

Designer's Guide to Flash-ADC Testing**3****Part 1****Flash ADCs Provide the Basis for High Speed Conversion**

by Walt Kester

Building high-performance circuits that take advantage of the high sampling rates of flash ADCs requires a knowledge of these converters' many intricacies. Part 1 of this 3-part series discusses the pitfalls of designing with flash ADCs, how to evaluate certain data-sheet specifications, and how to choose external components that complement your particular converter. Parts 2 and 3 will explore test and measurement methods that you can use to verify a converter's performance in your system.

To digitize analog signals whose bandwidths exceed 1 MHz, you'll probably need flash ADCs. Many flash converters with 4 to 10 bits of resolution are now available, thanks to recent advances in VLSI process technology and design techniques. However, to use these converters successfully at the high sampling rates that they provide, you must take into account and compensate for a variety of flash-converter characteristics.

The basic features of most flash converters are shown in Fig 1. A flash ADC simultaneously applies an analog-input signal to $2^N - 1$ latched comparators, where N is the number of the converter's output bits. A resistive voltage divider generates each comparator's reference voltage and sets each reference level 1

LSB higher than the level of the comparator immediately below it. Comparators that have a reference voltage below the input-signal level will assume a logic 1. The comparators with a reference voltage above the level of the input signal will produce a logic 0. A secondary logic stage decodes the thermometer code that results from the $2^N - 1$ comparisons. An optional output register latches the decoding stage's digital output for one clock cycle.

Timing is everything

One of the first difficulties you'll encounter when using flash converters is removing valid data from the converter. In practice, the comparator bank has two states controlled by a conversion-command signal. Various converters call this command the convert, the encode, or simply the clock command. When this signal is in its convert-command state, the comparators track the analog-input signal, and during this time the output data is invalid. When the command line changes state, it latches the comparator outputs. Valid output data is now available for transfer to an external register. You'll find most flash converters somewhat sensitive to the duty cycle and frequency of this command pulse. In other words, the performance of the converter, specifically its differential and integral nonlinearity performance, is related to the clock's duty cycle and frequency. Performance degradation is especially pronounced when you run the device at or near its maximum sampling rate.

Because of recent advances in VLSI processing and design techniques, many flash converters that have from 4 to 10 bits of resolution are available.

The way you handle the binary output depends on whether the converter has an internal output latch. Without a latch, the data will be invalid for a period equal to the sampling clock's pulse width. At high sample rates, the data-invalid time will impinge on the data-valid time, making it difficult for you to strobe the flash converter's output into an external register. For instance, if you operate a flash converter at 100M samples/sec with a 50%-duty-cycle sampling clock, the output data will be valid for only 5 nsec. When you consider the finite rise and fall time of the output binary bits, this short time doesn't leave you much leeway, even if you use the fastest external logic. In fact, you may ultimately lose data. The addition of an internal output latch simplifies clocking of the output data, because the output data is valid for approximately the entire clock cycle. In return for a longer data-valid time, you'll have to accept an inherent 1-cycle or more pipeline delay—an acceptable compromise in most systems applications.

Try to place an appropriate buffer register next to the flash converter. If you route the converter's digital

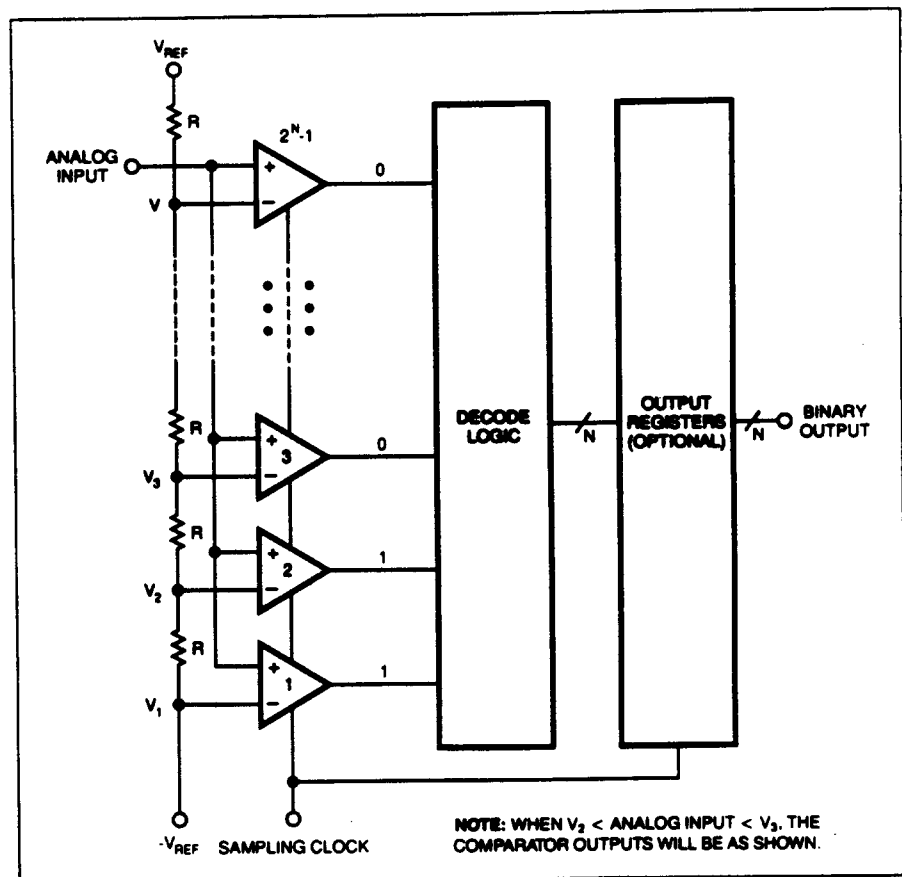
output directly to a backplane data bus through a card-edge connector, signal coupling between the digital-output signals and analog input will degrade the S/N ratio and harmonic performance.

In many high-speed data-acquisition designs, you'll need a large and fast buffer memory to store the output data. A 500M-sample/sec converter can fill 1M byte of memory in 2 msec. To reduce the required speed—and thus the cost—of the memory, you can demultiplex the high-speed data stream (Fig 2), which slows it to frequencies compatible with cost-efficient CMOS RAMs. Fig 2's circuit clocks the two output registers at half the sample rate, and it latches data in each register 180° out of phase from the other. Some flash converters that operate in excess of 200 MHz have onboard demultiplexing for added convenience.

All that sparkles isn't gold

So far, these timing difficulties refer to how you deal with the converter's output data. But flash converters can have internal problems as well. Low input frequencies can cause comparator metastability, and

Fig 1—Flash converters contain a bank of $2^N - 1$ comparators, where N is the number of output bits. Decoding logic transforms the comparator outputs into the appropriate N -bit result. Timing and input characteristics of the converter, such as mismatched comparator delays, mismatched ladder resistors, and nonlinear input capacitances, are just a few sources of converter errors.



high input frequencies can lead to errors caused by slew-rate and delay mismatches. All of these errors may manifest themselves as sparkle codes in a poorly designed flash converter.

Sparkle codes are random errors whose magnitude may approach the full-scale range of the converter. The term refers to the white dots or "sparkles" that appear against a gray background when the ADC output drives a video display. There are two sources of sparkle codes: comparator metastable states and thermometer-code bubbles.

A comparator metastable state occurs if the comparator output falls between the logic-0 and logic-1 threshold of the digital decoding logic. If the threshold uncertainty region has a width of ΔV_L and the comparator has a gain of A , then the error probability P_m is a uniformly distributed value that's equal to

$$P_m = \frac{\Delta V_L}{Aq}$$

where q is the weight of the LSB.

