Evaluating High Speed DAC Performance

by Walt Kester

INTRODUCTION

Unlike an ADC which requires an FFT processor to evaluate spectral purity, a DAC produces an analog output which can be examined directly using a traditional analog spectrum analyzer. A challenge in DAC evaluation is generating the digital input that can range from a single-tone sinewave to a complex wideband CDMA signal. Direct digital synthesis techniques can be used to generate digital sinewaves, but more sophisticated and expensive word generators are needed to produce the more complex digitial signals.

The ac specifications which are the most important in evaluating high speed DACs are *settling time, glitch impulse area, distortion, spurious free dynamic range (SFDR), and signal-to-noise ratio (SNR).* Time domain specifications will be addressed first, followed by frequency domain specifications.

DAC SETTLING TIME

The precise settling time of a DAC may or may not be of interest depending upon the application. It is especially important in high speed DACs used in video displays because of the high pixel rates associated with high resolution monitors. The DAC must be capable of making the transition from all "0"s (black level) to all "1"s (white level) in 5% to 10% of a pixel interval, which can be quite short. For instance, even the relatively common 1024 × 768, 60-Hz refresh-rate monitor has a pixel interval of only approximately 16 ns. This implies a required settling time of less than 2 ns to at least 8-bit accuracy (for an 8-bit system).

The fundamental definitions of full-scale settling time are shown in Figure 1. The definition is quite similar to that of the settling time of an op amp. Notice that settling time can be defined in two acceptable ways. The more traditional definition is the amount of time required for the output to settle with the specified error band measured with respect to the 50% point of either the data strobe to the DAC (if it has a parallel register driving the DAC switches) or the time when the input data to the switches changes (if there is no internal register). Another equally valid definition is to define the settling time with respect to the time the output leaves the initial error band. This effectively removes the "dead time" from the measurement. In video DAC applications, for instance, settling time with respect to the output is the key specification—the fixed delay (dead time) is of little interest.

The error band is usually defined in terms of an LSB or % full-scale. It is customary, but not mandatory, to define the error band as 1 LSB. However, measuring full-scale settling time to 1 LSB at the 12-bit level (0.025% FS) is possible with care, but measuring it to 1 LSB at the 16-bit level (0.0015% FS) presents a real instrumentation challenge. For this reason, high-speed DACs such as the TxDAC® family specify 14- and 16-bit settling time to the 12-bit level, 0.025% FS (typically less than 11 ns).

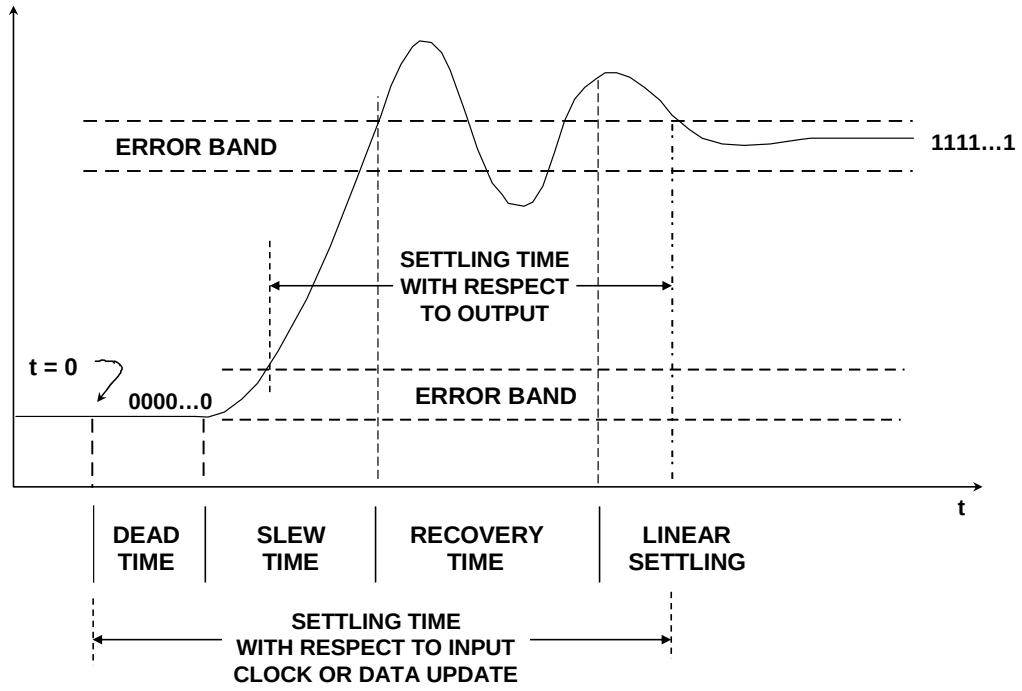

Figure 1: DAC Full-Scale Settling Time

Mid-scale settling time is also of interest, because in a binary-weighted DAC, the transition between the 0111…1 code and the 1000…0 code produces the largest transient. In fact, if there is significant bit skew, the transient amplitude can approach full-scale. Figure 2 shows a waveform along with the two acceptable definitions of mid-scale settling time. As in the case of full-scale settling time, mid-scale settling time can either be referred to the output or to the latch strobe (or the bit transitions if there is no internal latch).

