

digital den



**THIS HANDS-ON
PROJECT EXPLORES
LEADING COMPRESSION
TECHNOLOGIES; SOUND OFF
WITH YOUR FEEDBACK.**

Auditioning high-resolution surround-sound compression

By Brian Dipert, Technical Editor

This article and its companion Web-site addendum are the culmination of a several-month project that digs into the inner workings of today's most prevalent high-end audio compression technologies (references 1, 2, and 3). It focuses on audio containing more than two channels and having per-channel samples larger than 16 bits, sampling rates greater than 48 kHz, or both, and I investigate Dolby Labs' AC-3 (also known as Dolby Digital), DTS' (Digital Theater Sys-

tems') Coherent Acoustics, and Microsoft's WMA (Windows Media Audio) Professional algorithms. I intend *not* to assess how well the algorithms preserve the sonic qualities of real-life music and other complex sounds; only your ears can make that determination (see sidebar "WAV woes" on the Web version of this article at www.edn.com). I instead attempt to show how the algorithms respond to a set of carefully crafted test clips; from these results, you should be able to draw some conclusions about how the algorithms perform their bit-

slimming magic tricks.

Dolby Labs announced Dolby Digital in 1992 as the sound foundation for the movie *Batman Returns*; the AC-3 algorithm at the heart of Dolby Digital predates MP3! In 1996, DTS unveiled the compression technology for audio CDs and DVDs that I tested. (A different DTS-developed algorithm for movie theaters had first appeared three years earlier in the film *Jurassic Park*, and DTS was founded in 1990.) The commencement of both companies' technology-development efforts undoubtedly predated

AT A GLANCE


▶ Established leaders and up-and-coming contenders vie for dominance in high-definition-DVD and other emerging applications.

▶ Successful analysis requires a beefy computer and robust software tools; a high-quality sound card and speakers let you audition the results.

▶ Amplitude attenuation and boosting, lowpass-filtering, and signal distortion were some of the artifacts my preliminary examination uncovered.

▶ WMA Professional exemplifies the significant compression efficiency that modern algorithms deliver.

▶ Regularly visit the EDN Web-site addendum for sound files, screen shots, batch files, and updates.



(a)



(b)

Figure 1

A high-quality sound card (a, courtesy M-Audio) and speaker set (b, courtesy Creative Labs) combine to deliver outstanding surround sound on the PC.

Surf the Web-site WAVs.....	92
Damage control	94
For more information	98

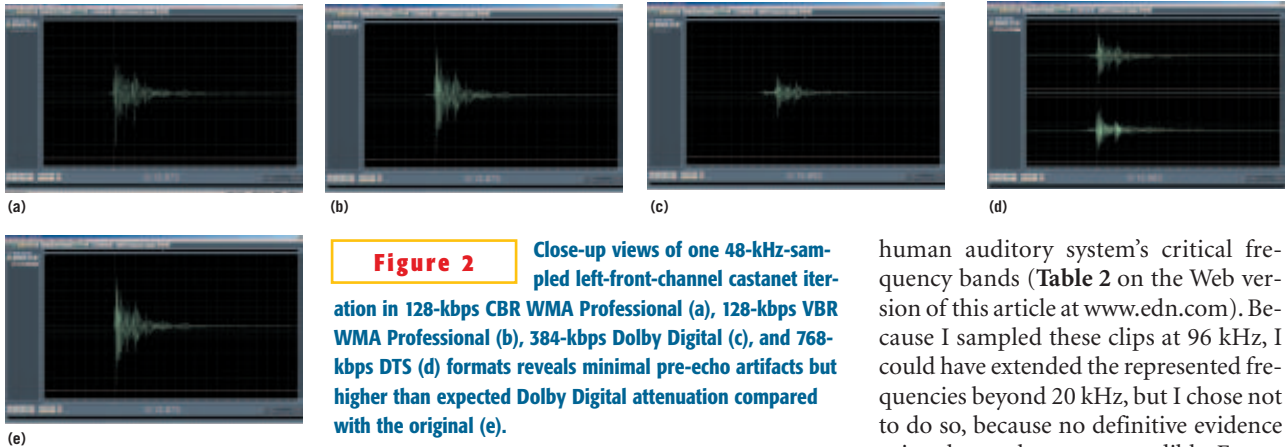


Figure 2 Close-up views of one 48-kHz-sampled left-front-channel castanet iteration in 128-kbps CBR WMA Professional (a), 128-kbps VBR WMA Professional (b), 384-kbps Dolby Digital (c), and 768-kbps DTS (d) formats reveals minimal pre-echo artifacts but higher than expected Dolby Digital attenuation compared with the original (e).

their products' public unveilings by several years. Conversely, the first iteration of WMA appeared in the spring of 1999, and Microsoft unveiled the final version of WMA Professional, which it built on a WMA foundation, in January. (The company released a public beta of what it then called Corona in September 2002.) WMA Professional is a newer technology than its competitors, benefiting both from feedback from Dolby Digital and DTS users, from additional audio theory and implementation research during the last decade-plus, and from advances in processing performance and reductions in computing costs. Keep the age discrepancies in mind as you peruse the results.