A latched comparator in a flash converter has a regenerative gain of

$$A = A_0 e^{v\tau}$$

when $t > 0$, and

$$A = A_0$$

when $t \leq 0$. τ equals the regeneration-time constant, and t is the time after the application of a latch command. The probability of a metastable state, P_m , for a regenerative comparator bank driving decoding logic is

$$P_m = \frac{\Delta V_L}{A_0 q} e^{-v\tau}$$

The magnitude of the sparkle code depends on the location of the metastable comparator in the comparator bank and on the logic-decoding scheme. For instance, the 128th comparator determines the MSB of an 8-bit flash converter with straight binary decoding logic. If this comparator's output is in a metastable state, the decoding logic may mistakenly convert the input voltage whose correct binary representation is 01111111 to 11111111, producing a full-scale error. If the comparator bank's thermometer-code output is first decoded into Gray code, latched, and then converted into binary code, the metastable error is reduced to 1 LSB, regardless of the comparator in error. Flash converters, however, rarely use this scheme because of the ripple-through time and the increased logic density of the Gray-to-binary circuitry. Instead, converter designers often use "pseudo-Gray" decoding techniques to eliminate the delay time associated with traditional Gray-to-binary circuits.

Note that the probability of a metastable-state error increases as the time-after-latch, t , decreases (assum-




Fig 2—Flash converters can operate at extremely high sample rates. Thus, you must have an output register that can handle this fast data. By demultiplexing the converter's outputs, you can slow the data to a rate that's compatible with standard CMOS memory.

One of the first difficulties you'll encounter when using flash converters is removing valid data from the converter.

ing a constant value of τ). This implies that the flash converter is more apt to produce metastable-state errors as you increase the sampling rate, because t must decrease correspondingly. Most manufacturers reduce metastable comparator states by minimizing the regeneration-time constant, τ . Lower regeneration-time constants result in higher power dissipation, which is one reason why many high-speed flash converters are power-hungry devices.

Thermometer-code bubbles are another potential source of sparkle codes. A well-behaved flash converter's comparator bank produces a specific sequence of ones up to a certain point in the input range of the converter, and it produces a sequence of zeros beyond that point. The decoding logic then assigns a binary number to the thermometer code. For low-frequency inputs, most flash converters' comparator banks are well behaved. At high speeds, however, delay mismatches among comparators may produce out-of-sequence ones and zeros in the thermometer code. The decoding logic then assigns an error binary code to these out-of-sequence points, or bubbles, which also result in sparkle codes. Again, proper comparator design and more sophisticated decoding-logic circuitry

within the ADC itself can reduce these errors to acceptable levels.

These comparator-timing errors degrade both the differential and integral linearity of a flash ADC as the input slew rate increases. In addition to static errors, such as missing codes and sparkle-code errors, slew-rate limitations manifest themselves as dynamic errors, such as increased harmonic distortion and degradation in the S/N ratio. Ideally, a flash converter should maintain its static performance specifications across the full Nyquist bandwidth, and certain applications demand full performance beyond the Nyquist bandwidth. The theoretical, rms S/N ratio for an N-bit ADC is given by the well-known equation

$$\text{S/N ratio} = 6.02N + 1.76 \text{ dB.}$$

However, as shown in a typical plot of S/N ratio versus input frequency for actual flash converters (Fig 3), the S/N ratio degrades as the input frequency increases. This degradation starts well below these converters' maximum sampling rates. The left vertical axis of Fig 3 shows the S/N ratio in another term: effective number of bits. The effective number of bits is simply the value of N when you solve the above equation using a specific value for the S/N ratio. Aperture jitter (sample-to-sample variations in the effective-sampling instant) can also cause degradation in the overall S/N ratio for high-slew-rate inputs. Jitter can be internal or external to the converter. Part 3 of this series will cover this topic in more detail. To minimize externally produced jitter components, you should always practice proper grounding, power-supply-decoupling, and pc-board-layout techniques.

Watch for dangerous data-sheet territory

Most of the timing difficulties and error sources discussed so far are common to all flash converters. However, each converter features its own unique design and specifications; thus the data sheets require scrutiny. Don't be fooled into believing that sampling rate and input bandwidth are interchangeable; they're different specifications. It wasn't until recently that you'd even find input bandwidth specified for an ADC. Even now, no accepted industry-wide definition exists for a flash converter's full-power-bandwidth specification. Scrutinize the data sheet carefully and make sure you understand the manufacturer's definition and test method. The full-power bandwidth of a traditional op amp is the maximum frequency at which the amplifier




Fig 3—Theoretically, a flash converter should maintain its performance across the full Nyquist bandwidth. But as the curves in this figure illustrate, the S/N ratio starts to degrade well below each device's maximum sampling rate. (Note that these curves are based on test sampling rates that are somewhat lower than each device's maximum possible rate.)

can produce the specified p-p output voltage at a specified level of distortion. Another commonly used definition calculates the full-power bandwidth by dividing the amplifier's slew rate by $2\pi V_o$, where the output-voltage range of the amplifier is $\pm V_o$.

When you apply traditional analog-bandwidth definitions to flash converters, the results can be misleading. The dynamic-error sources previously discussed may become predominant long before the comparator front end approaches its maximum bandwidth. If you use a common definition of full-power bandwidth as the frequency at which the p-p reconstructed-sine-wave output is reduced by 3 dB for a full-scale input, then the effective number of bits (S/N ratio) at this input frequency may render the flash converter useless in your system. Thus, to get a true idea of a converter's performance, you must consider both the full-power bandwidth and the effective number of bits (S/N ratio) at a specific sampling rate.

Another definition you'll encounter occasionally for full-power bandwidth is the maximum, full-scale input signal at a specified sampling rate that produces no missing codes. Using this definition always gives the most pessimistic number, so specifications based on this definition appear on only a few data sheets. The following is a recently proposed definition for full-power bandwidth (courtesy Chris Manglesdorf, a senior scientist at Analog Devices): the frequency at which the fundamental component of the reconstructed FFT output—excluding harmonics—is reduced by 3 dB from full scale.

Just when you think you've mastered the intricacies of flash ADCs, you'll realize that you have another component to worry about: the input buffer amplifier.

Fortunately (or unfortunately depending upon your perspective), the flash converter—not the amplifier—usually limits the converter's dynamic performance. A flash converter typically digitizes a signal from a 50, 75, or 93Ω bipolar or unipolar source. If the input range of the flash converter is incompatible with the signal, then you'll clearly need a wideband op amp to generate the required gain and offset (Fig 4). In addition, the input capacitance of some flash converters may vary as a function of the analog input's signal amplitude. Therefore, so that the nonlinear capacitance doesn't produce undesirable harmonics in the digitized signal, you'll need to use a buffer amplifier for isolation. For certain flash converters, the input capacitance is so high that a buffer amplifier is needed just to preserve the signal bandwidth.

A good choice for the buffer is a high-speed transimpedance amplifier. These amplifiers have high bandwidths and flat frequency responses over a broad range of input frequencies. Also, many transimpedance amplifiers exhibit extremely low distortion. Pairing the right amplifier with your converter is important. For instance, Fig 3 shows the S/N ratio of various converters plotted along with the harmonic distortion of the AD9617. Since the THD of the amplifier is better than the S/N ratio of the converters, the amplifier won't degrade the flash converters' performance over the major portion of their usable bandwidth.

Another factor to consider when driving the input of flash converters is the input-signal polarity. Positive input signals, which forward bias the substrate diode, can damage a converter that has a unipolar, negative input-voltage range. Installing an external Schottky diode provides effective protection.

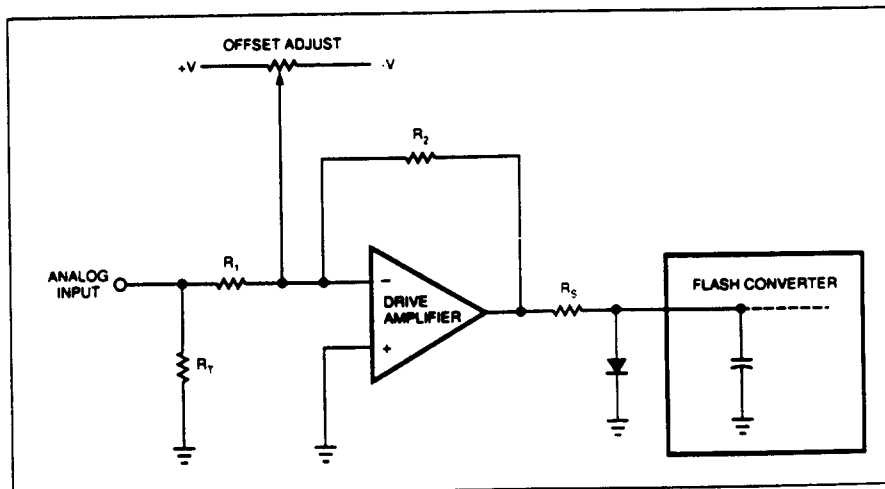


Fig 4—If the voltage range of the analog signal and the input range of the ADC aren't compatible, you'll have to adjust the gain using R_1 and R_2 , and also adjust the input signal's offset. Because of the substantial value and the often nonlinear nature of a flash converter's input capacitance, you must choose an appropriate drive amplifier and value for R_3 .