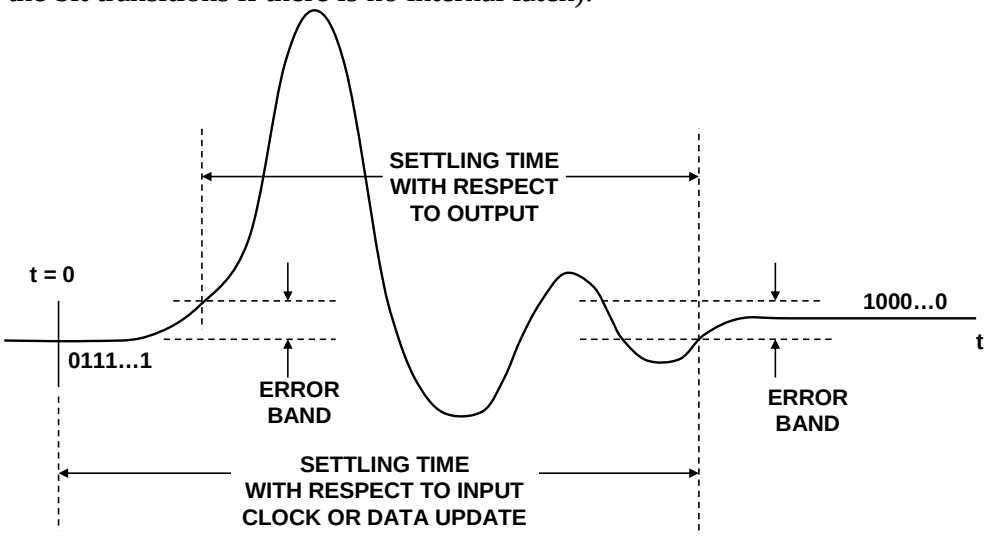

Figure 2: DAC Mid-Scale Settling Time

GLITCH IMPULSE AREA

Ideally, when a DAC output changes it should move from one value to its new one monotonically. In practice, the output is likely to overshoot, undershoot, or both. This uncontrolled movement of the DAC output during a transition is known as a *glitch*. It can arise from two mechanisms: capacitive coupling of digital transitions to the analog output, and the effects of some switches in the DAC operating more quickly than others and producing temporary spurious outputs.

Capacitive coupling frequently produces roughly equal positive and negative spikes (sometimes called a *doublet* glitch) which more or less cancel in the longer term. The glitch produced by switch timing differences is generally unipolar, much larger, and of greater concern.

Glitches can be characterized by measuring the *glitch impulse area*, sometimes inaccurately called *glitch energy*. The term *glitch energy* is a misnomer, since the unit for glitch impulse area is volt-seconds (or more probably µV-sec or pV-sec. The *peak glitch area* is the area of the largest of the positive or negative glitch areas.

Glitch impulse area is easily estimated from the mid-scale settling time waveform as shown in Figure 3. The areas of the four triangles are used to calculate the net glitch area. Recall that the area of a triangle is one-half the base times the height. If the total positive area equals the total negative glitch area, then the net area is zero. The specification given on most data sheets is the net glitch area, although in some cases, the peak area may specified instead.

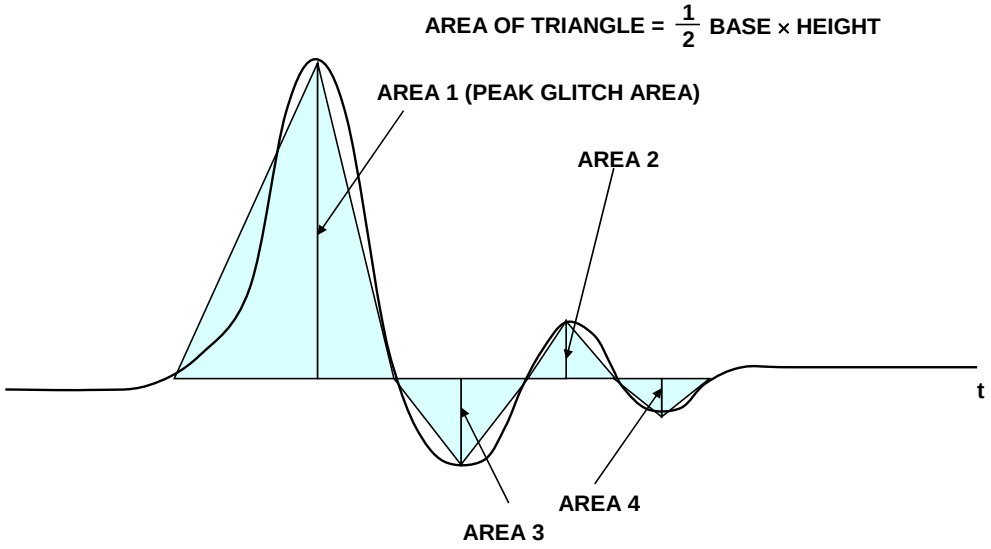

Figure 3: Glitch Impulse Area

OSCILLOSCOPE MEASUREMENT OF SETTLING TIME AND GLITCH IMPULSE AREA

A wideband fast-settling oscilloscope is crucial to accurate settling time measurements. There are several considerations in selecting the proper scope. The required bandwidth can be calculated based on the rise/falltime of the DAC output, for instance, a 1-ns output risetime and falltime corresponds to a bandwidth of $0.35/t_r$ = 350 MHz. A scope of at least 500-MHz bandwidth would be required. Preferably, the scope bandwidth should be at least three times the signal bandwidth to include the second and third harmonic components for a more accurate representation of the waveform.

Modern digital storage scopes (DSOs) and digital phosphor scopes (DPOs) are popular and offer an excellent solution for performing settling time measurements as well as many other waveform analysis functions (see Reference 3). These scopes offer real-time sampling rates of several GHz and are much less sensitive to overdrive than older analog scopes or traditional sampling scopes. Overdrive is a serious consideration in measuring settling time, because the scope is generally set to maximum sensitivity when measuring a full-scale DAC output change. For instance, measuring 12-bit settling for a 1-V output (20 mA into 50 Ω) requires the resolution of a signal within a 0.25-mV error band riding on the top of a 1-V step function.

From a historical perspective, older analog oscilloscopes were sensitive to overdrive and could not be used to make accurate step function settling time without adding additional circuitry. Quite a bit of work was done during the 1980s on circuits to cancel out portions of the step function using Schottky diodes, current sources, etc. References 4, 5, and 6 are good examples of various circuits which were used during that time to mitigate the oscilloscope overdrive problems.

Even with modern DSOs and DPOs, overdrive should still be checked by changing the scope sensitivity by a known factor and making sure that all portions of the waveform change proportionally. Measuring the mid-scale settling time can also subject the scope to considerable overdrive if there is a large glitch. The sensitivity of the scope should be sufficient to measure the desired error band. A sensitivity of 1-mV/division allows the measurement of a 0.25-mV error band if care is taken (one major vertical division is usually divided into five smaller ones, corresponding to 0.2 mV/small division). If the DAC has an on-chip op amp, the fullscale output voltage may be larger, perhaps 10 V, and the sensitivity required in the scope is relaxed proportionally.

Although there is a well-known relationship between the risetime and the settling time in a single-pole system, it is inadvisable to extrapolate DAC settling time using risetime alone. There are many higher order nonlinear effects involved in a DAC which dominate the actual settling time, especially for DACs of 12-bits or higher resolution.

