

Analog-to-Digital Converter Testing

Kent H. Lundberg

Analog-to-digital converters are essential building blocks in modern electronic systems. They form the critical link between front-end analog transducers and back-end digital computers that can efficiently implement a wide variety of signal-processing functions. The wide variety of digital-signal-processing applications leads to the availability of a wide variety of analog-to-digital (A/D) converters of varying price, performance, and quality.

Ideally, an A/D converter encodes a continuous-time analog input voltage, V_{IN} , into a series of discrete N -bit digital words that satisfy the relation

$$V_{IN} = V_{FS} \sum_{k=0}^{N-1} \frac{b_k}{2^{k+1}} + \epsilon$$

where V_{FS} is the full-scale voltage, b_k are the individual output bits, and ϵ is the quantization error. This relation can also be written in terms of the least significant bit (LSB) or quantum voltage level

$$V_Q = \frac{V_{FS}}{2^N} = 1 \text{ LSB}$$

as

$$V_{IN} = V_Q \sum_{k=0}^{N-1} b_k 2^k + \epsilon$$

A plot of this ideal characteristic for a 3-bit A/D converter is shown in Figure 1.

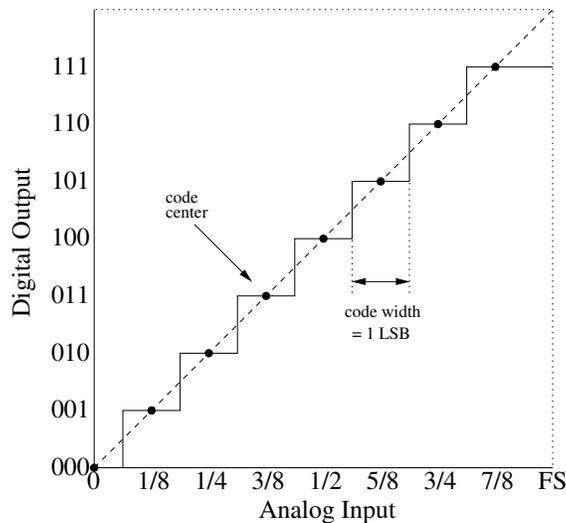


Figure 1: Ideal analog-to-digital converter characteristic.

Each unique digital code corresponds to a small range of analog input voltages. This range is 1 LSB wide (the “code width”) and is centered around the “code center.” All input voltages resolve to the digital code of the nearest code center. The difference between the analog input voltage and the corresponding voltage of the nearest code center (the difference between the solid and dashed lines in Figure 1) is the quantization error. Since the A/D converter has a finite number of output bits, even an ideal A/D converter produces some quantization error with every sample.

1 A/D Converter Figures of Merit

The number of output bits from an analog-to-digital converter do not fully specify its behavior. Real A/D converters can differ from ideal behavior in many ways. While static imperfections, such as gain and offset, are easy to quantify, the success of many signal-processing applications depends on the dynamic behavior of the A/D converter. Ultimately, the application determines the requirements, and A/D converter resolution may not be either necessary or sufficient to specify the required performance. In many cases, the quality of the A/D converter must be tested for the specific application.

The wide variety of analog-to-digital converter applications leads to a large number of figures of merit for specifying performance. These figures of merit include accuracy, resolution, dynamic range, offset, gain, differential nonlinearity, integral nonlinearity, signal-to-noise ratio, signal-to-noise-and-distortion ratio, effective number of bits, spurious-free dynamic range, intermodulation distortion, total harmonic distortion, effective resolution bandwidth, full-power bandwidth, full-linear bandwidth, aperture delay, aperture jitter, transient response, and overvoltage recovery.

These specifications can be loosely divided into three categories — static parameters, frequency-domain dynamic parameters, and time-domain dynamic parameters — and are defined in this section.

1.1 Static Parameters

Static parameters are the A/D converter specifications that can be tested at low speed, or even with constant voltages. These specifications include accuracy, resolution, dynamic range, offset, gain, differential nonlinearity, and integral nonlinearity.

1.1.1 Accuracy

Accuracy is the total error with which the A/D converter can convert a known voltage, including the effects of quantization error, gain error, offset error, and nonlinearities. Technically, accuracy should be traceable to known standards (for example, NIST), and is generally a “catch-all” term for all static errors.

1.1.2 Resolution

Resolution is the number of bits, N , out of the A/D converter. The characteristic in Figure 1 shows a 3-bit A/D converter. Probably the most noticeable specification, resolution determines the size of the least significant bit, and thus determines the dynamic range, the code widths, and the quantization error.

1.1.3 Dynamic Range

Dynamic range is the ratio of the smallest possible output (the least significant bit or quantum voltage) to the largest possible output (full-scale voltage), mathematically $20 \log_{10} 2^N \approx 6N$.

1.1.4 Offset Error

Offset error is the deviation in the A/D converter’s behavior at zero. The first transition voltage should be 1/2 LSB above analog ground. Offset error is the deviation of the actual transition voltage from the ideal 1/2 LSB. Offset error is easily trimmed by calibration. Compare the location of the first transitions in Figures 1 and 2.

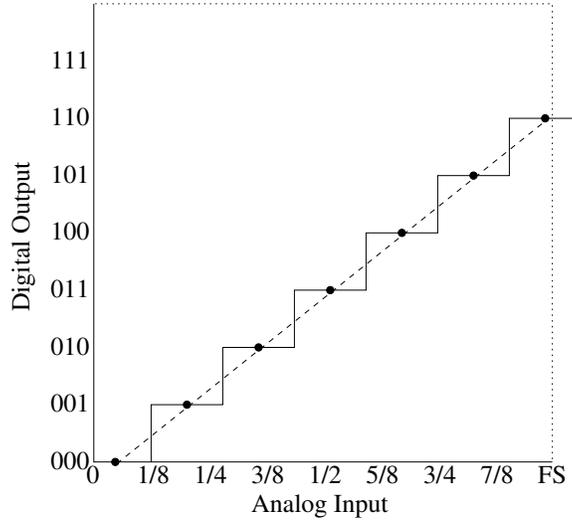


Figure 2: Analog-to-digital converter characteristic, showing offset and gain error.

