

SECTION 2

SAMPLED DATA SYSTEMS

- Discrete Time Sampling of Analog Signals
- ADC and DAC Static Transfer Functions and DC Errors
- AC Errors in Data Converters
- DAC Dynamic Performance

SAMPLED DATA SYSTEMS

SECTION 2

SAMPLED DATA SYSTEMS

Walt Kester, James Bryant

INTRODUCTION

A block diagram of a typical sampled data DSP system is shown in Figure 2.1. Prior to the actual analog-to-digital conversion, the analog signal usually passes through some sort of signal conditioning circuitry which performs such functions as amplification, attenuation, and filtering. The lowpass/bandpass filter is required to remove unwanted signals outside the bandwidth of interest and prevent aliasing.

FUNDAMENTAL SAMPLED DATA SYSTEM

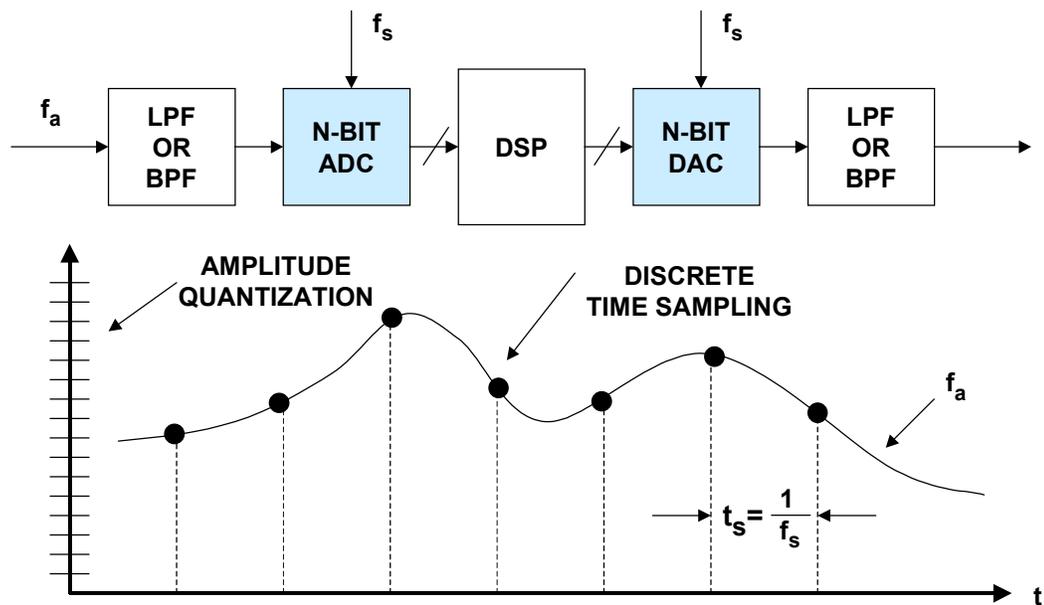


Figure 2.1

The system shown in Figure 2.1 is a real-time system, i.e., the signal to the ADC is continuously sampled at a rate equal to f_s , and the ADC presents a new sample to the DSP at this rate. In order to maintain real-time operation, the DSP must perform all its required computation within the sampling interval, $1/f_s$, and present an output sample to the DAC before arrival of the next sample from the ADC. An example of a typical DSP function would be a digital filter.

In the case of FFT analysis, a block of data is first transferred to the DSP memory. The FFT is calculated at the same time a new block of data is transferred into the memory, in order to maintain real-time operation. The DSP must calculate the FFT during the data transfer interval so it will be ready to process the next block of data.

Note that the DAC is required only if the DSP data must be converted back into an analog signal (as would be the case in a voiceband or audio application, for example). There are many applications where the signal remains entirely in digital format after the initial A/D conversion. Similarly, there are applications where the DSP is solely responsible for generating the signal to the DAC, such as in CD player electronics. If a DAC is used, it must be followed by an analog anti-imaging filter to remove the image frequencies.

There are two key concepts involved in the actual analog-to-digital and digital-to-analog conversion process: *discrete time sampling* and *finite amplitude resolution due to quantization*. An understanding of these concepts is vital to DSP applications.

DISCRETE TIME SAMPLING OF ANALOG SIGNALS

The concepts of *discrete time sampling* and *quantization* of an analog signal are shown in Figure 2.1. The continuous analog data must be sampled at discrete intervals, $t_s = 1/f_s$ which must be carefully chosen to insure an accurate representation of the original analog signal. It is clear that the more samples taken (faster sampling rates), the more accurate the digital representation, but if fewer samples are taken (lower sampling rates), a point is reached where critical information about the signal is actually lost. This leads us to the statement of Nyquist's criteria given in Figure 2.2.

NYQUIST'S CRITERIA

- A signal with a **bandwidth** f_a must be sampled at a rate $f_s > 2f_a$ or information about the signal will be lost.
- Aliasing occurs whenever $f_s < 2f_a$
- The concept of aliasing is widely used in communications applications such as direct IF-to-digital conversion.

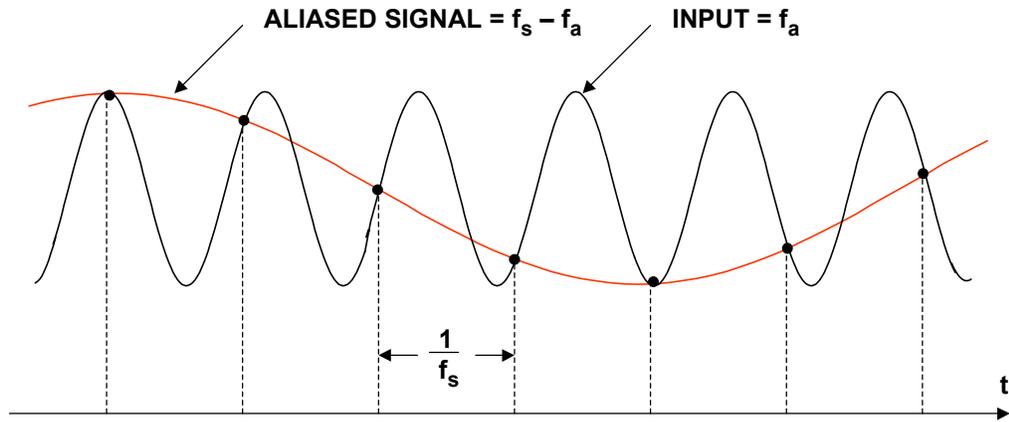
Figure 2.2

Simply stated, the Nyquist Criteria requires that the sampling frequency be at least twice the signal bandwidth, or information about the signal will be lost. If the sampling frequency is less than twice the analog signal bandwidth, a phenomena known as aliasing will occur.

