

DYNAMIC TESTS FOR A/D CONVERTER PERFORMANCE

This article describes useful theory and techniques for evaluating the dynamic performance of A/D converters. Four techniques are discussed: (1) beat frequency, (2) histogram analysis, (3) sine wave curve fitting, and (4) discrete finite Fourier transform.

The key to confidence in the quality of a waveform recorder is assurance that the analog-to-digital converter (ADC) encodes the signal without degrading it. Dynamic tests that cover the frequency range over which the converter is expected to operate can provide that assurance. The results of the dynamic tests give the user a model of resolution versus frequency for the recorder. More elaborate models of failure mechanisms can be obtained by varying the conditions of the tests.

All of the dynamic tests used for the 5180A Waveform Recorder use sine waves as stimulus. Sine waves were chosen primarily because they are the easiest to generate in practice at the frequencies of interest with adequate fidelity. While it may be possible to generate a square wave, for example, whose function is known to the 10-bit resolution of the 5180A, no square wave generators exist that can guarantee the same waveshape to 10-bit resolution at 10MHz from unit to unit. Another motivation for choosing a sine wave stimulus is the simple mathematical model a sine function provides for analysis. This benefit greatly simplifies the algorithms used for data analysis.

Four dynamic tests for waveform recorder characterizations are presented here: beat frequency testing⁽¹⁾ histogram analysis⁽²⁾ sine wave curve fitting,^(3,4) and discrete finite Fourier transform.⁽⁵⁾ The last three tests operate in the same way. A sine wave source is supplied to the waveform recorder and one or more records of data are taken. A computer is then used to analyze the data. The tests differ primarily in the analysis algorithms and consequently in the sort of errors brought to light. Critical to the success of these tests is the purity of the sine wave source. Synthesized sources are necessary to provide the short-term and long-term stability required by the dynamic range of the ADC. Passive filters (a six-pole elliptical filter is used for 5180A tests) are required to eliminate harmonic distortion from the source.

These tests provide the most stressful conditions for the ADC with the input signal amplitude at full scale. Generally speaking, nonlinear effects increase more quickly than the signal level increases because of the nonideal large-signal DC behavior of the ADC components and the higher slew rates large amplitudes imply.

BEAT FREQUENCY TESTING

The beat frequency and envelope tests are qualitative tests that provide a quick, simple visual demonstration of ADC dynamic failures. An input frequency is selected that provides worst-case range changes and maximal input slew rates that the ADC is expected to see in use. The output is then viewed on a display in real time.

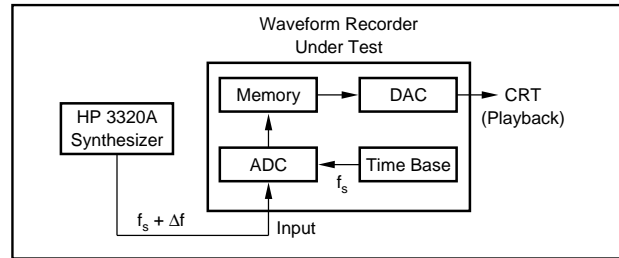


FIGURE 1. Beat Frequency Test Setup.




FIGURE 2. When the Input Frequency is Close to the Sample Rate f_s , the Encoded Result is Aliased to the Difference or Beat Frequency, Δf .

The name “beat frequency” describes the reasoning behind the test. The input sinusoid is chosen to be a multiple of the sample frequency plus a small incremental frequency (Figure 1). Successive samples of the waveform step slowly through the sine wave as a function of the small difference or beat frequency (Figure 2). Ideally, the multiplicative properties of sampling would yield a sine wave of the beat frequency displayed on the waveform recorder’s CRT. Errors can be seen as deviations from a smooth sine function. Missing codes, for example, appear as local discontinuities in the sine wave. The oversize codes that accompany missing codes are seen as widening in the individual codes appearing on the sine wave. By choosing an arbitrarily low beat frequency, a slow accurate DAC may be used for viewing the test output. For best results, the upper limit on the beat frequency choice is set by the speed with which the beat frequency walks through the codes. It is desirable to have one or more successive samples at each code. This

alleviates the settling constraint on the DAC and ensures that the display covers all possible code failures. For a 20MHz sample rate and a 10-bit ADC, this implies a 3kHz maximum beat frequency for a minimum of one sample per code bin.

Although the usual input frequency for a beat frequency test is near the sample rate, the analog bandwidth of the ADC may be measured by setting the carrier to a number of different multiples of the sample rate. The band limit is observed as a rolloff in amplitude as the carrier frequency is increased.

The envelope test differs from the beat frequency test in the choice of input frequency that the ADC encodes. Instead of a multiple of the sample frequency, an input frequency near one-half the sample rate is used. Now the ideal output is two out-of-phase sine waves at the beat frequency (Figure 3). This means that successive samples can be at the extreme ends of the ADC range, which is useful for examining slew problems that might not appear when successive samples are at adjacent codes. To avoid placing the same stress on the DAC used for display, a bank of D flip-flops removes every other sample before the data arrives at the DAC. Thus only one phase of the beat frequency remains.




FIGURE 3. When the Input Frequency is Near One-half the Sample Rate, the Envelope of the Difference Frequency Results.

Figure 4 shows 5180A beat frequency test results for a 10.0031MHz input sine wave sampled at 10MHz. For comparison, Figure 5 shows a 10.0031MHz sine wave being sampled at 10MHz by a commercially available 8-bit, 20MHz ADC.