1.1.5 Gain Error

Gain error is the deviation in the slope of the line through the A/D converter's end points at zero and full scale from the ideal slope of $2^N/V_{FS}$ codes-per-volt. Like offset error, gain error is easily corrected by calibration. Compare the slope of the dashed lines in Figures 1 and 2.

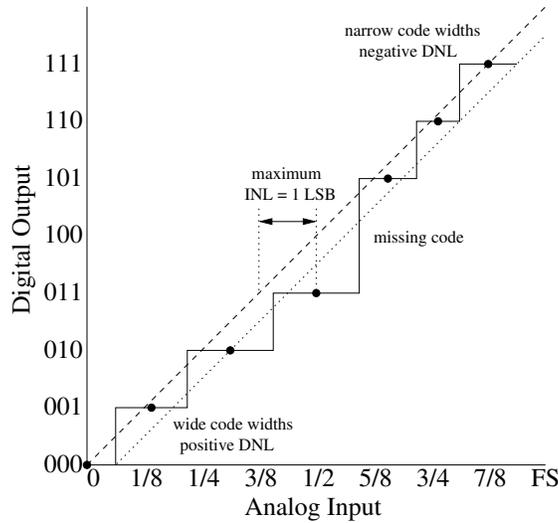


Figure 3: Analog-to-digital converter characteristic, showing nonlinearity errors and a missing code. The dashed line is the ideal characteristic, and the dotted line is the best fit.

1.1.6 Differential Nonlinearity

Differential nonlinearity (DNL) is the deviation of the code transition widths from the ideal width of 1 LSB. All code widths in the ideal A/D converter are 1 LSB wide, so the DNL would be zero

everywhere. Some datasheets list only the “maximum DNL.” See the wide range of code widths illustrated in Figure 3.

1.1.7 Integral Nonlinearity

Integral nonlinearity (INL) is the distance of the code centers in the A/D converter characteristic from the ideal line. If all code centers land on the ideal line, the INL is zero everywhere. Some datasheets list only the “maximum INL.” See the deviations of the code centers from the ideal line in Figure 3.

Note that there are two possible ways to express maximum INL, depending on your definition of the “ideal line.” Datasheet INL numbers can be decreased by quoting maximum INL from a “best fit” line instead of the ideal line. In Figure 3, the ideal line (shown dashed) exhibits an INL of 1 LSB, while the best-fit line (shown dotted) exhibits an INL of half that size. While this prevarication underestimates the effect of INL on total accuracy, it is probably a better reflection of the A/D converter’s linearity in AC-coupled applications.

1.1.8 Missing Codes

Missing codes are output digital codes that are not produced for any input voltage, usually due to large DNL. In some converters, missing codes can be caused by non-monotonicity of the internal D/A. The large DNL in Figure 3 causes code 100 to be “crowded out.”

1.2 Resolution and Quantization Noise

Quantization error due to the finite resolution N of the A/D converter limits the signal-to-noise ratio. Even in an ideal A/D converter, the quantization of the input signal creates errors that behave like noise. All inputs within $\pm 1/2$ LSB of a code center resolve to that digital code. Thus, there will be a small difference between the code center and the actual input voltage due to this quantization. If assume that this error voltage is uncorrelated and distributed uniformly, we can calculate the expected rms value of this “quantization noise.”

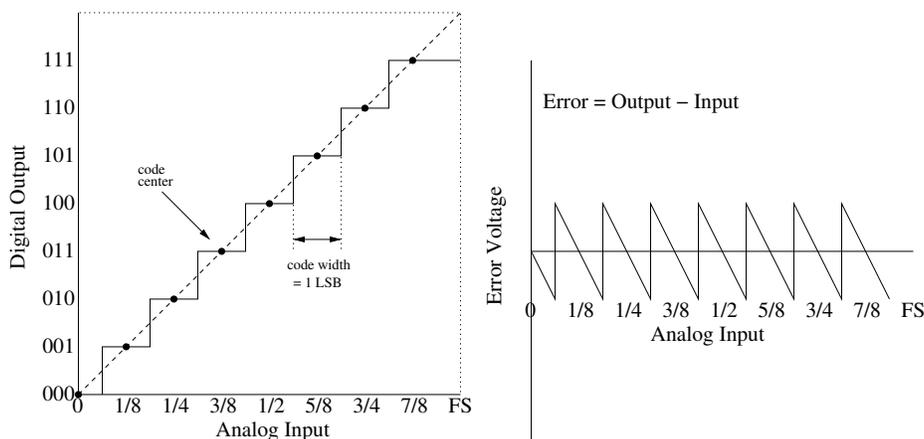


Figure 4: Quantization error voltage for ideal analog-to-digital converter.

The signal-to-noise ratio due to quantization can be directly calculated. The range of the error

voltage is the quantum voltage level (the least significant bit)

$$V_Q = \frac{V_{FS}}{2^N} = 1 \text{ LSB}$$

Finite amplitude resolution introduces a quantization error between the analog input voltage and the reconstructed output voltage. Assuming this quantization error voltage is uniformly distributed over the code width from $-1/2$ LSB to $+1/2$ LSB, the expectation value of the error voltage is

$$E\{\epsilon^2\} = \frac{1}{V_Q} \int_{-\frac{1}{2}V_Q}^{+\frac{1}{2}V_Q} \epsilon^2 d\epsilon = \frac{1}{V_Q} \left[\frac{\epsilon^3}{3} \right]_{-\frac{1}{2}V_Q}^{+\frac{1}{2}V_Q} = \frac{V_Q^2}{12}$$

This quantization noise is assumed to be uncorrelated and broadband. Using this result, the maximum signal-to-noise ratio for a full-scale input can be calculated. The rms value of a full-scale peak-to-peak amplitude V_F is

$$V_{\text{rms}} = \frac{V_{FS}}{2\sqrt{2}} = \frac{2^N V_Q}{2\sqrt{2}}$$

thus the signal-to-noise ratio is

$$\text{SNR} = 20 \log \left(\frac{V_{\text{rms}}}{\sqrt{E\{\epsilon^2\}}} \right) = 20 \log(2^N \sqrt{1.5}) = 6.02N + 1.76 \text{ dB}$$

when the noise is due only to quantization. Using this result, we can tabulate the ideal signal-to-noise ratio for several A/D converter resolutions in Figure 5.

resolution	signal-to-noise ratio
6 bits	37.9 dB
8 bits	49.9 dB
10 bits	62.0 dB
12 bits	74.0 dB
14 bits	86.0 dB
16 bits	98.1 dB

Figure 5: Ideal signal-to-noise ratio due to quantization versus resolution.