In order to understand the implications of *aliasing* in both the time and frequency domain, first consider case of a time domain representation of a single tone sinuswave sampled as shown in Figure 2.3. In this example, the sampling frequency f_s is only slightly more than the analog input frequency f_a , and the Nyquist criteria is violated. Notice that the pattern of the actual samples produces an *aliased* sinuswave at a lower frequency equal to $f_s - f_a$.

The corresponding frequency domain representation of this scenario is shown in Figure 2.4B. Now consider the case of a single frequency sinewave of frequency f_a sampled at a frequency f_s by an ideal impulse sampler (see Figure 2.4A). Also assume that $f_s > 2f_a$ as shown. The frequency-domain output of the sampler shows *aliases* or *images* of the original signal around every multiple of f_s , i.e. at frequencies equal to $|\pm Kf_s \pm f_a|$, $K = 1, 2, 3, 4, \dots$

ALIASING IN THE TIME DOMAIN



NOTE: f_a IS SLIGHTLY LESS THAN f_s

Figure 2.3

ANALOG SIGNAL f_a SAMPLED @ f_s USING IDEAL SAMPLER HAS IMAGES (ALIASSES) AT $|\pm Kf_s \pm f_a|$, $K = 1, 2, 3, \dots$

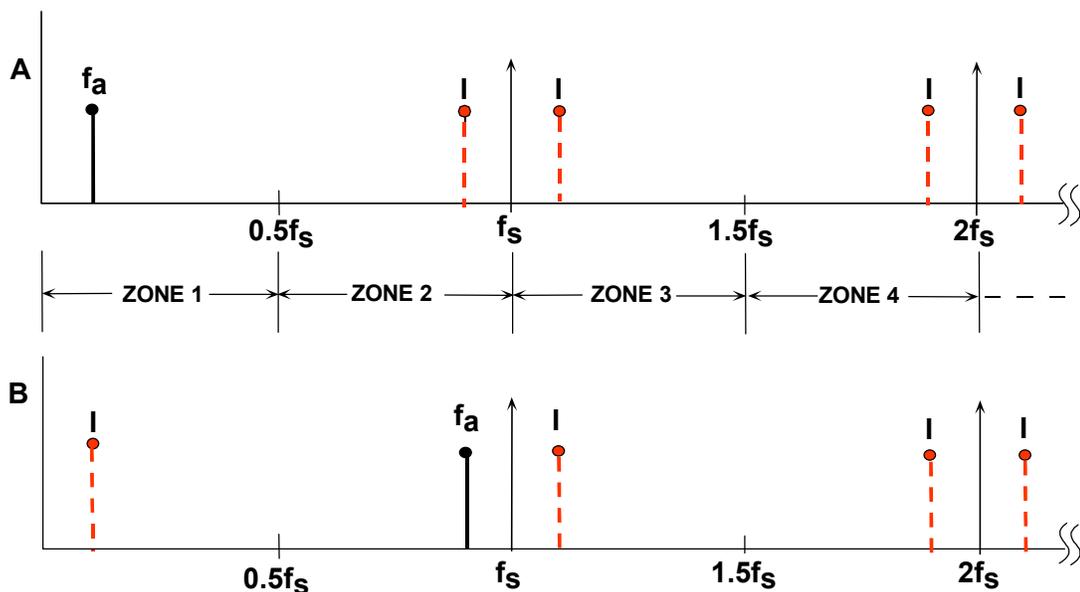


Figure 2.4

SAMPLED DATA SYSTEMS

The *Nyquist* bandwidth is defined to be the frequency spectrum from DC to $f_s/2$. The frequency spectrum is divided into an infinite number of *Nyquist zones*, each having a width equal to $0.5f_s$ as shown. In practice, the ideal sampler is replaced by an ADC followed by an FFT processor. The FFT processor only provides an output from DC to $f_s/2$, i.e., the signals or aliases which appear in the first Nyquist zone.

Now consider the case of a signal which is outside the first Nyquist zone (Figure 2.4B). The signal frequency is only slightly less than the sampling frequency, corresponding to the condition shown in the time domain representation in Figure 2.3. Notice that even though the signal is outside the first Nyquist zone, its image (or *alias*), $f_s - f_a$, falls inside. Returning to Figure 2.4A, it is clear that if an unwanted signal appears at any of the image frequencies of f_a , it will also occur at f_a , thereby producing a spurious frequency component in the first Nyquist zone.

This is similar to the analog mixing process and implies that some filtering ahead of the sampler (or ADC) is required to remove frequency components which are outside the Nyquist bandwidth, but whose aliased components fall inside it. The filter performance will depend on how close the out-of-band signal is to $f_s/2$ and the amount of attenuation required.

Baseband Antialiasing Filters

Baseband sampling implies that the signal to be sampled lies in the first Nyquist zone. It is important to note that with no input filtering at the input of the ideal sampler, *any frequency component (either signal or noise) that falls outside the Nyquist bandwidth in any Nyquist zone will be aliased back into the first Nyquist zone*. For this reason, an antialiasing filter is used in almost all sampling ADC applications to remove these unwanted signals.

Properly specifying the antialiasing filter is important. The first step is to know the characteristics of the signal being sampled. Assume that the highest frequency of interest is f_a . The antialiasing filter passes signals from DC to f_a while attenuating signals above f_a .

Assume that the corner frequency of the filter is chosen to be equal to f_a . The effect of the finite transition from minimum to maximum attenuation on system dynamic range is illustrated in Figure 2.5A.

Assume that the input signal has fullscale components well above the maximum frequency of interest, f_a . The diagram shows how fullscale frequency components above $f_s - f_a$ are aliased back into the bandwidth DC to f_a . These aliased components are indistinguishable from actual signals and therefore limit the dynamic range to the value on the diagram which is shown as *DR*.

Some texts recommend specifying the antialiasing filter with respect to the Nyquist frequency, $f_s/2$, but this assumes that the signal bandwidth of interest extends from DC to $f_s/2$ which is rarely the case. In the example shown in Figure 2.5A, the aliased components between f_a and $f_s/2$ are not of interest and do not limit the dynamic range.

The antialiasing filter transition band is therefore determined by the corner frequency f_a , the stopband frequency $f_s - f_a$, and the desired stopband attenuation, DR. The required system dynamic range is chosen based on the requirement for signal fidelity.