EVALUATION PLATFORM AND TEST SUITE

I used a 2-GHz Pentium 4-powered Dell OptiPlex GX260 as my encoding and decoding platform, as well as my listening station. The unit contained 256 Mbytes of DDR1-266 SDRAM and ran Windows XP Professional. To create my 96-kHz-sampled test clips and to resample both them and the 96-kHz clips I obtained elsewhere down to 48 kHz, I employed Syntrillium Software's Cool Edit Pro Version 2.1. In my past project, I had used both Cool Edit Pro and Sonic Foundry's Sound Forge. I preferred Sound Forge's frequency- and time-based displays to those in Cool Edit Pro, but its lack of greater-than-two-channel WAV support precluded its use this time. Ironically, as I wrapped up this print article, both companies announced that their products had been acquired; Sony now owns Sonic Foundry's Acid, Sound Forge, and Vegas products, and Adobe has purchased Syntrillium's technology assets.

I needed a sound card with six-channel, 24-bit, 96-kHz support; the PC's

built-in dual-channel Analog Devices codec, although acceptable for many computing tasks, was inadequate for my power-user needs. Both Creative Labs' Audigy 2 and M-Audio's Revolution 7.1 were available and would have worked; I chose the M-Audio card for its slightly better claimed SNR specifications, based on comparative reviews I'd read and because I had already installed an IEEE-1394 card in the PC, making that feature of the Audigy 2 sound card redundant. However, Creative Labs still factored into my setup; I listened to the uncompressed and compressed variants of my various audio clips through a set of Inspire 5.1 5300 speakers. The combination of high-quality PC, sound card, and speakers made for an enjoyable and highly recommended listening experience (Figure 1).

Those of you familiar with my earlier audio-compression study will undoubtedly note some similarities between that and this project's test clips (references 4 and 5 and Table 1 on the Web version of this article at www.edn.com). I first handcrafted six-channel, 24-bit, 96-kHz-sampled, pink- and white-noise files with every channel containing a unique noise pattern to complicate the compression algorithm's task. Each clip came in two variants: one in which all channels had near-0 dBFS (decibels relative to digital full-scale) peak levels and another in which the front-right, left-surround, and LFE (low-frequency-effects)-targeted channels were 20 dB below their front-left, right-surround, and center-channel peers. This version would enable me to quickly spot channel blending above a particular frequency threshold—a common bit-rate-reduction technique.

As before, I also created test clips that combined tones at the midpoints of the

human auditory system's critical frequency bands (Table 2 on the Web version of this article at www.edn.com). Because I sampled these clips at 96 kHz, I could have extended the represented frequencies beyond 20 kHz, but I chose not to do so, because no definitive evidence exists that such tones are audible. Front-right, left-surround, and LFE channels were 180% out of phase with their front-left, right-surround, and center-channel counterparts, again to complicate the compressor's job, and I created versions both with all channels at near-0-dBFS levels and with the front-right, left-surround, and LFE channels attenuated by 20 dB. I then added one- and three-quarter critical-band tones to each file, 20 dB down from the midpoint-tone counterparts, to create more complex multitone variants that would test for frequency masking. Because my earlier project's temporal masking tests had revealed no representative artifacts, I did not create the even *more* complex tonal combinations that would be necessary to look for them. I also didn't bother data-logging encoding and decoding delays; because some of the software I used predated both SSE instruction sets and the Windows 2000 and XP operating systems, I was concerned that the resulting performance impact would unfairly handicap the represented algorithm.

Last time, I successfully discovered pre-echo artifacts in the "castanet" clip, so I wanted this time to include some tests that also incorporated abrupt sound transients. My initial plan was to transform the previous project's two-channel, 16-bit, 44.1-kHz-sampled, EBU (European Broadcast Union) SQAM (Sound Quality Assessment Material) clips into six-channel, 24-bit, 96-kHz variants. Such resampling, though, wouldn't have created more meaningful audio data; at their heart, the clips would still deliver nothing more than Red Book quality. Eventually, I remembered that fellow audio aficionado Army Krueger's PCABX (www.pcabx.com) Web site offers two-channel, 24-bit, 96-kHz test clips, which