When making settling time measurements, is generally better to make a direct connection between the DAC output and the 50-Ω scope input and avoid the use of probes. FET probes are notorious for giving misleading settling time results. If probes must be used, compensated passive ones are preferable, but they should be used with care. Skin effect associated with even short lengths of properly terminated coaxial cable can give erroneous settling time results. In making the connection between the DAC and the scope, it is mandatory that a good low impedance ground be maintained. This can be accomplished by soldering the ground of a BNC connector to the ground plane on the DAC test board and using this BNC to connect to the scope's 50-Ω input. A manufacturer's evaluation board can be of great assistance in interfacing to the DAC and should be used if available.

Finally, if the DAC output is specifically designed to drive the virtual ground of an external current-to-voltage converter and does not have enough compliance to develop a measurable voltage across a load resistor, then an external op amp is required, and the test circuit measures the settling time of the DAC/op amp combination. In this case, select an op amp that has a settling time which is at least 3 to 5 times smaller than the DAC under test. If the settling time of the op amp is comparable to that of the DAC, the settling time of the DAC can be determined, because the total settling time of the combination is the root-sum-square of the DAC settling time and the op amp settling time. Solving the equation for the DAC settling time yields:

$$\text{DAC Settling Time} = \sqrt{(\text{Total Settling Time})^2 - (\text{Op Amp Settling Time})^2}\,. \qquad \text{Eq. 1}$$

DAC DISTORTION

If we consider the spectrum of a waveform reconstructed by a DAC from digital data, we find that in addition to the expected spectrum (which will contain one or more frequencies, depending on the nature of the reconstructed waveform), there will also be noise and distortion products.

Code-dependent glitches will produce both out-of-band and in-band harmonics when the DAC is reconstructing a digitally generated sinewave as in a Direct Digital Synthesis (DDS) system. For instance, the mid-scale glitch occurs twice during a single cycle of a reconstructed sinewave (at each mid-scale crossing), and will therefore produce a second harmonic of the sinewave, as shown in Figure 4. Note that the higher order harmonics of the sinewave, which also alias back into the Nyquist bandwidth (dc to $f_c/2$), cannot be filtered.

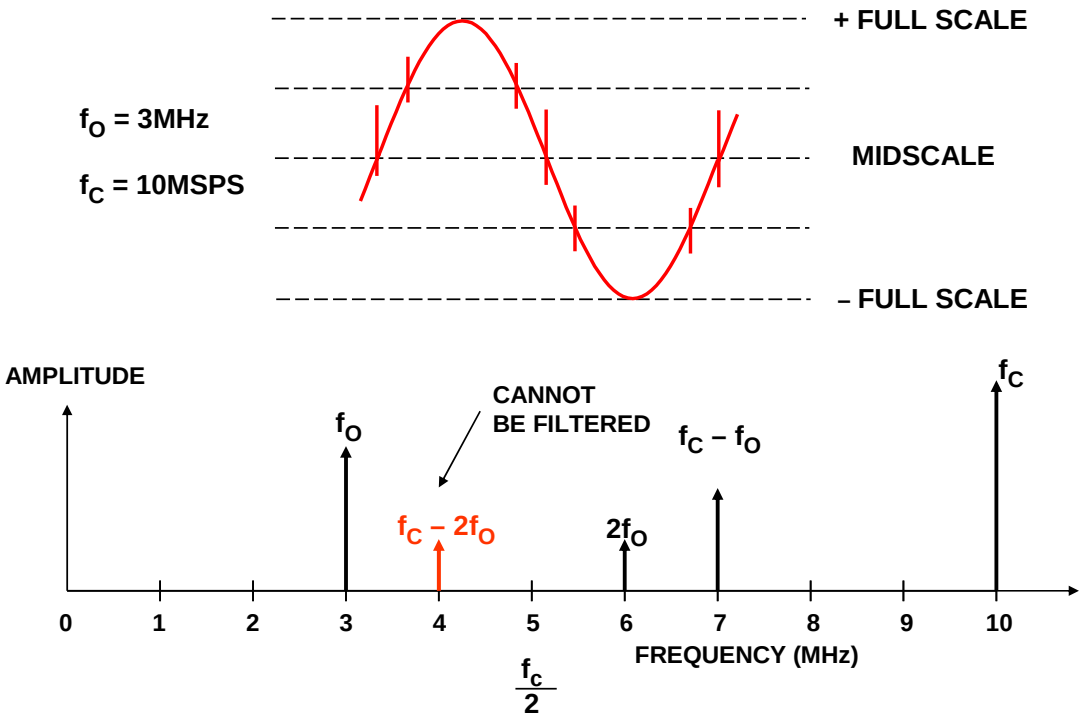

Figure 4: Effect of Code-Dependent Glitches on Spectral Output

Although segmented DAC architectures can be used to greatly minimize the distortion caused by code-dependent glitches, the distortion can never be completely eliminated.

It is difficult to predict the harmonic distortion or SFDR from the glitch area specification alone. Other factors, such as the overall linearity of the DAC, also contribute to distortion. In addition, integer ratios between the DAC sampling clock and the DAC output frequency and the cause the quantization noise to concentrate at harmonics of the fundamental thereby increasing the apparent distortion at these points.

Because so many DAC applications are in communications and frequency analysis systems, practically all modern DACs are now specified in the frequency domain. The basic ac specifications include *harmonic distortion*, *total harmonic distortion (THD)*, *signal-to-noise ratio (SNR)*, *total harmonic distortion plus noise (THD + N)*, *spurious free dynamic range (SFDR)*, etc. In order to test a DAC for these specifications, a proper digitally-synthesized signal must be generated to drive the DAC (for example, a single or multi-tone sinewave).

In the early 1970s, when ADC and DAC frequency domain performance first became important, "back-to-back" testing was popular. An ADC and its companion DAC were connected together, and the appropriate analog signal source was selected to drive the ADC. An analog spectrum analyzer was then used to measure the distortion and noise of the DAC output. This approach was logical, because ADCs and DACs were often used in conjunction with a digital signal processor placed between them to perform various functions. Obviously, it was impossible to

determine exactly how the total ac errors were divided between the ADC and the DAC. Today, however, ADCs and DACs are used quite independently of one another, so they must be completely tested on their own.

Figure 5 shows a typical test setup for measuring the distortion and noise of a DAC. The first consideration, of course, is the generation of the digital signal to drive the DAC. To achieve this, modern arbitrary waveform generators (for example Tektronix AWG2021 with Option 4) or word generators (Tektronix DG2020) allow almost any waveform to be synthesized digitally in software, and are mandatory in serious frequency domain testing of DACs (see Reference 3). In most cases, these generators have standard waveforms pre-programmed, such as sinewaves and triangle waves, for example. In many communications applications, however, more complex digital waveforms are required, such as two-tone or multi-tone sinewaves, QAM, GSM, and CDMA test signals, etc. In many cases, application-specific hardware and software exists for generating these types of signals and can greatly speed up the evaluation process.