OVERSAMPLING RELAXES REQUIREMENTS ON BASEBAND ANTIALIASING FILTER

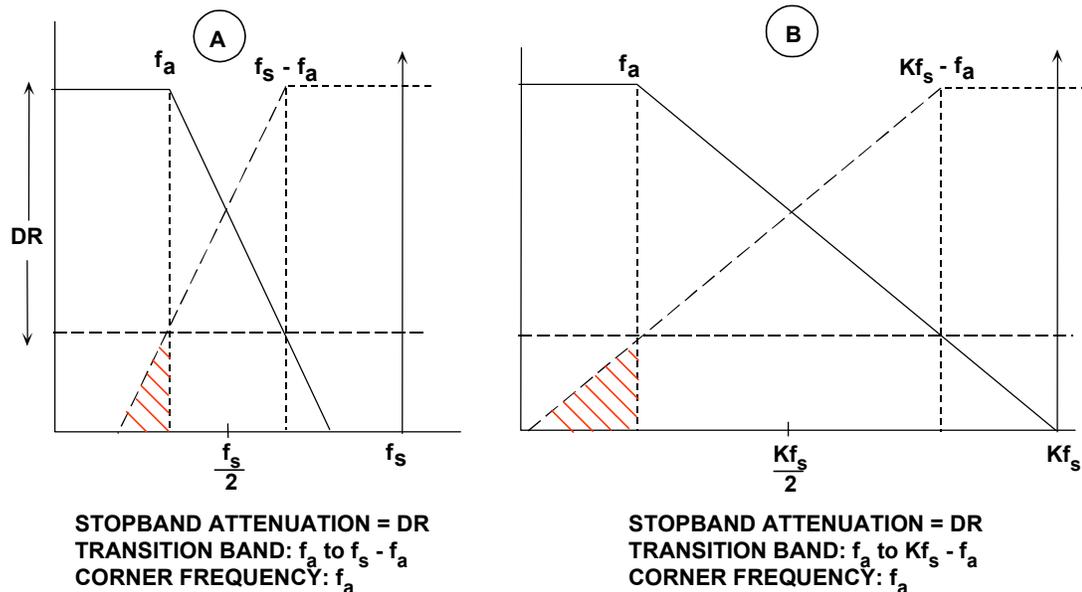


Figure 2.5

Filters become more complex as the transition band becomes sharper, all other things being equal. For instance, a Butterworth filter gives 6dB attenuation per octave for each filter pole. Achieving 60dB attenuation in a transition region between 1MHz and 2MHz (1 octave) requires a minimum of 10 poles - not a trivial filter, and definitely a design challenge.

Therefore, other filter types are generally more suited to high speed applications where the requirement is for a sharp transition band and in-band flatness coupled with linear phase response. Elliptic filters meet these criteria and are a popular choice. There are a number of companies which specialize in supplying custom analog filters. TTE is an example of such a company (Reference 1).

From this discussion, we can see how the sharpness of the antialiasing transition band can be traded off against the ADC sampling frequency. Choosing a higher sampling rate (oversampling) reduces the requirement on transition band sharpness (hence, the filter complexity) at the expense of using a faster ADC and processing data at a faster rate. This is illustrated in Figure 2.5B which shows the effects of increasing the sampling frequency by a factor of K, while maintaining the same analog corner frequency, f_a , and the same dynamic range, DR, requirement. The wider transition band (f_a to $Kf_s - f_a$) makes this filter easier to design than for the case of Figure 2.5A.

SAMPLED DATA SYSTEMS

The antialiasing filter design process is started by choosing an initial sampling rate of 2.5 to 4 times f_a . Determine the filter specifications based on the required dynamic range and see if such a filter is realizable within the constraints of the system cost and performance. If not, consider a higher sampling rate which may require using a faster ADC. It should be mentioned that sigma-delta ADCs are inherently oversampling converters, and the resulting relaxation in the analog anti-aliasing filter requirements is therefore an added benefit of this architecture.

The antialiasing filter requirements can also be relaxed somewhat if it is certain that there will never be a fullscale signal at the stopband frequency $f_s - f_a$. In many applications, it is improbable that fullscale signals will occur at this frequency. If the maximum signal at the frequency $f_s - f_a$ will never exceed X dB below fullscale, then the filter stopband attenuation requirement is reduced by that same amount. The new requirement for stopband attenuation at $f_s - f_a$ based on this knowledge of the signal is now only $DR - X$ dB. When making this type of assumption, be careful to treat any noise signals which may occur above the maximum signal frequency f_a as unwanted signals which will also alias back into the signal bandwidth.

Undersampling (Harmonic Sampling, Bandpass Sampling, IF Sampling, Direct IF to Digital Conversion)

Thus far we have considered the case of baseband sampling, i.e., all the signals of interest lie within the first Nyquist zone. Figure 2.6A shows such a case, where the band of sampled signals is limited to the first Nyquist zone, and images of the original band of frequencies appear in each of the other Nyquist zones.

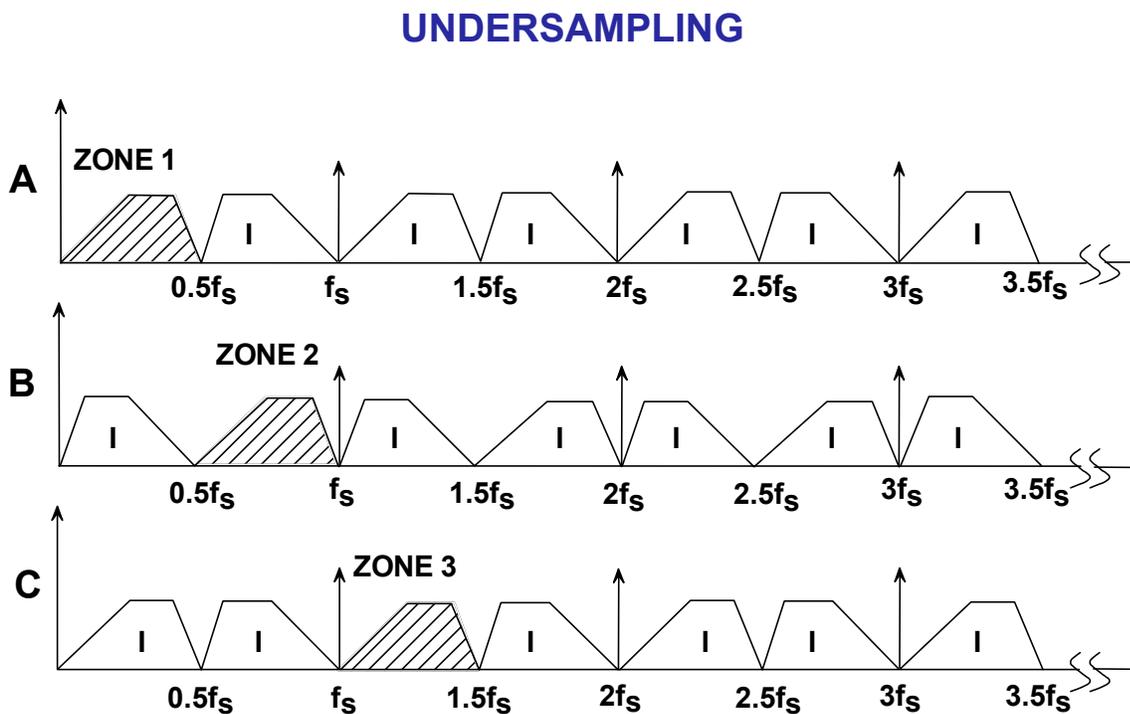


Figure 2.6

Consider the case shown in Figure 2.6B, where the sampled signal band lies entirely within the second Nyquist zone. The process of sampling a signal outside the first Nyquist zone is often referred to as *undersampling*, or *harmonic sampling*. Note that the first Nyquist zone image contains all the information in the original signal, with the exception of its original location (the order of the frequency components within the spectrum is reversed, but this is easily corrected by re-ordering the output of the FFT).

