ANALOG DEVICES

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Designer's Guide to Flash-ADC Testing

Part 3
Measure Flash-ADC Performance for Trouble-Free Operation

by Walt Kester

The first two parts of this series described the subtleties of flash A/D converters and the test methods used to evaluate these devices. Part 3 concludes the series with a discussion of the actual measurements you'll need to fully characterize flash A/D converters.

Although manufacturers have expanded the number of guaranteed specifications they put on their data sheets, the test conditions often won't match those of your system design. You can use the methods described in Part 2 of this series to test a flash A/D converter, but the measurements you need to perform depend on the converter's primary application. This final part of the series provides information on important measurements you'll need to characterize your converter's performance, including total harmonic distortion (THD), differential and integral nonlinearity, and noise power ratio. You'll probably want to start with the S/N ratio, a measurement that's common to most A/D converter applications.

The S/N ratio is the ratio of the rms fundamental to the rms quantization noise. As described in Part 2, you can measure this parameter by digitizing a pure sine wave and performing Fourier transformations on the data. The rms energy contained in the fundamental sine wave is equal to the square root of the sum of the squares of the peak value and the values of the appropriate number of samples, or bins, located on either side of the peak. The converter's resolution and its side-lobe roll-off characteristics determine the number of samples you'll need. For a detailed explanation of sampling requirements, see Part 2.

The rms energy in the remaining frequency bins represents the noise due to theoretical quantization, the converter's harmonic distortion and excess noise, and the FFT round-off error. Take the square root of the sum of the squares of the remaining samples (excluding the dc components) to determine the rms energy. The overall S/N ratio of the A/D converter is

S/N ratio = 20 log(rms signal level/rms noise level).

You can measure harmonic distortion in a similar manner. The test program (described in Part 2) examines the FFT frequency spectrum for the proper location of the desired harmonic (harmonics above $f_s/2$ will be aliased into the baseband) and determines the rms energy in that harmonic. The following equation calculates the harmonic distortion:

Harmonic distortion = 20 log(rms signal level/rms harmonic level).

Reprinted from EDN – February 1, 1990

The S/N ratio and harmonic distortion are key specifications in evaluating the performance of A/D converters.

The total harmonic distortion (THD) is the root-sum-square of the first five harmonics of the fundamental. Use this number in place of the rms harmonic level in the above formula.

Two-tone intermodulation tests using FFTs

In many applications, you don't have the simple case of a single input frequency. For example, in communication applications that multiplex several frequencies onto a single carrier, you need to measure intermodulation products. You determine this parameter by applying two sine waves of different frequencies (f_1 and f_2) to an A/D converter. You then measure the amplitudes of the third-order intermodulation products, which occur at frequencies $2f_1+f_2$, $2f_1-f_2$, $2f_2+f_1$, and $2f_2-f_1$.

Although it's possible to filter out most intermodulation distortion if the two tones are of similar frequencies, the third-order products will be very close to the fundamental frequencies and thus difficult to remove.

To avoid clipping-induced distortion, the amplitudes of the individual tones should be at least 6 dB below the full-scale range of the flash converter. In addition, the frequency separation of the two tones should be consistent with the resolution of the FFT. As discussed in Part 2, the spectral resolution of the FFT is a function of record length M, coherence vs noncoherence, and the properties of the windowing function that you choose.

In receiver applications, you often want to know the maximum ratio between the amplitude of a single-tone input signal and the amplitude of its maximum spurious component. For an ideal A/D converter, this ratio occurs for a full-scale input sinusoid. In a practical A/D converter, however, spurious content is a function of slew rate. Therefore, the maximum spurious-free dynamic range for a given input frequency will probably occur at a level somewhat below full scale. Because the spurious-free dynamic range is slew-rate dependent, it's a function of input frequency and amplitude.

Fig 1 is a plot of the typical maximum spurious level vs input signal level. Also shown is a plot of the corresponding spurious-free dynamic range. The plot demonstrates that the maximum spurious-free dynamic range of 38 dB occurs for an input signal that's about 3 dB below full scale.

The data you need to generate these plots is readily available from the family of FFTs calculated for the different input amplitudes. By knowing the input signal level that gives the highest spurious-free dynamic range at frequencies close to the Nyquist frequency, it's possible to set the gain of the system to take maximum advantage of the A/D converter's spectral characteristics.