Figure 6 shows the fast Fourier transform (FFT) spectrum of a sine wave sampled by an ideal 10-bit A/D converter. The noise floor in this figure is due only to quantization noise. For an M -point FFT, the average value of the noise contained in each frequency bin is $10 \log_2(M/2)$ dB below the rms value of the quantization noise. This “FFT process gain” is why the noise floor appears 30 dB lower than the quantization noise level listed in Figure 5.

1.3 Frequency-Domain Dynamic Parameters

All real analog-to-digital converters have additional noise sources and distortion processes that degrade the performance of the A/D converter from the ideal signal-to-noise ratio calculated above. These imperfections in the dynamic behavior of the A/D converter are quantified and reported in a variety of ways.

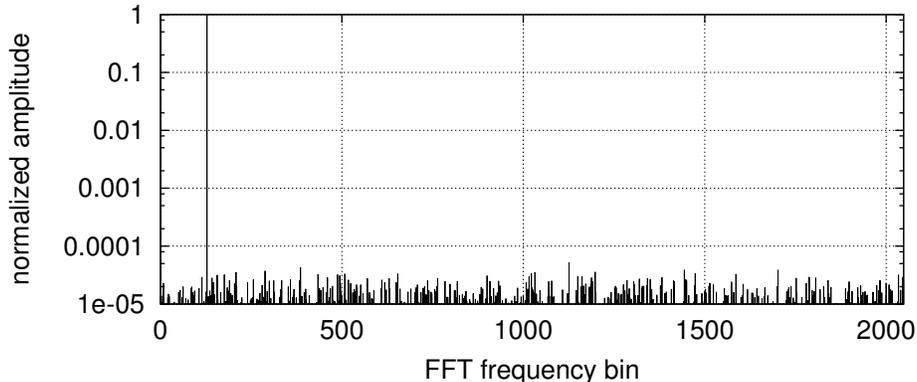


Figure 6: Quantization noise floor for an ideal 10-bit A/D converter (4096 point FFT).

1.3.1 Signal-to-Noise-and-Distortion Ratio

Signal-to-noise-and-distortion ratio (S/N+D, SINAD, or SNDR) is the ratio of the input signal amplitude to the rms sum of all other spectral components. For an M -point FFT of a sine wave test, if the fundamental is in frequency bin m (with amplitude A_m), the SNDR can be calculated from the FFT amplitudes

$$\text{SNDR} = 10 \log \left[A_m^2 \left(\sum_{k=1}^{m-1} A_k^2 + \sum_{k=m+1}^{M/2} A_k^2 \right)^{-1} \right]$$

To avoid any spectral leakage around the fundamental, often several bins around the fundamental are ignored. The SNDR is dependent on the input-signal frequency and amplitude, degrading at high frequency and power. Measured results are often presented in plots of SNDR versus frequency for a constant-amplitude input, or SNDR versus amplitude for a constant-frequency input.

1.3.2 Effective Number of Bits

Effective number of bits (ENOB) is simply the signal-to-noise-and-distortion ratio expressed in bits rather than decibels by solving the “ideal SNR” equation

$$\text{SNR} = 6.02N + 1.76 \text{ dB}$$

for the number of bits N , using the measured SNDR

$$\text{ENOB} = \frac{\text{SNDR} - 1.76 \text{ dB}}{6.02 \text{ dB/bit}}$$

In the presentation of measured results, ENOB is identical to SNDR, with a change in the scaling of the vertical axis.

1.3.3 Spurious-Free Dynamic Range

Spurious-free dynamic range (SFDR) is the ratio of the input signal to the peak spurious or peak harmonic component. Spurs can be created at harmonics of the input frequency due to nonlinearities in the A/D converter, or at subharmonics of the sampling frequency due to mismatch or clock

coupling in the circuit. The SFDR of an A/D converter can be larger than the SNDR. Measurement of SFDR can be facilitated by increasing the number of FFT points or by averaging several data sets. In both cases, the noise floor will improve, while the amplitude of the spurs will stay constant. The spectrum of an A/D converter with significant harmonic spurs is shown in Figure 7. Because SFDR is often slew-rate dependent, it will be a function of input frequency and magnitude [1]. The maximum SFDR often occurs at an amplitude below full scale.

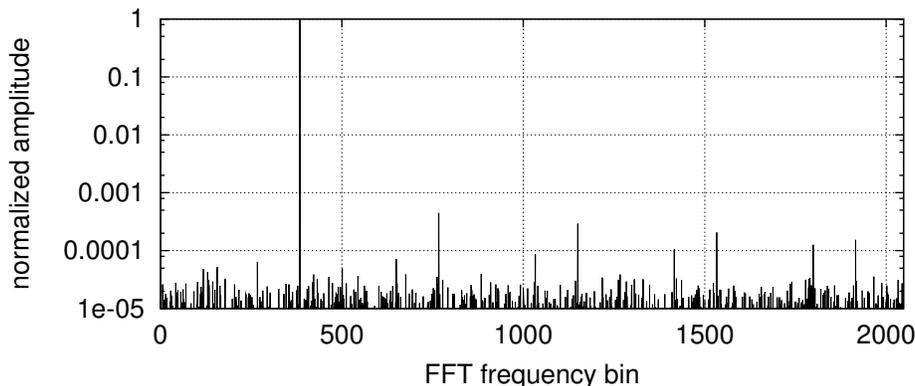


Figure 7: A/D converter with significant nonlinearity, showing poor SFDR and THD. Note that high frequency harmonics of the input signal are aliased down to frequencies below the Nyquist frequency.