Ideally, a flash converter should maintain its static-performance specifications across the full Nyquist bandwidth, but in reality ADCs fall far short of this ideal.

The flash converter's input capacitance and the drive amplifier's isolation resistor, R_s , form a lowpass filter. Typical series-resistor and input-capacitance values of 10Ω and 20 pF create a single-pole lowpass filter that has a bandwidth of 800 MHz . However, if the input capacitance changes from 20 to 15 pF over the input range of the converter, then an attenuation error of 1.4% occurs for a 50-MHz input signal. This 1.4% non-linearity will produce 37 dBc of harmonics. (The unit dBc refers to the number of dB between the signal you're measuring and the carrier frequency). If you minimize the value of R_s and still maintain op-amp stability, you can reduce the attenuation error caused by these lowpass filtering effects. The signal dependence of input capacitance is rarely specified, but as converters move toward higher bandwidths, you can expect to see this parameter on more data sheets.

Few, if any, flash converters contain an internal voltage reference, so in addition to an external drive amplifier, you must also design your own voltage-reference generator. Fig 5 illustrates a typical -2V , unipolar reference-voltage circuit for a flash converter. A buffer transistor is necessary because the resistance of the converter's ladder string is usually fairly low. The total reference-ladder resistance of a flash converter depends heavily on the fabrication process and may vary considerably from device to device. Also, the ladder's resistance may exhibit a large temperature coefficient.

If the flash converter allows bipolar operation, then you'll have to generate two reference voltages. The circuit in Fig 6 allows great flexibility in setting both the gain and the offset of a bipolar flash converter, and it operates on $\pm 5\text{V}$ power supplies. A few flash converters provide a sense pin for the voltage reference. You can use this pin to compensate for the voltage drop caused by the package's pin and bond-wire resistances, as shown in the bipolar-reference circuit in Fig 6. In addition, some ADCs give you access to one or more taps along the internal, reference-ladder resistor string. To achieve better integral linearity, you can drive these taps from low-impedance sources.

Improve dynamics with T/H amplifiers

As previously discussed, the effective-sample time-delay variations among the latched comparators degrade the S/N ratio and the harmonic performance. You can visualize the individual comparators within an array as having variable delay lines in series with their latch-strobe inputs. To understand the effect of this delay on performance, consider an 8-bit, 100M -sample/sec flash converter that's digitizing a full-scale, 50-MHz sine-wave input. You can express the sine wave as

$$v(t) = V_p \sin 2\pi ft.$$

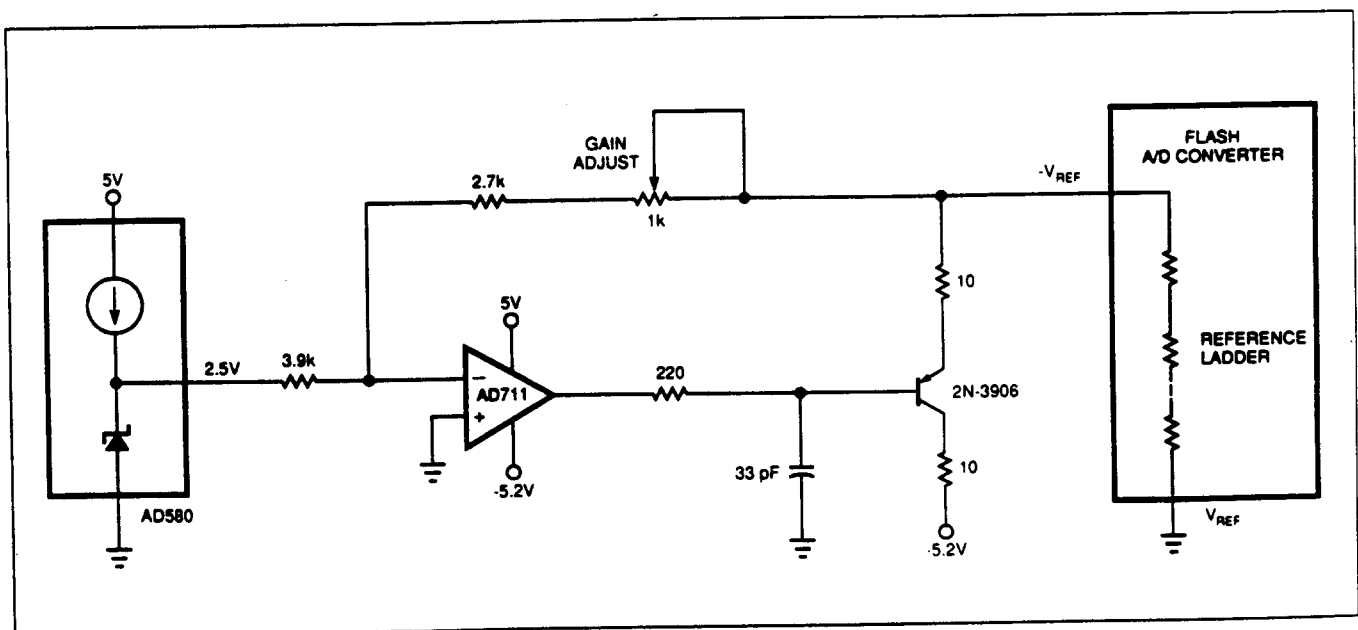


Fig 5—Flash converters don't have internal references, so you must design them externally. This particular circuit provides a stable -2V reference for a unipolar converter.

To improve flash-converter performance at sampling rates as high as 25 MHz, you can use front-end T/H amplifiers to implement a "track-and-slow-down" approach.

The maximum rate-of-change of this signal occurs at the zero-crossing point and is equal to

$$\left. \frac{dv}{dt} \right|_{\max} = 2\pi f V_P = \left. \frac{\Delta v}{\Delta t} \right|_{\max}$$

By solving this equation for Δt_{\max} , you obtain

$$\Delta t_{\max} = \frac{\Delta V}{2\pi f V_P}$$

If the input-voltage range of the flash converter is 2V, or $V_P = 1V$, then the LSB weight is 8 mV for an 8-bit ADC. For the flash converter's error to be less than 1 LSB, Δt_{\max} must equal 25 psec. The effective-sample delay mismatch between comparators can't exceed this value. If the mismatch is greater, a 50-MHz, full-scale sine-wave input will produce missing codes in the converter's output.

Placing an ideal track-and-hold (T/H) amplifier ahead of the flash converter theoretically would eliminate this problem, because the flash converter basically would

be digitizing a dc input. In actual practice, T/H amplifiers aren't ideal, especially at high speeds. The signal presented to the flash converter is still changing, although at a slower rate. Nevertheless, this "track-and-slow-down" approach can improve the flash-converter performance at sampling rates as high as approximately 25 MHz. At sampling rates above 25 MHz, the T/H circuit needs to be mounted on the same substrate as the flash converter in a suitable hybrid package. Monolithic T/H amplifiers in hybrid packages with 8-bit flash converters have successfully achieved 7 effective bits at Nyquist inputs and at sampling rates of 250 MHz. These hybrid packages exact a penalty of higher cost and power consumption, however.