FIGURE 4. A Beat Frequency Display Produced by the 5180A Waveform Recorder with a 10.0031MHz Input Frequency and a 10MHz Sample Rate. The smooth sine wave indicates freedom from dynamic errors.




FIGURE 5. A Beat Frequency Display for a Commercially Available 10MHz, 8-Bit ADC with a 10.0031MHz Input.

HISTOGRAM TESTING

A sine-wave-based histogram test provides both a localized error description and some global descriptions of the ADC. Using the histogram test, it is possible to obtain the differential nonlinearity of the ADC, to see whether any missing codes exist at the test frequency, and to get a measure of gain and offset at the test frequency. Of the sine-wave-based tests presented here, the histogram test yields the best information about individual code bin size at an arbitrary frequency.

A statistically significant number of samples of the input sinusoid are taken and stored as a record (Figure 6). The frequency of code occurrence in the record is then plotted as a function of code. For an ideal ADC, the shape of the plot would be the probability density function (PDF) of a sine wave (Figure 7) provided that the input and sample frequencies are relatively independent. The PDF of a sine wave is given by:

$$p(V) = \frac{1}{\pi\sqrt{A^2 - V^2}}$$

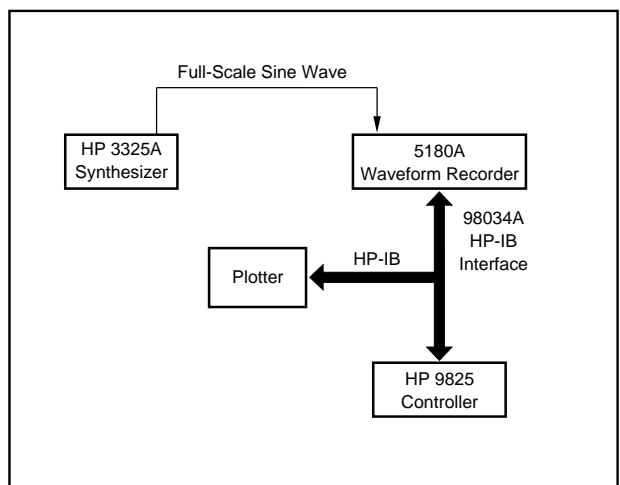


FIGURE 6. Setup for Histogram Test.




FIGURE 7. Sine Wave Probability Density Function.

Where A is the sine wave amplitude and V is the independent variable (voltage). For a real ADC, fewer than the expected number of occurrences for a given code bin indicates that the effective code bin width is smaller than ideal at the input frequency.⁽¹⁾ No occurrences indicate that the code bin width is zero for that input. A greater-than-expected number of occurrences implies a larger-than-ideal code bin width.

What is a statistically significant number of samples? We can determine significance from probability theory. For a given input PDF and record size, each bin of an ideal ADC has an expected number of occurrences and a standard deviation around that expectation. The confidence that the number of occurrences is close to the expectation is equal to the probability that the occurrences fall within the appropriate number of standard deviations. The ratio of the standard deviation to the expectation (and thus the error for a given confidence) decreases with more samples. To get the confidence for the entire range, the probabilities for all codes lying within the desired error are multiplied together.

For an ideal 10-bit ADC, 100,000 samples would give us a 12% confidence that the peak deviation from the input PDF is less than 0.3LSB and a 99.9% confidence that the peak deviation is less than 0.5LSB. The notion of confidence relies on the input's being a random process. We can model the sine wave input as random process only if the input and sample frequencies are relatively independent.

The specification of greatest interest that can be calculated using the histogram test is differential nonlinearity. Differential nonlinearity is a measure of how each code bin varies in size with respect to the ideal:

NOTE: (1) Histogram testing can be thought of as a process of sampling and digitizing the input signal and sorting the digitized samples into bins. Each bin represents a single output code and collects samples whose values fall in a specific range. The number of occurrences or samples collected in each bin varies according to the input signal. If N is the number of ADC bits, there are 2^N bins. Ideally, if B is the full-scale range of the ADC in volts, each bin corresponds to a range of sample sizes covering $B/2^N$ volts. In a real ADC, the bins may not all have the same width.

$$\text{Differential Nonlinearity} = \frac{\text{actual } P(\text{nth code})}{\text{ideal } P(\text{nth code})} - 1$$

Where actual $P(\text{nth code})$ is the measured probability of occurrence for code bin n , and ideal $P(\text{nth code})$ is the ideal probability of occurrence for code bin n . The code number n goes from 1 to 2^N , where N is the number of ADC bits. Using the probability of occurrence eliminates dependence on the number of samples taken. To calculate the probability for each code in the actual data record, the number of occurrences for each code is divided by the number of samples in the record. The ideal probability of occurrence is what an ideal ADC would generate with a sine wave input. For each code bin, this is the integral of the probability density function of a sine wave over the bin:

$$P(n) = \frac{1}{\pi} \left[\sin^{-1} \left(\frac{B(n - 2^{N-1})}{A2^N} \right) - \sin^{-1} \left(\frac{B(n - 1 - 2^{N-1})}{A2^N} \right) \right]$$

Where n is the code bin number, B is the full-scale range of the ADC, and N is the number of ADC bits. To avoid large differences in code probability caused by the sinusoid cusp, a sine wave amplitude A is chosen that slightly overdrives the ADC.

