 4 CPU completes the hardware picture.

I also hope to include some subjective, qualitative observations of the format alternatives, comparing them against each other and the original clips. In the living room, I'll be relying on

Sennheiser EH2200 headphones and a multivendor surround-speaker set, along with an Apex Digital AD-7701 SACD-plus-DVD-Audio player. In the office, I'll conduct auditions through a Creative Labs Inspire 5300 surround speaker set, driven by an Analog Devices-based AC'97 six-channel integrated audio subsystem, Creative Labs Audigy 2 Platinum and M-Audio Delta DI0 2496 PCI-based add-in cards, and a USB-based Creative Labs Extigy external audio peripheral (Figure A). The Audigy 2 comes with DVD-Audio playback software, and InterVideo has also promised me a WinDVD 4 upgrade with DVD-Audio capability.

Look for updates on my work over the next few months via a Web-site-addendum link in the online version of this sidebar. I'll likely also be summarizing my results in a print article later this year. As always, I welcome your suggestions for work I should make a top priority and DVD-Audio and SACD discs I might want to audition, and your feedback on the work I've done.



**Figure A**



**The Audigy 2 Platinum (a) and Extigy (b) are two of the many products I'll be testing this year (courtesy Creative Labs).**

high-frequency information that people may or may not hear.”

For most music, 192-kbps (variable-bit-rate) or 256-kbps (constant-bit-rate) encoding is sufficient, according to Microsoft. And for critical listening tests on challenging audio sources, the originally submitted estimates should suffice. Even 384 kbps is impressive, though, when you consider that it’s half the bit rate of the lowest of the two DTS 24-bit, 48-kHz, 5.1-channel encode options, and roughly one-fourth the bit rate of a two-channel, 16-bit, 44.1-kHz Red Book audio CD’s PCM stream. WMA Professional has obvious application in streaming-media delivery environments, but when you pair it with Microsoft’s version 9 video codec, it also opens the door to intriguing digital-cinema and red-laser-friendly high-resolution DVD scenarios (references 7 and 8). Microsoft acquired Pacific Microsonic’s HDCD technology in late 2000, adding yet another weapon to its high-resolution-audio arsenal.

**CONVERSION SEGMENTATION**

Audio, like any other human-sensory-input mechanism, is inherently an analog medium, both as it travels toward microphones during capture and as it travels out of transducers during playback. But, with the exception of the enduring cassette tape, all of today’s prevalent audio-storage and -distribution mediums are digital: CDs and DVDs, DAT (digital audio tape) and Minidiscs, along with AAC, MP3, RealAudio, WAV, WMA, and other file formats. Clearly, some conversion is going on, both to (with ADCs) and from (with DACs) the digital domain, as well as within (with SRCs, or sample rate converters) to bring all of the incoming digital data to a common sample rate prior to tackling mixing and other audio-processing functions.

Peruse the myriad ADC, DAC, ADC-plus-DAC (codec), and SRC options available from companies such as AKM Semiconductor, Analog Devices, Cirrus Logic, Texas Instruments, and Wolfson Microelectronics, and you may walk away with a severe headache. A tremendous diversity of alterna-

**TABLE 3—AVERAGE DVD-AUDIO PLAYBACK TIME WITH MLP COMPRESSION**

PCM format	Number of channels	Total playing time (min)
	6.0	89
	5.1	106
96-kHz/24-bit	5.1 and 2.0	74
	2.0	230
192-kHz/24-bit	2.0	125

tives exists; one obvious differentiator is the number of integrated channels. You’ll also discover that a few decibels’ difference in claimed dynamic range or THD (total harmonic distortion) plus noise, both inside and outside the audible frequency range, can significantly affect the price you’ll pay (Figure 4). First, though, you’ll need to ask yourself if you even believe the vendors’ specs (references 9 and 10)


An ADC might integrate one or multiple S/PDIF (Sony/Philips Digital Interface) transmitters. A DAC might include S/PDIF receivers, global or per-channel digital volume control, or the capability to directly interact with both multibit PCM and single-bit (SACD) inputs. A direct path from the SACD decoder to the DAC saves you from the added expense; the additional board space; and the audiophile purists’ wrath that a separate SACD-to-PCM transcode chip, such as Nippon Precision Circuits’ SM5816AF

would create. Keep in mind when evaluating your options that for DVD-Audio, only two channels have to support 192-kHz sampling rates.