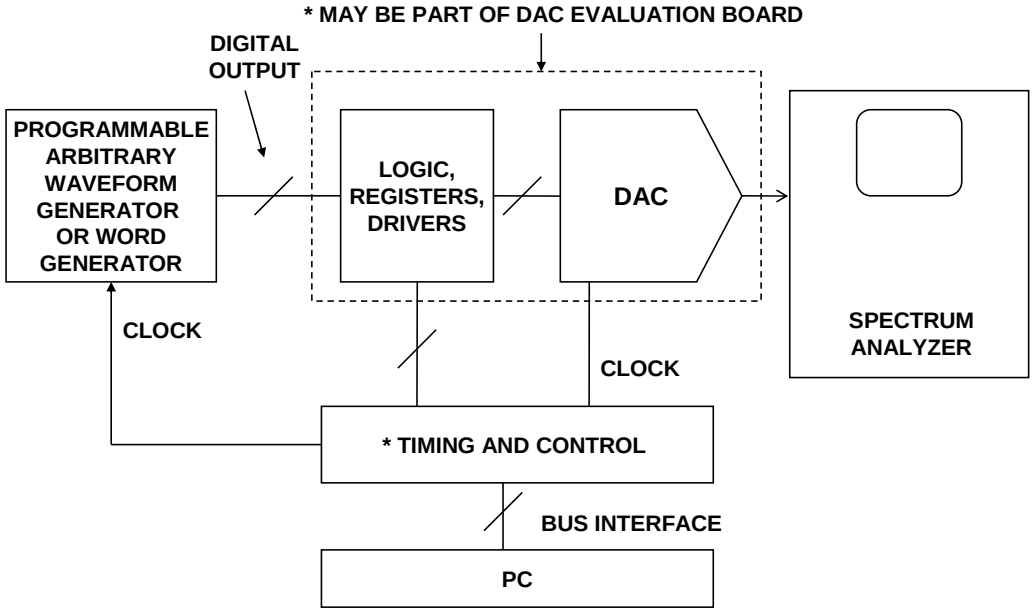

Figure 5: Test Setup for Measuring DAC Distortion and Noise

Analog Devices and other manufacturers of high performance DACs furnish evaluation boards which greatly simplify interfacing to the test equipment. Because many communications DACs (such as the TxDAC®-family) have quite a bit of on-chip control logic, their evaluation boards have interfaces to PCs via the SPI, USB, parallel, or serial ports, as well as Windows®-compatible software to facilitate setting the various DAC options and modes of operation.

Testing DACs which are part of a direct-digital-synthesis (DDS) system is somewhat easier because the DDS portion of the IC acts as the digital signal generator for the DAC. Testing these DACs often requires no more than the manufacturer's evaluation board, a PC, a stable clock source, and a high performance spectrum analyzer.

The spectrum analyzer chosen to measure the distortion and noise performance of the DAC should have at least 10-dB more dynamic range than the DAC being tested. The "maximum intermodulation-free range" specification of the spectrum analyzer is an excellent indicator of distortion performance (see Reference 7). However, spectrum analyzer manufacturers may specify distortion performance in other ways. Modern communications DACs such as the TxDAC®-series require high performance spectrum analyzers such as the Rhode and Schwartz FSEA30 (Reference 7).

As in the case of oscilloscopes, the spectrum analyzer must not be sensitive to overdrive. This can be easily verified by applying a signal corresponding to the full-scale DAC output, measuring the level of the harmonic distortion products, and then attenuating the signal by 6 dB or so and verifying that both the signal and the harmonics drop by the same amount. If the harmonics drop more than the fundamental signal drops, then the analyzer is distorting the signal.

In some cases, an analyzer with less than optimum overdrive performance can still be used by placing a bandstop filter in series with the analyzer input to remove the frequency of the fundamental signal being measured. The analyzer looks only at the remaining distortion products. This technique will generally work satisfactorily, provided the attenuation of the bandstop filter is taken into account when making the distortion measurements. Obviously, a separate bandstop filter is required for each individual output frequency tested, and therefore multi-tone testing is cumbersome.

Finally, there are a variety of application-specific analyzers for use in communications, video, and audio. In video, the Tektronix VM-700 and VM-5000 series are widely used (Reference 3). In measuring the performance of DACs designed for audio applications, special signal analyzers designed specifically for audio are preferred. The industry standard for audio analyzers is the Audio Precision, System Two (see Reference 8). There are, of course, many other application-specific analyzers available which may be preferred over the general-purpose types. In addition, software is usually available for generating the various digital test signals required for the applications.

Once the proper analyzer is selected, measuring the various distortion and noise-related specifications such as SFDR, THD, SNR, SINAD, etc., is relatively straightforward. The analyzer resolution bandwidth must be set low enough so that the harmonic products can be resolved above the noise floor. Figure 6 shows a typical spectral output where the SFDR is measured.

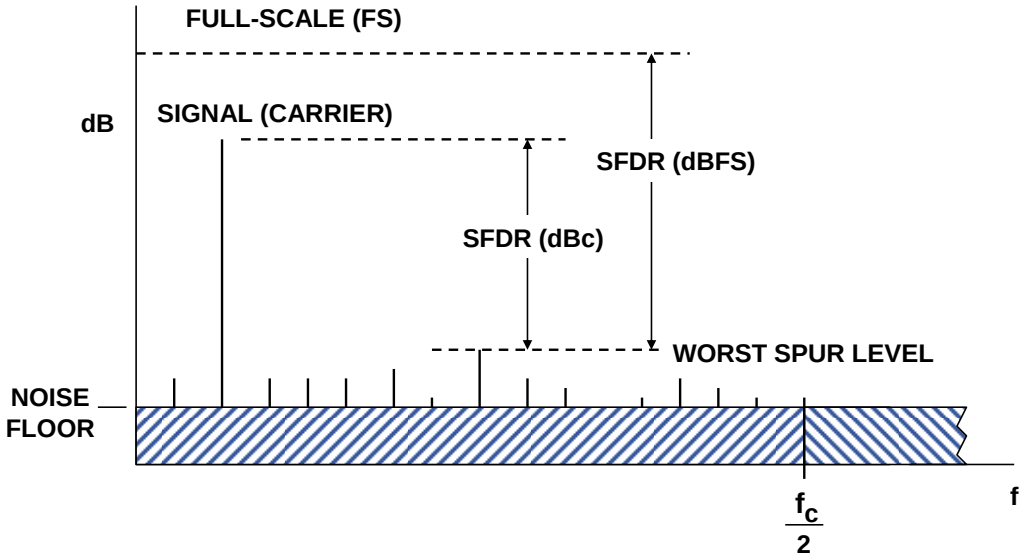

Figure 6: Measuring DAC Spurious Free Dynamic Range (SFDR)

Figure 7 shows how to measure the various harmonic distortion components with a spectrum analyzer. The first nine harmonics are shown. Notice that aliasing causes the 6^{th}, 7^{th}, 8^{th}, 9^{th}, and 10^{th} harmonic to fall back inside the $f_c/2$ Nyquist bandwidth.