A judicious choice of frequency for the input sinusoid in this test is necessary for realistic test results. An input frequency that is a submultiple of the sample frequency violates the relative independence criterion and will result in sampling of the same few codes each input cycle. Using an input frequency that has a large common divisor with the sample frequency generates similar problems since the codes repeat after each cycle of the divisor frequency. Ideally, the period of the greatest common divisor should be as long as the record length.

A 5180A histogram is shown in Figure 8 for an input sine wave at 9.85MHz. For comparison, Figure 9 shows data from a commercially available, 8-bit 20MHz ADC for an input sine wave at 9.85MHz, while Figure 10 shows data from an 8-bit, 100MHz ADC taken at 9.85MHz.

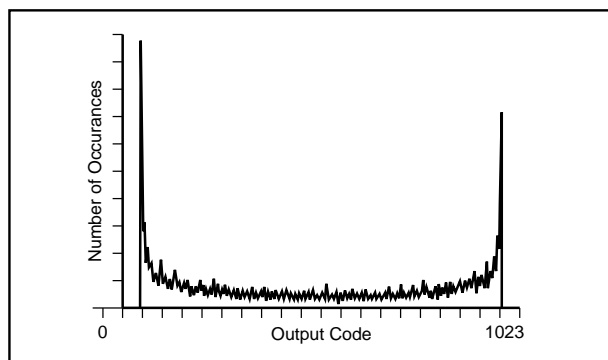


FIGURE 8. A 100,000-sample Histogram for a 5180A with a 9.85MHz Sine Wave Input. All Discontinuities are Less Than 1LSB.

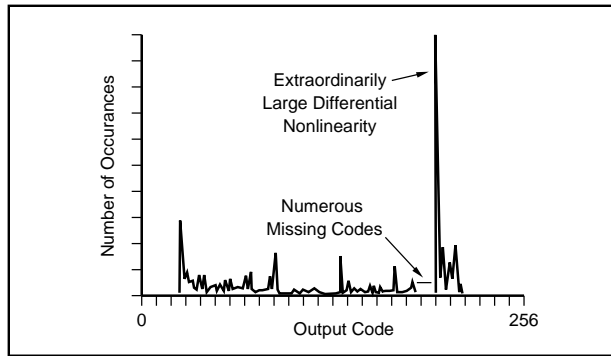


FIGURE 9. A 100,000-sample Histogram Plot for a Commercially Available 20MHz, 8-bit ADC with a 9.85MHz Input. Large differential nonlinearities and numerous missed codes are apparent.

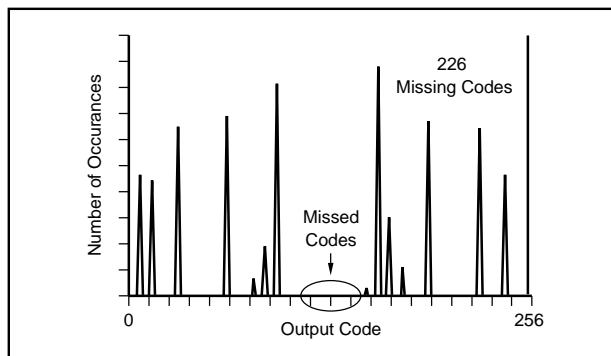


FIGURE 10. A 100,000-sample Histogram Plot for a 100MHz, 8-bit ADC with a 9.85MHz Input Sampled at 20MHz. Extremely large differential nonlinearities and numerous missed codes are apparent.

CURVE FITTING

The curve-fit test is a global description of the ADC. This means that the errors measured by the test are averaged to give a general measurement of the ADC transfer function. The result of this test is a figure of merit called the number of effective bits for the ADC. The effective bit number is a general measure of how much an ADC's nonlinearity has impaired its usefulness at a given frequency.

The number of effective bits is obtained by analyzing a record of data taken from a sine wave source (Figure 11). The analysis consists of generating a sine wave in software that is a best fit to the data record. Any difference between the data record and the best-fit sine wave is assumed to be error (Figure 12). The standard deviation of the error thus calculated is compared to the error an ideal ADC of the same number of bits might generate. If the error exceeds the ideal, the number of effective bits exhibited by the ADC is less than the number of bits it digitizes. Errors that cause degradation in this test are nonlinear effects such as harmonic distortion, noise, and aperture uncertainty. Gain, offset, and phase errors do not affect the results since they are ignored by the curve-fit process.

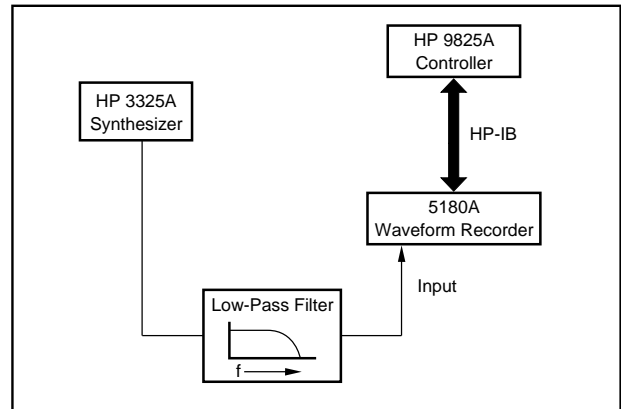


FIGURE 11. Setup for the Curve-fit Test and the Discrete Finite Fourier Transform (DFT) Test.

