SHARC® 2148x and 2147x Series Processors: Making High-Definition Audio Easier and More Affordable

High-definition audio is becoming pervasive, thanks to increases in digital signal processing power, the large data capacity of the Blu-ray Disc format, and the ever-increasing speed of Internet connections. Many consumer and professional products now offer high-definition audio capability, and market pressures will compel most consumer electronics OEMs to incorporate high-definition audio in upcoming products. This paper will define high-definition audio, describe applications in which it is used, and explain how Analog Devices' new SHARC 2148x and SHARC 2147x floating point digital signal processors (DSPs) can make high-definition audio easier and more affordable to implement.

What Is High-Definition Audio?

No industry standard exists to define "high-definition audio." In fact, there is even some issue with the term itself, with audiophiles preferring the term "high-resolution" audio. Regardless, a generally accepted meaning of "HD audio" is: audio that exceeds the resolution and/or number of channels available on commercial audio CDs.

The term can apply to any audio that exceeds CD's 16-bit word depth and 44.1-kilohertz sampling rate. Examples include:
1) The 20-bit/48-kHz (20/48) audio capability found in many professional audio recording devices
2) The 24-bit/96-kHz (24/96) audio used on many Blu-ray Discs and in the moribund DVD-Audio format
3) The 24-bit/192-kHz (24/192) audio used as the original recording format for many audiophile recordings (and occasionally released on DVD-Audio).

The term can also apply to audio with a greater number of channels than the two found on CDs, although this application is open to debate. For example, Dolby® Digital offers 5.1 channels, but most audio experts would argue that it cannot be considered high-definition because of the data reduction (or "lossy" compression) it employs. Few, though, would dispute that many greater-than-two-channel audio technologies can be considered high-definition. These include:
1) The lossless DTS™-HD Master Audio or Dolby TrueHD soundtracks found on most Blu-ray Discs
2) The uncompressed multichannel PCM tracks found on some Blu-ray Discs
3) The lossless 5.1 audio found on DVD-Audio and SACD discs
4) The uncompressed multichannel digital audio outputs of digital mixing consoles.

For the purposes of this paper, we will define high-definition audio as any audio that exceeds the 16-bit/44.1-kHz CD standard, or any audio that exceeds two channels without using lossy compression (we'll discuss a couple of possible exceptions to this rule though.) We will focus on devices that actually deliver a high-definition experience to the listener—after all, a flat-panel TV set that receives high-definition audio from a Blu-ray player and downconverts it to play through two 15mm x 50mm speakers cannot be considered a high-definition audio device.

Advantages of High-Definition Audio

For the consumer, the advantages of high-definition audio are twofold: the greater fidelity of high bit rates and sampling rates, and the greater realism possible through surround sound. For professionals, high-definition audio allows greater leeway in setting recording levels, more creative potential, and maximum flexibility in meeting the needs of various distribution formats.

Using a higher sampling rate of 88.2 or 96 kHz instead of CD's 44.1 kHz extends high-frequency response. A higher sampling rate also allows the anti-aliasing filter in an analog-to-digital converter to be raised to a frequency far above human hearing. Some audio engineers contend that the phase distortion caused by the 20-kHz anti-aliasing filter used in 44.1-kHz digital audio is audible at lower frequencies. Countless anecdotal accounts from audiophiles attest to the advantages of sampling rates greater than 44.1 kHz.

Word depths of 20 or 24 bits extend dynamic range past the maximum 96 dB of CD. Although a 96 dB dynamic range is theoretically adequate for consumer applications, it is seldom achieved in practice. When a recording engineer has only 16 bits of resolution to work with, input level must be set conservatively to avoid overloading the recording device. This practice often results in an effective dynamic range of perhaps 12 bits (72 dB). Alternatively, the recording engineer can use dynamic range compression to avoid overloads, but this practice reduces fidelity. A 24-bit system allows the recording engineer to get dynamic range of at least 20 bits (120 dB) with no risk of overloading the analog-to-digital converter, and without resorting to dynamic range compression.

These days, a consumer may experience an audio recording in 24/96 lossless audio through a high-quality 7.1-channel audio system—or in stereo sound compressed to 30-kilobits-per-second for Internet streaming, played through 20mm speakers in a laptop computer. And no one knows what distribution formats may be popular in the future. By producing audio in the maximum possible bit rate and highest possible sampling rate, audio professionals can be confident that their current productions will meet the needs of future distribution formats.

Using more than two channels of sound expands creative possibilities for recording professionals, and increases the realism of audio for consumers. With a two-channel format, a listener must sit equidistant from the two speakers in order to get a centered sonic image. A 5.1 or 7.1 format can deliver a consistent center image for a roomful of listeners, no matter where they're sitting. Likewise, two-channel sound generally requires an equidistant listener to produce an enveloping sound effect, but 5.1 and 7.1 systems can deliver a consistently enveloping sound effect over a wide variety of listening positions.