You'll find it difficult to select an appropriate discrete T/H amplifier, because the interaction between the T/H amplifier and the flash ADC is hard to predict. You should evaluate key T/H amplifier specifications such as acquisition time, full-power bandwidth, slew rate, and harmonic distortion. Harmonic-distortion specifications typically are provided for the track mode. The T/H amplifier's performance may be considerably

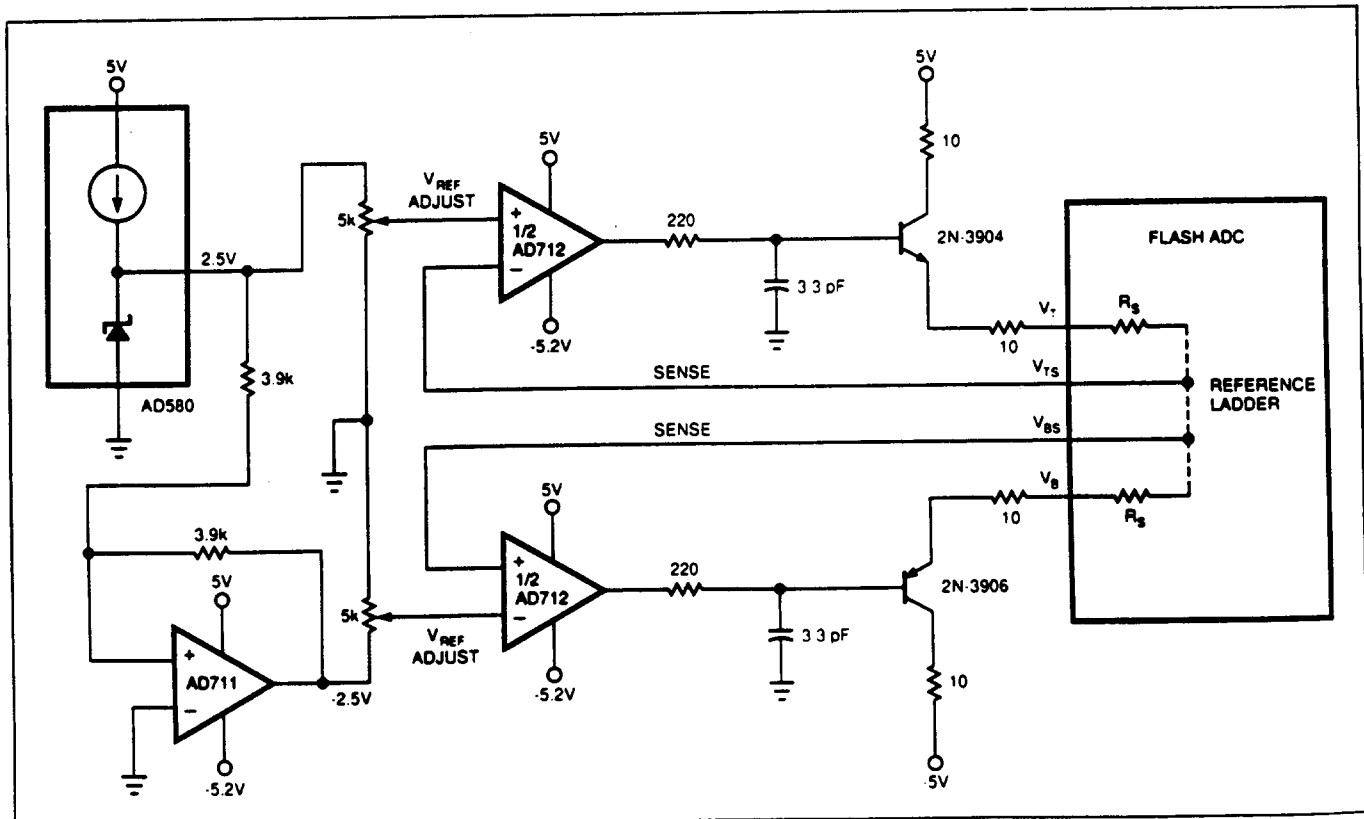


Fig 6—This reference generator for a bipolar flash converter uses the sense pins provided by the resistor ladder to compensate for the voltage drops in the package's pin and bond-wire resistances.

worse than the stated specifications when the amplifier is in the hold mode and actually driving a flash converter. Also, the loading effects of the converter may degrade the T/H amplifier's performance. In addition, obtaining the optimum relationship between the various timing pulses that drive the T/H amplifier and the flash converter may require considerable experimentation.

Because of these many difficulties, the current goal of many ADC manufacturers is either to provide flash converters whose dynamic performance is acceptable without a T/H function, or to integrate the T/H function and converter on the same chip. In either case, manufacturers can fully specify the ADC for dynamic performance and spare you the somewhat difficult design problems associated with interfacing the T/H amplifier to the converter.

The knowledge of internal ADC features and the requirements these ADCs place on external circuits should guide your initial design efforts. But once you finish the design and build your circuit, you'll want to ensure that the combined performance of your converter and its support circuits meets your requirements. Part 2 will discuss various DSP-based test methods that are particularly effective in flash-converter testing.

Reference

1. Sheingold, Dan, *Analog-Digital Conversion Handbook*, 3rd ed, Prentice-Hall, Englewood Cliffs, NJ, 1986.
-

Designer's Guide to Flash-ADC Testing**3****Part 2
DSP Test Techniques Keep Flash ADCs in Check**by **Walt Kester**

By testing your flash A/D converter, you can ensure that it's faithful to all the specifications listed on its data sheet. Part 2 of this 3-part series presents a number of methods, including sine-wave curve fitting and the FFT, that you can use to test flash converters. Readily available benchtop instruments or personal computers are the only equipment that you'll need to use these methods.

It's important to know how your flash A/D converter will perform in real-world applications. Therefore, you may want to perform any one of a variety of tests on your converter to determine its deviation from ideal performance. As Part 1 of this series discussed, flash ADCs exhibit errors due to static and dynamic nonlinearities, and these errors increase as the input signal's slew rate increases. Thus, the actual S/N ratio will fall short of the converter's theoretical value. Even if you don't apply these tests yourself, becoming familiar with them will help you evaluate data-sheet specifications more accurately, because many manufacturers use these same methods.

Another reason you may want to test your flash converter is to gain information that the manufacturer doesn't provide. Specifications such as S/N ratio and its related effective number of bits are key in all applications and are normally specified, but other specifications that are more important for your particular application may not be included on the data sheet. For

example, video designs typically require that you know a converter's differential phase and gain (Ref 1). Communications systems may even depend on esoteric specifications such as the spurious-free dynamic range, which isn't available on many data sheets.

For a full-scale sine-wave input, the theoretical rms-signal to rms-quantization noise ratio is

$$\text{S/N RATIO} = 6.02N + 1.76 \text{ dB,}$$

where N equals the number of bits (Ref 2). The rms quantization-noise voltage for an ideal ADC within the Nyquist bandwidth is $q/\sqrt{12}$, where q is the weight of the LSB expressed in volts.

The most popular method for extracting a flash converter's S/N ratio and effective number of bits is through discrete Fourier transforms (DFTs). Today, you can perform sophisticated DSP tests with PC-based test systems and standard software packages. The test system in Fig 1, for example, can execute a 1024-point FFT in less than one second. Most of the hardware you'll need is available as plug-in boards for the PC. However, you'll have to do a fair amount of work before you can begin to use a PC-based test system. First, you'll need to design a high-speed buffer-memory board to capture the data from the flash ADC. Typically, you'll need to use high-speed static CMOS or ECL RAMs. Second, plan to design an appropriate logic interface to connect this buffer memory to the digital I/O card of the PC.

Another hardware feature you might consider is an evaluation board, which certain manufacturers of video-speed ADCs supply to ease design testing. Many evaluation boards contain reference voltages, power-supply decoupling, timing circuits, output registers, and connectors. The evaluation boards usually have a

DSP test techniques determine your converter's deviation from ideal performance, and they even tell you certain specifications that the ADC's data sheet doesn't.

matching reconstruction DAC. In most cases, the manufacturer has optimized the design of these boards so that your ADC test won't be corrupted by faulty or poorly designed support circuits.

Your software must include a program to capture the data and then load it into the memory of the PC. If you plan to use FFT analysis, you must link a standard FFT software package to your test program. You may also have to generate a look-up table to store any special weighting functions required by your particular sampling scheme. Also, adding a coprocessor card will speed up the thousands of multiplications that FFT-based analysis requires.

If you don't have the time or the energy to build your own test system, consider one of the benchtop instruments available from a number of instrumentation manufacturers. These turnkey systems typically utilize a high-speed logic analyzer to capture data. Because menu-driven software allows you to select from a variety of tests, you incur practically no hardware or software development time.