Figure 2.6C shows the sampled signal restricted to the third Nyquist zone. Note that the first Nyquist zone image has no frequency reversal. In fact, the sampled signal frequencies may lie in *any* unique Nyquist zone, and the first Nyquist zone image is still an accurate representation (with the exception of the frequency reversal which occurs when the signals are located in even Nyquist zones). At this point we can clearly restate the Nyquist criteria:

*A signal must be sampled at a rate equal to or greater than twice its **bandwidth** in order to preserve all the signal information.*

Notice that there is no mention of the absolute *location* of the band of sampled signals within the frequency spectrum relative to the sampling frequency. The only constraint is that the band of sampled signals be restricted to a *single* Nyquist zone, i.e., the signals must not overlap any multiple of $f_s/2$ (this, in fact, is the primary function of the antialiasing filter).

Sampling signals above the first Nyquist zone has become popular in communications because the process is equivalent to analog demodulation. It is becoming common practice to sample IF signals directly and then use digital techniques to process the signal, thereby eliminating the need for the IF demodulator. Clearly, however, as the IF frequencies become higher, the dynamic performance requirements on the ADC become more critical. The ADC input bandwidth and distortion performance must be adequate at the IF frequency, rather than only baseband. This presents a problem for most ADCs designed to process signals in the first Nyquist zone, therefore an ADC suitable for undersampling applications must maintain dynamic performance into the higher order Nyquist zones.

ADC AND DAC STATIC TRANSFER FUNCTIONS AND DC ERRORS

The most important thing to remember about both DACs and ADCs is that either the input or output is digital, and therefore the signal is quantized. That is, an N-bit word represents one of 2^N possible states, and therefore an N-bit DAC (with a fixed reference) can have only 2^N possible analog outputs, and an N-bit ADC can have only 2^N possible digital outputs. The analog signals will generally be voltages or currents.

The resolution of data converters may be expressed in several different ways: the weight of the Least Significant Bit (LSB), parts per million of full scale (ppm FS), millivolts (mV), etc. Different devices (even from the same manufacturer) will be specified differently, so converter users must learn to translate between the different types of specifications if they are to compare devices successfully. The size of the least significant bit for various resolutions is shown in Figure 2.7.

QUANTIZATION: THE SIZE OF A LEAST SIGNIFICANT BIT (LSB)

RESOLUTION N	2^N	VOLTAGE (10V FS)	ppm FS	% FS	dB FS
2-bit	4	2.5 V	250,000	25	-12
4-bit	16	625 mV	62,500	6.25	-24
6-bit	64	156 mV	15,625	1.56	-36
8-bit	256	39.1 mV	3,906	0.39	-48
10-bit	1,024	9.77 mV (10 mV)	977	0.098	-60
12-bit	4,096	2.44 mV	244	0.024	-72
14-bit	16,384	610 μ V	61	0.0061	-84
16-bit	65,536	153 μ V	15	0.0015	-96
18-bit	262,144	38 μ V	4	0.0004	-108
20-bit	1,048,576	9.54 μ V (10 μ V)	1	0.0001	-120
22-bit	4,194,304	2.38 μ V	0.24	0.000024	-132
24-bit	16,777,216	596 nV*	0.06	0.000006	-144

*600nV is the Johnson Noise in a 10kHz BW of a 2.2k Ω Resistor @ 25°C

Remember: 10-bits and 10V FS yields an LSB of 10mV, 1000ppm, or 0.1%.
All other values may be calculated by powers of 2.

Figure 2.7

Before we can consider the various architectures used in data converters, it is necessary to consider the performance to be expected, and the specifications which are important. The following sections will consider the definition of errors and specifications used for data converters. This is important in understanding the strengths and weaknesses of different ADC/DAC architectures.

The first applications of data converters were in measurement and control where the exact timing of the conversion was usually unimportant, and the data rate was slow. In such applications, the DC specifications of converters are important, but timing and AC specifications are not. Today many, if not most, converters are used in *sampling* and *reconstruction* systems where AC specifications are critical (and DC ones may not be) - these will be considered in the next part of this section.

Figure 2.8 shows the ideal transfer characteristics for a 3-bit unipolar DAC, and Figure 2.9 a 3-bit unipolar ADC. In a DAC, both the input and the output are quantized, and the graph consists of eight points - while it is reasonable to discuss the line through these points, it is very important to remember that the actual transfer characteristic is *not* a line, but a number of discrete points.

