ANALOG DEVICES

AN-215B
APPLICATION NOTE

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

Designer's Guide to Flash-ADC Testing

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

Reprinted from EDN – January 18, 1990

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-

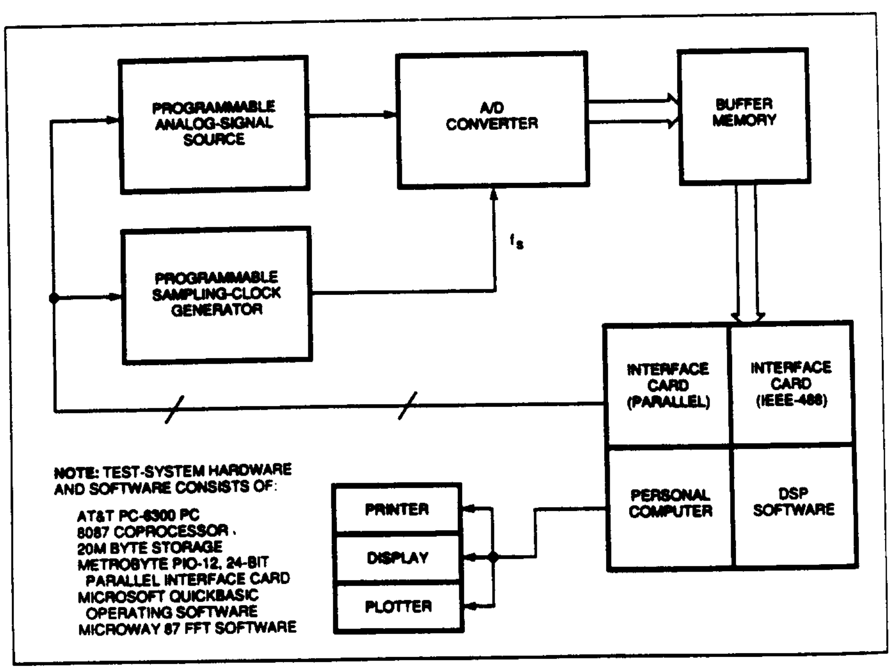

Fig. 1—This DSP test system for a flash ADC can execute a 1024-point FFT in less than one second, but the system requires a significant design effort. Hardware requirements include a high-speed buffer memory and logic interface between this memory and your PC. Software requirements include a program 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 distributed across the Nyquist bandwidth. Ultimately, this condition leads to incorrect harmonic-distortion test results.

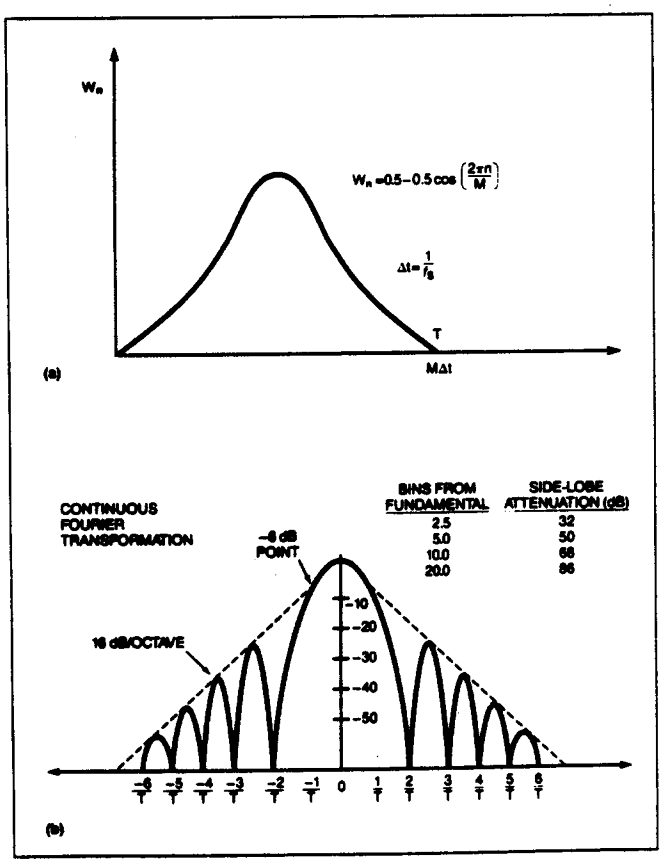

Fig 2—When the record of samples doesn't contain an integer number of input cycles—that is, when you're using noncoherent sampling—you must precondition the data with a weighting function. The Hanning window shown here in the time domain (a) and the sampled-frequency domain (b) is a popular weighting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

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

3

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

where W_n is the weighting coefficient for the nth 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

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:

$$\frac{f_{in}}{f_S}=\frac{M_C}{M}.$$

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_{in} from two locked frequency synthesizers.

