## DAC ICs: How Many Bits is Enough?

### by Robert Adams

**The DACs designed for use in digital audio playback circuitry are available in a diverse choice of architectures, and each have their own performance tradeoffs.**

At the heart of every digital audio playback system lies the single-most critical component for high-fidelity audio: the digital-to-analog converter (DAC). These converters handle the delicate task of translating the 16-bit binary words encoded on the disc or tape into corresponding analog signals worthy of amplification and, ultimately, of the human ear.

Magnavox's first CD player (circa 1983) employed dual 14-bit converters, and by 1989 many models heralded 20-bit converters. Today however, the words "bitstream" and "MASH" — displayed prominently on the face of many CD players and other digital reproduction devices — refer to a 1-bit converter architecture that is revolutionizing the digital audio industry and attracting the attention of many.

Just how many bits are necessary? Given that CDs are encoded in a 16-bit PCM format, why are 18- and 20-bit converters showing up?

This article examines the DACs designed and optimized for use in digital audio playback circuitry. It compares the diverse choice of architectures available for the job and discusses the design and performance tradeoffs of each.

### Basic theoretical concepts

Although you don't have to be an expert to understand audio D/A conversion, there are two basic concepts that are fundamental to the understanding of this topic: *sampling* and *quantization*.

Sampling is the act of taking an analog signal waveform, which has a value at every instant of time, and sampling it at regular intervals. This supplies us with information about the value of the waveform only at the sampling instants. The Nyquist Theorem states that, as long as the original analog signal contains no frequencies higher than half of the sampling frequen-

Adams is with the Semiconductor Division of Analog Devices, Wilmington, MA.

cy, the signal can be re-constructed completely from the sampled values.

For most digital audio applications, we have two common sampling frequencies: 44.1kHz (the consumer standard) and 48kHz (the professional standard). Both of these frequencies are a little above twice the highest frequency of interest (20kHz).

If we imagine that we have a "perfect" sampler, free of real-world non-idealities, we can say that the signal can be recovered completely, without any added noise or distortion. We assume that the output of this perfect sampler is either a noiseless analog voltage (held on a capacitor, for example) or a digital word with infinite precision. To store the sampled values for later playback, one must use a finite number of bits to represent the sampled value. This is where *quantization* comes into play.

By using a finite digital word-length to represent the sampled value, we have added an uncertainty to the sampled value, one which can no longer be completely recovered on playback. Quantization can be thought of as a *many-to-one mapping* (a continuous range of input values results in a single output value). For consumer digital audio equipment, the sampled values are stored in a 16-bit format, with a theoretical dynamic range of 98.1dB. (It's more complicated than just 6dB/bit×16 bits=96dB, and it is commonly misunderstood.)

Errors introduced by quantization are heavily dependent on the signal being quantized. High-amplitude signals with a complex spectrum (pop music, for example) generally cause quantization errors that are like white noise. Very soft signals that are not much larger than one quantization level can be severely distorted unless dither noise (a random noise uncorrelated with the signal) is added to the signal before quantization. The importance of using dither to de-correlate the quantization error from the input signal

has received wide attention, and most manufacturers now use it where appropriate in their designs.

### A typical DAC circuit

Over the years, many modifications in D/A circuitry have improved performance and reduced cost. For example, early CD players generally used a single, parallel-input DAC, alternated between both right and left channels. They employed two output sample-and-hold circuits — one per channel — to reduce the signal-dependent "glitch" that they produced when switching from one sample to the next, and to divide the output samples into separate left and right signals. The sample-and-hold amplifier's distortion and the inter-channel time delay drew criticism from audiophiles, who were convinced that such degradations were audible. Modern designs usually employ one DAC for each channel, and special design techniques are used in the DAC itself to reduce the glitch to the point where de-glitching circuitry is no longer necessary.

Figure 1 shows a typical D/A circuit, which might be found inside a modern CD player. The most important IC is, of course, the D/A chip itself. In this case, our AD1860, an 18-bit DAC, is being used. Like most modern DACs, this device uses a serial interface to reduce pin count. This serial interface consists of three signals: serial data input, bit clock input and word clock input. This serial bit stream is converted internally to a parallel word, where it drives the individual current-steering switches. One advantage of using an internal latch is that the timing of when each bit actually switches can be carefully controlled, which reduces the signal-dependent glitch problem found in older designs.