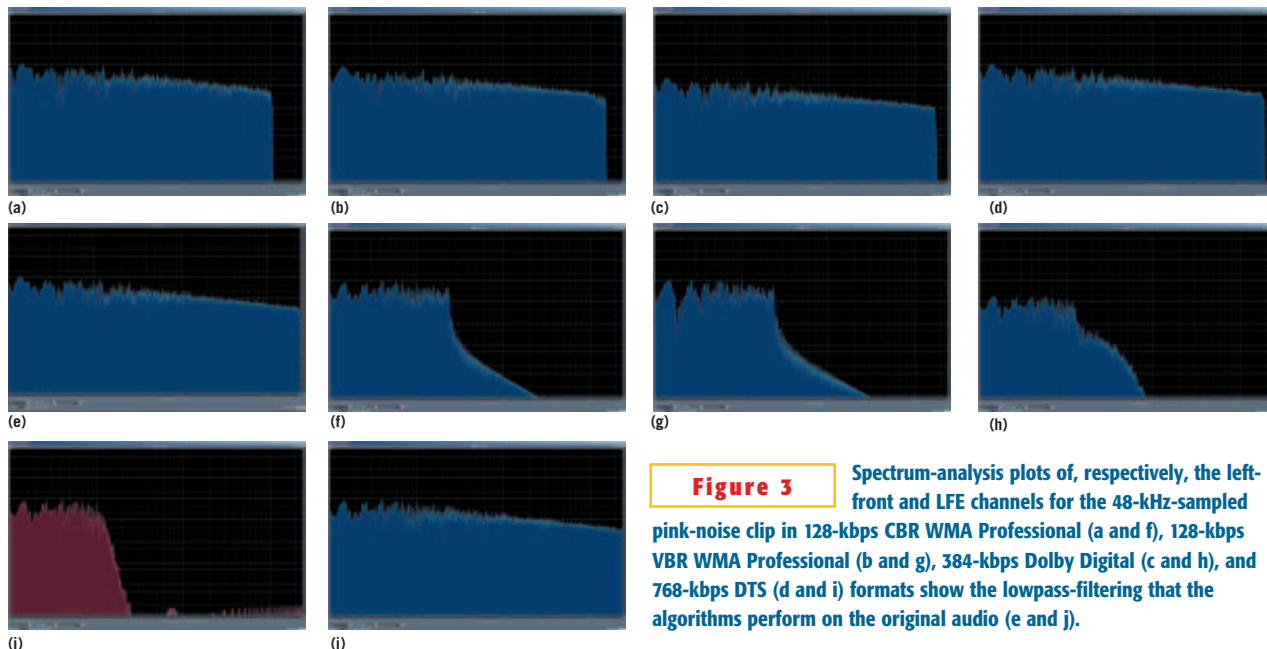


Figure 3 Spectrum-analysis plots of, respectively, the left-front and LFE channels for the 48-kHz-sampled pink-noise clip in 128-kbps CBR WMA Professional (a and f), 128-kbps VBR WMA Professional (b and g), 384-kbps Dolby Digital (c and h), and 768-kbps DTS (d and i) formats show the lowpass-filtering that the algorithms perform on the original audio (e and j).

he recorded using high-quality audio-capture equipment. I extended these clips to 30-sec playbacks through copy-and-paste waveform repetitions within Cool Edit Pro. I then duplicated the clips' front-left and -right channels to create, respectively, the left- and right-surround channels and blended the left and right channels to create the center and LFE channels. Because my Dolby Digital encoder requires 24-bit, 48-kHz source

files, I concluded the clip-creation portion of the project by downsampling (with dither on, 0.5-bit dither depth, a triangular probability-distribution function, and no noise shaping) *all* of my test clips from 96 kHz in Cool Edit Pro to generate 48-kHz variants.

ENCODING AND DECODING TOOLS

Microsoft's Windows Media 9 Encoder, which you can download for free from the

vendor's Web site, fortunately supports optional script-driven operation. Considering the number of files I was planning to compress and the number of combinations of bit rate and other configuration settings with which I hoped to encode each file, the ability to run the encoder in batch mode made for *much* more efficient use of my time. Microsoft also gave me a command-line-driven WMA-to-WAV decoding utility that supported as many

SURF THE WEB-SITE WAVS

My work so far answers some questions, but it predictably generates others. I've focused my attention on a few test clips and on low bit rates; how do results differ with other clips and at higher bit rate values? For WMA Professional, how does the presence and extent of compression artifacts vary if I encode to other bit rates, two-pass CBR (constant-bit-rate) mode; single-pass, quality-based VBR (variable-bit-rate) mode; or two-pass, peak-bit-rate-based VBR mode? DTS is moving into European digital-television broadcast and other applications; how will the Coherent Acoustics algorithm hold up at lower-than-768-kbps bit rates? I've also thus far con-

centrated on the front-left and -right and LFE (low-frequency-effects) channels. Will I see more or different artifacts when I turn my focus to the center and surround channels that, compression-algorithm developers might assume, respectively contain less critical dialogue and ambiance effects?