Histograms are helpful

Differential and integral nonlinearity are also important measurements of converter performance. Try a histogram test to obtain these measurements. To make a histogram analysis, digitize a known periodic input at a rate that's asynchronous relative to the input signal. To gather the sample data for the histogram, you'll need a buffer memory and a test system, as described

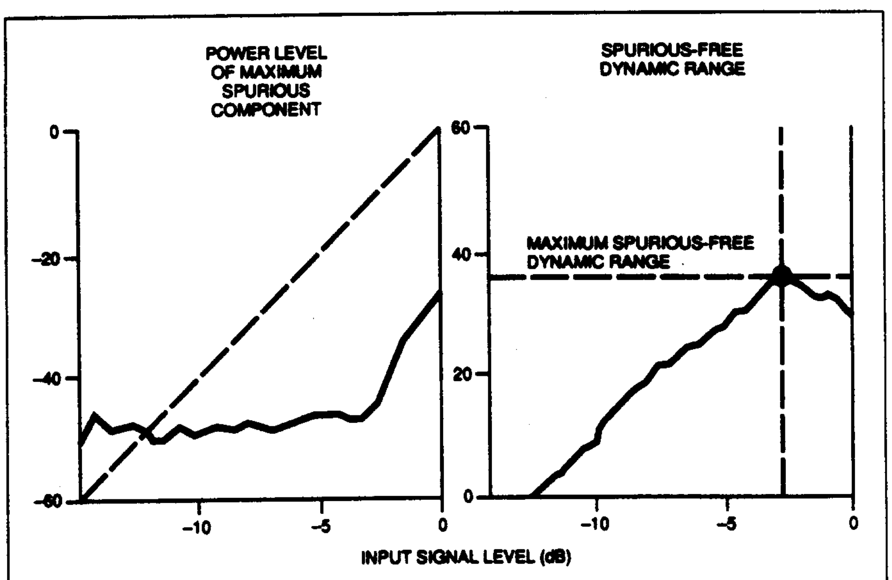

Fig 1—These dynamic-range plots show the power levels of spurious frequencies and the maximum spurious-free dynamic range. In this example, the maximum spurious-free dynamic range occurs at an input signal level that's 3 dB below full scale.

in Part 2. The buffer memory will probably be too small to hold a statistically significant number of samples from a single run (several hundred thousand are usually required). For this reason, run several tests to acquire the data and load the contents of the buffer into the main memory of your test system after each run. Benchtop test systems from Hewlett Packard and Tektronix also provide histogram test capability.

After the test system accumulates a statistically significant number of samples, it can determine the relative number of occurrences of each digital code (the code density). This test routine then normalizes the data based upon the input signal and analyzes the results for linearity errors.

For an ideal A/D converter with a full-scale triangular-wave input, you'd expect an equal number of codes in each bin. The number of counts in the nth bin, H(n), divided by the total number of samples taken, M, is the bin width as a fraction of full scale. The ratio of the actual bin width to the ideal bin width, P(n), is the differential linearity. Ideally, this ratio should be unity. Subtracting 1 LSB gives you the differential nonlinearity.

You can determine integral nonlinearity with a cumulative histogram; the cumulative bin widths are the transition levels. However, the cumulative effects of errors can make the integral-nonlinearity measurement inaccurate. Histograms are used more often in evaluating differential nonlinearity.

High-speed, high-accuracy triangular waves are difficult to generate, so use a sine wave. All codes aren't equally probable with a sine-wave input, however, and you should normalize the histogram data using the probability density function for a sine wave, as shown in **Fig 2**.

To obtain accurate results, you need to take a large number of samples. For example, to determine the differential nonlinearity for an 8-bit flash converter to within 0.1 bit with 99-percent confidence, you'll need 268,000 samples. You can use hardware to count these samples, thus speeding up the software processing time. For high-speed sampling, decimate the output data to clock rates that are compatible with a slower-speed memory.

Using noise-power-ratio tests

You can use noise-power-ratio (NPR) tests to measure the transmission characteristics of frequency-division-multiplexed (FDM) communications links. In a typical FDM system, 4-kHz-wide voice channels are "stacked" in frequency for transmission over coaxial, microwave, or satellite equipment. At the receiving end, the FDM equipment demultiplexes the data and returns it to individual, 4-kHz baseband channels. In an FDM system that has 100 channels or more, Gaussian noise with the appropriate bandwidth approximates the FDM signal.

The test setup of **Fig 3** measures an individual 4-kHz channel for quietness by using a narrow-band notch (bandstop) filter and a tuned receiver (**Ref 4**), both of which measure the noise power inside this 4-kHz notch. The NPR measurements are straightforward. With the