TRANSFER FUNCTION FOR IDEAL 3-BIT DAC

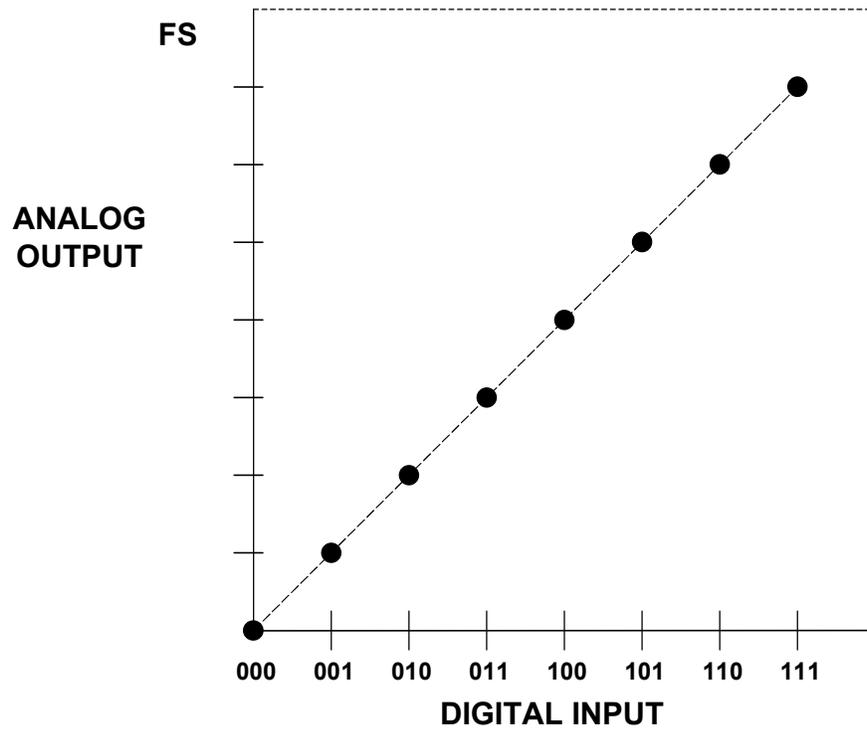


Figure 2.8

TRANSFER FUNCTION FOR IDEAL 3-BIT ADC

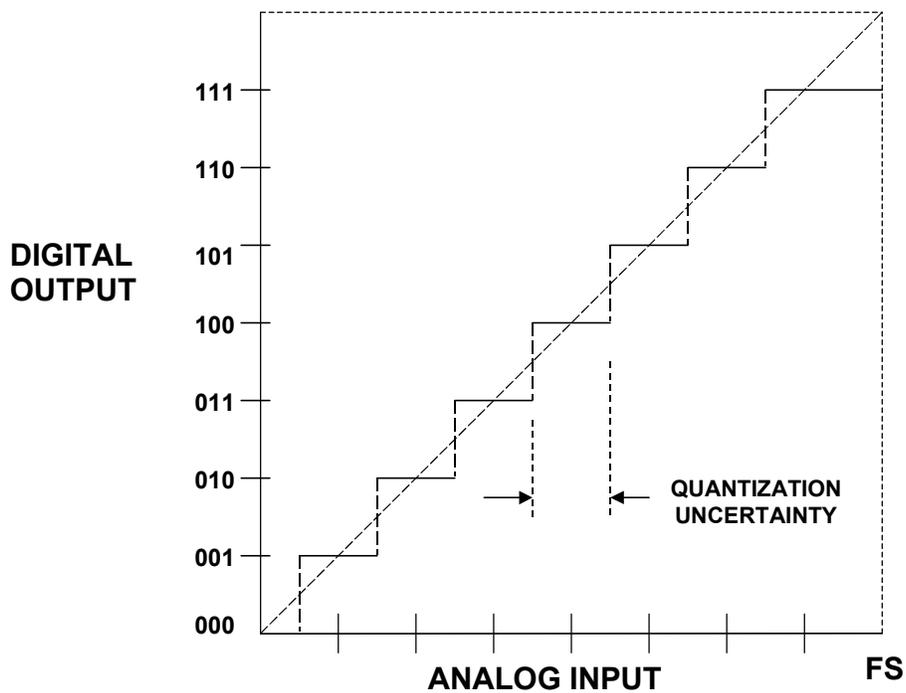


Figure 2.9

SAMPLED DATA SYSTEMS

The input to an ADC is analog and is not quantized, but its output is quantized. The transfer characteristic therefore consists of eight horizontal steps (when considering the offset, gain and linearity of an ADC we consider the line joining the midpoints of these steps).

In both cases, digital full scale (all "1"s) corresponds to 1 LSB below the analog full scale (the reference, or some multiple thereof). This is because, as mentioned above, the digital code represents the *normalized* ratio of the analog signal to the reference.

The (ideal) ADC transitions take place at $\frac{1}{2}$ LSB above zero, and thereafter every LSB, until $1\frac{1}{2}$ LSB below analog full scale. Since the analog input to an ADC can take any value, but the digital output is quantized, there may be a difference of up to $\frac{1}{2}$ LSB between the actual analog input and the exact value of the digital output. This is known as the *quantization error* or *quantization uncertainty* as shown in Figure 2.9. In AC (sampling) applications this quantization error gives rise to *quantization noise* which will be discussed in the next section.

There are many possible digital coding schemes for data converters: *binary*, *offset binary*, *1's complement*, *2's complement*, *gray code*, *BCD* and others. This section, being devoted mainly to the *analog* issues surrounding data converters, will use simple *binary* and *offset binary* in its examples and will not consider the merits and disadvantages of these, or any other forms of digital code.

The examples in Figures 2.8 and 2.9 use *unipolar* converters, whose analog port has only a single polarity. These are the simplest type, but *bipolar* converters are generally more useful in real-world applications. There are two types of bipolar converters: the simpler is merely a unipolar converter with an accurate 1 MSB of negative offset (and many converters are arranged so that this offset may be switched in and out so that they can be used as either unipolar or bipolar converters at will), but the other, known as a *sign-magnitude* converter is more complex, and has N bits of magnitude information and an additional bit which corresponds to the sign of the analog signal. Sign-magnitude DACs are quite rare, and sign-magnitude ADCs are found mostly in digital voltmeters (DVMs).

The four DC errors in a data converter are *offset error*, *gain error*, and two types of *linearity error*. Offset and gain errors are analogous to offset and gain errors in amplifiers as shown in Figure 2.10 for a bipolar input range. (Though offset error and zero error, which are identical in amplifiers and unipolar data converters, are not identical in bipolar converters and should be carefully distinguished.) The transfer characteristics of both DACs and ADCs may be expressed as $D = K + GA$, where D is the digital code, A is the analog signal, and K and G are constants. In a unipolar converter, K is zero, and in an offset bipolar converter, it is -1 MSB. The offset error is the amount by which the actual value of K differs from its ideal value. The gain error is the amount by which G differs from its ideal value, and is generally expressed as the percentage difference between the two, although it may be defined as the gain error contribution (in mV or LSB) to the total error at full scale. These errors can usually be trimmed by the data converter user. Note, however, that amplifier offset is trimmed at zero input, and then the gain is trimmed near to full scale. The trim algorithm for a bipolar data converter is not so straightforward.