1.3.4 Total Harmonic Distortion

Total harmonic distortion (THD) is the ratio of the rms sum of the first five harmonic components (or their aliased versions, as in Figure 7) to the input signal

$$\text{THD} = 10 \log \left(\frac{V_2^2 + V_3^2 + V_4^2 + V_5^2 + V_6^2}{V_1^2} \right)$$

where V_1 is the amplitude of the fundamental, and V_n is the amplitude of the n -th harmonic.

1.3.5 Intermodulation Distortion

Intermodulation distortion (IMD) is the ratio of the amplitudes of the sum and difference frequencies to the input signals for a two-tone test, sometimes expressed as “intermod-free dynamic range (IFDR)” See the FFT spectrum in Figure 8. For second-order distortion, the IMD would be

$$\text{IMD} = 10 \log \left(\frac{V_+^2 + V_-^2}{V_1^2 + V_2^2} \right)$$

where V_1 and V_2 are the rms amplitudes of the input signals, and V_+ and V_- are the rms amplitudes of the sum and difference intermodulation products. See Figure 8.

1.3.6 Effective Resolution Bandwidth

Effective resolution bandwidth (ERBW) is the input-signal frequency where the SNDR of the A/D converter has fallen by 3 dB (0.5 bit) from its value for low-frequency input signals.

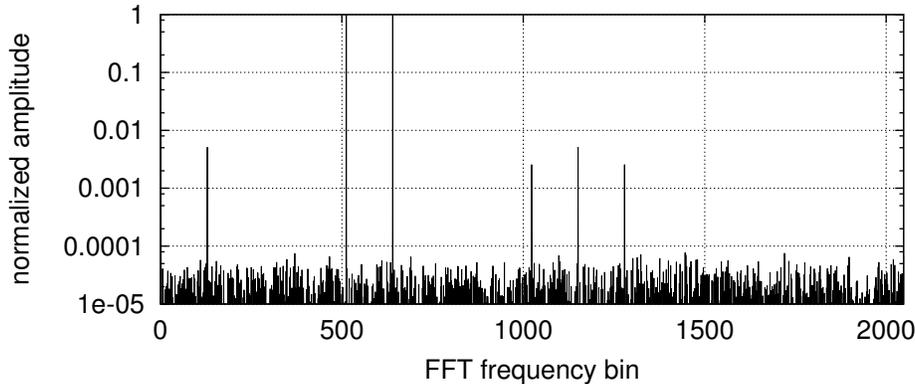


Figure 8: Two-tone IMD test with second-order nonlinearity, showing (from left) $F_2 - F_1$ product, F_1 input, F_2 input, $2F_1$ product, $F_1 + F_2$ product, and $2F_2$ product.

1.3.7 Full-Power Bandwidth

In amplifiers, full-power bandwidth (FPBW) is the maximum frequency at which the amplifier can reproduce a full-scale sinusoidal output without distortion (sometimes calculated as slew-rate divided by $\pm 2\pi V_{\max}$), or where the amplitude of full-scale sinusoid is reduced by 3 dB. Using this definition for A/D converters can result in optimistic frequencies where the SNDR is severely degraded. Some manufacturer report the full-power bandwidth as the frequency where the amplitude of the *reconstructed* input signal is reduced by 3 dB [2].

1.3.8 Full-Linear Bandwidth

Full-linear bandwidth is the frequency where the slew-rate limit of the input sample-and-hold begins distorting the input signal by some specified amount ([3] uses 0.1 dB)

1.4 Time-Domain Dynamic Parameters

1.4.1 Aperture Delay

Aperture delay is the delay from when the A/D converter is triggered (perhaps the rising edge of the sampling clock) to when it actually converts the input voltage into the appropriate digital code. Aperture delay is also sometimes called aperture time.

1.4.2 Aperture Jitter

Aperture jitter is the sample-to-sample variation in the aperture delay. The rms voltage error caused by rms aperture jitter decreases the overall signal-to-noise ratio, and is a significant limiting factor in the performance of high-speed A/D converters [4].

If we assume that the input waveform is a sinusoid

$$V_{IN} = V_{FS} \sin \omega t$$

then the maximum slope of the input waveform is

$$\left. \frac{dV_{IN}}{dt} \right|_{\max} = \omega V_{FS}$$

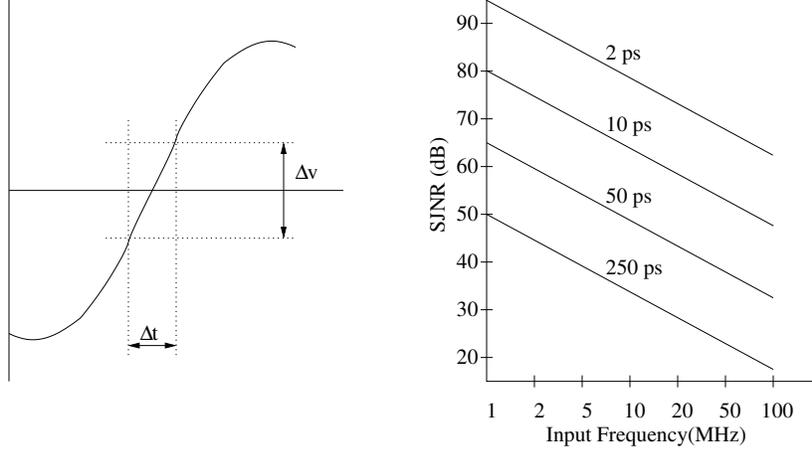


Figure 9: Effects of aperture jitter.

which occurs at the zero crossings. If there is an rms error in the time at which we sample (aperture jitter, t_a) during this maximum slope, then there will be an rms voltage error of

$$V_{\text{rms}} = \omega V_{FS} t_a = 2\pi f V_{FS} t_a$$

Since the aperture time variations are random, these voltage errors will behave like a random noise source. Thus the signal-to-jitter-noise ratio

$$\text{SJNR} = 20 \log \left(\frac{V_{FS}}{V_{\text{rms}}} \right) = 20 \log \left(\frac{1}{2\pi f t_a} \right)$$

The SJNR for several values of the jitter t_a is shown in Figure 9.