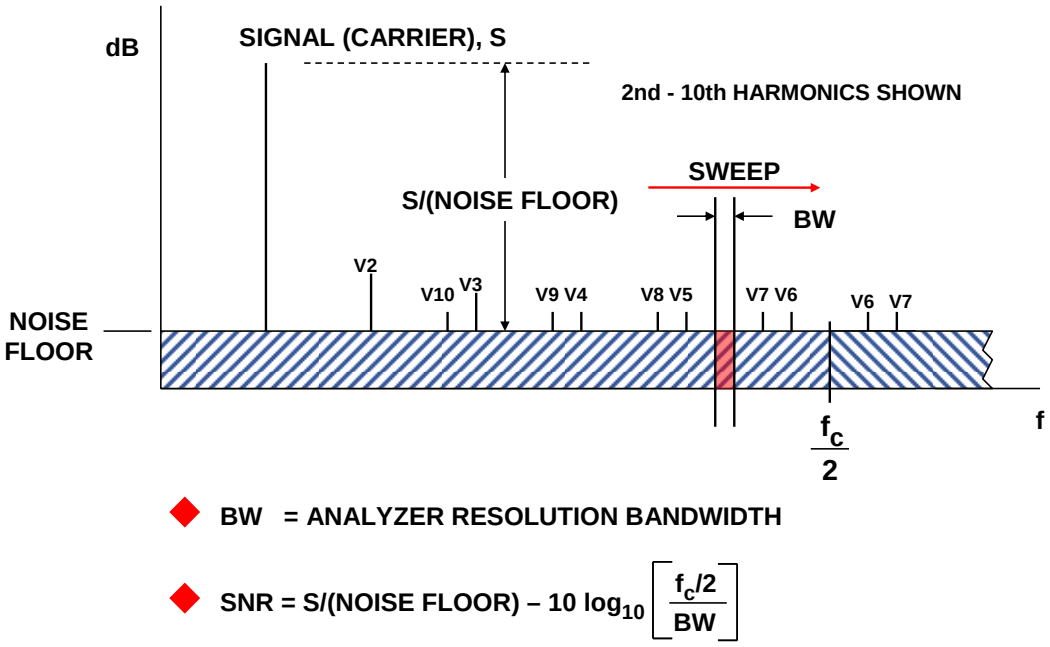

Figure 7: Measuring DAC Distortion and SNR with an Analog Spectrum Analyzer

The harmonics of the input signal can be distinguished from other distortion products by their location in the frequency spectrum. Figure 8 shows a 7-MHz input signal sampled at 20 MSPS and the location of the first 9 harmonics. Aliased harmonics of f_o fall at frequencies equal to $|\pm Kf_c \pm nf_o|$, where n is the order of the harmonic, and K = 0, 1, 2, 3,.... The second and third

harmonics are generally the only ones specified on a data sheet because they tend to be the largest, although some data sheets may specify the value of the *worst* harmonic. An interactive Harmonic Image Calculator applet is available on the Analog Devices' Design Center website which shows the locations of the second and third harmonics as a function of output frequency and DAC update rate. In addition, the tool shows the attenuation effects of the sin x/x rolloff and the output anti-imaging filter.

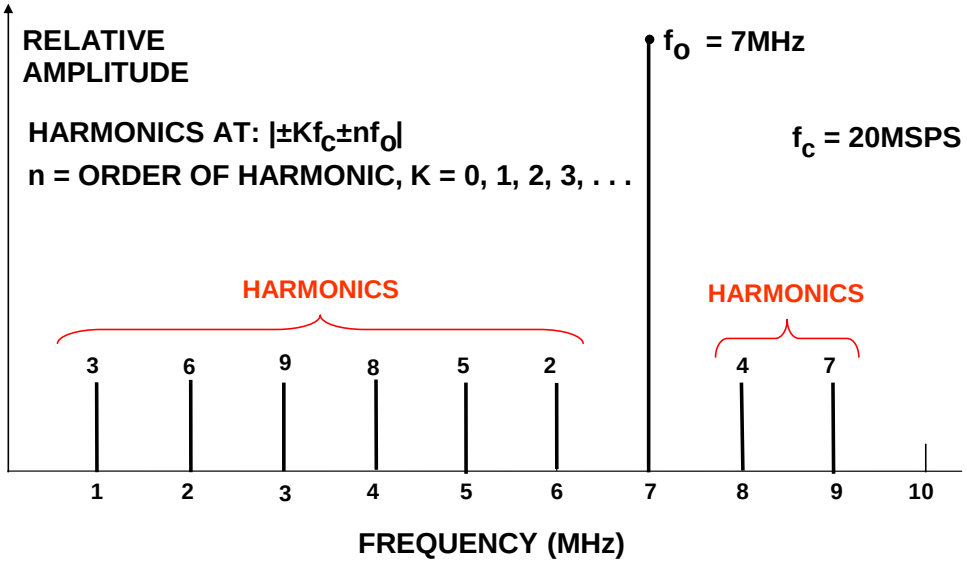

Figure 8: Location of First 9 Harmonic Products: Output Signal = 7 MHz, DAC Update Rate = 20 MSPS

DAC NOISE

The spectrum analyzer can also be used to measure SNR if the proper correction factors are taken into account. Figure 7 shows the analyzer sweep bandwidth, BW, which in most cases will be considerably less than $f_c/2$. First, measure the noise floor level with respect to the signal level at a point in the frequency spectrum which is relatively free of harmonics. This corresponds to the value "S/(NOISE FLOOR)" in the diagram. The actual SNR over the dc to $f_c/2$ bandwidth is obtained by subtracting the process gain, $10\log_{10}(f_c/2 \cdot BW)$, from the S/(NOISE FLOOR).

$$SNR = S/(NOISE\ FLOOR) - 10\log_{10}(f_c/2 \cdot BW). \qquad \text{Eq. 2}$$

In order for this SNR result to be accurate, one must precisely know the analyzer bandwidth. The bandwidth characteristics of the analyzer should be given out in the manufacturer's documentation. Also, if there is any signal averaging used in the analyzer, that may affect the net correction factor.