CONVERTER OFFSET AND GAIN ERROR

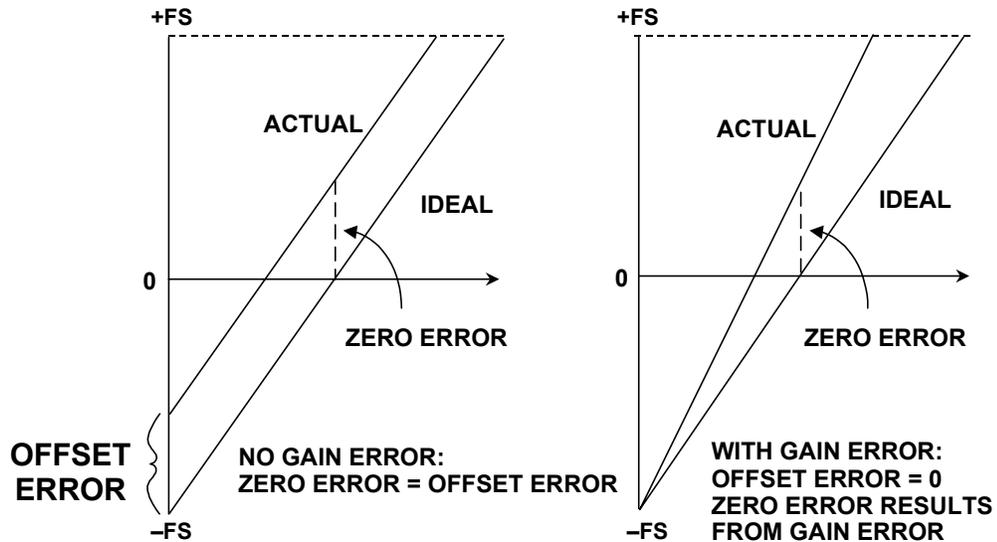


Figure 2.10

The integral linearity error of a converter is also analogous to the linearity error of an amplifier, and is defined as the maximum deviation of the actual transfer characteristic of the converter from a straight line, and is generally expressed as a percentage of full scale (but may be given in LSBs). There are two common ways of choosing the straight line: *end point* and *best straight line* (see Figure 2.11).

METHOD OF MEASURING INTEGRAL LINEARITY ERRORS (SAME CONVERTER ON BOTH GRAPHS)

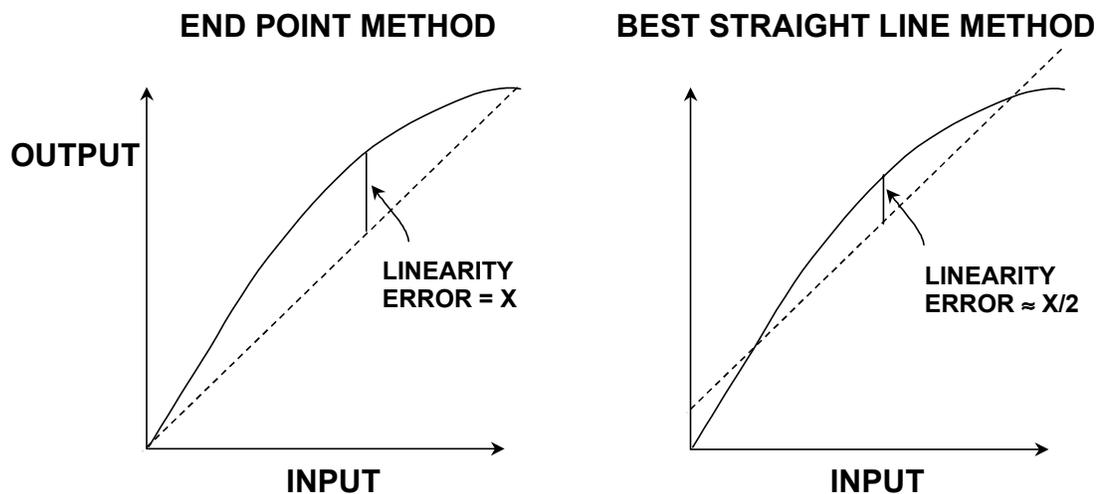


Figure 2.11

SAMPLED DATA SYSTEMS

In the end point system, the deviation is measured from the straight line through the origin and the full scale point (after gain adjustment). This is the most useful integral linearity measurement for measurement and control applications of data converters (since error budgets depend on deviation from the ideal transfer characteristic, not from some arbitrary "best fit"), and is the one normally adopted by Analog Devices, Inc.

The best straight line, however, does give a better prediction of distortion in AC applications, and also gives a lower value of "linearity error" on a data sheet. The best fit straight line is drawn through the transfer characteristic of the device using standard curve fitting techniques, and the maximum deviation is measured from this line. In general, the integral linearity error measured in this way is only 50% of the value measured by end point methods. This makes the method good for producing impressive data sheets, but it is less useful for error budget analysis. For AC applications, it is even better to specify distortion than DC linearity, so it is rarely necessary to use the best straight line method to define converter linearity.

The other type of converter non-linearity is *differential non-linearity* (DNL). This relates to the linearity of the code transitions of the converter. In the ideal case, a change of 1 LSB in digital code corresponds to a change of exactly 1 LSB of analog signal. In a DAC, a change of 1 LSB in digital code produces exactly 1 LSB change of analog output, while in an ADC there should be exactly 1 LSB change of analog input to move from one digital transition to the next.

Where the change in analog signal corresponding to 1 LSB digital change is more or less than 1 LSB, there is said to be a DNL error. The DNL error of a converter is normally defined as the maximum value of DNL to be found at any transition.

If the DNL of a DAC is less than -1 LSB at any transition (see Figure 2.12), the DAC is *non-monotonic* i.e., its transfer characteristic contains one or more localized maxima or minima. A DNL greater than $+1$ LSB does not cause non-monotonicity, but is still undesirable. In many DAC applications (especially closed-loop systems where non-monotonicity can change negative feedback to positive feedback), it is critically important that DACs are monotonic. DAC monotonicity is often explicitly specified on data sheets, although if the DNL is guaranteed to be less than 1 LSB (i.e., $|DNL| \leq 1\text{LSB}$) then the device must be monotonic, even without an explicit guarantee.

ADCs can be non-monotonic, but a more common result of excess DNL in ADCs is *missing codes* (see Figure 2.13). Missing codes (or non-monotonicity) in an ADC are as objectionable as non-monotonicity in a DAC. Again, they result from $DNL > 1$ LSB.