If Dolby Labs cooperates, I'd like to get to the root cause of the significant attenuation that the SoftEncode compression-plus-decompression combination did on the source files. Similarly, I hope to chat with representatives at DTS and Microsoft to find the basis of the amplitude boosting I found with their encoders and decoders and

to rerun my batch-compression sessions with reduced-amplitude source files that avert the clipped peaks I encountered. Although little to no evidence so far exists of amplitude-based channel mixing, I haven't yet looked for any indications that would suggest whether an algorithm has eliminated phase differences between channels.

Marc Nutter from Sonic Sense is sending me additional 24-bit, 96-kHz-sampled surround-sound files he's recorded, which I'd like to add to my test suite. Visit the version of this article sidebar on the EDN Web site for a link to a regularly updated addendum, in which I hope to answer these and other questions that you or I

can think of. There, you'll also find Cool Edit Pro screen shots of every test clip, for all six audio channels, in both frequency- and time-based modes, through all three compression algorithms, and at all tested parameter combinations. I'll document the exact settings I used with each encoder and provide copies of the batch files I ran. You'll also find available for download the test clips I used, so that you can conduct your own analyses. And, if I can secure sufficient server space, you'll be able to peruse and listen to some or all of the files that Dolby Digital, DTS, and WMA Professional compression created. I welcome and appreciate your feedback.

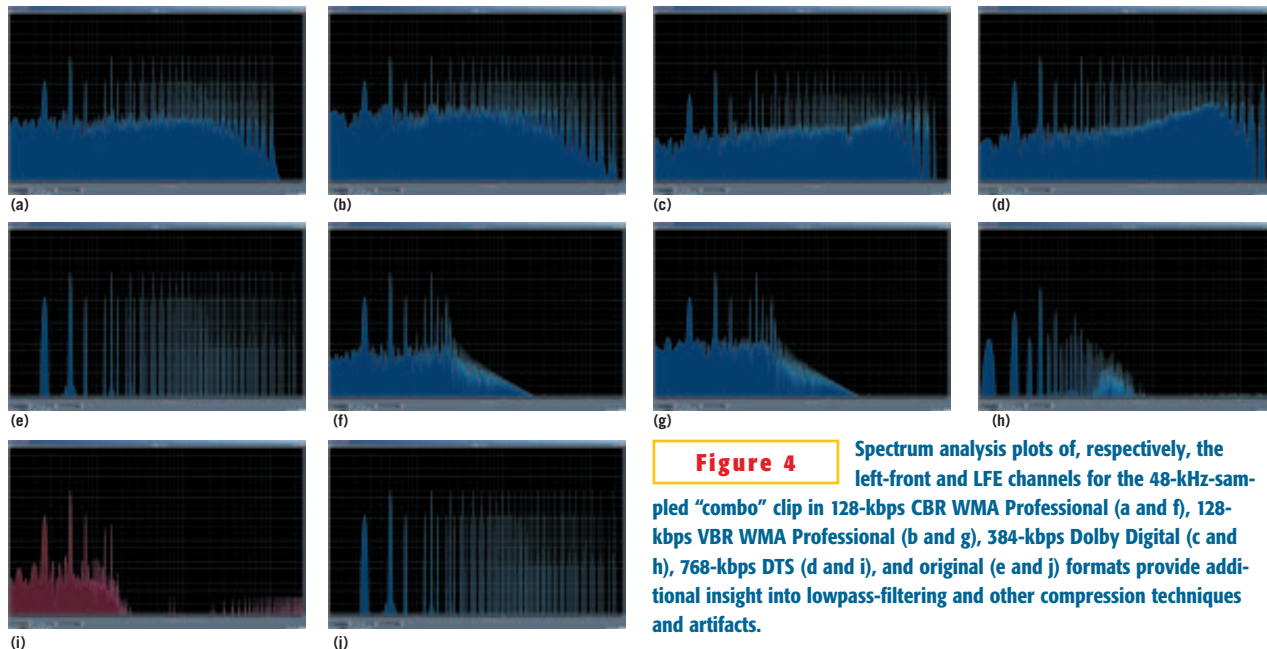


Figure 4 Spectrum analysis plots of, respectively, the left-front and LFE channels for the 48-kHz-sampled “combo” clip in 128-kbps CBR WMA Professional (a and f), 128-kbps VBR WMA Professional (b and g), 384-kbps Dolby Digital (c and h), 768-kbps DTS (d and i), and original (e and j) formats provide additional insight into lowpass-filtering and other compression techniques and artifacts.

as eight channels. I encoded each test clip at 128-, 192-, and 256-kbps bit rates to both one-pass CBR (constant-bit-rate) and two-pass, bit-rate-based VBR (variable-bit-rate) WMA files. The 384- and 448-kbps WMA files were CBR-only, because the similar-rate Dolby Digital algorithm is also CBR. For similar reasons, I encoded the 768-kbps WMA files, delivering DTS-like rates, only in CBR mode.