For practical purposes, ask yourself just how wide the dynamic range and how low the THD plus noise really need to be, given that the codec is just one piece of an audio-processing chain that’s constrained by its weakest link. How high is the quality of the speakers that consumers will likely hook up, directly or indirectly, to this piece of equipment? How much degradation will occur through equipment interconnect? What kind of music will the average user listen to? And, perhaps most importantly, what are the characteristics of the anticipated listening environment?

Will your target customer be auditioning audio while sitting still in a pin-drop quiet anechoic chamber, or inside his or her car in the garage late at night with the engine off? Or will the system you design be playing background music at dinner parties with attendees milling about, be reproducing mostly dialogue and explosions for home theater, compete with an open window and jabbering front- and back-seat passengers on the open road, or inhabit the hazardous environment endured by a PC internal sound system?

Some examples from recent Cirrus Logic press releases may enlighten you to the many trade-offs you face. In October 2001, Cirrus quoted the CS4362 six-channel DAC at \$5.35 (10,000), and priced the CS4382 8-channel DAC at \$6.50. In May 2002, Cirrus priced the CS5361 ADC, with differential inputs, 114-dB dynamic range, and 105-dB THD plus noise at \$4.95 (10,000). The pin-compatible CS5351, with single-ended inputs, 108-dB dynamic range, and 100-dB THD noise was \$3.95. Combine channel and spec choices, and the decision becomes even more complicated, as Cirrus’s latest product-family announcement reveals (Table 4). All of the vendors slice and dice their product lines in a similar manner. It’s up to you to balance trade-offs and pick an option that works best in your situation.



**Figure 3**



(b)


**SACD’s PDM (pulse-density modulation) employs single-bit coding at a very high sampling rate (a, courtesy Sony), therefore the DSD (direct-stream digital) moniker. Hybrid discs are backward-compatible with CD players although in a two-channel, 16-bit, 44.1-kHz-sampled fashion (b, courtesy Philips).**

Larger chunks of audio data, flowing into and out of the system at faster rates, make correspondingly greater demands on the processing subsystem. Estimates of the number-crunching needed to decode DTS 96/24, for example, start at 25 MIPS (according to DTS, on the Analog Devices 21065L 32-bit floating-point SHARC DSP) and can rise far above that figure, reflecting DSP-architecture variations (such as the 32-bit integer processing in Analog Devices' Melody 32 DSPs), the use of high-level, inefficient languages to code the algorithms, and other factors. Factors *other* than a DSP's clock

rate are also critical in determining its performance; the amount of embedded memory, for example, is also key. Every time the DSP exceeds the capacities of internal RAM and ROM and must access much slower external memory, sustained performance will suffer.

The emergence of 24-bit audio has fueled the long-running debate over 24-versus 32-bit processors. Pragmatically, as long as you have enough time to do the necessary multipass calculations and enough memory to hold interim data, *any* processor data width will adequately perform the job you require of it without creating an adverse amount of truncation- and overflow-caused rounding error. Time, though, is of course the limiting factor; listeners won't long tolerate nonreal-time audio processing! The wider the internal data path (extending to the adoption of floating-point capability), the fewer the passes the algorithm must make through the DSP. Similarly, the greater the amount of onboard hardware acceleration, the more work the processor can do in each clock.