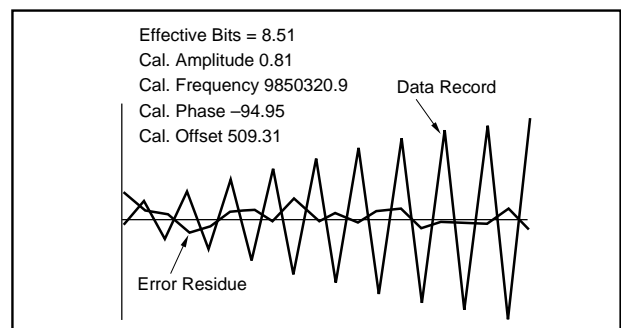


FIGURE 12. The First 20 Points of the Curve-fit Data Record and the Error Residue from a Fitted Sine Wave.

The number of effective bits is computed using expressions for average errors as follows:

$$\text{Effective bits} = N - \log_2 \left(\frac{\text{actual rms error}}{\text{ideal rms error}} \right)$$

where N is the number of ADC bits. The ideal rms error is not actually computed for the input waveform, but is assumed to be the quantization noise exhibited by an ideal ADC with a uniform-probability-density (UPD) input such as a perfect triangle wave. The ideal error is found from the expectation of squared error for a rectangular distribution. A rectangular distribution is used since that represents a UPD taken over an ideal code bin. The result thus obtained is:

$$\text{Ideal rms error} = \frac{Q}{\sqrt{12}}$$

where Q is the ideal code bin width. Although the input sine wave is not a UPD function, the UPD assumption is still valid since it is locally applied over each code bin. The deviation from a UPD over each code bin is very small, so the errors in using sine waves to approximate UPD inputs are negligible.

The actual rms error is simply the square root of the sum of the squared errors of the data points from the fitted sine wave. The actual rms error is given by:

$$E = \sum_{k=1}^m [x_k - A \cos(\omega t_k + P) - C]^2 \quad (1)$$

where E is the actual rms error, X_k and t_k are the data points, m is the number of data points in the record, and the fitted sine wave parameters are amplitude A , frequency ω , phase P , and offset C .

Equation 1 is also used to find the best-fit sine wave by minimizing the error E . The error is minimized by adjusting the fit parameters: frequency, phase, gain, and offset. This is done by taking the partial derivative of E in Equation 1 with respect to each of the four parameters. The error minimum occurs when all of the derivatives are equal to zero. This gives the four simultaneous equations:

$$\sum_{k=1}^m x_k \cos(\omega t_k + P) = A \sum_{k=1}^m \cos^2(\omega t_k + P) + C \sum_{k=1}^m \cos(\omega t_k + P) \quad (2)$$

$$\sum_{k=1}^m x_k = A \sum_{k=1}^m \cos(\omega t_k + P) + nC \quad (3)$$

$$\begin{aligned} \sum_{k=1}^m x_k t_k \sin(\omega t_k + P) = \\ A \sum_{k=1}^m t_k \cos(\omega t_k + P) \sin(\omega t_k + P) + C \sum_{k=1}^m t_k \sin(\omega t_k + P) \end{aligned} \quad (4)$$

$$\begin{aligned} \sum_{k=1}^m x_k \sin(\omega t_k + P) = \\ A \sum_{k=1}^m \cos(\omega t_k + P) \sin(\omega t_k + P) + C \sum_{k=1}^m \sin(\omega t_k + P) \end{aligned} \quad (5)$$

Equations 2 and 3 result from gain and offset adjustments. These are substituted into the other two equations, 4 and 5, giving two nonlinear equations:

$$\frac{\sum_{k=1}^m (x_k - \bar{x}) t_k \sin(\omega t_k + P)}{\sum_{k=1}^m (x_k - \bar{x}) \cos(\omega t_k + P)} = \frac{\sum_{k=1}^m [\cos(\omega t_k + P) - \bar{a}] t_k \sin(\omega t_k + P)}{\sum_{k=1}^m [\cos(\omega t_k + P) - \bar{a}] \cos(\omega t_k + P)} \quad (6)$$

$$\frac{\sum_{k=1}^m (x_k - \bar{x}) \sin(\omega t_k + P)}{\sum_{k=1}^m (x_k - \bar{x}) \cos(\omega t_k + P)} = \frac{\sum_{k=1}^m [\cos(\omega t_k + P) - \bar{a}] \sin(\omega t_k + P)}{\sum_{k=1}^m [\cos(\omega t_k + P) - \bar{a}] \cos(\omega t_k + P)} \quad (7)$$

$$\text{Where } \bar{a} = \sum_{k=1}^m \cos(\omega t_k + P)$$

These are solved iteratively to give values for the parameters. The difference between the right and left sides of Equation 6 is defined as error parameter R and the difference between the right and left sides of Equation 7 is defined as error parameter S . An approximation algorithm using a first-order Taylor series expansion drives R and S to zero. This approximation algorithm requires an initial guess for frequency and phase close to the solution to ensure conver-

gence to the best-fit sine wave. For frequency, the frequency of the generator output in Figure 11 is used as a guess. For phase, a guess is based on an examination of the data record by a software routine.

Although the result of this process is a single figure of merit, some enlightenment can be gained about the error components in the ADC by varying the test conditions. White noise produces the same degradation regardless of input frequency or amplitude. That is, the error term in Equation 1 is independent of test conditions for this sort of error. Another way of identifying noise in this test is by the randomness in the error residue, or the difference between the best-fit sine wave and the data taken.