1.4.3 Transient Response

Transient response is the settling time for the A/D converter to full accuracy (to within $\pm 1/2$ LSB) after a step in input voltage from zero to full scale

1.4.4 Overvoltage Recovery

Overvoltage recovery is the settling time for the A/D converter to full accuracy after a step in input voltage from outside the full scale voltage (for example, from $1.5V_F$ to $0.5V_F$)

2 Dynamic A/D Converter Testing Methods

A wide variety of tests have been developed to measure dynamic specifications. Many of these tests rely on Fourier analysis using the discrete Fourier transform (DFT) and the fast Fourier transform (FFT), as well as other mathematical models.

2.1 Testing with the Fast Fourier Transform

The simplest frequency-domain tests use the direct application of the fast Fourier transform. Taking the FFT of the output data while driving the A/D converter with a single, low distortion sine wave, the SNDR, ENOB, SFDR, and THD can easily be calculated. It is useful to take these measurements at several input amplitudes and frequencies, and plot the results. Taking data for high input frequencies allows the full-power and full-linear bandwidths to be calculated.

Two more tests are completed while driving the A/D converter with an input composed of two sine waves of different frequencies. The FFT of this test result is used to calculate the IMD (for second-order and third-order products) and the two-tone SFDR.

2.2 Signal Coherence for FFT Tests

One potential problem in using Fourier analysis is solved by paying close attention to the coherence of the sampled waveform in the data record. Unless the sampled data contains a whole number of periods of the input waveform, spectral leakage of the input frequency can obscure the results [5]. Figure 10 shows the FFT of a 4096-point data record containing 127.5 periods of a perfect sine wave. This spectrum should contain a single impulse in frequency, but the incomplete cycle of the sine wave at the end of the data record causes the spectrum to be broadly smeared.¹

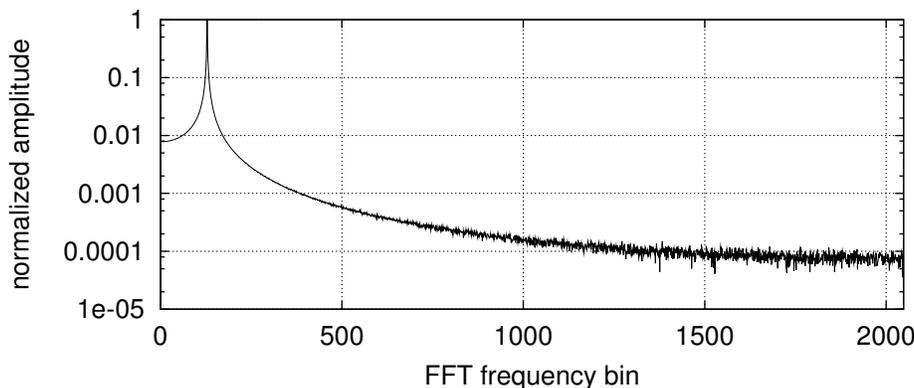


Figure 10: FFT of non-integer number of cycles (4096 points, 127.5 cycles).

In addition, the number of periods of the input waveform in the sample record should not be a non-prime integer sub-multiple of the record length. For example, using a power-of-two for both the number of periods and the number of samples results in repetitive data. Figure 11 shows the FFT of a 4096-point data record that contains 128 periods of a sine wave. In this case, the first 32 samples of the data record are simply repeated 128 times. This input vector does a poor job of exercising the device under test, and the quantization noise is concentrated in the harmonics of the input frequency rather than uniformly distributed across the Nyquist bandwidth [7].

¹This smearing can be solved by windowing [6], but it is easier to simply use a whole number of cycles.

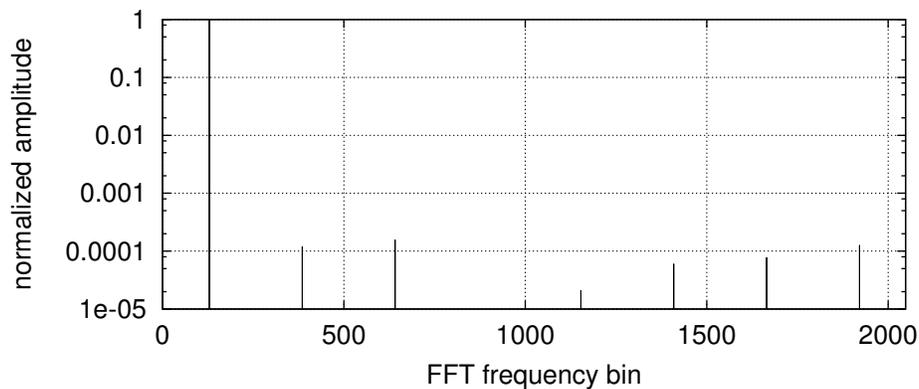


Figure 11: FFT of even divisor number of cycles (4096 points, 128 cycles).

Figure 12 shows the FFT of a 4096-point data record that contains 127 periods of a perfect sine wave. Since 127 is both odd and prime, there are no common factors between the number of input periods and the number of samples in the data record. Because of these choices, there are no repeating patterns and every sample in the data record is unique. This FFT correctly shows the frequency impulse due to the input sine wave and the noise floor due to quantization, without any measurement-induced artifacts.

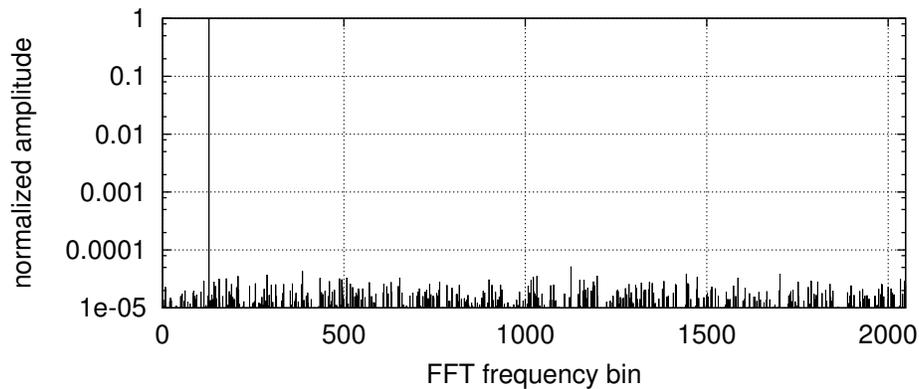


Figure 12: FFT of non-divisor prime number of cycles (4096 points, 127 cycles).