This trend reaches its peak (of impressiveness or ludicrousness—take your pick) with Texas Instruments' TMS-320DA610 audio DSP, the first member in the company's Aureus line. TI claims that the DA610, which runs at 225 MHz, delivers a mind-boggling 1800 MIPS and 1350 Mflops of performance. How did TI come up with those astronomical numbers? Under certain conditions,



**Figure 4**

**When comparing specifications, whether on speakers for yourself or DACs for your design, look beyond the frequency-range claims and examine the plus-and-minus tolerances, which will separate the good (a) from the bad (b) (courtesy AKM Semiconductor).**

the DA610 can execute as many as eight instructions within a single clock. TI doesn't have an exclusive on this feature, though. As one of its competitors points out, it's possible for *many* vendors' DSPs to simultaneously execute multiple operations. For example, the following three parallel operations:

1. mac x0,y0,b
2. x:(r0)+,x0
3. y:(r4)+,y1

break down into the following seven "atomic" instructions:

1. multiply x0 and y0
2. add result to b
3. move x:(r0) to x0
4. increment r0
5. move y:(r4) to y1
6. increment r4
7. round result of multiplication

Will you always achieve 1800-MIPS performance in the DA610? Of course not. You probably won't even consistently achieve 225 MIPS of processing muscle, for reasons such as the earlier-described external memory bottleneck. Regardless of the vendors' claims, you need to make sure there's *always* enough overhead to handle not only the audio-decoding functions but also various

postprocessing tasks, including bass management and other types of speaker compensation, THX adjustments, and surround-sound speaker virtualization. If you run out of gas and you're using a traditional single-core DSP architecture, such as Analog Devices' SHARC line, Motorola's DSP5636x DSPs, or TI's DA610, you'll need to incorporate a second DSP in your design and allocate functions between the two processors, which is never a simple task.

Alternatively, Cirrus Logic's CS49400 is, all by itself, a dual-core DSP. The CS49400 walks a middle path in the 24- versus 32-bit debate. A 24-bit processor handles decoding, a separate 32-bit DSP handles post-decode functions, and the partitioned DTS 96/24 algorithm splits between the two processors. Motorola recommends the DSP56311 or DSP56321, with their EFCOP (Enhanced Filter Coprocessor) cores, as companion chips for its main DSP5636x audio DSPs.

An upcoming revision of the DSP56367, code-named Onyx, will boost both the amount of on-chip memory and the clock speed of today's Motorola chips, the latter to a 180-MHz target. Motorola also plans to increase the speed of its EFCOP-inclusive DSPs to 180 MHz.

All of the aforementioned processors tend to find homes in home-theater receivers and other high-end gear. Eventually, these DSPs will probably be directly decoding DVD-Audio and SACD bit streams, just as they decode Dolby Digital and DTS streams today. Until the interconnect quagmire—which this article will later explore—gets straightened out, though, the optical disc players are responsible for this task. DVD chip sets from companies such as Cirrus Logic, STMicroelectronics, and Zoran are already handling DVD-Audio-targeted functions, such as watermark detection,

digital-rights-management decryption, and MLP decoding. SACD processing currently takes place in a separate Sony-sourced chip. But if the format achieves widespread popularity, and the IC vendors obtain licenses from Philips and Sony, SACD functions will also become integrated as another in a long line of simplifica-

**TABLE 4—CS425XX FEATURES AND PRICES**

Dynamic range (dB)	DAC channels	
	Six	Eight
109	CS42516 (\$3.93)	CS42526 (\$5.29)
114	CS42518 (\$4.65)	CS42528 (\$6.97)

**Notes:**

1. 10,000-unit quantities.
2. All devices also integrate a two-channel ADC, S/PDIF receivers, a S/PDIF transmitter, and other features.

tion steps so critical in cost-sensitive, consumer-electronics applications.