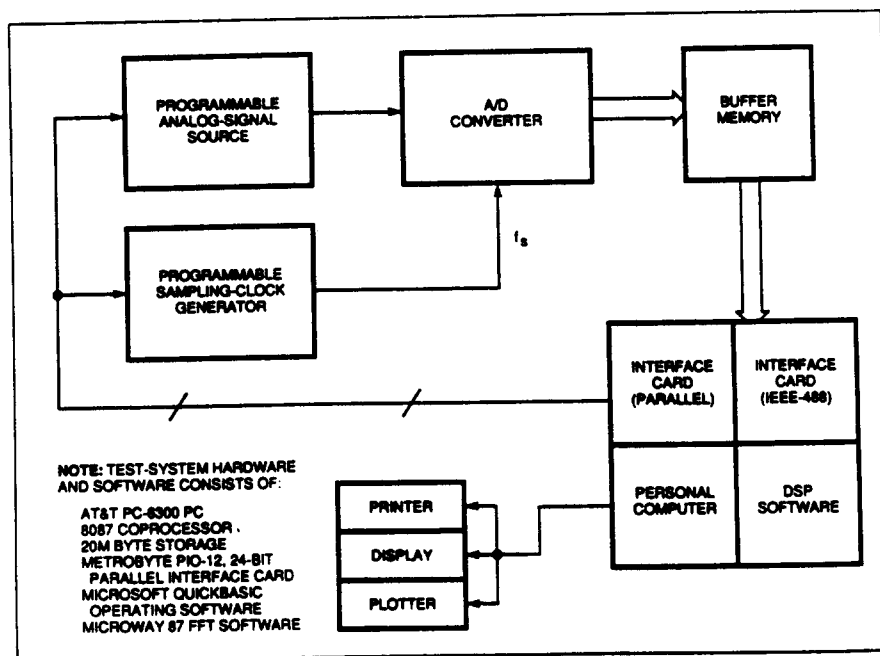
To test a flash ADC using Fourier analysis, you must apply a spectrally pure sine wave to the converter and store a number of contiguous output data samples. Then, using DFT techniques, your test program calculates the rms-signal and rms-noise content and determines the ratio of the two. Noise calculations using DFT techniques include not only the converter's quantization noise but also the harmonics of the input sine wave. In addition, harmonics that fall outside the

Nyquist bandwidth are aliased back into the Nyquist bandwidth because of the sampling process. Thus, to achieve accurate and repeatable results, the purity of the sampling clock and the input sine wave is critical.

You can use either coherent or noncoherent sampling to evaluate the ADC performance. Coherent sampling simply means that your record of samples contains an integer number of sine-wave input cycles. Alternatively, noncoherent sampling produces a record that contains noninteger multiples of the input. You must choose between these sampling schemes based on the type of input data you expect. Coherent testing is more suited to a laboratory environment when you know the precise frequency content of an input signal, and it requires careful attention in the selection of the input and sampling frequencies. Noncoherent testing yields a better representation of ADC performance in a real-world application such as spectral analysis, because the precise frequency content of the signal being digitized is a mystery.

However, whenever the number of time samples doesn't contain an integer number of input cycles (noncoherent testing), you'll have to time-weight the samples to reduce frequency side lobes. Without weighting, discontinuities will cause the main lobe's energy—the fundamental—to leak into many other frequency bins. The term "bins" refers to the spaces between spectral lines or spectral peaks. The number of bins for a particular spectrum equals the sampling fre-

Fig. 1—This DSP test system for a flash ADC can execute a 1024-point FFT in less than one second, but the system requires a significant design effort. Hardware requirements include a high-speed buffer memory and logic interface between this memory and your PC. Software requirements include a program to capture the data and load it into the memory of the PC, as well as a link between your test program and a standard FFT software package.



quency divided by the record length, or f_s/M . The leakage of the signal from the central bin to side-lobe bins makes accurate spectral measurements impossible—you simply can't distinguish the frequency bins that contain actual signal information from those that contain noise. Another reason to time-weight the samples is that the end user of your A/D-conversion system may be interested in the performance of the ADC using an identical or similar window.

Noncoherent sampling involves fewer input- and sampling-frequency restrictions than coherent sampling does, but it requires careful attention in the selection and use of the weighting function. Also, to prevent masking out harmonics of the fundamental, avoid using inputs that are integer submultiples of the sampling frequency. If your input frequency is an integer submultiple of the sampling frequency, the quantization noise, $q/\sqrt{12}$, will be concentrated in the harmonics of the input frequency rather than uniformly distrib-

uted across the Nyquist bandwidth. Ultimately, this condition leads to incorrect harmonic-distortion test results.

One popular weighting function is the Hanning window (Ref 3), which is described by the equation

$$W_n = 0.5 - 0.5 \cos\left(\frac{2\pi n}{M}\right),$$

where W_n is the weighting coefficient for the n th data sample, and M is the total number of samples. Fig 2 graphically depicts the Hanning window in both the time and the frequency domains.

To calculate the S/N ratio, you have to decide the number of frequency bins to include in the fundamental and the number of bins to consider as noise. As Fig 2 shows, you can correlate the amount of side-lobe attenuation with the lobe's distance, in terms of bins, from the fundamental bin. Fig 2b includes a table that lists some of these values. You'll have to make your decision based on the theoretical S/N ratio of the converter you're testing.

For example, an 8-bit converter has a theoretical, maximum S/N ratio of approximately 50 dB. In order to ensure that the side-lobe energy doesn't cause an artificially high noise measurement (and hence an artificially low S/N-ratio measurement), you should include at least 10 frequency bins on either side of the fundamental when calculating the signal level. (Simply take the square root of the sum of the squares of all 21 bins as your signal level.) Now, any side-lobe energy outside this region will be at least 68 dB below the fundamental signal level (18 dB below the theoretical, 8-bit quantization noise floor of 50 dB), and side-lobe leakage won't significantly affect the accuracy of your S/N-ratio measurement.

Other weighting functions may better suit your application. For example, Fig 3 compares the popular Hanning window's spectral representation with the more sophisticated, minimum 4-term, Blackman-Harris type. For the same record length, the Blackman-Harris window provides better spectral resolution than the Hanning window, making it more suitable for critical spectral analysis, such as measuring 2-tone, third-order intermodulation-distortion products. The extra computations for the Blackman-Harris window don't lengthen processing time, because you calculate them only once and store them in a look-up table.

As previously stated, you can use coherent sampling if you know the characteristics of your input signal and if you choose the sampling rate accordingly. Coherent




Fig 2—When the record of samples doesn't contain an integer number of input cycles—that is, when you're using noncoherent sampling—you must precondition the data with a weighting function. The Hanning window shown here in the time domain (a) and the sampled-frequency domain (b) is a popular weighting function.

Today, you can perform sophisticated DSP tests with PC-based test systems and standard software packages.

sampling eliminates leakage and the need for windowing (Ref 4); the spectral result of a coherently sampled signal is simply a single-frequency peak. Certain restrictions apply to the choice of the sampling rate and the sine-wave frequency, however. First, you must observe the following ratio:

$$\frac{f_m}{f_s} = \frac{M_C}{M}$$

M_C equals the number of integer cycles of the sine wave during the record period. For a whole number of cycles, M_C must be an integer. To ensure that you don't take repetitive data, M_C should also be a prime number: 1, 3, 5, 7, 11, 13, 17, etc. By using prime numbers, you ensure that all samples during the record period are unique. When using coherent sampling, it's mandatory that the ratio M_C/M be constant. This requirement implies that you derive f_s and f_m from two locked frequency synthesizers.