TRANSFER FUNCTION OF NON-IDEAL 3-BIT DAC

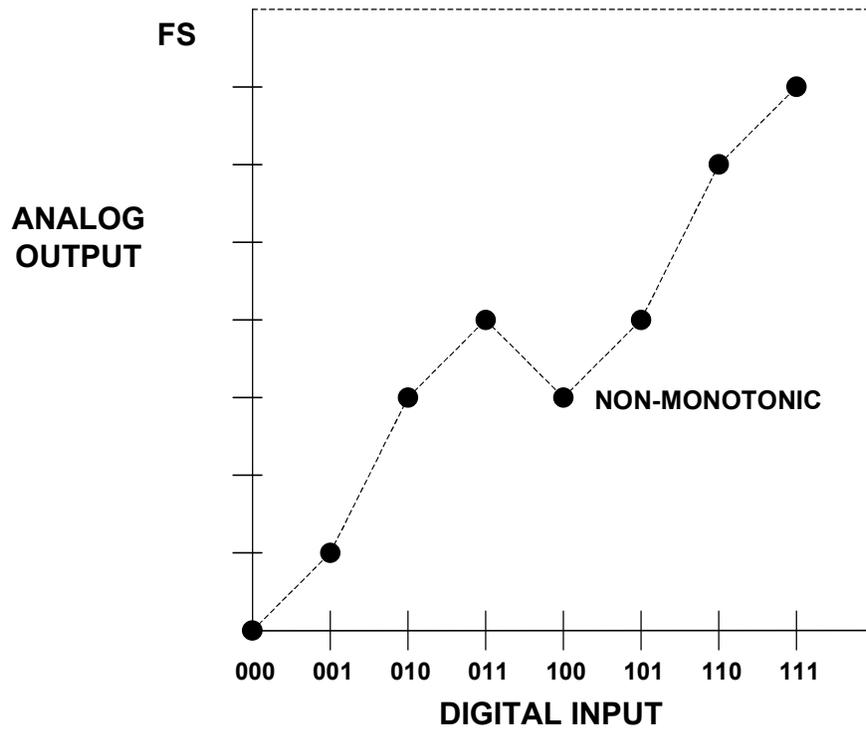


Figure 2.12

TRANSFER FUNCTION OF NON-IDEAL 3-BIT ADC

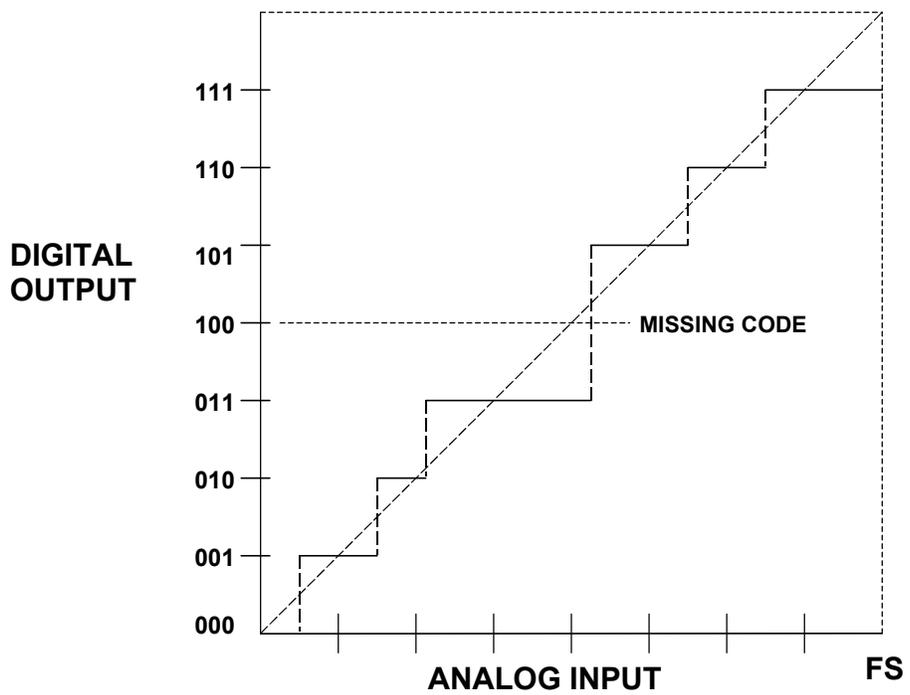


Figure 2.13

Defining missing codes is more difficult than defining non-monotonicity. All ADCs suffer from some transition noise as shown in Figure 2.14 (think of it as the flicker between adjacent values of the last digit of a DVM). As resolutions become higher, the range of input over which transition noise occurs may approach, or even exceed, 1 LSB. In such a case, especially if combined with a negative DNL error, it may be that there are some (or even all) codes where transition noise is present for the whole range of inputs. There are therefore some codes for which there is *no* input which will *guarantee* that code as an output, although there may be a range of inputs which will *sometimes* produce that code.

COMBINED EFFECTS OF ADC CODE TRANSITION NOISE AND DNL

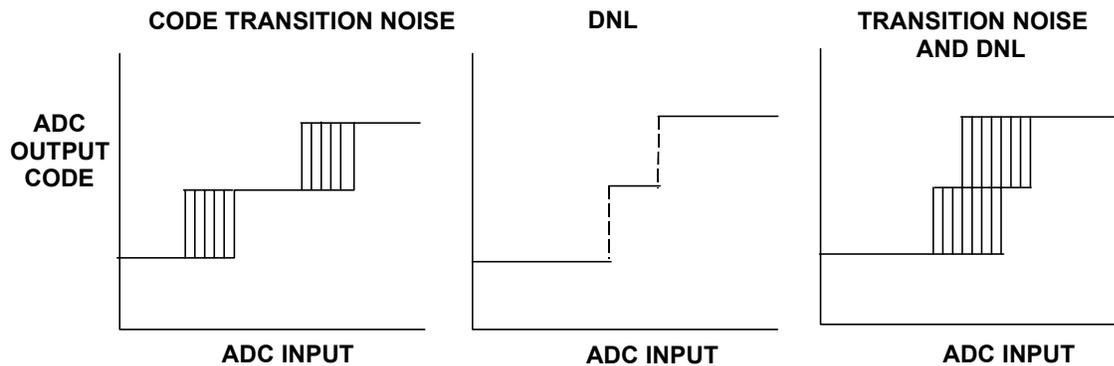


Figure 2.14

For lower resolution ADCs, it may be reasonable to define *no missing codes* as a combination of transition noise and DNL which guarantees some level (perhaps 0.2 LSB) of noise-free code for all codes. However, this is impossible to achieve at the very high resolutions achieved by modern sigma-delta ADCs, or even at lower resolutions in wide bandwidth sampling ADCs. In these cases, the manufacturer must define noise levels and resolution in some other way. Which method is used is less important, but the data sheet should contain a clear definition of the method used and the performance to be expected.

AC ERRORS IN DATA CONVERTERS

Over the last decade, a major application of data converters is in AC sampling and reconstruction. In very simple terms, a *sampled data system* is a system where the instantaneous value of an AC waveform is sampled at regular intervals. The resulting digital codes may be used to store the waveform (as in CDs and DATs), or intensive computation on the samples (Digital Signal Processing, or DSP) may be used to perform filtering, compression, and other operations. The inverse operation, reconstruction, occurs when a series of digital codes are fed to a DAC to reconstruct

an AC waveform - an obvious example of this is a CD or DAT player, but the technique is very widely used indeed in telecommunications, radio, synthesizers, and many other applications.

The data converters used in these applications must have good performance with AC signals, but may not require good DC specifications. The first high performance converters to be designed for such applications were often manufactured with good AC specifications but poor, or unspecified, DC performance. Today the design tradeoffs are better understood, and most converters will have good, and guaranteed, AC and DC specifications. DACs for digital audio, however, which must be extremely competitive in price, are generally sold with comparatively poor DC specifications - not because their DC performance is poor, but because it is not tested during manufacture.