The other key function is the digital oversampling filter chip, which takes in data at 44.1kHz (or 48kHz) and produces an output at a higher rate (two, four, eight
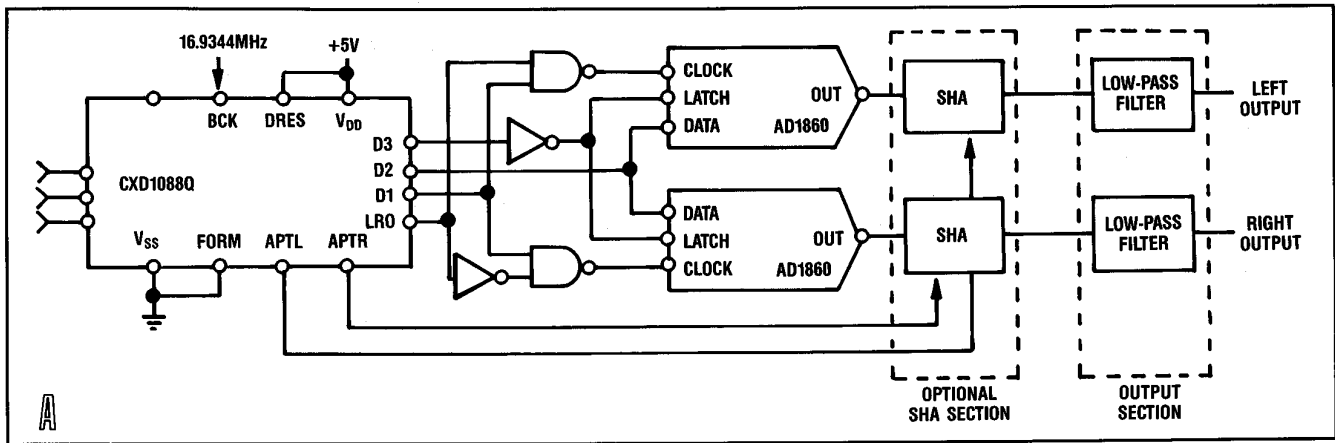
Figure 1. This typical D/A circuit uses an 18-bit DAC, which uses a serial interface to reduce pin count.

or even 16 times the input sample rate). Next we will look into the theory behind DAC filters, and why analog brick-wall filters have almost universally given way to digital interpolation filters.

### D/A reconstruction filters: Is digital better?

On the playback side, sampling theory predicts that there will be "images" of the original signal around multiples of the sampling frequency. For example, a 1kHz signal that was originally sampled at 44.1kHz will have components at 43.1kHz, 45.1kHz, 87.2kHz, 89.2kHz, etc., as shown in Figure 2.

If we want perfect reconstruction of the original continuous waveform, we must filter out or otherwise eliminate these unwanted look-alike or image components
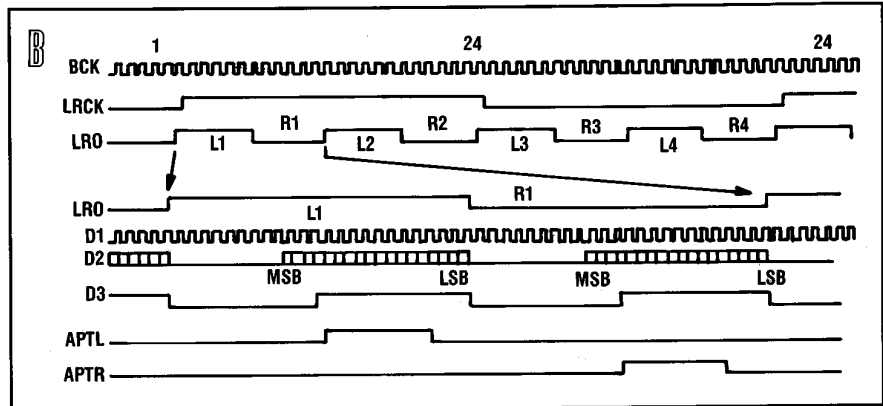


before they are fed to the rest of our audio system. However, this filtering is, strictly speaking, not necessary. All of the unwanted components are in the ultrasonic range, and the only reason to remove them is to ensure that the analog circuitry in a typical power amp or pre-amp will not become non-linear in the presence of large ultra-high-frequency signals.