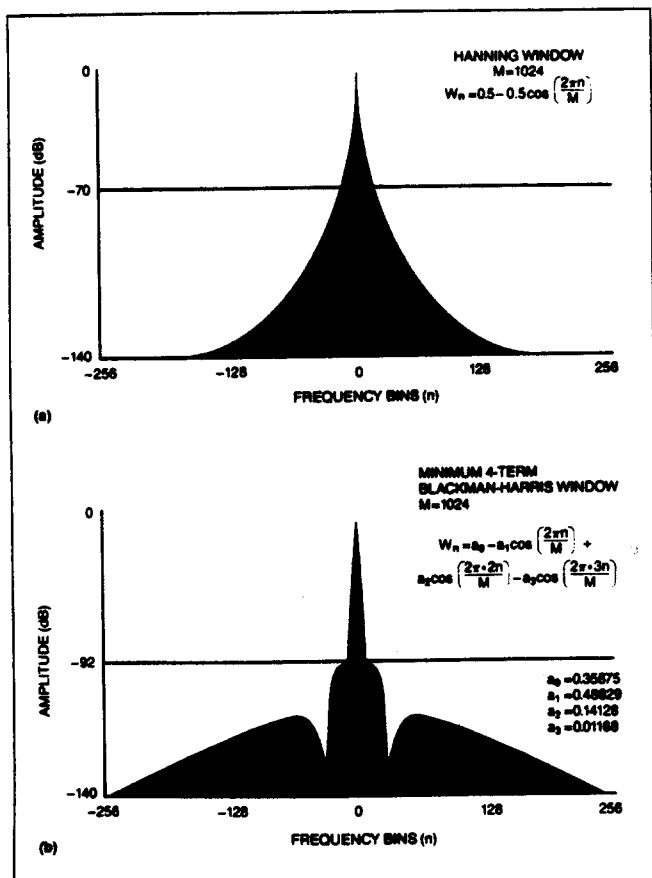


Fig 3—The Blackman-Harris windowing function (b) resolves closer peaks in a frequency spectrum than does the Hanning window (a). The mathematical expression for the Blackman-Harris window is more complex, but you only need to calculate the terms once and then store them in a look-up table.

Calculate the DFT

After selecting the record length and determining the weighting function (for noncoherent sampling), you must write your DFT test program. Your program must find the DFT of the sequence of weighted data samples for $M/2$ frequencies (the Nyquist frequency). Thus, the program should solve the following two equations for the k th frequency:

$$A_k = \frac{1}{M} \sum_{n=1}^M W_n D_n \cos \left[\frac{2\pi k(n-1)}{M} \right]$$

$$B_k = \frac{1}{M} \sum_{n=1}^M W_n D_n \sin \left[\frac{2\pi k(n-1)}{M} \right]$$

In these equations, A_k and B_k represent the magnitudes of the cosine and sine parts of the k th spectral line; n is the number of time samples; W_n is the weighting function; D_n is the amplitude of the time-function data point; and k is the number of a spectral line. The total magnitude of the k th spectral line is

$$\text{MAGNITUDE}_k = \sqrt{A_k^2 + B_k^2}$$

The program's results yield $M/2$ components, which are the frequency-domain representation of the M time samples. The resolution or spacing between the spectral lines, Δf , equals f_s/M and is the bin size or bin width.

Typically, you should select the number of time samples (M) to be between 256 and 4096, depending on the desired resolution and the size of the buffer memory. M must be equal to an integer that's a power of two. If you're using noncoherent sampling, you can compress leakage around the main lobe by using a larger record length, thereby leaving a larger percentage of the Nyquist spectrum uncontaminated. For example, the Hanning weighting leakage is ± 10 bins from the fundamental for 68-dB side-lobe suppression. If the record length is 256, then the leaky fundamental occupies 20 bins out of 128 spectral components, or 16% of the digital spectrum. When M equals 1024, the percentage reduces to 20/512, or 4%.

In practice, you can use one of the many FFT algorithms to simplify and speed the DFT calculations (Ref 5). An FFT algorithm will produce the same results as the DFT equations above, and the computation time is much faster.

The discrete Fourier transform is the most popular method for determining a converter's true S/N ratio and effective number of bits.

Verify the FFT

Consider the noise floor when verifying the FFT. Assuming that the round-off error contributed by the DSP-noise calculation (the error caused by using a finite number of bits in the FFT multiplications and additions) is negligible, the rms-signal to rms-noise level in a single frequency bin of width Δf is

$$\text{S/N RATIO}_{\text{FFT}} = 6.02N + 1.76 \text{ dB} + 10 \text{ LOG}_{10} \left(\frac{M}{2} \right).$$

This equation represents the FFT noise floor. You should choose M so that any spurious components you want to resolve lie at least 10 dB above this floor.

Basic software can easily generate an ideal N -bit sine wave by using the Integer (INT) function to truncate the value to the proper resolution. For instance, an input signal of frequency f_{in} is equal to

$$V_q = V_o \sin \left(\frac{2\pi n f_{in}}{f_s} \right),$$

where n is the n th time sample for an ADC that has infinite resolution. You can calculate the corresponding quantized value using

$$V_{q(n)} = \text{INT} \left(\frac{V_o \sin \left(\frac{2\pi n f_{in}}{f_s} \right)}{q} \right),$$

where $q = 2V_o/2^N$. Substituting this expression for q in the above equation yields

$$V_{q(n)} = \text{INT} \left[2^{N-1} \sin \left(\frac{2\pi n f_{in}}{f_s} \right) \right].$$

The INT function simply truncates the fractional portion of $V_{q(n)}$.

To check the dynamic range of the FFT, calculate the S/N ratio by using $6.02N + 1.76$ dB for increasing values of N and observing the point at which the S/N ratio no longer increases by 6.02 dB/bit. The sine-wave input to the weighting function and the FFT are more ideal as N approaches infinity. By making N arbitrarily large, you can greatly reduce quantization-error effects and analyze the true noise floor of the FFT. You can also examine the characteristics of the weighting function.

Match the sine wave to a curve

Another test method to use with flash ADCs is sine-wave curve fitting. You perform this test after the ADC digitizes the sine wave and after your test system stores the data in its memory. A record length of 1024 samples is usually sufficient. The software then calcu-

lates the best-fit, ideal N -bit sine wave to match the data points, based on the sine wave's amplitude, offset, frequency, and phase required to minimize the rms error between the actual and the ideal sine wave (Refs 6 and 7). This method also requires that the input sine-wave frequency contains no subharmonics of the sampling rate. If you know the precise sine-wave frequency, the curve-fit algorithm is much simpler than the FFT method, and the probability that the algorithm will converge is higher.

After the software computes the rms error, Q_A , between the ideal sine wave and the actual sine wave, you can calculate the effective number of bits by using

$$\text{EFFECTIVE NUMBER OF BITS} = N - \text{LOG}_2 \left(\frac{Q_A}{Q_T} \right),$$

where Q_T is the theoretical rms quantization error, $q/\sqrt{12}$. This measurement includes errors due to differential nonlinearity, integral nonlinearity, missing codes, aperture jitter, and noise, in addition to the quantization noise.

The effective number of bits that you calculate using the sine-wave curve-fitting method correlates with the value of the full-scale, FFT S/N-ratio measurement obtained using the equation

$$\text{EFFECTIVE NUMBER OF BITS} = \frac{\text{S/N RATIO}_{\text{ACTUAL}} - 1.76 \text{ dB}}{6.02}.$$

However, if you measure the effective bits of a sine-wave input signal whose amplitude is less than full scale, you must include the following correction factor in the above equation to achieve correlation between the two methods:

$$\text{EFFECTIVE NUMBER OF BITS} = \frac{\text{S/N RATIO}_{\text{ACTUAL}} - 1.76 \text{ dB}}{6.02}$$

$$+ \frac{\text{LEVEL OF SIGNAL BELOW FULL SCALE (dB)}}{6.02}$$

One useful method for reducing the effects of the D/A converter in making gross back-to-back measurements on an ADC is the beat-frequency method. Fig 4 illustrates a basic test setup. This test method stresses the converter with a near-Nyquist signal and drives the converter at its maximum sampling rate. Thus, the analog-input sine wave should be slightly lower in frequency than half the sampling frequency. The test system updates the registers that drive the DAC at an even submultiple of the sampling rate, f_s/N ,

Coherent testing is more suited to a laboratory environment; noncoherent testing more closely represents ADC performance in the real world.

where N is a power of 2. (N is not the ADC's resolution.) The resulting signal from the DAC is a low-frequency sine wave whose exact frequency equals the difference between half the sampling rate and the analog-input frequency. As Fig 4 shows, you should clock the DAC at a much lower rate, f_s/N —known as the decimation rate—thereby reducing the effects of glitches and other dynamic errors.