Applications for High-Definition Audio

In the consumer market, high-definition audio's most prominent application is in Blu-ray Disc players, which can at least pass DTS-HD Master Audio and Dolby TrueHD lossless audio and uncompressed multichannel PCM through their HDMI outputs. Many of these players incorporate complete decoding of these formats for analog output. Most home-theater-in-a-box (HTiB) systems that include an integrated Blu-ray player/receiver also decode these formats.

High-definition audio may be found in devices that play audio streamed or downloaded from the Internet. These may be computers, video streaming devices such as the Vudu™ and Roku™ set-top boxes, or newer Blu-ray players with Internet streaming built in. Many of these devices are capable of reproducing better-than-CD audio, and most are capable of passing 5.1 audio through HDMI or SPDIF.

A few Internet download sites now offer high-definition audio. The most prominent is HDTracks.com™, which at the time this paper was written offered downloads of more than 130 albums in 24/96 audio, compressed using the Free Lossless Audio Codec (FLAC) algorithm.

High-definition audio is often decoded by an audio/video receiver or surround-sound processor. At the time of this writing, almost all receivers and surround processors priced above $500 offer internal decoding of DTS-HD Master Audio and Dolby TrueHD. Most of these products offer full 7.1-channel output, and some higher-priced models offer technologies that expand a 5.1 or 7.1 signal to as many as 11.1 channels.

In the professional world, digital audio products that record or process 24/96 audio are now the norm. These may include digital mixing consoles, sound effects processors (reverb, compression, etc.), and digital equalization and crossover processors for PA systems. Multichannel capability is also common; most digital mixing boards now offer at least eight output channels.

Important Technologies for High-Definition Audio

High-definition audio signals are typically compressed for storage and transmission because they are so data-intensive. A 24/96 eight-channel signal may consume as much as 12 times the data required for a 16/44.1 two-channel signal.

Two technologies, DTS-HD Master Audio and Dolby TrueHD, are currently used to compress multichannel high-definition audio for distribution on Blu-ray Discs. Both of these are lossless codecs, meaning they deliver bit-for-bit reproduction of an original master recording. DTS-HD Master Audio is capable of 24/192 resolution in two-channel mode, and 24/96 resolution in up to eight (7.1) channels. Dolby TrueHD is capable of as many as 14 full-range channels in up to 24/192 resolution.

Although we stated that only lossless codecs should be considered high-definition audio, there are a couple of lossy compression technologies that might make the grade. The most advanced of these are Dolby Digital Plus and DTS-HD High Resolution Audio. Dolby Digital Plus offers resolution up to 24/48 and as many as 13.1 channels, while DTS-HD High Resolution Audio offers resolution up to 24/96 and as many as eight (7.1) channels. However, these formats are seldom used at present. It's worth noting that the basic DTS codec can deliver up to 6.1 channels in 24/96, and Dolby Digital format can produce 5.1 channels in 20/48.

Some codecs used to distribute audio over the Internet also offer high-resolution capability. The FLAC codec previously mentioned achieves resolution up to 32/655; while it is commonly used to distribute two-channel content, it can also be used for 5.1. The Windows® Media Audio Lossless and Apple® Lossless codecs also have the technical capability to support 5.1 and 24/96, although most applications and devices using these codecs support only 16/44.1 two-channel audio.

It's important to note that various post-processing technologies are often used along with high-resolution audio technologies such as DTS-HD Master Audio and Dolby TrueHD. These technologies include volume management algorithms such as Audyssey Dynamic Volume, Dolby Volume, SRS® TruVolume, and THX® Loudness Plus™. Room equalization technologies that automatically compensate for flaws in room acoustics and speaker positioning and performance are also popular; these include Audyssey MultEQ, Dirac® Live, and Trinnov® Optimizer. At least two post-processing technologies expand the number of channels beyond 7.1—Dolby Pro Logic IIz adds two height speakers to create 9.1 sound, and Audyssey DSX adds height and/or width speakers for a maximum total of 11.1 channels.

New Analog Devices DSPs for High-Definition Audio

Analog Devices has introduced two new series of SHARC DSPs designed to make incorporation of high-definition audio easier and less expensive for OEMs. One of these DSPs can be used where two or three were previously needed. The new DSPs are available in two series. The SHARC 2148x series offers performance up to 400 MHz. The SHARC 2147x series offers performance up to 266 MHz and is designed for applications requiring low power consumption, such as portable products.

Because high-definition audio is so processor-intensive, Analog Devices included several features that free the core processor from having to perform simple tasks that can be better handled by separate, dedicated components within the DSP. In the past, the task of decoding high-resolution audio formats used up most or all of the processing power of a single DSP. Any post-processing such as room equalization, volume management, or creating extra channels of sound would have to be performed in additional DSPs. With the SHARC 2148x and 2147x series processors, a single chip can handle both the high-definition audio decoding and practically all of the post-processing options available today.

A key feature of the SHARC 2148x and 2147x series processors is the built-in accelerator for Fast Fourier Transform (FFT), Finite Impulse Response (FIR) filters, and Infinite Impulse Response (IIR) filters. Room EQ technologies such as Audyssey MultEQ use many long FIR filters, which consume a great deal of processing power. The simplicity of these filters makes it possible to offload this task to a separate piece of silicon within the DSP. The FIR/IIR/FFT accelerators can perform most of the processing required for room EQ, speaker crossovers, and tonal adjustments, so the core processor can concentrate on more complex tasks such as high-definition audio decoding. In terms of multiply accumulates (MACs) per second, the FIR/IIR/FFT accelerators roughly equal the speed of the core processor—thus the combination of the core and the accelerator double the overall performance of the system.