Figure 3 shows two methods for removing these unwanted image signals. The first method is to follow the DAC with a *brick-wall* analog filter, which has flat response over audio frequencies and sharply cuts off just above 20kHz. The second method is to do most of this filtering digitally and use a simple analog filter at the DAC output.
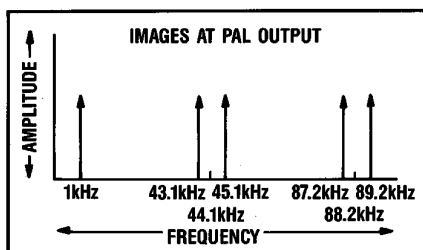


Figure 2. Sampling theory predicts that there will be "images" of the original signal around multiples of the sampling frequency. Perfect reconstruction of the original continuous waveform requires that these unwanted look-alike or image components be eliminated.
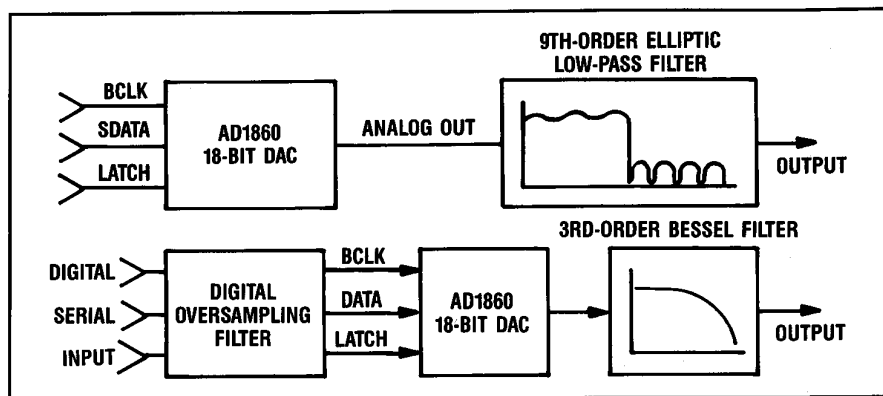


Figure 3. There are two methods for removing unwanted image signals: Follow the DAC with a brick-wall analog filter, which has flat response over audio frequencies and sharply cuts off just above 20kHz, or do most of this filtering digitally and use a simple analog filter at the DAC output.

## The analog reconstruction filter: Fast disappearing

In the early days of digital audio, the reconstruction filter was typically a complex, analog, active (or passive) filter. This filter had to have very flat frequency response over the audio band and then drop precipitously just above 20kHz, with an ultimate attenuation of at least 80dB.

These filters became known as *brick-wall* filters, and they required exceptionally high-tolerance components to meet audio standards. These filters also introduced large phase shifts near the edge of the audio band. Whether this phase shift introduced audible artifacts became a controversial subject, and early detractors of digital audio often blamed the so-called "gritty" sound of digital recordings on the high phase shift introduced by this filter, when in fact they were probably hearing the effects of poor low-level linearity in the ADC or DAC (more on this later).

In any event, the combination of high cost, high power consumption and large phase shifts drove the development of digital oversampling filters that have effectively replaced brick-wall analog filters in most applications.

## To the rescue ...

Digital oversampling filters take the burden off the analog filters by filtering out most of the undesired information (images) in the digital domain before the signal is applied to the DAC. There must still be an analog output filter, but it can be a simple third-order design with minimal phase shift.

Figure 4 shows how oversampling filters work in both the time and frequency domains. We start with a sampled signal at 44.1kHz (Figure 4a), which has images in the frequency domain (Figure 4b). The next step in the process is to increase the sample rate of the digital signal by inserting zero-valued samples between the existing samples (Figure 4c), resulting in the spectrum shown in Figure 4d.