2.3 Histogram Test for Linearity

The linearity (INL and DNL) of the A/D converter can be determined with a histogram test [8, 9]. A histogram of output digital codes is recorded for a large number of samples for an input sine wave. The results are compared to the number of samples expected from the theoretical sine-wave probability-density function. If the input sine wave is

$$V = A \sin \omega t$$

then the probability-density function is

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

The probability of a sample being in the range (V_a, V_b) is found by integration

$$P(V_a, V_b) = \int_{V_b}^{V_a} p(V) dV = \frac{1}{\pi} \left(\arcsin \frac{V_b}{A} - \arcsin \frac{V_a}{A} \right)$$

The difference between the measured and expected probability of observing a specific output code is a function of that specific code width, which can be used to calculate the differential nonlinearity.

Due to the statistical nature of this test, the length of the data record for the histogram test can be quite large [9]. The number of samples, M , required for a high-confidence measurement is

$$M = \frac{\pi 2^{N-1} Z_{\alpha/2}^2}{\beta^2}$$

where N is the number of bits, β is the DNL-measurement resolution, and for a 99% confidence level, $Z_{\alpha/2} = 2.576$. Thus for a 6-bit A/D converter, and a DNL-measurement resolution of 0.1 LSB

$$M = \frac{\pi 2^{N-1} Z_{\alpha/2}^2}{\beta^2} = \frac{\pi 2^5 (2.576)^2}{(0.1)^2} = 67,000$$

If the sample size is large enough, this test will work with an asynchronous input signal, however a coherent input signal with unique samples (odd and prime frequency ratio, as explained above) is also possible. Sampling an input signal harmonically related to the sampling frequency would create the same problems that occur with the FFT tests. Figure 13 shows the histogram for a 6-bit A/D converter using an insufficient number of points, while Figure 14 shows a properly smooth histogram result from using enough data.

Once the histogram data is collected, the experimental code widths are from the measured data H_k using the expected probability [5]

$$P_k = \frac{1}{\pi} \left[\arcsin \left(\frac{(k+1)V_Q}{A} \right) - \arcsin \left(\frac{kV_Q}{A} \right) \right]$$

The differential nonlinearity is thus

$$\text{DNL}_k = 1 \text{ LSB} \left(\frac{H_k}{P_k} - 1 \right)$$

However, this method is sensitive to errors in the measurement of the input sine wave amplitude A .

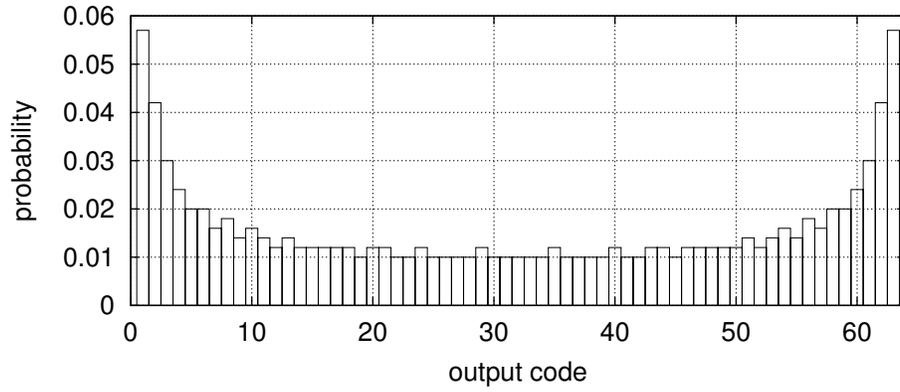


Figure 13: Histogram of 1000 points for 6-bit A/D converter (insufficient).

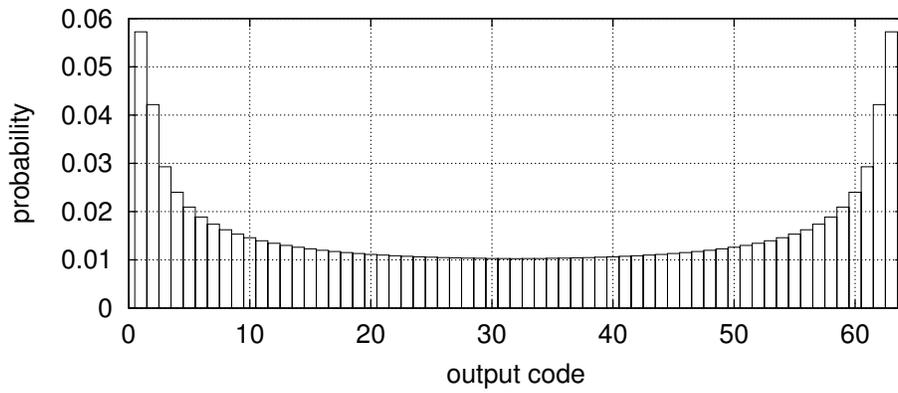


Figure 14: Histogram of 100,000 points for 6-bit A/D converter.

A better method is to calculate the INL and DNL from a cumulative histogram [9]. First, the offset is found by equating the number of positive samples M_p and the number of negative samples M_n

$$M_n = \sum_{k=1}^{2^{N-1}} H_k \quad M_p = \sum_{k=2^{N-1}+1}^{2^N} H_k$$

$$V_{OS} = \frac{A\pi}{2} \sin\left(\frac{M_p - M_n}{M_p + M_n}\right)$$

Second, the transition voltages are found from

$$V_j = -A \cos\left(\frac{\pi}{M} \sum_{k=0}^j H_k\right)$$

Once the transition voltages are known, the INL and DNL follow

$$\text{INL}_j = \frac{V_j - V_1}{1 \text{ LSB}} \quad \text{DNL}_j = \frac{V_{j+1} - V_j}{1 \text{ LSB}}$$

This method makes the amplitude A a linear factor in the calculations, which reduces the sensitivity to errors in A , and makes the final result easily normalizable.