Dolby Labs declined to directly participate in this project, and, as a result its partner Minnetonka Audio also reluctantly turned down my request for a copy of the SurCode Dolby Digital-encoding software. Fortunately, I also had a copy of Sonic Foundry’s SoftEncode 5.1 handy. This several-years-old product, which the company’s Dolby Digital encoding plugin for Acid and Vegas now supersedes, runs under Windows XP and includes both a WAV-to-AC3 encoder and an AC3-

to-PCM-formatted WAV decoder. Alas, the decoder doesn’t run in batch mode, and the batch-mode encoder is GUI-driven, not controlled from a command-line prompt. I encoded 384- and 480-kbps AC3 files, which represent the bit rates on DVDs and in DTV; Dolby Digital in theaters runs at a film-sprocket-spacing- and frame-rate-defined 320-kbps rate, but, because Windows Media Professional didn’t support a comparable bit rate, I decided to forgo this encoding option. Again note that Dolby Digital lacks support for 96-kHz sampling; I could encode only the 48-kHz variants of my test clips.

DTS supplied me with command-line-controlled encoder and decoder utilities that, for 48-kHz sampling, enabled both 768-kbps and 1.536-Mbps encoding and, for 96-kHz sampling, supported only 1.536-Mbps encoding. As was the case with Dolby Digital, the DTS-compression

algorithm does only single-pass encoding and generates CBR files. I was initially flabbergasted when I discovered that the 768-kbps and 1.536-Mbps variants of each 48-kHz, DTS-encoded WAV file were the same size! I eventually discovered that both cases employed zero padding. DTS documentation states that the 768-kbps DTS stream is packed in the first 1008 bytes of 2048 bytes available in the S/PDIF carrier for one DTS frame (512 samples at 48 kHz); the remaining 1048 bytes are filled with zeros. The 1536-kbps DTS stream is packed in the first 2013 bytes of 2048 bytes available in the S/PDIF carrier for one DTS frame; the remaining 35 bytes are filled with zeros.

PRELIMINARY RESULTS

After several weeks of creating test clips and constructing and running batch files, I began looking for compression artifacts. I first searched for pre-echo. I began my examination with Dolby Digital, which I deemed most likely to exhibit the phenomenon due both to its being older than the other two algorithms and to its low bit rate versus DTS. The pre-echo I did find was of insufficient duration or amplitude to cause much disquiet. However, of more concern was roughly 8 dB of attenuation compared with the source clip (Figure 2). Broadening my inspection to additional clips, I saw attenuation values of 8 to 12 dB.

With the home-theater environment in mind, Dolby Digital’s developers built

DAMAGE CONTROL

Intel is now shipping its i865 chip sets, which include Sonic Focus’ DSP algorithms and Analog Devices codecs. Sonic Focus claims that these algorithms compensate for a significant amount of the signal distortion that two-channel lossy-compression schemes such as MP3 cause. Jim Barber, vice president of R&D for Sonic Focus, points out that the company’s highly configurable technology supports multichannel operation and that you can parameterize it on the fly to operate with any popular compression scheme and transport bit rate. He also says that it handles samples as large as 32 bits using either 32- or 64-bit processing cores and that it currently works with sampling rates as high as 48 kHz on AC-97/Win32-driver platforms. Although Barber remains mum on future-product specifics, he’s not immune to dropping hints. For example, he admits that the company’s engineers have to date probably played the DTS track of *Seven Bridges Road* on The Eagles’ *Hell Freezes Over* DVD approximately 5000 times!

into the algorithm format the ability for content providers, through a “dialnorm” variable, to tell the decoder to attenuate the signal by a specific amount in the process of playing it back. This dialogue-normalization feature saves viewers from constantly readjusting the volume as they switch between types of content, such as a television program and its accompanying commercials or multiple television channels.

Dolby Digital’s target loudness value is -31 dBFS, and the attenuation amount is $31 + \text{dialnorm}$ dB. I had expected a 4-dB level drop, because I left the Dolby Digital encoder at its default -27 -dB dialnorm setting, but I saw two to three times more attenuation than that amount. I expected some additional attenuation with wideband-frequency source material, such as pink and white noise, due to the lowpass filtering that’s part of all lossy-compression schemes. But for simple clips, such as castanets, this filter-induced attenuation shouldn’t be a significant factor.

Conversely, I occasionally saw amplitude boosting with DTS and WMA Professional that resulted in signal clipping. This phenomenon occurred only with 96 kHz-sampled complex patterns, such as pink and white noise, and at WMA Professional’s 768-kbps bit rate and DTS’ 1.536-Mbps bit-rate settings. With simpler patterns, with 48-kHz-converted variants of the pink- and the white-noise clips (which, as part of resampling, were

lowpass-filtered and therefore attenuated) and at lower encoded bit rates, the DTS- and WMA Professional-compressed files more closely approximated the amplitude of the source clips. I suspect that cumulative rounding errors through the numerous steps in the compression and decompression algorithms are the root causes of this boosted amplitude.