This may seem like a trivial operation, because we have not really changed the spectrum of the signal. But it is, in fact, an important conceptual step; the signal in Figure 4a is sampled at 44.1kHz and has images around multiples of 44.1kHz, but the signal in Figure 4c is sampled at four times 44.1kHz and has images around four times 44.1kHz.

Even though the spectrum looks the same, we now regard the 44.1kHz images as the actual signal itself, which just happens to have repeated regions at multiples of 44.1kHz. This conceptual step is necessary because any digital filter that we build will have its own image response at multiples of the input sample rate. Because we want our filter response to attenuate frequencies beyond 44.1kHz, it has to run at
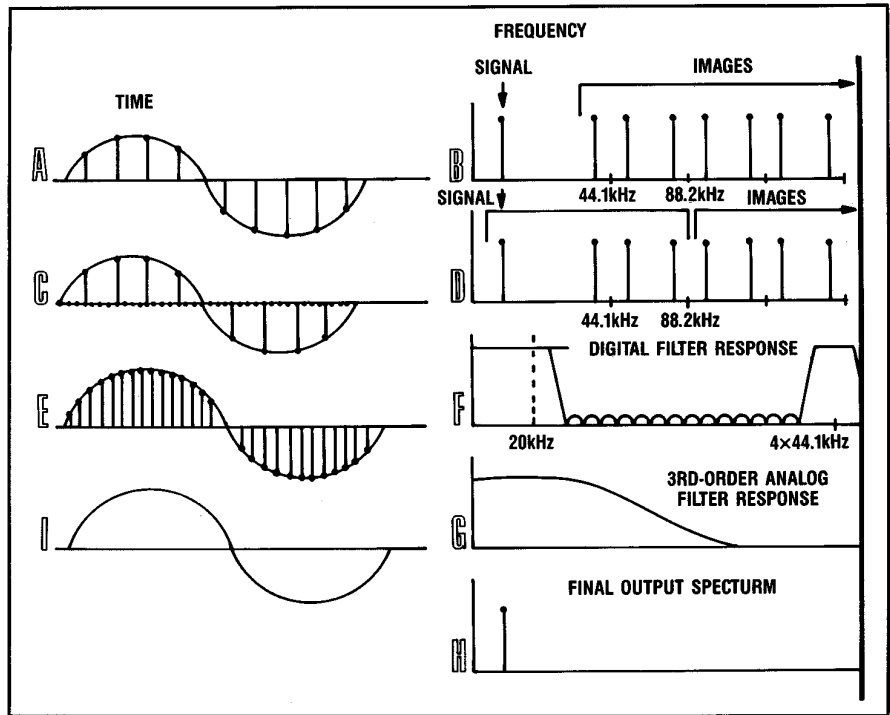


*Figure 4.* For digital oversampling, we start with a sampled signal at 44.1kHz (Figure 4a), which has images in the frequency domain (Figure 4b). The sample rate of the digital signal is increased by zero-valued samples inserted between the existing samples (Figure 4c), resulting in the spectrum shown in Figure 4d. This zero-stuffed signal is applied to a digital low-pass filter. The filter response for a 4× interpolator is shown in Figure 4f, and the smoothed time-domain signal is shown in Figure 4e. A simple analog filter finishes the job by removing the high-frequency image, as shown in Figures 4g and 4h, resulting in the reconstructed waveform shown in Figure 4i.

a faster sampling rate than 44.1kHz.

The third step in the process is to apply this zero-stuffed signal to a digital low-pass filter. The filter response for a 4× interpolator is shown in Figure 4f, and the smoothed time-domain signal is shown in Figure 4e. Note that the filter response has an image around the 4× oversampling rate. Therefore, an image of the baseband signal appears at the DAC output around this frequency and must be removed by the external analog filter. But the digital filter has done most of the hard work, and a simple analog filter can now finish the job by removing this high-frequency image, as shown in Figures 4g and 4h, resulting in the reconstructed waveform shown in Figure 4i.

One advantage of using a finite impulse response (FIR) digital filter is that the phase response can be made linear. (See "The Performance Aspects of Digital Oversampling" and "The Saori Signal Processor" in the October 1989 and September 1990 issues, respectively for more on FIR filters.) This solves the problem (if it is a problem) of too much phase shift in the audio band, although we should bear in mind that the digital filter is every bit as much of a "brick-wall" filter as the old analog filter was. It does not provide relief for those who can hear beyond 20kHz.