2.4 Sine Wave Curve Fit for ENOB

Another way to calculate the effective number of bits is by using a sine wave curve fit [7]. A least-squared-error sine wave is fit to the measured data, and compared to the input waveform. The resulting rms error of the curve fit, E , is a measure of the effective number of bits lost due to A/D converter error sources. The effective number of bits (ENOB) can be calculated from

$$\text{ENOB} = N - \log_2 \left(\frac{E_{\text{rms}}}{V_Q/\sqrt{12}} \right)$$

Since the frequency of the input and the output must be the same, only the amplitude, offset, and phase of the output sinusoid must be found.

2.4.1 General Form

The fixed-frequency sine-wave curve-fit method [10] is used to find the best fit. The sinusoid to fit to the output data is

$$x_n = A \cos(\omega t_n) + B \sin(\omega t_n) + C$$

where t_n are the sample times, ω is the input frequency (in radians/second) and A , B , and C are the fit parameters. The squared error between the samples y_n and the curve fit is

$$E = \sum_{k=1}^M (y_k - x_k)^2 = \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C]^2$$

The error is minimized by setting the partial derivatives with respect to the fit parameters to zero

$$\begin{aligned} 0 = \frac{\partial E}{\partial A} &= -2 \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C] \cos(\omega t_k) \\ 0 = \frac{\partial E}{\partial B} &= -2 \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C] \sin(\omega t_k) \\ 0 = \frac{\partial E}{\partial C} &= -2 \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C] \end{aligned}$$

Defining $\alpha_k = \cos(\omega t_k)$ and $\beta_k = \sin(\omega t_k)$ and rearranging terms gives a set of linear equations

$$\begin{aligned} \sum_{k=1}^M y_k \alpha_k &= A \sum_{k=1}^M \alpha_k^2 + B \sum_{k=1}^M \alpha_k \beta_k + C \sum_{k=1}^M \alpha_k \\ \sum_{k=1}^M y_k \beta_k &= A \sum_{k=1}^M \alpha_k \beta_k + B \sum_{k=1}^M \beta_k^2 + C \sum_{k=1}^M \beta_k \\ \sum_{k=1}^M y_k &= A \sum_{k=1}^M \alpha_k + B \sum_{k=1}^M \beta_k + CM \end{aligned}$$

These equations can be expressed as a single linear equation $Y = UX$

$$\sum_{k=1}^M \begin{bmatrix} y_k \alpha_k \\ y_k \beta_k \\ y_k \end{bmatrix} = \left(\sum_{k=1}^M \begin{bmatrix} \alpha_k^2 & \alpha_k \beta_k & \alpha_k \\ \alpha_k \beta_k & \beta_k^2 & \beta_k \\ \alpha_k & \beta_k & 1 \end{bmatrix} \right) \begin{bmatrix} A \\ B \\ C \end{bmatrix}$$

This linear equation has a solution $X = U^{-1}Y$. The total squared error from above is

$$E = \sum_{k=1}^M [y_k - A\alpha_k - B\beta_k - C]^2$$

and can be rewritten as

$$E = \sum_{k=1}^M [y_k^2 - 2Ay_k\alpha_k - 2By_k\beta_k - 2Cy_k + A^2\alpha_k^2 + 2AB\alpha_k\beta_k + 2AC\alpha_k + B^2\beta_k^2 + 2BC\beta_k + C^2]$$

$$E = \left(\sum_{k=1}^M y_k^2 \right) - 2X^T Y + X^T U X$$

The rms error of the curve fit is then

$$E_{\text{rms}} = \sqrt{\frac{E}{M}}$$

So the effective number of bits (ENOB) is

$$\text{ENOB} = N - \frac{1}{2} \log_2 \left(\frac{12E}{V_Q^2 M} \right)$$

2.4.2 Simplified Coherent Form

One advantage of the curve-fit method of finding the effective number of bits is, unlike the methods that use the FFT and SNDR, it does not require windowing or coherent sampling. However, if coherent sampling is used, such that the expected data record has an integer number of sine-wave cycles, the linear equation above simplifies significantly. The sums over α_k , β_k , and $\alpha_k\beta_k$ vanish, and

$$\sum_{k=1}^M \alpha_k^2 = \sum_{k=1}^M \beta_k^2 = \frac{M}{2}$$

The linear equation becomes

$$\sum_{k=1}^M \begin{bmatrix} y_k \alpha_k \\ y_k \beta_k \\ y_k \end{bmatrix} = \begin{bmatrix} M/2 & 0 & 0 \\ 0 & M/2 & 0 \\ 0 & 0 & M \end{bmatrix} \begin{bmatrix} A \\ B \\ C \end{bmatrix}$$

which has simple, closed-form solutions

$$A = \frac{2}{M} \sum_{k=1}^M y_k \alpha_k \quad B = \frac{2}{M} \sum_{k=1}^M y_k \beta_k \quad C = \frac{1}{M} \sum_{k=1}^M y_k$$

The least squared error simplifies to

$$\begin{aligned} E &= \sum_{k=1}^M \left[y_k^2 - 2Ay_k\alpha_k - 2By_k\beta_k - 2Cy_k + A^2\alpha_k^2 + 2AB\alpha_k\beta_k + 2AC\alpha_k + B^2\beta_k^2 + 2BC\beta_k + C^2 \right] \\ &= \left(\sum_{k=1}^M y_k^2 \right) - MA^2 - MB^2 - 2MC^2 + \frac{M}{2}A^2 + 0 + 0 + \frac{M}{2}B^2 + 0 + MC^2 \\ &= \left(\sum_{k=1}^M y_k^2 \right) - \frac{M}{2}(A^2 + B^2 + 2C^2) \end{aligned}$$

The effective number of bits (ENOB) remains

$$\text{ENOB} = N - \frac{1}{2} \log_2 \left(\frac{12E}{V_Q^2 M} \right)$$

2.4.3 Comparison of Methods

Clearly, calculating the effective number of bits using an FFT and the signal-to-noise-and-distortion ratio should produce the same result as the sine-wave curve-fit method. Figure 15 shows part of a 4096-point quantized sine wave with added noise, plotted with the best-fit sine wave, calculated as described above. The effective number of bits from the curve fit is 3.02. Figure 16 shows the FFT of the same data. The effective number of bits calculated from SNDR is 3.01, which matches nicely with the curve-fit result.