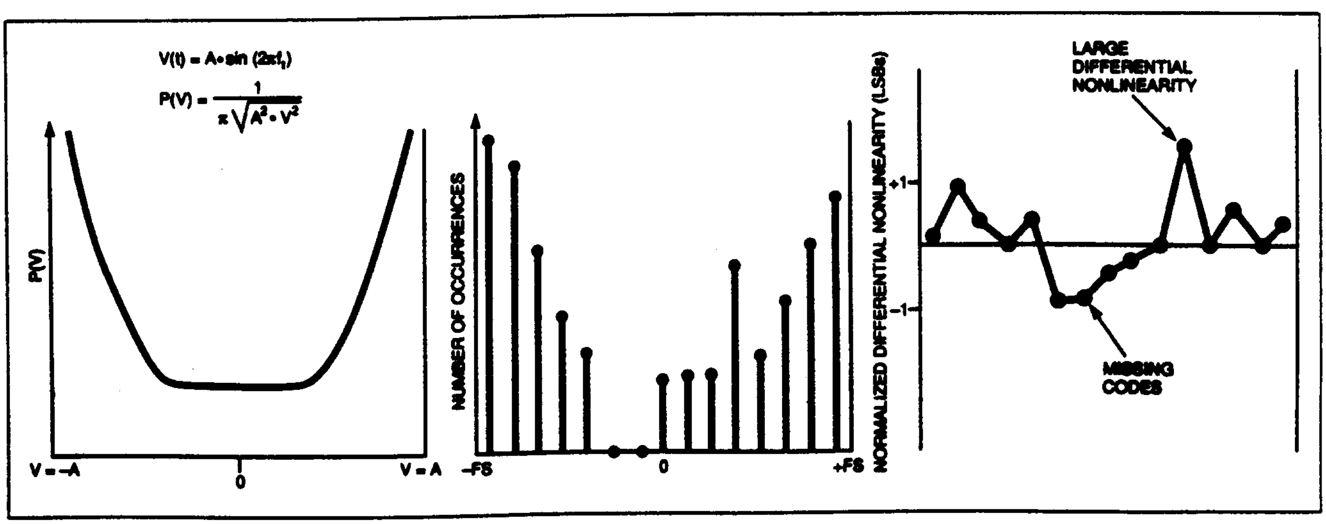

***Fig 2—Histograms are often used to plot differential nonlinearity.** Shown here is a curve for the probability density function of a sine wave, which is used to normalize histogram data to produce a plot of differential nonlinearity.*

Where multiple frequencies exist on a single carrier, you need to measure intermodulation distortion as well as harmonic distortion.

notch filter out, the receiver determines the rms noise power of the signal inside the notch. The notch filter is then switched in, and the receiver determines the residual noise inside the 4-kHz slot. The ratio of the two readings, expressed in dB, is the NPR. You should test several slot frequencies across the noise bandwidth—low, midband, and high.

The NPR is usually plotted on an NPR curve as a function of rms noise level referred to the peak range of the system. For very low noise levels, the undesired noise is primarily thermal noise and is independent of the input noise level. Over this region of the curve, a 1-dB increase in the noise level causes a 1-dB increase in the NPR. As the noise level increases, the amplifiers in the system begin to overload, creating intermodulation products that cause the noise floor of the system to rise. As the input noise increases further, the effects of overload noise predominate, reducing the NPR dramatically. FDM systems are usually operated at a noise-loading level a few decibels below the point of maximum NPR.

In a digital system containing an A/D converter, the noise within the slot is primarily quantizing noise when low values of noise input signals are applied. The NPR curve is linear in this region. As the noise input level increases, the hard-limiting action of the converter causes clipping noise to dominate.

In a practical A/D converter, any dc or ac nonlinearities cause a departure from the theoretical NPR. Although the peak value of NPR occurs at a fairly low input noise level (rms noise $= 1/4\ V_O$, where $\pm V_O$ is the range of the A/D converter), the broadband nature of the noise signal stresses the device, and the test provides a good indication of its dynamic performance.

Theoretically, NPR readings should be independent of any particular slot frequency. However, because of increased nonlinearities for the higher input frequencies, the NPR readings in the higher slots tend to be lower.

NPR testing using DSP techniques

Using FFT analysis techniques, you'll find NPR measurements a real challenge. Consider the case where the record length is 1024 and the sampling rate is 20 MHz. The FFT of 1024 contiguous time samples would place a spectral component every 19.53 kHz (20 MHz/1024). Because the notch-filter slot width is approximately 4 kHz, the probability of a spectral component falling within the notch is very low.

To achieve reasonable data stability in the FFT NPR analysis, a number of samples must fall within the notch. If ten samples are within the 4-kHz notch, then the resolution of the FFT would need to be 400 Hz, necessitating a record length of 50,000 for a sampling