In order to verify the process gain calculation, several LSBs can be disabled—under these conditions, the SNR performance of the DAC should approach ideal. For instance, measuring the 8-bit SNR of a low distortion, low noise 12-, 14-, or 16-bit DAC should produce near theoretical

results. The theoretical 8-bit SNR, calculated using the formula SNR = 6.02N + 1.76 dB, is 50 dB. The process gain can then be calculated using the formula:

$$\text{PROCESS GAIN} = \text{S/(NOISE FLOOR)} - \text{SNR}. \qquad \text{Eq. 3}$$

The accuracy of this measurement should be verified by enabling the 9th bit of the DAC and ensuring that the analyzer noise floor drops by 6 dB. If the noise floor does not drop by 6 dB, the measurement should be repeated using only the first 6 bits of the DAC. If near theoretical SNR is not achieved at the 6-bit level, the DAC under consideration is probably not suitable for ac applications where noise and distortion are important.

The relationship between SINAD, SNR, and THD can be derived as follows. THD is defined as the ratio of the signal to the root-sum-square (rss) of a specified number harmonics of the fundamental signal. IEEE Std. 1241-2000 (Reference 9) suggests that the first 10 harmonics be included. Various manufacturers may choose to include fewer than 10 harmonics in the calculation. Analog Devices defines THD to be the root-sum-square of the first 6 harmonics (2nd, 3rd, 4th, 5th, and 6th) for example. In practice, the difference in dB between THD measured with 10 versus 6 harmonics is less than a few tenths of a dB, unless there is an extreme amount of distortion. The various harmonics, V2 through V6, are measured with respect to the signal level, S, in dBc. They are then converted into a ratio, combined on an rss basis, and converted back into dB to obtain the THD.

The signal-to-noise-and-distortion, SINAD, can then be calculated by combining SNR and THD as a root-sum-square:

$$\text{SINAD} = 20\log_{10}\sqrt{\left(10^{-\text{SNR}/20}\right)^2 + \left(10^{-\text{THD}/20}\right)^2}\,. \qquad \text{Eq. 4}$$

An SNR/THD/SINAD Calculator applet is available on the Analog Devices' Design Center website to assist in these conversions.

One of the most important factors in obtaining accurate distortion measurements is to ensure that the DAC output frequency, f_o is not a sub-harmonic of the update rate, f_c. If f_c/f_o is an integer, then the quantization error is not random, but is correlated with the output frequency. This causes the quantization noise energy to be concentrated at harmonics of the fundamental output frequency, thereby producing distortion which is an artifact of the sampling process rather than nonlinearity in the DAC. It should be noted that these same artifacts can occur in evaluating ADCs.

To illustrate this point, Figure 9 shows simulated results for an ideal 12-bit DAC where 9A shows the output frequency spectrum for the case of f_c/f_o = 40. Notice that the SFDR is approximately 77 dBc. The right-hand spectral output (9B) shows the case where the f_c/f_o ratio is a non-integer—the quantization noise is now random, and the SFDR is 93 dBc.

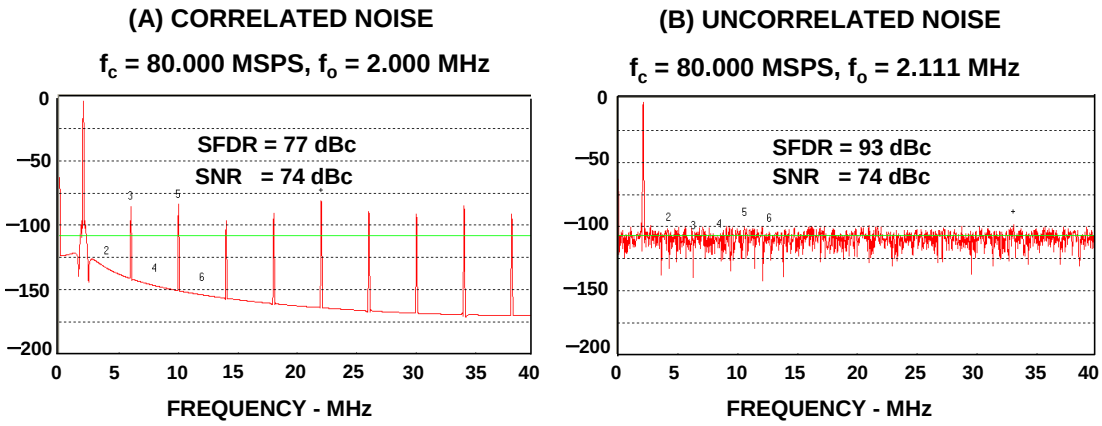

Figure 9: Correlated (A) and Uncorrelated (B) Quantization Noise Spectrum of an Ideal 12-Bit DAC

Because of the wide range of possible clock and output frequencies, Analog Devices offers special fast-turnaround measurements on TxDACs for specific customer test vectors. This important service allows system designers to do advance frequency planning to ensure optimum distortion performance for their application.

In lieu of specific frequency measurements, the SFDR performance of a DAC is often plotted as a function of the output frequency at fixed clock rates. This data is usually taken for sinewave outputs of various amplitudes as shown in Figure 10 for the AD9777 16-bit TxDAC. Note that this plot does not include data points where there is strong correlation between the quantization noise and the signal (i.e., where the ratio of the clock frequency to the output frequency is an integer number).

