

MUST A "16-BIT" CONVERTER BE 16-BIT MONOTONIC AND SETTLE TO 16 PPM?

by Dave Robertson and Steve Ruscak

Q. *I recently saw a data sheet for a low-cost 16-bit, 30 MSPS D/A converter. On examination, its differential nonlinearity (DNL) was only at the 14-bit level, and it took 35 ns (1/28.6 MHz) to settle to 0.025% (12 bits) of a full scale step. Isn't this at best a 14-bit, 28 MHz converter? And if the converter is only 14-bit monotonic, the last two bits don't seem very effective; why bother to keep them? Can I be sure they're even connected?*

A. That's a lot of questions. Let's take them one at a time, starting with the last one. You can verify that the 15th and 16th bits are connected by exercising them and observing that 0..00, 0..01, 0..10, and 0..11 give a very nice 4-level output staircase, with each step of the order of 1/65,536 of full scale. You can see that they would be especially useful in following a waveform that spent some of its time swinging between 0..00 and 0..11, or providing important detail to one swinging through a somewhat wider range. This is the crux of the *resolution* spec, the ability of the DAC to output 2^{16} individual voltage levels in response to the 65,536 codes possible with a 16-bit digital word.

Systems that must handle both strong and weak signals require large dynamic range. A notable example of this is the DACs used in early CD player designs. These converters offered 16-20 bits of dynamic range but only about 14 bits of differential linearity. The somewhat inaccurate representation of the digital input was far less important than the fact that the dynamic range was much wider than that of LP records and allowed both loud and soft sounds to be reproduced with barely audible noise—and that the converters' low cost made CD players affordable.

The resolution is what makes a 16-bit DAC a "16-bit DAC". Resolution is closely associated with *dynamic range*, the ratio of the largest signal to the smallest that can be resolved. So dynamic range also depends on the noise level; the irreducible "noise" level in ideal ADCs or DACs is *quantization noise*.

Q. *What is quantization noise?*

A. The sawtooth-wave-shaped quantization noise of an ideal n -bit converter is the difference between a linearly increasing analog value and the stepwise-increasing digital value. It has an rms value of $1/(2^{n+1}\sqrt{3})$ of span, or $-(6.02n + 10.79)$ dB (below p-p full scale). For a sine wave, with peak-to-peak amplitude equal to the converter's span, rms is $\sqrt{2}/4$, or -9.03 dB, of span, so the full-scale signal-to-noise ratio of an ideal n -bit converter, expressed in dB, becomes the classical

$$6.02n + 1.76 \text{ dB.} \quad (1)$$

As the analog signal varies through a number of quantization levels, the associated quantization noise resembles superimposed "white" noise. In a real converter, the circuit noise produced by the devices that constitute it adds to quantization noise in root-sum-of-squares fashion, to set a limit on the amplitude of the minimum detectable signal.

Q. *But I still worry about that differential nonlinearity spec. Doesn't 14-bit differential nonlinearity mean that the converter may be non-monotonic at the 16-bit level, i.e., that those last two bits have little influence on overall accuracy?*

*The AD768 is an example of such a DAC. Circle 22

A. That's true, but whether to worry about it depends on the application. If you have an instrumentation application that really requires 16-bit resolution, 1/2-LSB accuracy for all codes, and 1-LSB full-scale settling in 31.25 ns (we'll get to that discussion shortly), this isn't the right converter. But perhaps you really need 16-bit dynamic range to handle fine structure over small ranges, as in the above example, while high overall accuracy is not needed—and is actually a burden if cost is critical.

What you need to consider in regard to DNL in signal-processing applications is 1) the noise power generated by the DNL errors and 2) the types of signals that the D/A will be generating. Let's consider how these might affect performance.

In many cases, DNL errors occur only at specific places along the converter's transfer function. These errors appear as spurious components in the converter's output spectrum and degrade the signal-to-noise ratio. If the power in these spurs makes it impossible to distinguish the desired signal, the DNL errors are too large. Another way to think about it is as a ratio of the quantity of good codes to bad codes (those having large DNL errors). This is where the type of signal is important.

The various applications may concentrate in differing portions of the converter's transfer function. For example, assume that the D/A converter must be able to produce very large signals and very small signals. When the signals are large, there is a high proportion of DNL errors. But, in many applications, the signal-to-noise ratio will be acceptable because the signal is large.

Now consider the case where the signal is very small. The proportion of DNL errors that occur in the region of the transfer function exercised by the signal may be quite small. In fact, in this particular region, the spurs produced by the DNL errors could be at a level comparable to the converter's quantization noise. When the quantization noise becomes the limiting factor in determining signal-to-noise ratio, 16 bits of resolution will really make a difference (12 dB!) when compared to 14 bits.

Q. *OK, I understand. That's why there's such a variety of converters out there, and why I have to be careful to interpret the specs in terms of my application. In fact, maybe data sheets that have a great number of "typical" plots of parameters that are hard to spec are providing really useful information. Now, how about the settling-time question?*

A. *Update rate* for a D/A converter refers to the rate at which the digital input circuitry can accept new inputs, while *settling time* is the time the analog output requires to achieve a specified level of accuracy, usually with full-scale steps.

As with accuracy, time-domain performance requirements differ widely between applications. If full accuracy and full-scale steps are required between conversions, the settling requirements will be quite demanding (as in the case of offset correction with CCD image digitizers). On the other hand, waveform synthesis typically requires relatively small steps from sample to sample. The solid practical ground is that full-scale steps in consecutive samples mean operation at the Nyquist rate (half the sampling frequency), which makes it extremely difficult (how about "impossible"?) to design an effective anti-imaging filter.

Thus, DACs used for waveform reconstruction and many other applications* inevitably oversample. For such operation, full-scale settling is not required; and in general, smaller transitions require less time to settle to a given accuracy. Oversampled waveforms, taking advantage of this fact, achieve accuracy and speed greater than are implied by the full-scale specification. 