High-resolution audio support is increasingly appearing in mainstream PCs, not just in high-end machines for professional use. Until the AC'97 specification and silicon undergo another revision, 20-bit, 48-kHz, six-channel audio and 20-bit, 96-kHz, two-channel audio define the upper-end limit for codecs, such as those by SigmaTel. Higher bandwidth PCI-, USB-, and IEEE-1394-based boards and external peripherals fulfill today's ultrasonic audio needs. Creative Labs' THX-certified Audigy 2 line, based on the company's Emu processors, handles the generation of 6.1-channel audio (DTS-ES, Dolby Digital EX, and others); the decoding of DVD-Audio's full range of sample sizes and rates; and recording at 24-bit, 96-kHz quality.

Via Technologies' acquisition of IC Ensemble in late 2000 gave it 24-bit, 96-kHz (with the Envy24 processor) and 24-bit, 192-kHz (with Envy24HT) capability, of which numerous third-party sound-card manufacturers are taking advantage. Note that while 6x and faster DVD-ROM drives will play DVD-Audio discs when partnered with an appropriate audio subsystem and software, such as InterVideo's upcoming upgrade of WinDVD 4, no solution currently exists for playing SACDs on the PC, aside from the

legacy audio-CD layer on "hybrid" discs. Perhaps that topic will be the next one that Project Bar-B-Q tackles ([www.projectbarbq.com](http://www.projectbarbq.com)).

For more on high-resolution audio, including a discussion of various analog and digital interconnection approaches available for your designs, go to the Web version of this article at [www.edn.com](http://www.edn.com). □

#### REFERENCES

1. Dipert, Brian, "Expanding options bring surround sound to the forefront...and the back...and the sides...and the ceiling...and the floor..." *EDN*, Jan 10, 2002, pg 34.
2. Dipert, Brian, "Now hear this," *EDN*, Feb 3, 2000, pg 50.
3. Dipert, Brian, "Digital audio breaks the sound barrier," *EDN*, July 20, 2000, pg 72.
4. Dipert, Brian, "Digital audio gets an audition: part two, lossy compression," *EDN*, Jan 18, 2001, pg 87.
5. Dipert, Brian, "Digital audio gets an audition: part one, lossless compression," *EDN*, Jan 4, 2001, pg 48.
6. Dipert, Brian, "Security scheme doesn't hold water (marking)," *EDN*, Dec 21, 2000, pg 35.
7. Dipert, Brian, "Video characterization creates hands-on headaches: part one," *EDN*, July 25, 2002, pg 53.
8. Dipert, Brian, "Video characteriza-

tion creates hands-on headaches: part two," *EDN*, Aug 8, 2002, pg 81.

9. Travis, Bill, "Demystifying ADCs," *EDN*, March 27, 1997, pg 26.

10. Travis, Bill, "Remystifying ADCs," *EDN*, Oct 9, 1997, pg 87.

#### ACKNOWLEDGMENTS

Thanks to Analog Devices, Cirrus Logic, and Via Technologies for their useful research materials. Thanks, too, to the companies who provided hardware and software for my hands-on work. I'd also like to acknowledge the members of the AES and participants in the various rec.audio newsgroups for their thought-provoking insights and opinions.

#### AUTHOR'S BIOGRAPHY

Technical editor Brian Dipert's auditory system may be unable to detect ultrasonic jams, but he's sure his four dogs and three



cats have keen-enough hearing, and that they would want their tunes to be as high quality as possible. Do you think his pet-sensitive argument will secure spousal approval

for the equipment he's always eyeing? Send other suggestions for how to persuade her to 1-916-454-5242, fax 1-916-454-5101, [bdipert@edn.com](mailto:bdipert@edn.com).

## FOR MORE INFORMATION...

For more information on products such as those discussed in this article, go to [www.edn.com/info](http://www.edn.com/info) and enter the reader service number. When you contact any of the following manufacturers directly, please let them know you read about their products in *EDN*.