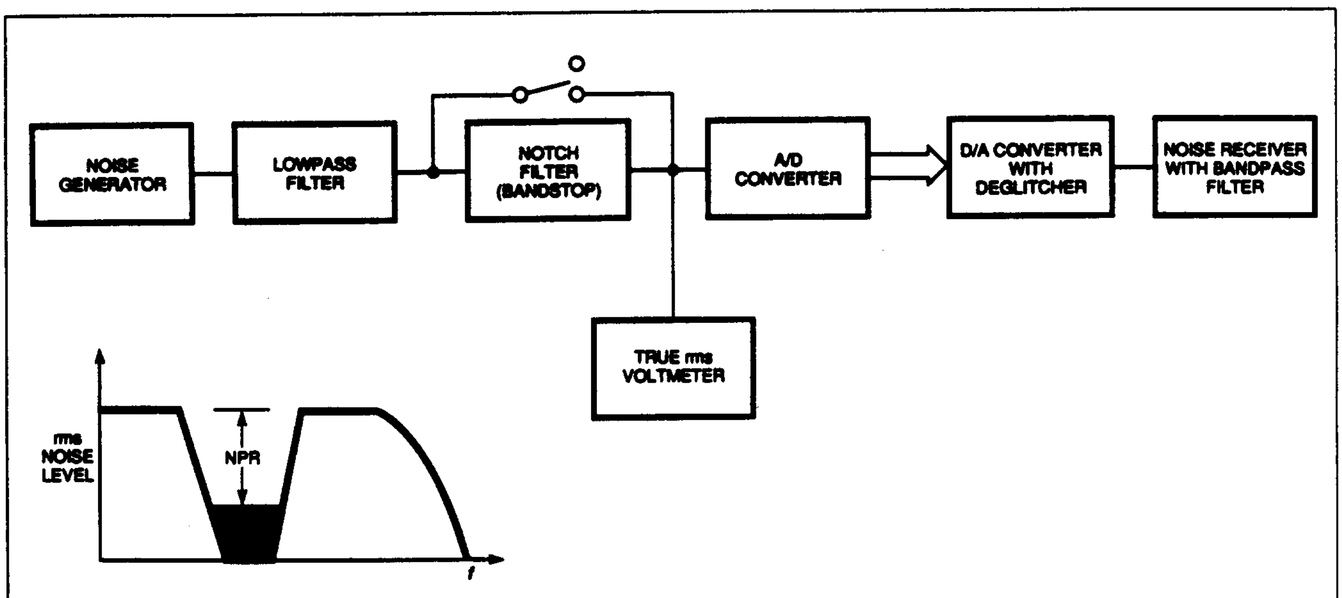

Fig 3—You can use this test setup to measure noise power ratio (NPR). With the notch filter out, the receiver determines the noise power of the signal inside the notch. With the notch filter switched in, the receiver measures the residual noise inside the typical 4-kHz slot. The ratio of the two readings (in decibels) is the NPR.

rate of 20 MHz. To avoid an extremely large buffer memory (and hence more demands on the FFT processor), you need to make the notch filter wider. For 20-MHz sampling and a 1024-word buffer memory, a notch filter that has a width of 200 kHz will provide ten frequency bins inside the notch. Even under these conditions, however, you should average the NPR calculations for several records to provide reasonable data stability.

Transient-response testing

The response of a flash converter to a transient input such as a square wave is often critical in radar applications. The major difficulty in implementing this test is obtaining a flat pulse that's commensurate with the converter's resolution.

A test setup for measuring the transient response of an A/D converter is shown in **Fig 4**. If you mount the Schottky-diode flat-pulse generator as close as possible to the analog input of the A/D converter, you can apply a signal to the A/D converter that's flat to at least 10-bit accuracy a few nanoseconds after it reverse biases the Schottky diodes.

You can use the same test setup to measure overvoltage recovery time. The amount of overvoltage is generally specified as a percentage of the A/D converter's range. For a converter with a 2V input range, 50% overvoltage corresponds to 1V above or below the nominal 2V input range. You make the starting point of the flat pulse correspond to the desired overvoltage condition. The actual recovery time is referenced to the time the input signal re-enters the A/D-converter input range. As in the transient-response test, you must consider the sampling (aperture) time delay when making this measurement.

The aperture-time and -jitter specifications of video A/D converters have probably been the least understood and most misused specifications in the entire field. The original concept of aperture time is centered around the classic S/H circuit of **Fig 5**. In an ideal S/H circuit, the switch has zero resistance when closed and opens instantly on receipt of an encode command. In practice, the sampling switch changes from a low to a high resistance over a certain finite time interval. An error occurs because the circuit tends to average the input signal over the finite time interval required to open the switch. As a result, the sampled voltage varies from the voltage at the instant the switch starts to open. The time required to open the switch is the aperture time. The error is determined by $E_a = t_a\,dV/dt$, where E_a is the aperture error, t_a is the aperture time, and dV/dt is the rate at which the input signal changes.

A simple first-order analysis, which neglects nonlinear effects, shows that no real error exists for such a switch. As long as the switch opens in a repeatable fashion, there is an effective sampling time that will cause an ideal S/H amplifier to produce the same hold voltage. The difference between this effective sampling point and the leading edge of the sampling clock is a fixed delay, which doesn't constitute an error. This effective aperture delay is the period from the leading edge of the sampling clock to the instant when the input signal equals the hold value. This specification