The onboard sample rate converter found in the SHARC 2148x and 2147x series processors offers additional independent processing capability. This sample rate converter (a duplicate of the AD1896 standalone sample rate converter) can be used to convert low sample rates such as 44.1 kHz up to higher rates like 96 kHz, or vice-versa. It can also be used for jitter reduction, which creates a more precise digital audio signal by removing the clock from the incoming signal and replacing it with an internally generated high-precision clock. The sample rate converter comprises four independent two-channel sections, which can be combined to deliver as many as eight channels—all precisely timed with zero interchannel phase error. A built-in SPDIF interface makes it easy for external devices to use these sample rate converters. No matter how you use the sample rate conversion capability, the task runs independently, without accessing the processing power of the core.

More Memory, Greater Efficiency

SHARC 2148x and 2147x series processors are available with 3 or 5 megabits of on-chip RAM, which is up to two and a half times the RAM found on competitors' DSPs, and substantially more than the 1 Mb on the original SHARCs. The onboard memory saves money because it allows many memory-intensive functions, such as room EQ and reverb, to be performed without the need for external RAM. The extra memory is a big advantage for small consumer audio equipment manufacturers and professional audio equipment manufacturers because it cuts programming time. Coding is simpler because the programmer doesn't have to be so conscious of memory limitations.

The SHARCs' Direct Memory Access (DMA) feature further lessens the load on the core processor by managing the DSP's internal memory. Outside devices can access the internal memory straight through the DMA, without having to go through the DSP core. The DMA allows the core processor to receive data in blocks rather than in single samples, thus dramatically reducing the number of interrupts and increasing speed in the process.

SHARC 2148x and 2147x series processors natively support 32-bit floating-point arithmetic. This design simplifies algorithm development, enabling engineers to focus on the audio aspects of the design rather than being distracted by arcane numerical issues that must be considered when using fixed-point integer arithmetic. The standard 32-bit floating-point format sets aside 23 bits for the mantissa, eight for the exponent, and one for the sign bit. This format is sufficient for storing 24-bit high-precision audio samples, but even more precision is required when performing arithmetic operations. To this end, the internal registers of these DSPs are actually 40 bits wide, allowing an additional 8 bits of precision in the mantissa without impacting the speed of the DSP. Once data is loaded into the registers, 40-bit arithmetic can be performed at the same speed as 32-bit arithmetic while preserving the 24-bit fidelity of the format.

The new processors' Variable Instruction Set Architecture (VISA) makes more efficient use of memory than past SHARCs, saving as much as 30 percent on code space. The original SHARCs required 48-bit words for all data and instructions, so even the simplest instructions used up 48 bits of memory. With VISA, a given instruction can be encoded in 16, 32, or 48 bits, thus reducing the number of unused bits.

The Quad Flat Package (QFP) versions of the SHARC 2148x and 2147x series processors makes them easier to use in low-cost manufacturing operations. Relatively common and inexpensive surface-mount equipment can be used to install these DSPs, so there’s no need for OEMs to invest in new manufacturing equipment and processes.

Conclusion
With high-definition audio becoming the norm in both consumer and professional applications, and with the increasing popularity of processing-intensive technologies such as Audyssey DSX and Dolby Volume, the demands on audio DSPs have grown tremendously. Yet market pressures force OEMs to deliver these new features at minimal additional cost. Analog Devices’ SHARC 2148x and 2147x series processors make implementing these advanced technologies easy and cost-effective for practically any OEM, and for practically any audio product.

Brent Butterworth is a writer who specializes in audio and video technology. He is former editor-in-chief of Home Entertainment *and* Home Theater *magazines, and former director of consumer marketing for Dolby Laboratories.*

ADI SHARC® Floating-Point Digital Signal Processers
Analog Devices' 32-bit floating-point SHARC digital signal processors are based on a Super Harvard architecture that balances exceptional core and memory performance with outstanding I/O throughput capabilities. This Super Harvard architecture extends the original concepts of separate program and data memory busses by adding an I/O processor with its associated dedicated busses. In addition to satisfying the demands of the most computationally intensive, real-time signal-processing applications, SHARC processors integrate large memory arrays and application-specific peripherals designed to simplify product development and reduce time to market.

Comprehensive Development and Support Ecosystem
Analog Devices software and hardware development tools are designed to provide easier and more robust methods for engineers to develop and optimize systems, simplifying product development processes and reducing time to market. The SHARC processor family leverages familiar development tools including the VisualDSP++® integrated development and debug environment (IDDE) and the EZ-Kit Lite® evaluation and application prototyping platform. A rich third-party software support network further enables developers to design more intelligent and efficient solutions for a wide range of applications.

For more information about ADI’s full portfolio of signal processing components, software, development tools and support, visit www.analog.com