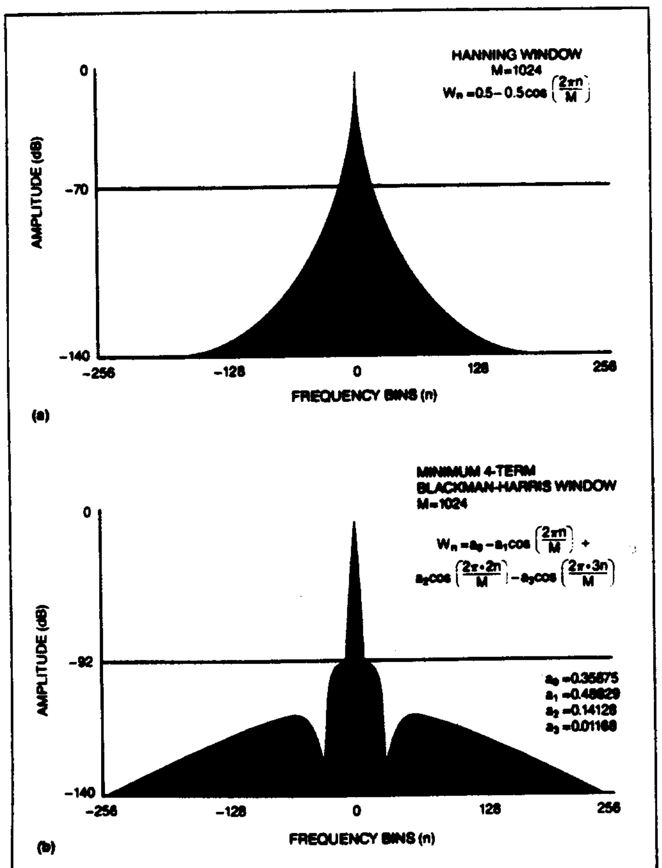

***Fig 3—The Blackman-Harris windowing function** (b) resolves closer peaks in a frequency spectrum than does the Hanning window (a). The mathematical expression for the Blackman-Harris window is more complex, but you only need to calculate the terms once and then store them in a look-up table.*

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 kth frequency:

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

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

In these equations, A_k and B_k represent the magnitudes of the cosine and sine parts of the kth 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 kth 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

$$V_q = V_0 \sin\left(\frac{2\pi n f_{in}}{f_S}\right),$$

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

$$V_q(n) = \text{INT}\left(\frac{V_0 \sin \frac{2\pi f_{in}}{f_S}}{q}\right),$$

where $q = 2V_0/2^N$. 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 calculates 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,

3

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 frequencies. 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.

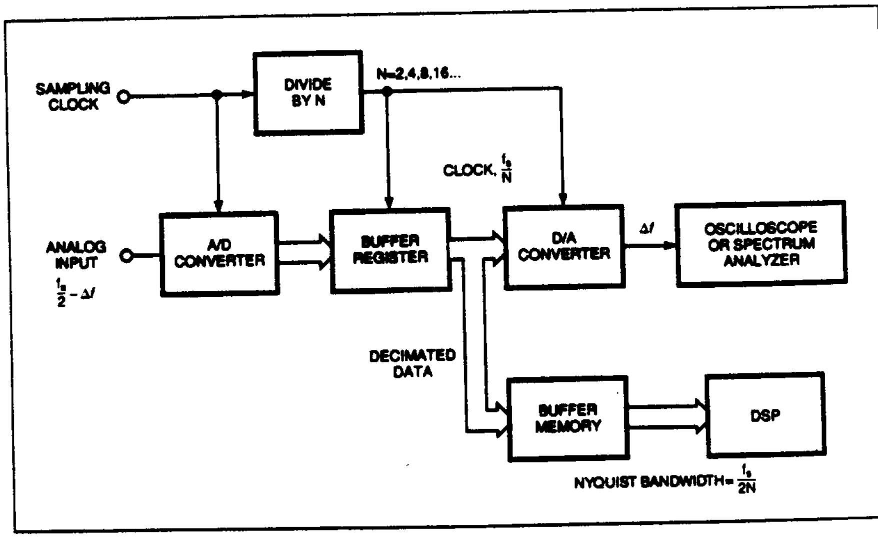

Fig 4—The beat-frequency test stresses the converter with a near-Nyquist input signal, *and it drives the converter at its maximum sampling rate. You can then examine the low-frequency beat on an oscilloscope and search for missing codes and other nonlinearities.*