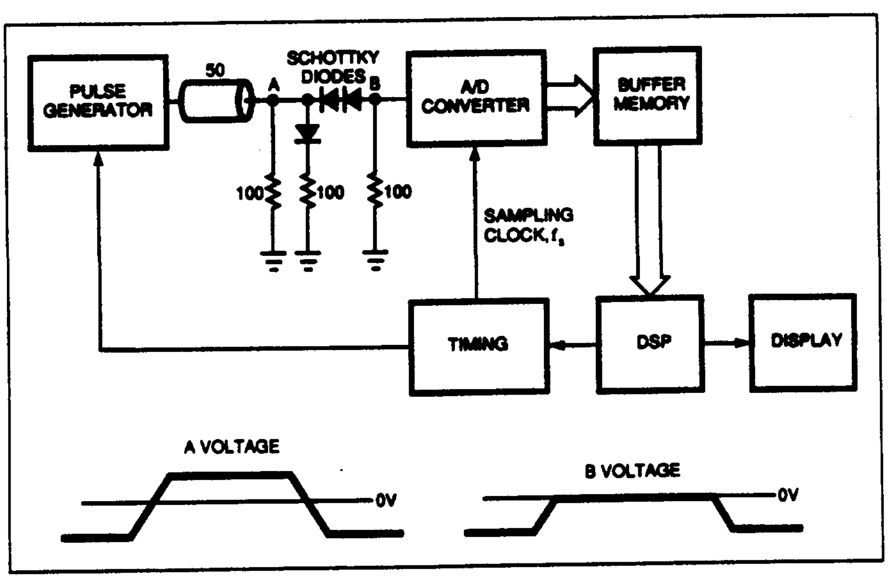

***Fig 4—This test setup measures the transient response of an A/D converter.** The Schottky-diode network, located between points A and B in the circuit, generates a flat pulse for the input of the converter.*

In a practical A/D converter, the spurious-free dynamic range is a function of the converter's slew rate and can occur at a level below full scale.

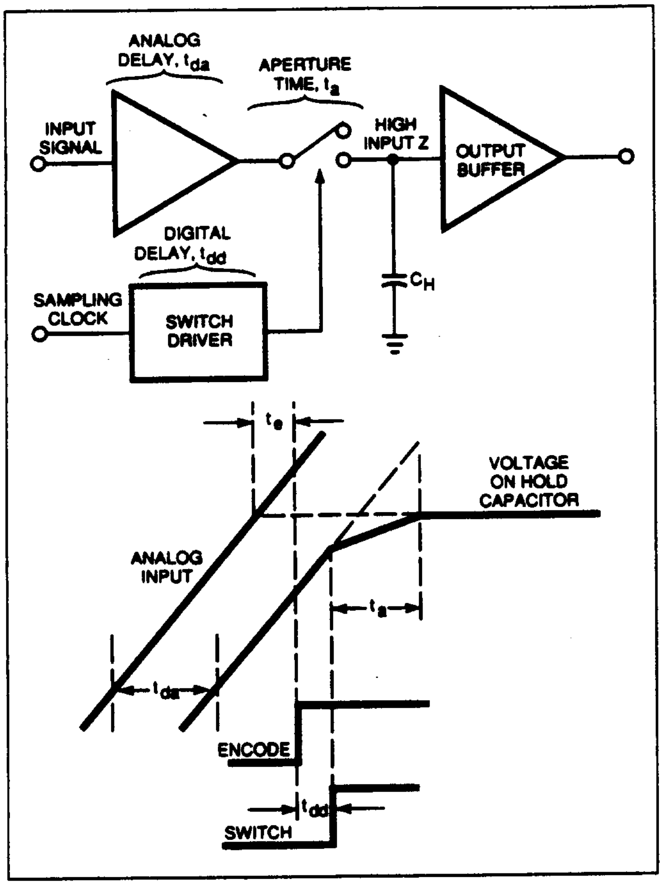

Fig 5—The concept of aperture time centers around the S/H circuit. In practice, the sampling switch generates an error because of input-signal averaging over the finite time interval needed to open the switch. The aperture time is the time needed to open the switch.

is important because it helps you determine when to apply the sampling clock with respect to the input signal timing.

The variation in effective aperture delay is important in simultaneous S/H applications. For example, in both I (in-phase) and Q (quadrature) radar receivers you may have to provide adjustable delays in the sampling clock to match the effective aperture delay times of several A/D converters. You should also consider delay-time tracking over a range of temperatures, especially in military systems where the specified operating temperature ranges from −55 to +125°C.

True aperture errors, however, do result from variable time delays. In a practical A/D converter, the sampling clock is often phase-modulated by some unwanted source; the source can be wideband random noise, power-line frequency, or digital noise due to poor grounding techniques. Phase jitter on the input sine wave can produce the same effect as jitter on the sampling clock. The resulting error is called aperture jitter. The corresponding rms voltage error caused by the rms aperture jitter qualifies as a valid aperture error.

The aperture-jitter specification is sometimes interpreted as a measure of the converter's ability to accurately digitize rapidly changing input signals. An A/D converter with an impressive aperture-jitter specification still may lose effective bits when digitizing a sine wave that has a maximum slew rate calculated from the aperture formula $E_a = t_a\,dV/dt$.