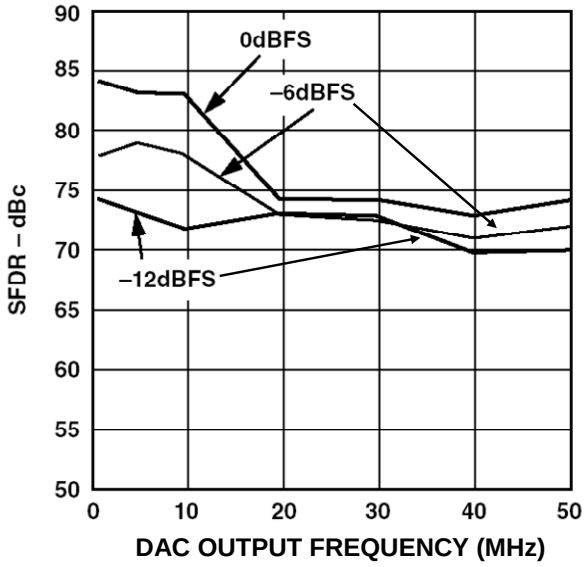

Figure 10: AD9777 16-bit TxDAC™ SFDR, Data Update Rate = 160 MSPS

There is another useful test method that gives a good overall indicator of the DAC performance at various combinations of output and clock frequencies. Specifically, this involves testing distortion for output frequencies, f_o, equal to $f_c/3$ and $f_c/4$. In practice, the output frequency is slightly offset by a small amount, Δf, where Δf is a non-integer fraction of f_c, i.e., $\Delta f = kf_c$, where $k << 1$. For an output frequency of $f_c/3 - \Delta f$, the even-order harmonics are spaced at intervals of Δf around the fundamental f_o output frequency as shown in Figure 11. The worst even-order harmonic is measured at various clock frequencies up to the maximum allowable while maintaining this same ratio. The same procedure should be repeated for an output frequency $f_c/4 - \Delta f$, in which case the odd-order harmonics are uniformly spaced around the output frequency as shown in Figure 12.

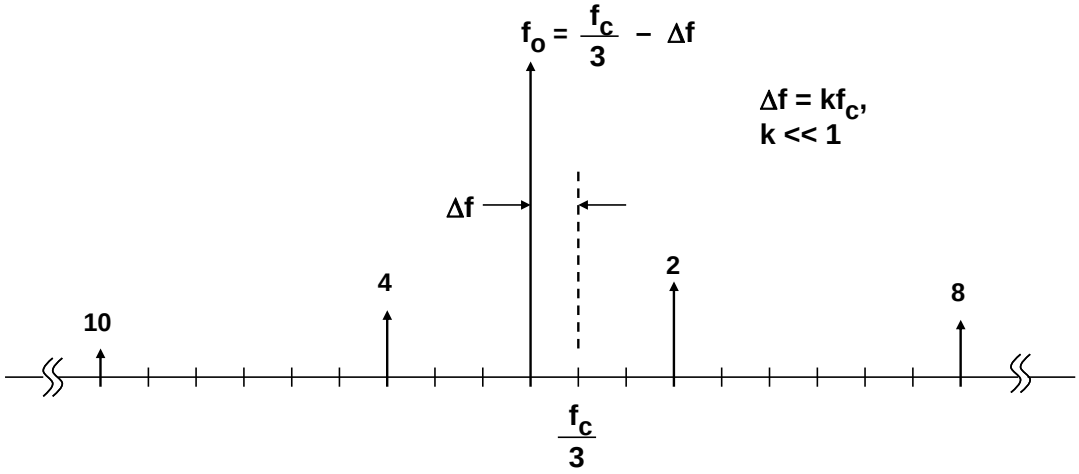

Figure 11: Location of Even Harmonics for $f_o = f_c/3 - \Delta f$

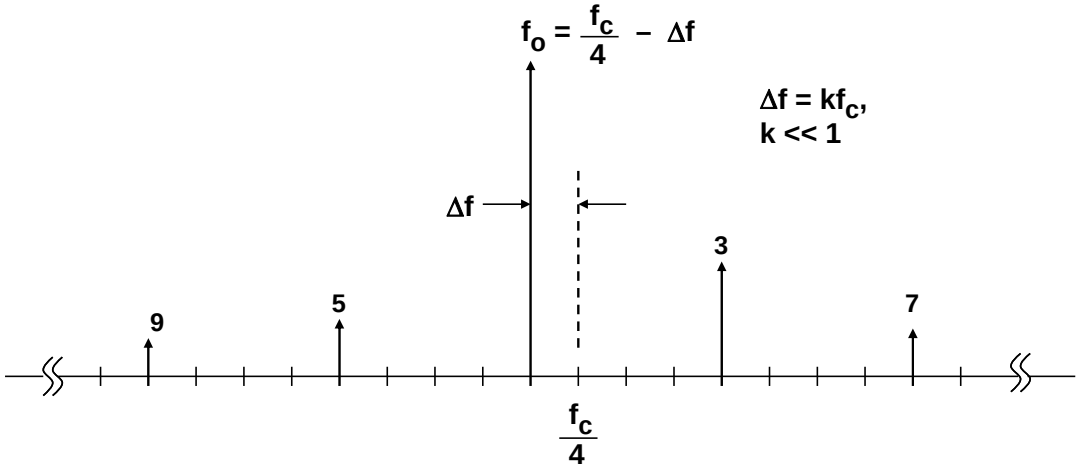

Figure 12: Location of Odd Harmonics for $f_o = f_c/4 - \Delta f$

These measurements are relatively easy to make, since once the ratio of f_o to f_c is established by the DDS or digital waveform generator, it is preserved as the master clock frequency is changed. Figure 13 shows a typical plot of SFDR versus clock frequency for a low distortion DAC with two output frequencies $f_c/3$ and $f_c/4$. In most cases, the $f_c/3$ distortion represents a worst case condition and is good for comparing various DACs.