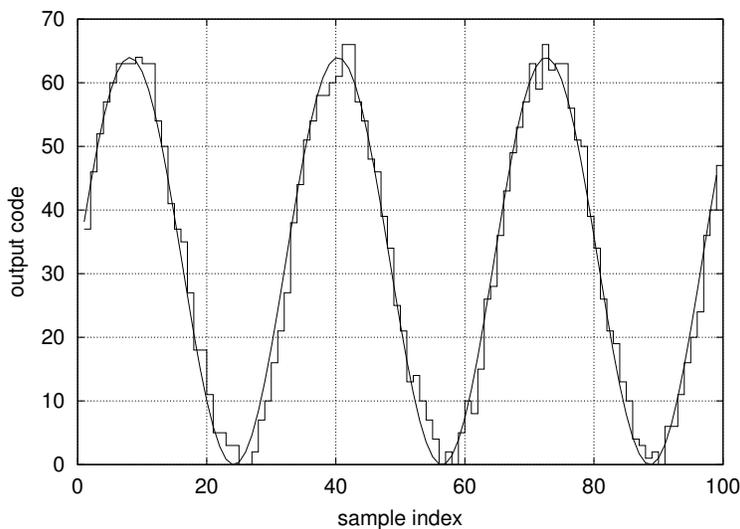


Figure 15: Partial data plot showing sine wave and added noise, along with best-fit sine wave. ENOB calculated from curve fit is 3.01.

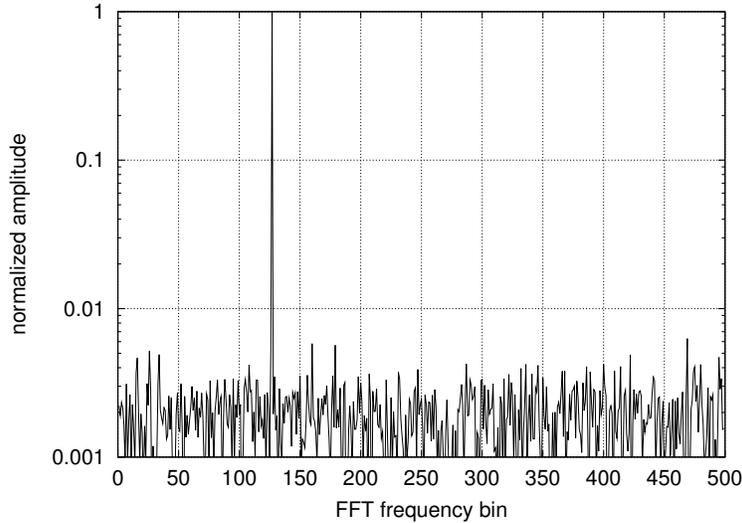


Figure 16: Partial FFT showing sine wave and added noise. ENOB calculated from SNDR is 3.02.

2.5 Overall Noise and Aperture Jitter

Overall noise of an A/D converter can be measured by grounding the input of the A/D converter (in a unipolar A/D converter, the input should be connected to a mid-scale constant voltage source) and accumulating a histogram. Only the center code bin should have counts in it. Any spread in the histogram around the center code bin is caused by noise in the A/D converter. This test can also be used to determine the offset of the A/D converter as in the histogram test,

Aperture jitter is measured by repeatedly sampling the same voltage of the input waveform. For example, a sine wave input is used, and the A/D converter is triggered to repeatedly sample the positive-slope zero crossing. If the input sine wave and the sampling clock are generated from phase-locked sources, there should be no spread in the output digital codes from this measurement. However, a real A/D converter will produce a spread in output codes due to aperture jitter.

The aperture jitter is calculated from a histogram of output codes produced from this measurement. For an input sine wave sampling at the zero crossings, the aperture jitter is

$$t_a = \frac{V_{\text{rms}}}{2\pi f A}$$

where A is the amplitude and f is the frequency of the input sine wave. As the amplitude of the input increases, the slope at the zero crossings increases, and the spread of output codes should proportionally increase due to aperture jitter.

References

- [1] Walter Kester. Measure flash-ADC performance for trouble-free operation. *EDN*, pages 103–114, February 1, 1990.
- [2] Walter Kester. Flash ADCs provide the basis for high-speed conversion. *EDN*, pages 101–110, January 4, 1990.
- [3] Analog Devices Inc. AD678 12-bit 200 kSPS complete sampling ADC. *Datasheet*, 2000.
- [4] Robert H. Walden. Analog-to-digital converter survey and analysis. *IEEE Journal on Selected Areas in Communication*, 17(4):539–550, April 1999.
- [5] Brendan Coleman, Pat Meehan, John Reidy, and Pat Weeks. Coherent sampling helps when specifying DSP A/D converters. *EDN*, pages 145–152, October 15, 1987.
- [6] Fredric J. Harris. On the use of windows for harmonic analysis with the discrete fourier transform. *Proceedings of the IEEE*, 66(1):51–83, January 1978.
- [7] Walter Kester. DSP test techniques keep flash ADCs in check. *EDN*, pages 133–142, January 18, 1990.
- [8] Walter Kester. Test video A/D converters under dynamic conditions. *EDN*, pages 103–112, August 18, 1982.
- [9] Joey Doernberg, Hae-Seung Lee, and David A. Hodges. Full speed testing of A/D converters. *IEEE Journal of Solid State Circuits*, 19(6):820–827, December 1984.
- [10] Brian J. Frohring, Bruce E. Peetz, Mark A. Unkrich, and Steven C. Bird. Waveform recorder design for dynamic performance. *Hewlett-Packard Journal*, 39(1):39–48, February 1988.