DTS has developed an infamous reputation in some audiophile cliques for using volume boosting as a short cut to achieving better perceived quality than Dolby Digital. As a result of my findings, I’m now inclined to believe that DTS doesn’t deserve this shady status; any amplitude differences that listeners experience when comparing Dolby Digital- and DTS-encoded soundtracks on their DVDs, for example, may be the result only of Dolby Digital dialogue normalization, a feature that DTS chose to not support (references 6 through 9). When you visit the EDN Web-site-based addendum to this article, you’ll see waveforms that expose another Dolby Digital-unique phenomenon: The left- and right-surround channels are 90° out of phase (see sidebar “Surf the Web-site WAVs”). This behavior finds use if the Dolby Digital decoder needs to downmix the audio to two-channel Dolby Surround matrix format (Reference 10).

At first, I was surprised to encounter no multichannel-to-monochannel amplitude mixing above a specific frequency threshold. Recall that collapsing the

multichannel presentation at high frequencies is a common lossy-compression technique (Reference 11). Listeners are least likely to perceive the channel-to-channel differences at these frequencies, and these differences are most likely to create sample-to-sample randomness. Revisiting my early 2001 project, I recalled that I saw significant evidence of a two- to one-channel combination only at MP3’s lowest 64-kbps bit rate. Dolby Digital at 384 kbps, especially in conjunction with the LFE channel’s bit-rate-slimming 120-Hz cutoff filter, may have enough bits available to forgo the need for significant multichannel “image” collapse. Tool limitations might also partly explain why I couldn’t easily observe channel blending. When I imported the six-channel WAV files that the Dolby Digital and WMA Professional decoders generate into Cool Edit Pro, the program split them into six tracks. I couldn’t simultaneously view multiple channels’ frequency spectrums.

Lowpass filtering, the other common lossy-compression bit-slimming technique, was present but less significantly than I had predicted. Examining the lowest bit-rate settings for each algorithm with the 48-kHz-sampled pink-noise test clip reveals the lowpass filter’s cutoff point and attenuation slope for both front-left and LFE channels (Figure 3). DTS has a higher frequency cutoff than Dolby Digital, but, at more than twice the bit rate of its competitor, it should. Conversely, WMA Professional delivers an impressive frequency range even at one-third the bit rate of Dolby Digital. Perusal of the “combo” test clip gives additional insight on lowpass filter operation, clearly showing which of the original clip’s tones didn’t outlive the encoder’s transformations and to what degree it disfigured the survivors (Figure 4). Note, too, the injected noise between the tones, present to some degree with all of the algorithms, and the Dolby Digital LFE-channel additional-tone artifacts, which, I’m guessing, are harmonics created during the compression process.

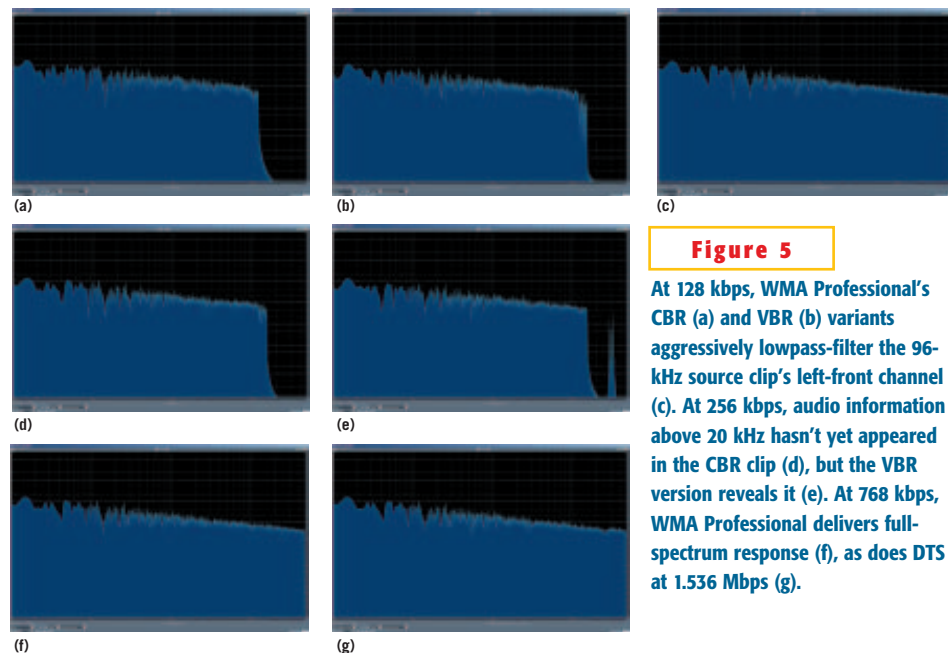


Figure 5

At 128 kbps, WMA Professional’s CBR (a) and VBR (b) variants aggressively lowpass-filter the 96-kHz source clip’s left-front channel (c). At 256 kbps, audio information above 20 kHz hasn’t yet appeared in the CBR clip (d), but the VBR version reveals it (e). At 768 kbps, WMA Professional delivers full-spectrum response (f), as does DTS at 1.536 Mbps (g).