Aperture uncertainty is identifiable because it generates an error that is a function of input slew rate. When this is the dominant error causing a low number of effective bits, the number of effective bits will vary linearly with both input frequency and amplitude. If the input waveform is sampled only at points of constant slew rate, such as zero crossings, then the aperture uncertainty corresponds to the amount that the effective bits decline as a function of slew rate.

Harmonic distortion is usually a nonlinear function of amplitude and frequency. Its distinguishing characteristic is the presence of the harmonics (or aliased harmonics if the fundamental is close to the Nyquist frequency) in the error residue. The amplitudes of the harmonics can be extracted by fitting the error residue with best-fit sine waves of the important harmonic frequencies. The impact of noise and aperture uncertainty in the presence of large distortion errors can be assessed by effective bit values and error residues with the fitted harmonics removed.

The greatest pitfall in the curve-fit test is using an input frequency that is a submultiple of the sample frequency. Since the same codes are sampled at exactly the same voltage each cycle, the locally uniform probability distribution assumption is violated. In the worst case, a submultiple of one-half, the quantization error would not be measurable at all. From a practical standpoint, this also defeats the global description of the test by sampling only a handful of codes.

Another potential pitfall is lack of convergence of the curve-fit algorithm. There are a few occasions where this can become a problem, such as when the data is very poor or the computational resolution is inadequate.

Figure 12 shows the error plot for a 5180A curve-fit test taken at a 9.85MHz input frequency. The number of effective bits associated with this error is 8.51.

FFT TESTING

The fast Fourier transform (FFT) is used to characterize an ADC in the frequency domain in much the same way that a spectrum analyzer is used to determine the linearity of an analog circuit. The data output for both techniques is a presentation of the magnitude of the Fourier spectrum for the circuit under test. Ideally, the spectrum is a single line

that represents the pure sine wave input and is devoid of distortion components generated by the circuit under test. There are, however, significant differences between the spectrum analyzer and ADC spectra because of the sampling operation of the ADC.

The Fourier transform of a signal $x(t)$ that is continuous for all time is defined as:

$$X(f) = \int_{-\infty}^{\infty} x(t)e^{-i2\pi ft} dt$$

and includes the amplitude and phase of every frequency in $x(t)$. The Fourier transform cannot be used in this form for an ADC, however, because $x(t)$ is only digitized at a finite number of points, M , spaced Δt apart. Instead, the discrete finite transform (DFT) must be used. It is defined as:

$$XD(f) = \sum_{m=0}^{M-1} x(m\Delta t)e^{-i2\pi f(m\Delta t)} \Delta t$$

There are significant differences between $X(f)$ and $XD(f)$. While $X(f)$ has infinite spectral resolution, $XD(f)$ has a discrete frequency resolution of $\Delta f = 1/M\Delta t$ because of the finite number of points in the data record. The finite record size also accounts for another difference between $X(f)$ and $XD(f)$ whenever a nonintegral number of cycles of $X(t)$ is contained in the record. Since the DFT assumes that the record repeats with a period of $M\Delta t$ (to satisfy the Fourier transform condition that $x(t)$ be continuous for all time) sharp discontinuities at the points where the start of one record joins the end of the preceding record cause the spectral components of $X(f)$ to be spread or smeared in $XD(f)$.

The smearing, called leakage, can be explained as follows. The finite record size of $x(t)$ can be considered the consequence of multiplying $x(t)$ by a rectangular function having unity amplitude during the time period $M\Delta t$ that the record is acquired and zero amplitude elsewhere. Since multiplication of two functions in one domain (time, in this case) is equivalent to convolution in the other, the spectrum of $XD(f)$ is derived by convolving $X(f)$ with $W(f)$, the Fourier trans-

form of the rectangular function. $W(f)$ is the familiar sinc/x function (see Figure 13 for $|W(f)|$), consisting of a main lobe surrounded by a series of sidelobes whose amplitudes decay at a 6dB-per-octave rate. It is these sidelobes that are responsible for leakage. Even if the spectrum of $X(f)$ is a single line, the sidelobes of $W(f)$ during the convolution smear the energy in the single line into a series of spectral lines spaced $1/M\Delta t$ apart whenever the frequency of $x(t)$ is not an integral multiple of $1/M\Delta t$.

Leakage can be reduced by multiplying the data in the record by a windowing function that weights the points in the center of the record heavily while smoothly suppressing the points near the ends. Many different windowing functions exist that offer various tradeoffs of amplitude resolution versus frequency resolution. A function commonly used with sine waves is the Hanning window, defined by $|(1/2)(1 - \cos 2\pi t/M\Delta t)|$. Notice in Figure 13 that both the window and its derivative approach zero at the two ends of the record and that the transform's main lobe is twice as wide as that of the rectangular function, while the amplitudes of the sidelobes decay by an additional 12dB per octave. The reduced level of the sidelobes reduces leakage, but the wider main lobe limits the ability to resolve closely spaced frequencies. Furthermore, the shape of the main lobe can attenuate the spectral amplitudes of $X(f)$ by as much as 1.5dB. However, for the DFT testing to be described here, the Hanning window was selected as a good compromise between frequency and amplitude resolution.

The third difference between the spectra of $X(f)$ and $XD(f)$ is the limited range of frequencies displayed for $XD(f)$. The sampling process causes the two-sided spectrum of $X(f)$, symmetrical about the origin, to be replicated as the sampling frequency L and at all of its harmonics. If $X(f)$ contains components that exceed $f_s/2$, then these components are folded back, or aliased, onto spectral lines below f_s , causing aliasing errors. The frequency $f_s/2$ is sometimes called the Nyquist frequency, referring to the Nyquist criterion, which requires the sampling rate to be twice the highest frequency present in the input signal to define the waveform uniquely.