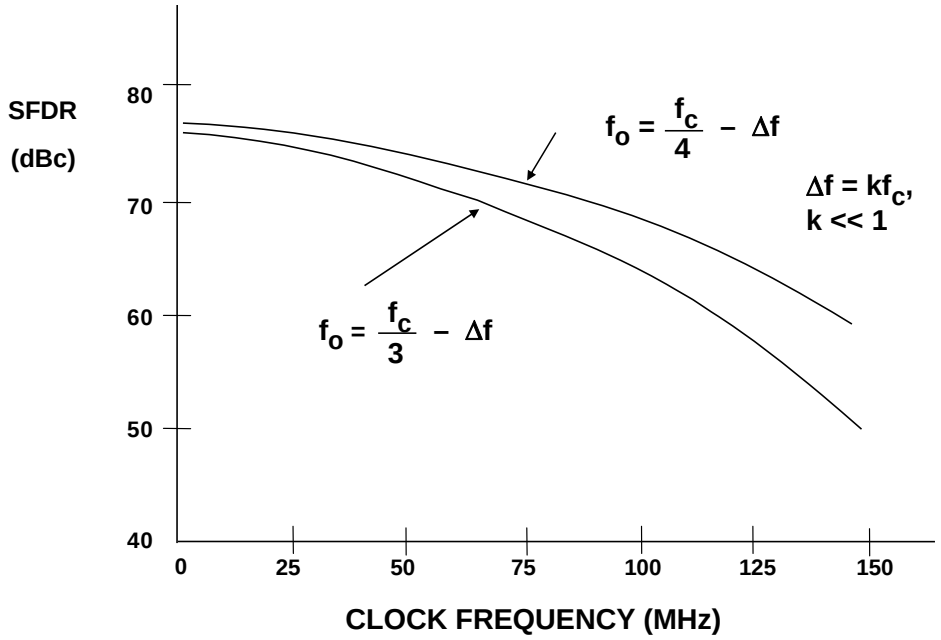

Figure 13: Worst Harmonic vs. Clock Frequency for $f_o = f_c/3 - \Delta f$ and $f_o = f_c/4 - \Delta f$

DAC OUTPUT SPECTRUM AND SIN (X)/X FREQUENCY ROLLOFF

The output of a reconstruction DAC can be represented as a series of rectangular pulses whose width is equal to the reciprocal of the clock rate as shown in Figure 14. Note that the reconstructed signal amplitude is down 3.92 dB at the Nyquist frequency, $f_c/2$. An inverse sin(x)/x filter can be used to compensate for this effect in most cases and is usually designed as part of the anti-imaging filter. The images of the fundamental signal occur as a result of the sampling function and are also attenuated by the sin(x)/x function.

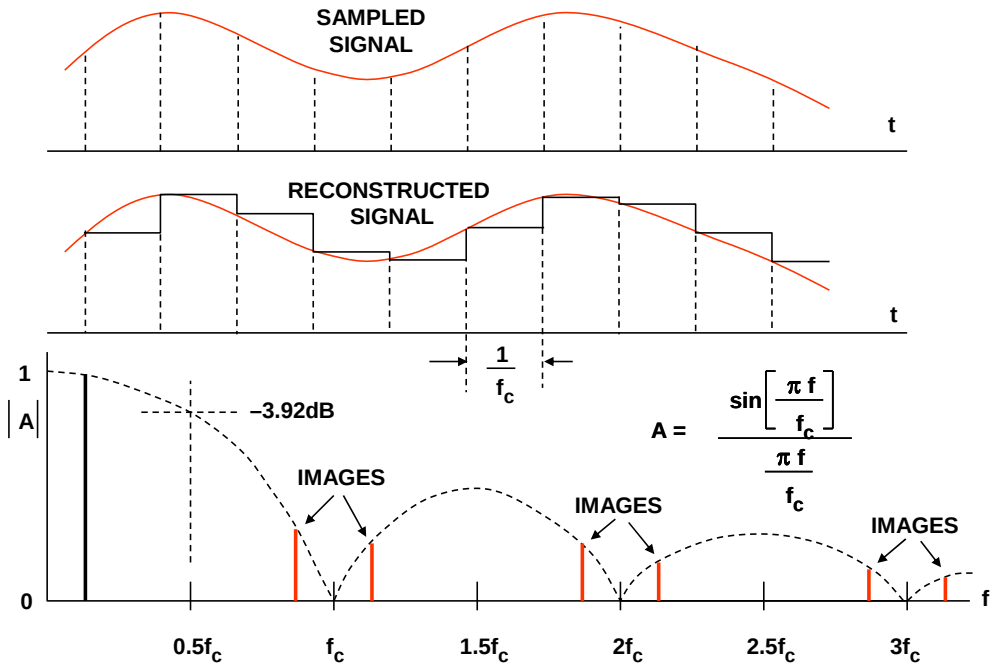

Figure 14: DAC sin(x)/x Roll Off (Amplitude Normalized)

If there is no compensation for the sin(x)/x rolloff, it must be considered when making bandwidth measurements on the DAC output. The effect of the rolloff on distortion and SNR measurements is negligible over the Nyquist bandwidth, dc to $f_c/2$.

An interactive Harmonic Image Calculator applet is available on the Analog Devices' Design Center website which shows the locations of the second and third harmonics as a function of output frequency and DAC update rate. In addition, the tool shows the attenuation effects of both the sin(x)/x rolloff and the output anti-imaging filter.

REFERENCES

1. Jim R. Naylor, "Testing Digital/Analog and Analog/Digital Converters," *IEEE Transactions on Circuits and Systems*, Vol. CAS-25, July 1978, pp. 526-538.

2. Dan Sheingold, *Analog-Digital Conversion Handbook, 3rd Edition*, Analog Devices and Prentice-Hall, 1986, ISBN-0-13-032848-0. *(the defining and classic book on data conversion).*

3. Tektronix, Inc., 14200 SW Karl Braun Drive, P. O. Box 500, Beaverton, OR 97077, Phone: (800) 835-9433, http://www.tek.com. *(the website contains a wealth of information on oscilloscopes, measurement techniques, probing, etc., as well as complete specifications on products).*

4. Howard K. Schoenwetter, "High Accuracy Settling Time Measurements," *IEEE Transactions on Instrumentation and Measurement*, Vol. IM-32, No. 1, March 1983, pp. 22-27.

5. James R. Andrews, Barry A. Bell, Norris S. Nahman, and Eugene E. Baldwin, "Reference Waveform Flat Pulse Generator," *IEEE Transactions on Instrumentation and Measurement*, Vol. IM-32, No. 1, March 1983, pp. 27-32.

6. Barry Harvey, "Take the Guesswork out of Settling-Time Measurements," *EDN*, September 19 1985, pp. 177-189.

7. Rohde & Schwarz, Inc., 8661A Robert Fulton Dr., Columbia, MD 21046-2265, Phone: (410) 910-7800, http://www.rohde-schwarz.com. *(a premier manufacturer of spectrum analyzers, the website contains tutorials on frequency analysis as well as product specifications).*

8. Audio Precision, 5750 S.W. Arctic Drive, Beaverton, Oregon 97005, http://www.audioprecision.com. *(the recognized industry standard for professional audio measurement equipment).*

9. *IEEE Std. 1241-2000, IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters*, IEEE, 2001, ISBN 0-7381-2724-8.

10. Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN 0-916550-27-3, Chapter 2 and 5. Also available as *The Data Conversion Handbook*, Elsevier/Newnes, 2005, ISBN 0-7506-7841-0, Chapter 2 and 5.