You can use the beat-frequency method to make signal-to-noise measurements over the Nyquist bandwidth, $f_s/2N$. You also can examine the low-frequency beat on an oscilloscope for missing codes and other nonlinearities. To measure the harmonic content of the beat frequency, you can use a low-frequency spectrum analyzer. The harmonics of the low-frequency beat are directly related to the harmonics of the analog-input frequency. A beat frequency of a few hundred kilohertz works well. To prevent jitter on the low-frequency beat signal, you must derive both the analog-input sine wave and the sampling frequency from frequency synthesizers or crystal oscillators.

This beat-frequency test is also effective in measuring the flash converter's performance for input signals near the sampling frequency. The performance under these conditions is useful for radar in-phase and quadrature-phase systems and in IF-to-digital conversion. To perform this test, set the ADC's analog-input frequency to slightly less than the sampling rate. The circuit generates a low-frequency beat even if the DAC updates at the sampling rate. However, updating the DAC at f_s/N reduces the effects of DAC dynamic errors on the measurements.

You can use DSP techniques and FFTs to analyze Fig 4's decimated data for a wide range of input fre-

quencies. You do have to remember the rules of aliasing, however, to know where to expect the fundamental signal to show up in the FFT output spectrum. You may think your FFT is sampling your signal at a rate of f_s/N , but the converter is actually sampling at a rate of f_s .

Once you understand how to use these various techniques, you can start to probe your particular converter to measure its real performance. Part 3 will discuss how you apply these techniques to actually test an ADC in your system and determine a number of static and dynamic specifications.

References

1. Kester, W A, "PCM Signal Codecs for Video Applications," *SMPTE Journal*, No. 88, November 1979, pg 770.
2. Bennett, W R, "Spectra of Quantized Signals," *Bell System Technical Journal*, No. 27, July 1948, pg 446.
3. Harris, Frederic J, "On the use of windows for harmonic analysis with the discrete Fourier transform," *IEEE Proceedings*, Vol 66, No. 1, January 1978, pg 51.
4. Coleman, Brendan, Pat Meehan, John Reidy, and Pat Weeks, "Coherent sampling helps when specifying DSP A/D converters," *EDN*, October 15, 1987, pg 145.
5. Ramirez, Robert W, *The FFT: Fundamentals and Concepts*, Prentice Hall, Englewood Cliffs, NJ, 1985.
6. Peetz, B E, A S Muto, and J M Neil, "Measuring Waveform Recorder Performance," *HP Journal*, Vol 33, No. 11, November 1982, pg 21.
7. Frohring, B J, B E Peetz, M A Unkrich, and S C Bird, "Waveform Recorder Design for Dynamic Performance," *HP Journal*, Vol 39, February 1988, pg 39.
8. "Dynamic performance testing of A to D converters," HP Product Note 5180A-2, 1982.

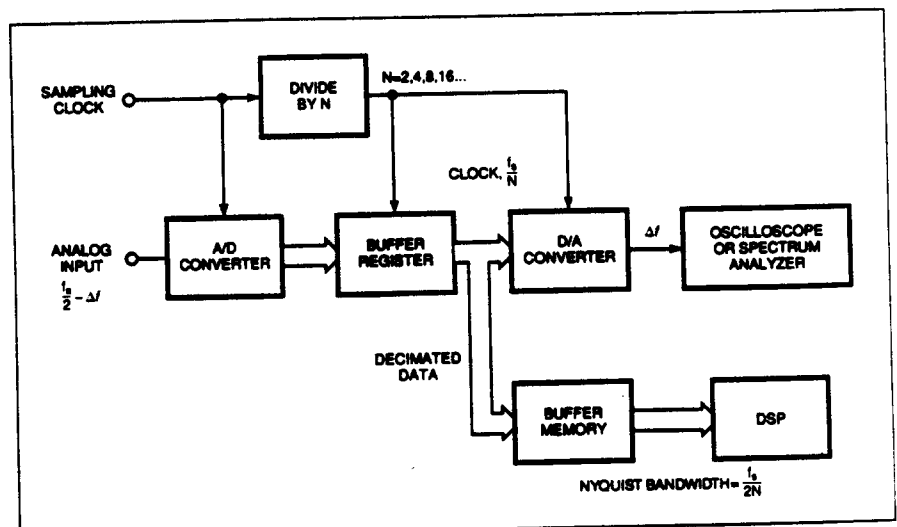


Fig 4—The beat-frequency test stresses the converter with a near-Nyquist input signal, and it drives the converter at its maximum sampling rate. You can then examine the low-frequency beat on an oscilloscope and search for missing codes and other nonlinearities.

Designer's Guide to Flash-ADC Testing**3****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

$$\text{S/N ratio} = 20 \log(\text{rms signal level}/\text{rms noise level}).$$

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:

$$\text{Harmonic distortion} = 20 \log(\text{rms signal level}/\text{rms harmonic level}).$$

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

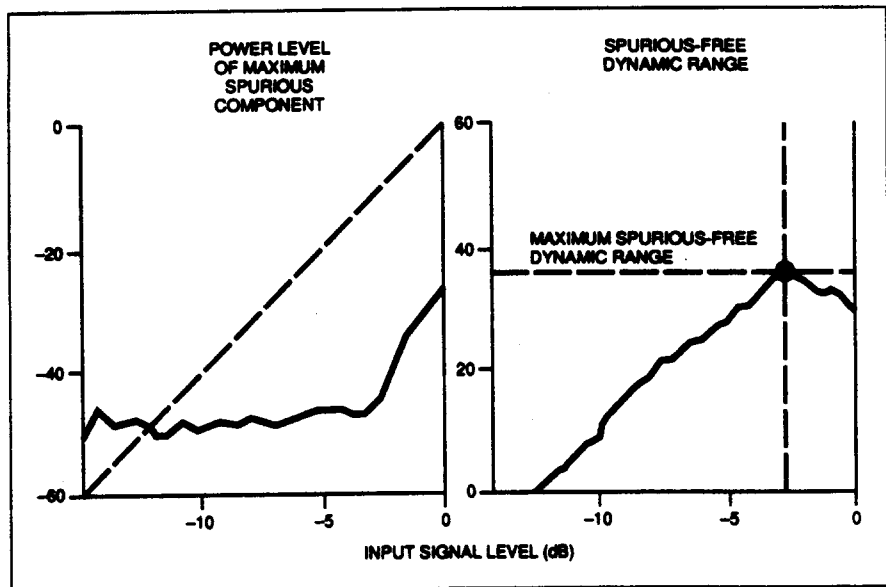


Fig 1—These dynamic-range plots show the power levels of spurious frequencies and the maximum spurious-free dynamic range. In this example, the maximum spurious-free dynamic range occurs at an input signal level that's 3 dB below full scale.

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 *n*th 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 non-linearity.

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 = $1/4 V_0$, where $\pm V_0$ 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