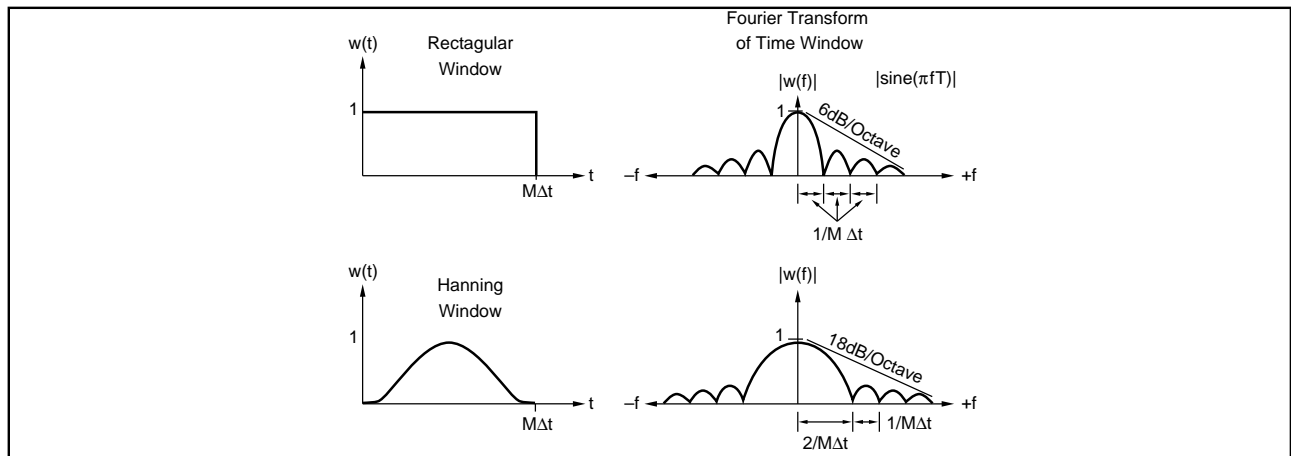


FIGURE 13. Time-domain and Frequency-domain Representations of Rectangular and Hanning Windows.

The result is that the spectrum of $XD(f)$ is displayed only from DC to $f_s/2$ and the maximum input frequency must be limited to less than $f_s/2$ to avoid aliasing.

Figure 14 presents the magnitude of the spectra derived from the DFT for perfect 10-bit and 6-bit ADC's given a pure sinusoidal input. Useful information about the ADC's performance can be derived from three features of the spectra: the noise floor, the harmonic level, and the spurious level.

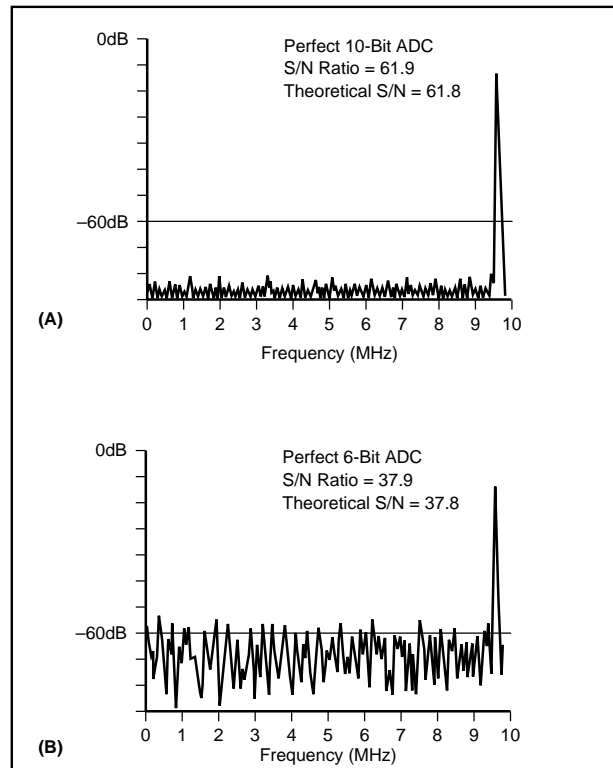


FIGURE 14. FFT Plots for 0.85MHz Data Quantized by Perfect 10-bit (A) and 6-bit (B) ADCs. The signal-to-noise ratio computed in each case agrees closely with the theoretical value of $6N + 1.8\text{dB}$ where N is the number of ADS bits.

Two classes of noise sources determine the level of the noise floor. The first is called quantization noise. This is the error, bounded by $\pm 1/2\text{LSB}$, that is inherent in the quantization of the input amplitude into discrete levels. As can be seen in Figure 14, even ideal ADCs have noise floors determined by quantization noise. The higher the number of bits, the smaller the error bound and, therefore, the lower the noise floor.

All real-life ADCs have noise floors that are higher than that solely from quantization noise. The second class of noise source includes wideband noise generated within the ADC, along with other sources. In a parallel-ripple ADC, for example, such things as misadjustment between the first-pass and second-pass ranges (exceeding the redundancy range) or inadequate DAC settling can cause localized code errors or differential nonlinearities in the ADC's status

transfer function. Furthermore, localized code errors can increase in amplitude and in the number of codes affected under dynamic input conditions. Aperture jitter is another major source of dynamic error; the magnitude of this localized code error is dependent upon the slew rate of the input at the time of sampling. Each of these localized code errors can be modeled as a sharp discontinuity in the time domain that when transformed into the frequency domain results in a broad spectrum that raises the height of the noise floor above that caused by quantization noise alone.