For example, assume that a 20-MHz, 8-bit flash converter has a bipolar input range of $\pm V_0$ ($2V_0$ p-p) and an aperture jitter specification of 20 psec rms. To calculate the maximum aperture-jitter error, convert the rms aperture jitter into a maximum value. If you consider that aperture jitter follows a Gaussian distribution similar to white noise, the rms aperture jitter, t_a, corresponds to the sigma (σ) of the distribution. The 2σ point on the distribution is a good place to set the maximum value, and the maximum aperture jitter becomes is $2t_a$.

If the corresponding maximum voltage error (ΔV) at the zero crossing of a full-scale sine wave is set to ½ LSB (½ LSB$=2V_0/2^{N+1}$, where N equals the resolution of the A/D converter), then you can calculate the maximum full-scale sine-wave frequency, f_{max}, which will produce the ½ LSB aperture error, by using the following equations:

$$V(t) = V_0 \cdot \sin(2\pi f_t),$$

$$\frac{dV}{dt} = 2\pi f V_0 \cdot \cos(2\pi f_t),$$

$$\left.\frac{dV}{dt}\right|\max = \frac{\Delta V}{2t_a} = 2\pi V_0 f_{max}\Big|, \text{ and}$$

$$f_{max} = \frac{\Delta V}{4\pi V_0 t_a} = 2\pi t_a \cdot 2^{N+1}.$$

For $t_a = 20$ psec rms and $N = 8$, f_{max} is 16 MHz. These calculations imply that a 20-MHz flash converter can accurately digitize a full-scale sine wave of 16 MHz. In actual practice, however, the device may begin to suffer from skipped codes, decreased effective bits and S/N ratio, and ac nonlinearities at much lower frequencies.

You can calculate the effects of aperture jitter on

Histograms are useful in evaluating the differential nonlinearity of an A/D converter.

the full-scale sine-wave S/N ratio as follows:

$$V(t) = V_0 \cdot \sin(2\pi f_t),$$

$$\frac{dV}{dt} = 2\pi f V_0 \cdot \cos(2\pi f_t), \text{ and}$$

$$\frac{dV}{dt}_{rms} = \frac{2\pi f V_0}{\sqrt{2}}.$$

For an rms error voltage, ΔV_{rms}, and an rms aperture jitter of t_a,

$$\frac{\Delta V_{rms}}{t_a} = \frac{2\pi f V_0}{\sqrt{2}}, \text{ and}$$

$$\Delta V_{rms} = \frac{2\pi f V_0 t_a}{\sqrt{2}}.$$

The rms-signal to rms-noise ratio, expressed in decibels, is

$$\text{S/N ratio} = 20 \log \left[\frac{V_0/\sqrt{2}}{\Delta V_{rms}}\right]$$

$$= 20 \log \left[\frac{1}{2\pi f t_a}\right] \text{dB}.$$

The S/N ratio that's due exclusively to aperture jitter in the above equation is plotted in Fig 6 as a function of the full-scale input-sine-wave frequency for various values of aperture jitter.

Consider an 8-bit, 20-MHz A/D converter with an rms aperture jitter of 20 psec. For an 8-MHz full-scale input, the S/N ratio due only to aperture jitter is 60 dB, as calculated from the equation. The theoretical S/N ratio due to quantizing noise in an 8-bit flash converter is 50 dB. When you combine the S/N ratio of 60 dB with the S/N ratio of 50 dB, you obtain a theoretical S/N ratio of 49.6 dB, which encompasses both the ideal quantizing noise and the noise due to aperture jitter. A practical 8-bit device that has an rms aperture-jitter specification of 20 psec may, however, only achieve an S/N ratio of 40 dB under these conditions.

Therefore, to accurately evaluate the A/D converter's dynamic performance, you must carefully examine the S/N ratio, effective number of bits, and aperture-jitter specifications.

Try measuring the aperture jitter of an A/D converter using the test setup shown in Fig 7. The low-jitter pulse generator produces both the sampling clock and the analog input signal to minimize the phase jitter between them. Adjust the phase shifter until the A/D converter repetitively samples the sine wave at its point of maximum slew rate at midscale. Then take a histogram on the digitized A/D-converter output data.

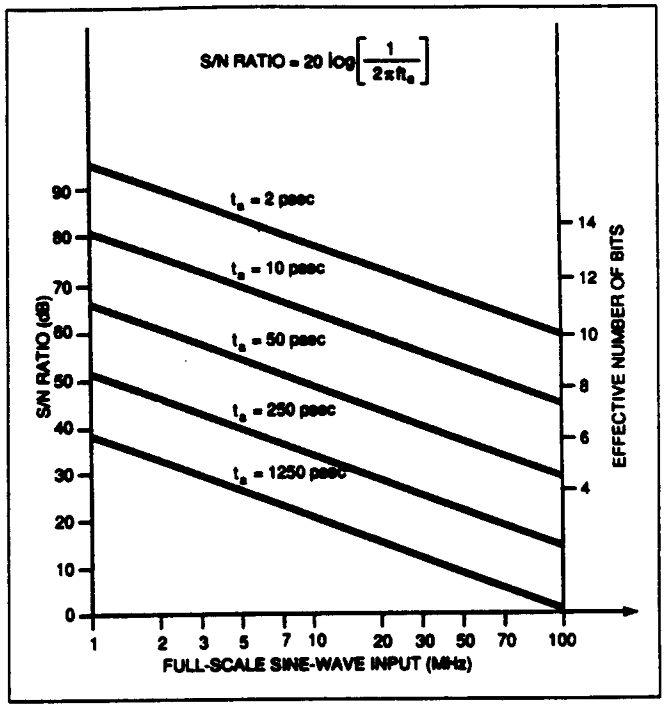

Fig 6—This plot compares the S/N ratio to the full-scale sine-wave input frequency for various values of aperture jitter.