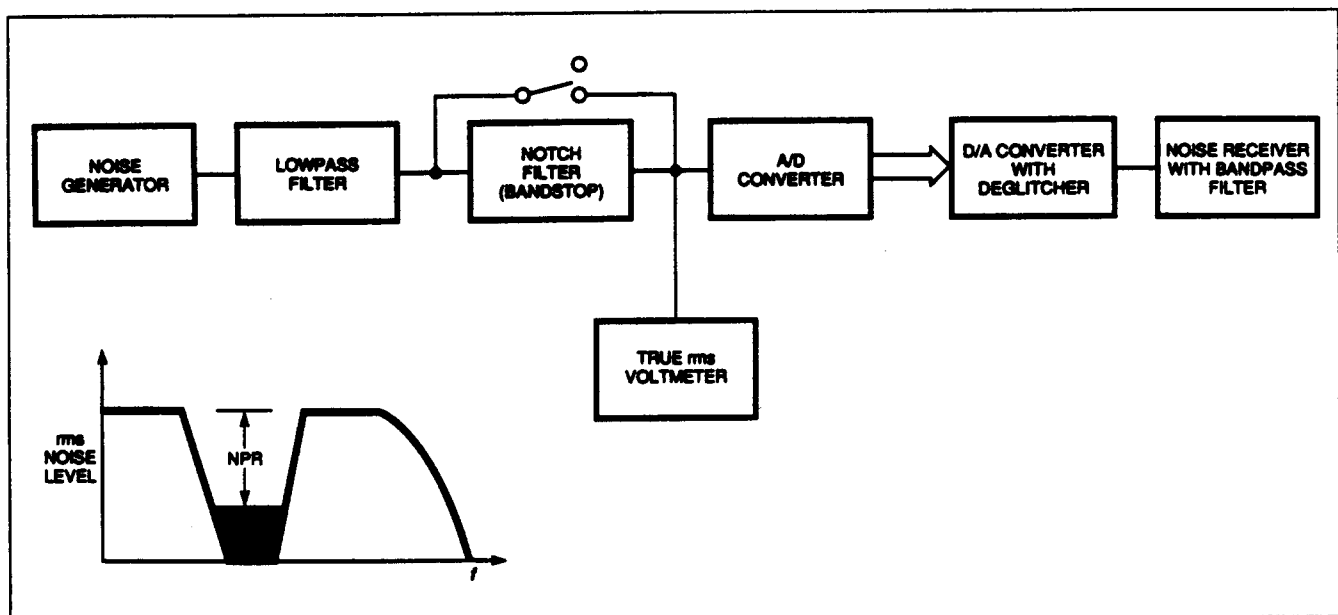


Fig 3—You can use this test setup to measure noise power ratio (NPR). With the notch filter out, the receiver determines the noise power of the signal inside the notch. With the notch filter switched in, the receiver measures the residual noise inside the typical 4-kHz slot. The ratio of the two readings (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 $E_a = t_a \cdot dV/dt$, 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 non-linear 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

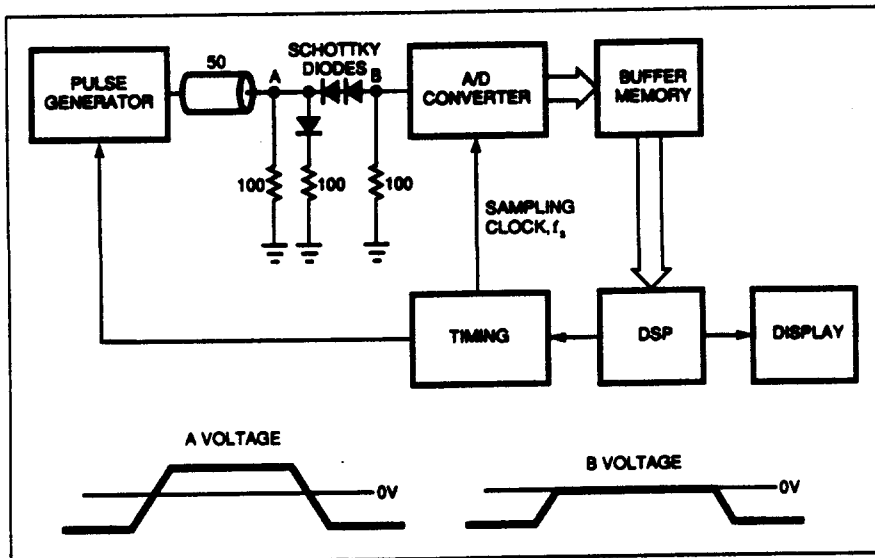


Fig 4—This test setup measures the transient response of an A/D converter. The Schottky-diode network, located between points A and B in the circuit, generates a flat pulse for the input of the converter.

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 $+125^{\circ}\text{C}$.

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 $E_a = t_a \cdot dV/dt$.

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 $2t_a$.

If the corresponding maximum voltage error (ΔV) at the zero crossing of a full-scale sine wave is set to $\frac{1}{2}$ LSB ($\frac{1}{2}$ 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 $\frac{1}{2}$ LSB aperture error, by using the following equations:

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

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

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

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

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:

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

$$\frac{dV}{dt} = 2\pi fV_0 \cdot \cos(2\pi ft), \text{ and}$$

$$\frac{dV}{dt}_{\text{rms}} = \frac{2\pi fV_0}{\sqrt{2}}.$$

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

$$\frac{\Delta V_{\text{rms}}}{t_a} = \frac{2\pi fV_0}{\sqrt{2}}, \text{ and}$$

$$\Delta V_{\text{rms}} = \frac{2\pi fV_0 t_a}{\sqrt{2}}.$$

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

$$\begin{aligned} \text{S/N ratio} &= 20 \log \left[\frac{V_0/\sqrt{2}}{\Delta V_{\text{rms}}} \right] \\ &= 20 \log \left[\frac{1}{2\pi f t_a} \right] \text{ dB.} \end{aligned}$$

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-




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

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.

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

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

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.

tribution 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.

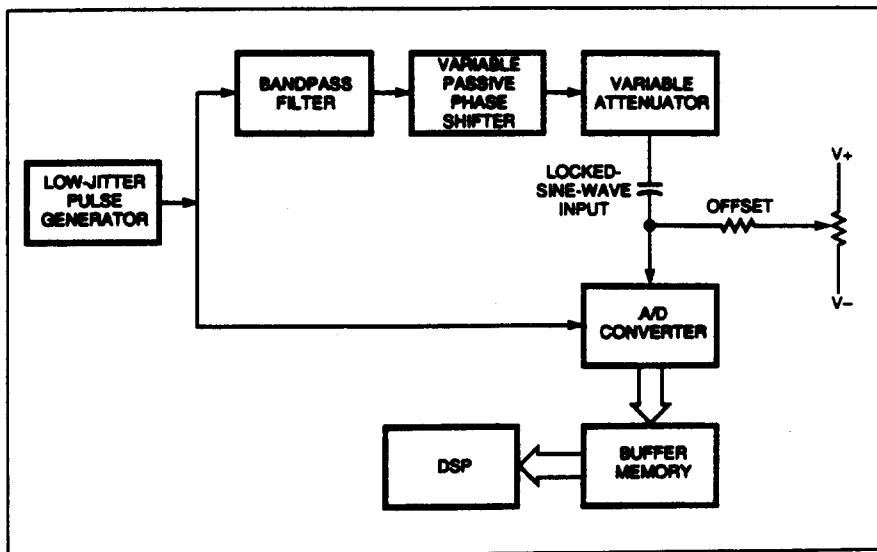
Use the circuit in Fig 8 to measure the error rate for low-frequency input signals. Apply a low-frequency, full-scale sine wave (or triangle wave) to the A/D converter so that its rate of change is less than 1 LSB/sample. This step ensures that the transition zones between codes are all adequately exercised. 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 $10 \times 256 = 2560$ samples. The frequency of the triangle wave is

$$f_t = \frac{1}{2560 \cdot 2 \cdot 10 \text{ nsec}} = 19.5 \text{ kHz.}$$

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.

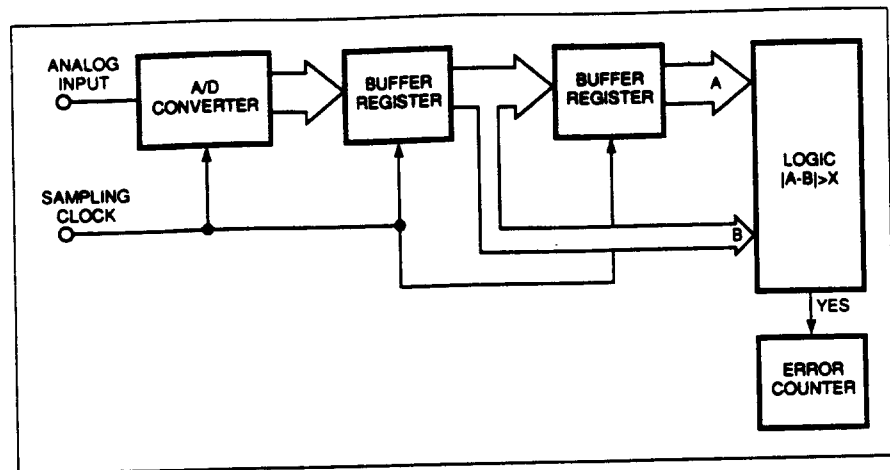


Fig 8—The effective aperture delay is the time difference between the leading edge of the sampling-clock pulse and the actual zero crossing of the sine-wave input.

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.