The second feature of the DFT-derived spectrum that indicates an ADC's level of dynamic performance is the harmonic content. Static and dynamic integral nonlinearities cause curvature in the ADC's transfer function. If the input frequency f_{IN} is much lower than the Nyquist frequency ($f_s/2$), then the harmonic components will be in the expected locations: $2f_{IN}$, $3f_{IN}$, etc. If, on the other hand, the harmonics of f_{IN} are greater than $f_s/2$, then these frequencies will be aliased onto components below $f_s/2$. Take, for instance, a 20-megasample-per-second (f_s) ADC with an input of 9.85MHz. The second harmonic at 19.7MHz is aliased to 0.3MHz, the third harmonic at 29.55MHz is aliased to 9.55MHz, the fourth at 39.4MHz is aliased to 0.6MHz, and so on.

Care must be exercised in selecting the input frequency for the DFT test. An incorrectly chosen frequency can alias one of its harmonic components on to the fundamental and thereby understate the harmonic distortion (in the example above, an input of exactly 5MHz would place the third harmonic at the fundamental frequency). The input frequency should be chosen so that the harmonics are far enough away to be easily resolvable from the fundamental, whose energy has been spread into several adjacent bins ($1/M\Delta t$ locations) by the Hanning window. This accounts for the 0.15MHz offset from 10MHz used in the example of Figure 14.

The third feature of the DFT-based spectrum that is indicative of the ADC's level of dynamic performance is the spurious content. Spurious components are spectral components that are not harmonically related to the input. For example, a strong signal near the ADC may contaminate the ADC's analog ground somehow and thereby appear in the spectrum. The nearby signal will not only appear as itself, but because of nonlinearities within the ADC, can combine with the input signal to form sum and difference terms resulting in intermodulation distortion.

The combined effects of noise floor, harmonic distortion and spurious errors are reflected in the ADC's rms signal-to-noise ratio, which can be derived from the DFT magnitude spectrum. The signal energy is determined by summing the energy in all the bins associated with the fundamental. The noise energy is the sum of the energy in all other bins. By taking the logarithm of the ratio of signal energy to noise energy and multiplying by 20, the signal-to-noise ratio for the ADC can be calculated. An ideal N -bit ADC having quantization noise only is theoretically known to have a signal-to-noise ratio equal to $(6N + 1.8)\text{dB}$, which sets an

upper bound. A signal-to-noise ratio below this ideal limit is indicative of errors of all types that the ADC produces.

The FFT test setup is presented in Figure 11. A full-scale sine wave of a properly chosen frequency is applied to the ADC under test. The low-pass filter ensures a spectrally pure input. A 1024-point record sampled at the maximum sampling rate is then taken and given to the computer, which calculates the DFT using an FFT algorithm. The spectral magnitude is plotted as a function of frequency.

Figure 15 shows the graphical outputs for the 5180 for full-scale sine wave input at 0.95MHz and 9.85MHz. As might be expected, the distortion increases with increasing frequency. Harmonic and spurious components are typically better than -60dBc below 1MHz and -54dBc at 9.85MHz. The latter spectrum at 9.85MHz is the frequency-domain representation for one of the most demanding tests of an ADC, called the envelope test which was described earlier.

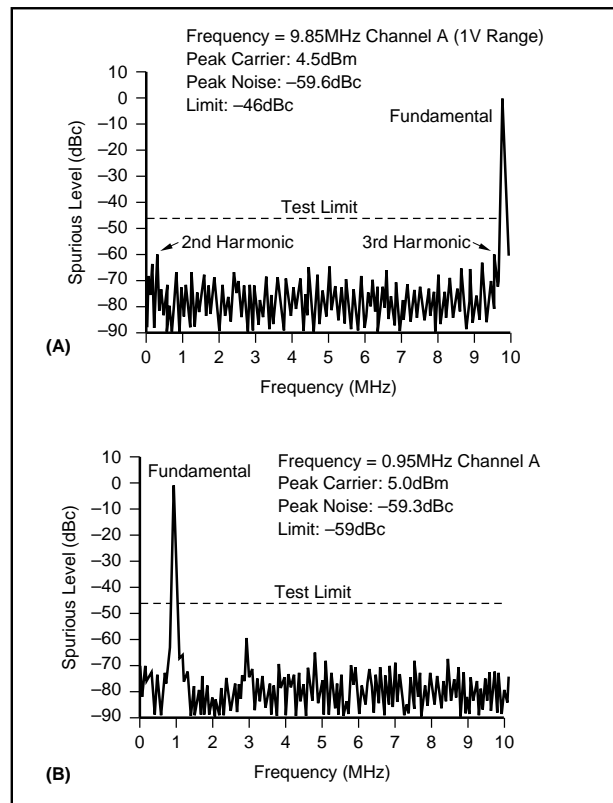


FIGURE 15. DFT Plots for the 5180A with Input Frequencies of 9.85MHz (a) and 0.95MHz (b). The low harmonic distortion indicates very low integral nonlinearity.