An ideal A/D converter with no aperture jitter would have only one code present on the histogram. A practical converter gives a distribution of codes that you can fit to the normal distribution. The sigma (Σ) of the distribution corresponds to the rms error voltage, ΔV_{rms}, produced by the rms aperture jitter. Calculate the aperture jitter, t_a, from the formula

$$t_a = \frac{\Delta V_{rms}}{\frac{dV}{dt}}$$

where dV/dt is the rate-of-change of the sine wave at zero crossing.

If you sufficiently attenuate the input sine wave, any spreading of the distribution around the nominal code is due to intrinsic A/D-converter noise. As the input sine wave increases in amplitude, the slew rate, dV/dt, becomes proportionally greater, and the distri-

Noise-power-ratio tests are useful in determining the transmission characteristics of frequency-division-multiplexed communications links.

bution begins to spread because of the aperture jitter. Because high slew rates can affect the ac differential linearity of the converter, you should exercise caution when interpreting the histogram for high slew-rate inputs.

The offset adjustment shown in **Fig 7** lets you position the sine wave at different points on the A/D-converter range. In this way, you can see variations attributed to range-dependent differential-linearity characteristics. When offsetting the sine wave, make sure you don't exceed the A/D converter's input range.

It's also possible to measure effective aperture delay by using the locked-sine-wave technique. Adjust the phase shifter until the output reads midscale. Use a dual-trace scope to determine the difference between the leading edge of the sampling-clock pulse and the actual zero crossing of the sine-wave input. This difference is the effective aperture delay, which can be either negative or positive, depending on the values of the internal analog and digital delays in the S/H portion of the A/D converter.

At present, no industry standard exists for either the definition or the test for A/D-converter error rates. In flash converters, comparator metastable states can occur for low- or high-frequency input signals. At high frequencies, bubbles in the thermometer code of the comparator-bank output can also produce erroneous output codes.

Because error rates less than 1×10^{-16} are typical for well-behaved A/D converters, you need to take a large number of samples to properly measure the error rate. You must also take great care in the test-set layout, grounding, shielding, and power-supply decoupling so that 60-Hz, EMI, or RFI glitches don't create erroneous errors.

Use the circuit in **Fig 8** to measure the error rate for low-frequency input signals. Apply a low-frequency, full-scale sine wave (or triangle wave) to the A/D converter so that its rate of change is less than 1 LSB/sample. This step ensures that the transition zones between codes are all adequately exercised. An error amplitude of X LSBs is established as the lower limit for the definition of a qualified error. Usually, you select X to be several LSBs so that random noise doesn't produce errors. The software or hardware then examines the difference between each adjacent sample and records the number of times this difference exceeds the error threshold, X. If NQ is the number of qualified errors that occur, and NT is the total number of samples taken, then the error rate, ER, is given by the equation $ER = NQ/2\cdot NT$.

As an example, consider an 8-bit, 100M-sample/sec flash converter designed to take at least ten samples at each code level. For one slope of the triangle-wave input, the number of samples required is $10\times256=2560$ samples. The frequency of the triangle wave is

$$f_t = \frac{1}{2560\cdot2\cdot10\ \text{nsec}} = 19.5\ \text{kHz}.$$

At a 100-MHz sampling rate, the average time required to make an error for an error rate of 1×10^{-9} is 10 seconds.