#### AKM Semiconductor

[www.akm.com](http://www.akm.com)  
Enter No. 346

#### Analog Devices

[www.analog.com](http://www.analog.com)  
Enter No. 347

#### Cirrus Logic

[www.cirrus.com](http://www.cirrus.com)  
Enter No. 348

#### Creative Labs

[www.creative.com](http://www.creative.com)  
Enter No. 349

#### Digital Theater Systems (DTS)

[www.dtsonline.com](http://www.dtsonline.com)  
Enter No. 350

#### Dolby Labs

[www.dolby.com](http://www.dolby.com)  
Enter No. 351

#### InterVideo

[www.intervideo.com](http://www.intervideo.com)  
Enter No. 352

#### Meridian Audio

[www.meridian-audio.com](http://www.meridian-audio.com)  
Enter No. 353

#### Microsoft

[www.microsoft.com](http://www.microsoft.com)  
Enter No. 354

#### Motorola

[www.motorola.com](http://www.motorola.com)  
Enter No. 355

#### SigmaTel

[www.sigmatel.com](http://www.sigmatel.com)  
Enter No. 356

#### Silicon Image

[www.siimage.com](http://www.siimage.com)  
Enter No. 357

#### Sony

[www.sony.com](http://www.sony.com)  
Enter No. 358

#### STMicroelectronics

[www.st.com](http://www.st.com)  
Enter No. 359

#### Texas Instruments

[www.ti.com](http://www.ti.com)  
Enter No. 360

#### Via Technologies

[www.via.com.tw](http://www.via.com.tw)  
Enter No. 361

#### Wolfson Microelectronics

[www.wolfsonmicro.com](http://www.wolfsonmicro.com)  
Enter No. 362

#### Zoran

[www.zoran.com](http://www.zoran.com)  
Enter No. 363

#### OTHERS MENTIONED IN THIS ARTICLE

##### Apex Digital

[www.apexdigitalinc.com](http://www.apexdigitalinc.com)

##### Audio Engineering Society

[www.aes.org](http://www.aes.org)

##### Broadcom

[www.broadcom.com](http://www.broadcom.com)

##### Denon

[www.denon.com](http://www.denon.com)

##### Fraunhofer Institute

[www.fraunhofer.de](http://www.fraunhofer.de)

##### Genesis Microchip

[www.genesis-microchip.com](http://www.genesis-microchip.com)

##### Hitachi

[www.hitachi.com](http://www.hitachi.com)

##### Intel

[www.intel.com](http://www.intel.com)

##### Matsushita/Panasonic

[www.matsushita.co.jp](http://www.matsushita.co.jp)

#### M-Audio/Midiman

[www.m-audio.net](http://www.m-audio.net)

#### Minnetonka Software

[www.minnetonka-software.com](http://www.minnetonka-software.com)

#### Nippon Precision Circuits

[www.npc.co.jp](http://www.npc.co.jp)

#### Philips

[www.philips.com](http://www.philips.com)

#### Pioneer

[www.pioneer-electronics.com](http://www.pioneer-electronics.com)

#### Sennheiser

[www.sennheiser.com](http://www.sennheiser.com)

#### Sonic Foundry

[www.sonicfoundry.com](http://www.sonicfoundry.com)

#### Syntrillium Software

[www.syntrillium.com](http://www.syntrillium.com)

#### Thomson/RCA

[www.thomson.net](http://www.thomson.net)

#### THX

[www.thx.com](http://www.thx.com)

#### Toshiba

[www.toshiba.com](http://www.toshiba.com)

#### University of Essex

[www.essex.ac.uk](http://www.essex.ac.uk)

#### University of Salford

[www.salford.ac.uk](http://www.salford.ac.uk)

#### University of Waterloo

[www.uwaterloo.ca](http://www.uwaterloo.ca)

#### Verance

[www.verance.com](http://www.verance.com)

#### Warner Brothers

[www.warnerbros.com](http://www.warnerbros.com)

#### SUPER INFO NUMBER

For more information on the products available from all of the vendors listed in this box, enter no. 364 at [www.edn.com/info](http://www.edn.com/info).