Finally, consider some 96-kHz-sampled results (Figure

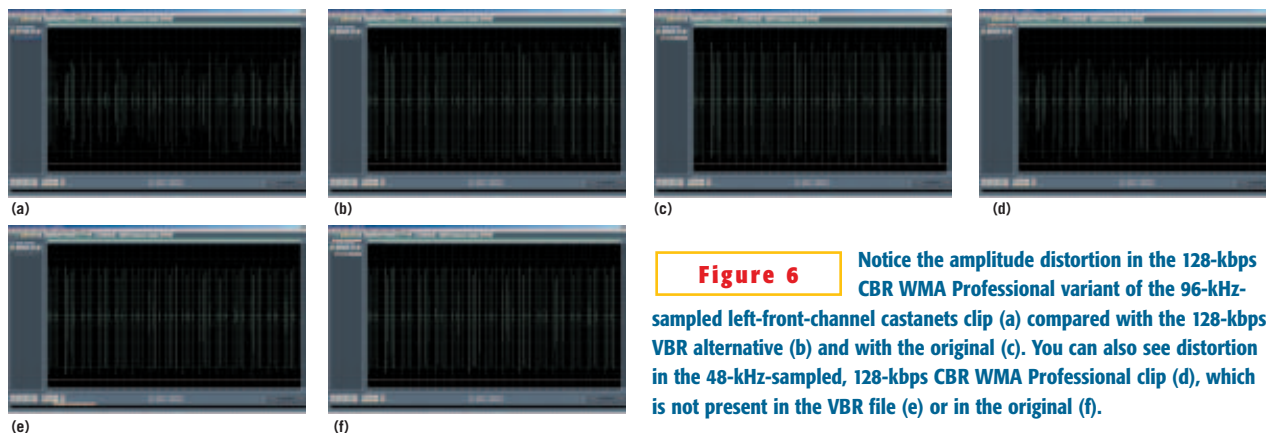


Figure 6 Notice the amplitude distortion in the 128-kbps CBR WMA Professional variant of the 96-kHz-sampled left-front-channel castanets clip (a) compared with the 128-kbps VBR alternative (b) and with the original (c). You can also see distortion in the 48-kHz-sampled, 128-kbps CBR WMA Professional clip (d), which is not present in the VBR file (e) or in the original (f).

5). At 128 kbps, WMA Professional aggressively lowpass-filters the pink-noise source clips. The spectral response differs from that of the 48-kHz example; Microsoft appears to employ different compression techniques for 48- and 96-kHz sources, even at WMA Professional's lowest bit rates. Also, note that two-pass VBR compression delivers wider spectral response than does CBR. I generally found that, at comparable bit-rate values and especially at low bit rates, two-pass VBR encoding produced fewer and less conspicuous artifacts than did single-pass CBR encoding. The VBR files were also almost always *smaller* than their CBR peers, suggesting that VBR mode delivers better quality at a lower average bit rate than does CBR mode.

Compare, for example, the CBR and VBR WMA Professional variants in **figures 2** through **4**. Look, too, at the strange, sinusoidal carrier-wave pattern of the 128-kbps CBR version of the 96 kHz-sampled castanets clip compared with both the original source file and the 128-kbps VBR version of the WMA Professional clip (**Figure 6**). Other channels' amplitudes oscillated with a similar periodicity, and the audible result was an unstable surround presentation; the 48 kHz-sampled, 128-kbps CBR castanets clip

also exhibited signal distortion that was not present in either the original or the VBR version. Alas, your application may not support the two-pass VBR format; live streaming presentations don't allow for two-pass encoding, and fluctuating VBR bit rates may overwhelm the transmission channel's available bandwidth.

At 256-kbps, audio information above 20 kHz begins to emerge in the WMA Professional VBR version of the pink-noise clip. (Notice the odd notch-filter phenomenon.) And, at 768 kbps, the WMA Professional encoder delivers a uniform spectral response all the way up to the Nyquist-defined 48-kHz limit. This bit rate may at first glance may seem high, but it's half the bit rate that DTS supports and less than 6% of the original, nearly 14-Mbps bit rate. Remember, too, that I neither generated VBR-compression files for 320 or 480 kbps nor tested the 640-kbps bit rate that WMA Professional also supports. WMA Professional may be able to achieve full-spectrum response at even lower bit rates than 768 kbps. And, if you believe, as I do, that full-frequency response to 48 kHz represents an extreme example of engineering overkill, more moderate WMA Professional bit rates will serve you just as well as 768 kbps does.