Figure 16 presents, for comparison, the test results for commercially available digitizers: a 20-megasample-per-second, 8-bit ADC and a 100-megasample-per-second, 10-bit ADC with a full-scale, 9.85MHz sine wave input, sampled at 20 megasamples-per-second. The numerous large harmonic components, both odd and even, are indicative of

severe harmonic distortion errors resulting from integral nonlinearity in the transfer functions of both of these ADCs.

A rule of thumb has evolved that uses the DFT-based spectrum as a quick overview of an N-bit ADC's dynamic performance. If all harmonic and spurious components are at least 6N dB below the full-scale amplitude of the fundamental, then the ADC is performing satisfactorily, since each error component has a peak-to-peak amplitude smaller than an LSB. If, on the other hand, harmonic or spurious components are less than 6N dB down, or if the noise floor is elevated, then other tests can be performed that are better at isolating the particular integral and differential nonlinearity errors. In particular, the FFT test can be followed by the histogram test or the beat frequency test (or envelope test), as conditions warrant.

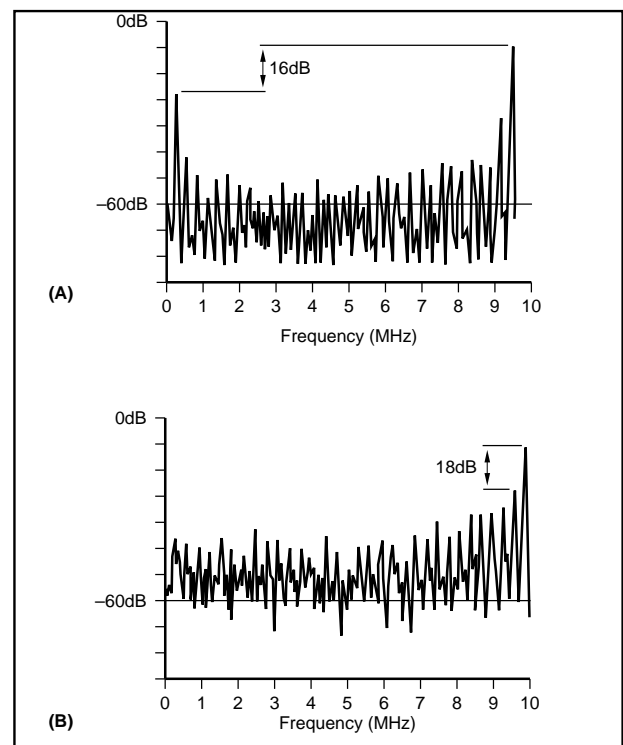


FIGURE 16. DFT Plots for a 20MHz, 8-Bit ADC (a) and a 100MHz, 8-Bit ADC (b). Full-scale input sine waves at 9.85MHz were sampled at a rate of 20MHz. The high levels of harmonic distortion indicate severe integral nonlinearities.

CONCLUSION

The four sine-wave-based ADC tests described provide information about the quality of any recorder. The tests may be used to isolate specific failures, even at high-speed and

fine resolution (Figure 17). The tests are simple to run, requiring only a synthesized generator and an HP-1B computer.

ERROR	HISTOGRAM	DFT	SINE WAVE CURVE-FIT	BEST FREQUENCY TEST
Differential Nonlinearity	Yes—shows up as spikes	Yes—shows up as elevated noise floor	Yes—part of rms error	Yes
Missing Codes	Yes—shows up as bins with 0 counts	Yes—shows up as elevated noise floor	Yes—part of rms error	Yes
Integral Nonlinearity	Yes—(could be measured directly with a highly linear ramp waveform)	Yes—shows up as harmonics on fundamental aliased into baseband	Yes—part of rms error	Yes
Aperture Uncertainty	No—averaged out. Can be measured with "locked" histogram	Yes—shows up as elevated noise floor	Yes—part of rms error	No
Noise	No—averaged out. Can be measured with "locked" histogram	Yes—shows up as elevated noise floor	Yes—part of rms error	No
Bandwidth Errors	No	No	No	Yes—used to measure analog bandwidth
Gain Errors	Yes—shows up in peak-to-peak spread of distribution	No	No	No
Other Errors	Yes—shows up in offset of distribution average	No	No	No

FIGURE 17. Summary of the Errors Exposed by the Dynamic Tests.

REFERENCES

1. D.J. Packard, "Beat Frequency Testing of Real Time A/D Converters," Workshop on High Speed A/D Conversion, Portland, Oregon, February, 1980.
 2. H.U. Koller, "New Criterion for Testing Analog-to-Digital Converters for Statistical Evaluation," *IEEE Transactions on Instrumentation and Measurement*, Vol. 1M-22, pp. 214-17, September, 1973.
 3. R. Potter, "Least-Squares Estimation of Sinusoidal Parameters from Measured Data," Hewlett-Packard internal memo, June, 1974.
 4. L. Ochs, "Measurement and Enhancement of Waveform Digitizer Performance," IEEE International Convention, Boston, May, 1976.
 5. R.N. Bracewell, *The Fourier Transform and Its Applications*, McGraw-Hill, New York, 1965.
-

The information provided herein is believed to be reliable; however, BURR-BROWN assumes no responsibility for inaccuracies or omissions. BURR-BROWN assumes no responsibility for the use of this information, and all use of such information shall be entirely at the user's own risk. Prices and specifications are subject to change without notice. No patent rights or licenses to any of the circuits described herein are implied or granted to any third party. BURR-BROWN does not authorize or warrant any BURR-BROWN product for use in life support devices and/or systems.