Another more tangible advantage of the

digital approach is that we have now replaced a fussy analog circuit requiring high-precision components and lots of circuitry with a monolithic IC that is inexpensive and whose performance does not change significantly from one unit to the next. It is this combination of economic incentive and improved performance that has lead virtually all manufacturers to incorporate digital oversampling filters into their designs.

Today it is common to find digital audio devices that use 18- or even 20-bit DACs. The number of bits used in the DAC is often prominently displayed on the front panel, with more expensive units using more bits. The technically savvy person will often ask the obvious question: If there are only 16 bits in the original signal, why would you need an 18- or 20-bit converter?

There are several answers to this question, some of which are technical in nature and some that have more to do with marketing issues. (How does a manufacturer differentiate a low-end CD player from a high-end unit if they both sound the same?) The technical answer is that the digital oversampling filter produces a digital output with much more than 16 bits. This does *not* mean that the digital filter has somehow magically removed noise from the signal or increased the dy-

namic range of the recording. It simply means that the interpolated values which the filter computes do not usually coincide with 16-bit integer values. They usually fall in between 16-bit values, and, therefore, a longer digital word is required to represent the interpolated value. So now we face the question of what to do with this long digital word (often more than 30 bits).

If we simply drop the lower bits and feed the upper 16 bits into a 16-bit DAC, we have effectively re-quantized the signal, and some quantization noise will be added to the signal that was not present in the original digital signal. If we postulate that the original recording was made with a perfect 16-bit ADC with a theoretical 16-bit dynamic range of 98.1dB, then we have added another noise source of the same magnitude to the signal.

Because this re-quantization is occurring at the interpolated rate and not at the original 44.1kHz rate, its noise is spread out over a wider frequency range, and less of it appears at audio frequencies. For an 8:1 interpolation ratio, the total noise will be about 1.5dB higher than theoretical 16-bit performance, resulting in a noise floor of 96.6dB. This assumes that the truncation noise is not correlated with the signal, which may be an optimistic assumption for low-level signals. A more pessimistic value for the noise degradation might be 3dB, resulting in a S/N ratio of 95.1dB.

If instead we save the upper 18 bits of the digital filter output signal and feed it to an 18-bit DAC, we are adding a noise source that is 12dB below the original 16-bit theoretic noise produced by the ADC. When we add these two noise sources, the result is only a fraction of a decibel higher than the theoretic 16-bit noise floor, resulting in almost no degradation of the signal. This is the primary technical reason that 18- and 20-bit converters are often used in digital audio playback systems.

But one should bear in mind that this slim 1.5dB (or 3dB worst case) noise advantage is mostly lost if the quantization noise of the converters is not the dominant noise source in the signal. In most consumer applications, noise on the original recording and in the playback environment will exceed the quantization noise by a fair margin, and it is highly questionable whether this improvement could realistically be audible.

Another reason for using converters of higher resolution is that the distortion performance (linearity) of higher-resolution converters is typically better than that of lower-resolution products. For example, a 16-bit DAC, such as our AD1856, typically exhibits full-scale distortion of 0.002%; a 20-bit part (such as the AD1862J) has distortion numbers of less than 0.0012%. Yet another advantage of high-resolution converters is that the thermal noise performance (noise with no digital codes changing) is typically better than the noise performance of their 16-bit counterparts.

The requirements for the professional audio market are quite different from those of the consumer market. When you are recording live, uncompressed material, there is a legitimate need for very-high-resolution ADCs and DACs. The dynamic range of a closely miked percussion instrument using a low-noise condenser microphone can exceed 120dB without much trouble!

The professional recording industry is in the midst of a major push toward 18- and even 20-bit converters to meet these needs. At present, few product offerings provide this level of performance, and most of these products are expensive hybrids. But in the near future, the same performance will be available in inexpensive monolithic ICs that use a radically different design approach known as sigma-delta conversion.

In Part II, we'll cover design approaches for D/A conversion and discuss the specifications in product literature.