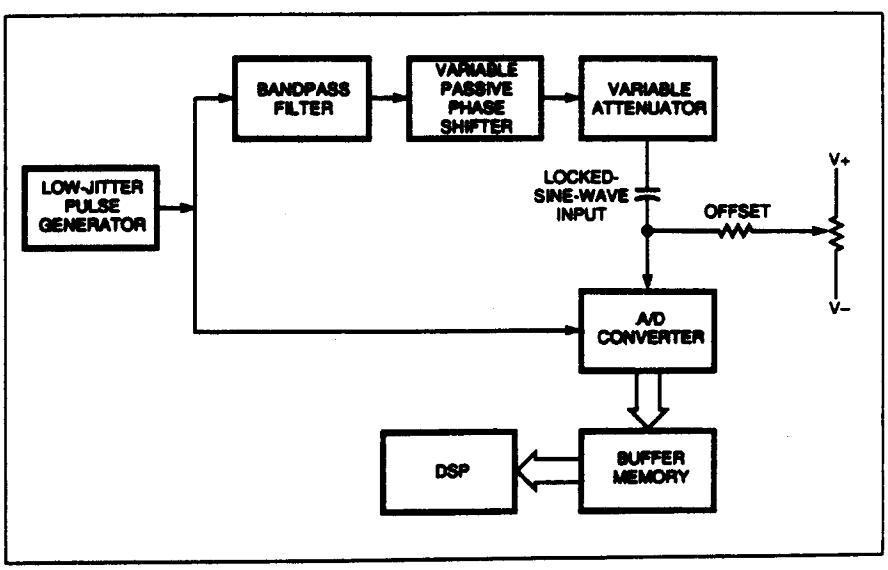

Fig 7—In this test setup for measuring aperture jitter, you adjust the phase shifter until the A/D converter repetitively samples the sine wave at its point of maximum slew rate. You then take a histogram of the digitized A/D-converter output data. The offset adjustment lets you position the sine wave at different points on the converter's range.

Aperture time and aperture jitter for A/D converters are probably the most misunderstood and misused specifications.

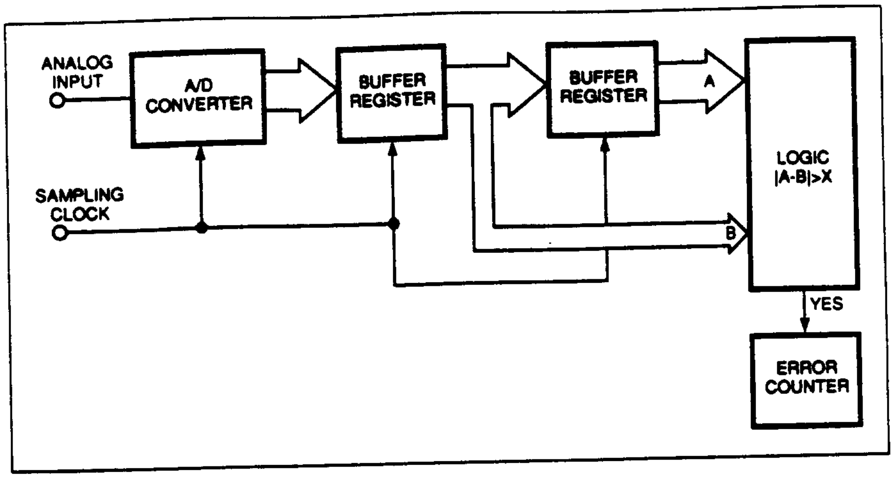

Fig 8—The effective aperture delay is the time difference between the leading edge of the sampling-clock pulse and the actual zero crossing of the sine-wave input.

3

In a similar manner, you can measure dynamic errors caused by fast input signals by using the beat-frequency approach. You choose the low-frequency beat frequency to give the proper number of samples per code level, and then you examine the decimated digital outputs for adjacent sample differences that exceed the allowable error amplitude.

In summary, determining appropriate error-rate criteria for an A/D converter depends upon both the application and the characteristics of the converter under consideration. Flash converters that use straight binary decoding with no additional correction logic are most subject to large metastable errors at midscale. For this situation, a low-amplitude dither signal centered on the midscale code transition might be an appropriate stimulus. In a more well-behaved flash converter, a full-scale signal that exercises all codes might be desirable.

If you plan to digitize composite video signals, you'll need to measure the differential-gain and -phase performance of the flash A/D converter. Differential gain is the percentage difference between the digitized amplitudes of two signals. Likewise, differential phase is the phase difference between the digitized values of the same two input signals. The input signals are typically a high-frequency low-level sine wave representing the color subcarrier frequency, superimposed on a low-frequency sine wave. Distortion-free processing of the color signal requires that the flash converter alters neither the amplitude nor the phase of the chrominance signal as a function of the luminance-signal level.

The best method for performing composite video tests is to use an A/D converter back-to-back with a D/A converter. Connect a TV test signal to the A/D converter and use the output of the D/A converter to drive a vectorscope. To ensure that the test accurately measures the A/D converter's performance, use a low-glitch D/A converter followed by a track-and-hold deglitcher. In addition, the dc accuracy of the D/A converter should exceed that of the A/D converter. When testing an 8-bit flash converter, use a D/A converter with at least 10 bits of accuracy.

References

1. Andrews, James R, 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, pg 27.

2. Schoenwetter, Howard K, "A Programmable Voltage Step Generator for Testing Waveform Recorders," *IEEE Transactions on Instrumentation and Measurement*, Vol IM-33, No. 3, September 1984, pg 196.

3. Gray, G A, and G W Zeoli, "Quantization and Saturation Noise Due to Analog-Digital Conversion," *IEEE Transactions on Aerospace and Electronic Systems*, January 1971, pg 222.

4. Tant, M J, *The White Noise Book*, Marconi Instruments, July 1974.