I was impressed with the advancements

in audio-compression technology that WMA Professional exemplifies and that its peers, such as AAC and Ogg Vorbis, also exhibit (see **sidebar** "Damage control"). Impressed, yes, but not surprised, given my experiences with the WMA algorithm both in the lab and off-hours as my chosen format for the "ripped" archive of music CDs residing on our home-media server. With Windows Media 9 reportedly one of the formats under consideration as the compression foundation for red-laser-based high-definition DVD, WMA Professional has convinced me that it can supply its piece of the puzzle that solves the apparent contradiction of delivering high-quality multimedia at low transmission bit rates and storage capacities. □

REFERENCES

You can find the references to this article on the Web version at www.edn.com.

ACKNOWLEDGMENTS

Special thanks to Army Krueger; Lorr Kramer and Kristin Thomson of DTS; and David Caulton, Ming-Chieh Lee, and Amir Majidimehr of Microsoft for their assistance and insights.

AUTHOR'S BIOGRAPHY



Technical editor Brian Dipert is patiently waiting for someone to start shipping PC-based DVD Audio-playback software that's compatible with his M-

Audio Revolution 7.1 sound card. Reach him at 1-916-454-5242, fax 1-617-558-4470, bdipert@edn.com, and www.bdipert.com.

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, contact any of the following manufacturers directly, and please let them know you read about their products in *EDN*.

Creative Labs
www.creative.com

M-Audio
www.m-audio.com

Syntrillium Software
www.syntrillium.com

Intel, www.intel.com
Minnetonka Audio

www.minnetonkaaudio.com

DTS (Digital Theater Systems)
www.dtsonline.com

Microsoft
www.microsoft.com

OTHER COMPANIES:

Adobe Systems
www.adobe.com

Sonic Focus
www.sonicfocus.com

Sonic Sense
www.sonicsense.com

Dolby Labs
www.dolby.com

Sonic Foundry
www.sonicfoundry.com

Analog Devices
www.analog.com
Dell, www.dell.com

Sony, www.sony.com

WAV WOES

In addition to the test clips, I also encoded a series of 24-bit, 48-kHz-sampled surround music files that Microsoft supplied. They are the source tracks for the music examples you can stream from Microsoft's Windows Media site. With these files, along with the others, I hoped to supplement my Cool Edit Pro-based objective analysis results with subjective auditory observations derived from playing the original and lossy-compressed clips on my PC and its high-quality sound card and speakers. Format inconsistencies unfortunately conspired to dash my plans.

Perhaps not surprisingly, my best results came when I employed the Microsoft WMA

Professional encoder-and-decoder set. I could directly play the WMA files in Windows Media Player 9; it also successfully handled the six-channel WAV clips that the decoder utility generated.

Alas, however, the Dolby Digital decoder scrambled the channel order in the six-channel WAV files it created. Windows Media Player downright refused to play these WAVs as well as any others, regardless of the number of channels they contained, which File->Properties->Summary explicitly reported as having a PCM-audio format. WAVs without this header information played in Windows Media Player. Cool Edit Pro's multichannel-playback facilities

are complicated and time-consuming to access and employ. WinAmp 2.9 easily played all of the WAVs, but the scrambled Dolby Digital channel assignments made any attempt at comparison with WMA Professional an exercise in futility.

The DTS decoder doesn't output a six-channel WAV; it instead generates three two-channel WAVs—one for the front channels, another for the surround channels, and one combining the LFE (low-frequency-effects) and the center channels. Theoretically, I could have tediously reassembled six-channel WAVs with Cool Edit Pro's multichannel encoder, but Cool Edit Pro assumed that the LFE

and center channels were in an opposite sequence within the WAV from the way the DTS decoder placed them, and I couldn't any way to override this mapping.

For both Dolby Digital and DTS, I had another playback option: I could have output a bit stream over the Revolution 7.1 sound card's S/PDIF (Sony/Philips digital-interface) connection for my home stereo receiver or (for Dolby Digital) Creative Labs Extigy to subsequently decode. But that playback scenario would have required that I alter more variables, such as the amplifier and speakers, than my controlled experiment permitted.

DAMAGE CONTROL

Intel is now shipping its i865 chip sets, which include Sonic Focus' DSP algorithms and Analog Devices codecs. Sonic Focus claims that these algorithms compensate for a significant amount of the signal distortion that two-channel lossy-compression schemes such as MP3 cause. Jim Barber, vice president of R&D for Sonic Focus, points out that the company's highly configurable technology supports multichannel operation and that you can parameterize it on the fly to operate with any popular compression scheme and transport bit rate. He also says that it handles samples as large as 32 bits using either 32- or 64-bit processing cores and that it currently works with sampling rates as high as 48 kHz on AC-97/Win32-driver platforms. Although Barber remains mum on future-product specifics, he's not immune to dropping hints. For example, he admits that the company's engineers have to date probably played the DTS track of *Seven Bridges Road* on The Eagles' *Hell Freezes Over* DVD approximately 5000 times!