While it is easier to discuss the DC parameters of both DACs and ADCs together, their AC specifications are sufficiently different to deserve separate consideration.

Distortion and Noise in an Ideal N-Bit ADC

Thus far we have looked at the implications of the sampling process without considering the effects of ADC quantization. We will now treat the ADC as an ideal sampler, but include the effects of quantization.

The only errors (DC or AC) associated with an ideal N-bit ADC are those related to the sampling and quantization processes. The maximum error an ideal ADC makes when digitizing a DC input signal is $\pm 1/2\text{LSB}$. Any AC signal applied to an ideal N-bit ADC will produce quantization noise whose rms value (measured over the Nyquist bandwidth, DC to $f_s/2$) is approximately equal to the weight of the least significant bit (LSB), q , divided by $\sqrt{12}$. (See Reference 2). This assumes that the signal is at least a few LSBs in amplitude so that the ADC output always changes state. The quantization error signal from a linear ramp input is approximated as a sawtooth waveform with a peak-to-peak amplitude equal to q , and its rms value is therefore $q/\sqrt{12}$ (see Figure 2.15).

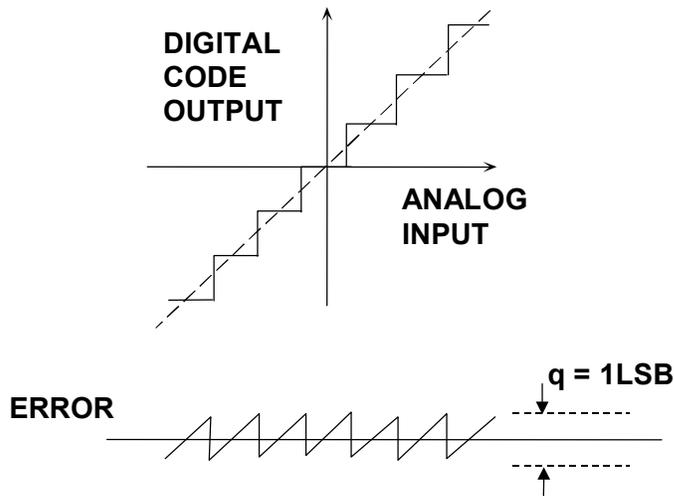
It can be shown that the ratio of the rms value of a full scale sinewave to the rms value of the quantization noise (expressed in dB) is:

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

where N is the number of bits in the ideal ADC. *This equation is only valid if the noise is measured over the entire Nyquist bandwidth from DC to $f_s/2$ as shown in Figure 2.16. If the signal bandwidth, BW , is less than $f_s/2$, then the SNR within the signal bandwidth BW is increased because the amount of quantization noise within the signal bandwidth is smaller. The correct expression for this condition is given by:*

$$\text{SNR} = 6.02N + 1.76\text{dB} + 10 \log\left(\frac{f_s}{2 \cdot BW}\right).$$

IDEAL N-BIT ADC QUANTIZATION NOISE



$$\text{RMS ERROR} = q/\sqrt{12}$$

$$\text{SNR} = 6.02N + 1.76\text{dB} + 10\log \left[\frac{f_s}{2 \cdot \text{BW}} \right] \text{ FOR FS SINEWAVE}$$

Figure 2.15

QUANTIZATION NOISE SPECTRUM

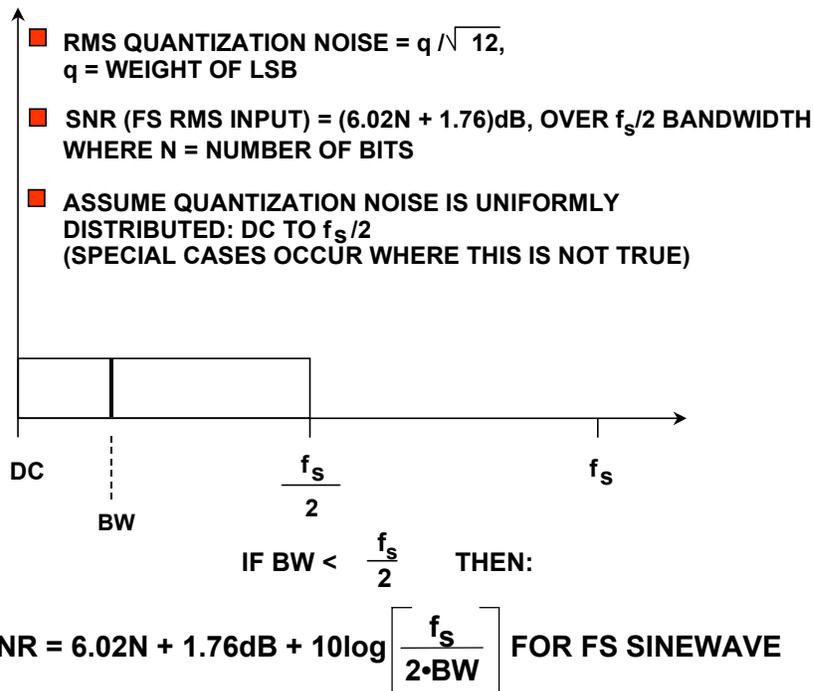


Figure 2.16

The above equation reflects the condition called *oversampling*, where the sampling frequency is higher than twice the signal bandwidth. The correction term is often called *processing gain*. Notice that for a given signal bandwidth, doubling the sampling frequency increases the SNR by 3dB.

Although the rms value of the noise is accurately approximated by $q/\sqrt{12}$, its frequency domain content may be highly correlated to the AC input signal. For instance, there is greater correlation for low amplitude periodic signals than for large amplitude random signals. Quite often, the assumption is made that the theoretical quantization noise appears as white noise, spread uniformly over the Nyquist bandwidth DC to $f_s/2$. Unfortunately, this is not true. In the case of strong correlation, the quantization noise appears concentrated at the various harmonics of the input signal, just where you don't want them.

In most applications, the input to the ADC is a band of frequencies (usually summed with some noise), so the quantization noise tends to be random. In spectral analysis applications (or in performing FFTs on ADCs using spectrally pure sinewaves - see Figure 2.17), however, the correlation between the quantization noise and the signal depends upon the ratio of the sampling frequency to the input signal. This is demonstrated in Figure 2.18, where an ideal 12-bit ADCs output is analyzed using a 4096-point FFT. In the left-hand FFT plot, the ratio of the sampling frequency to the input frequency was chosen to be exactly 32, and the worst harmonic is about 76dB below the fundamental. The right hand diagram shows the effects of slightly offsetting the ratio, showing a relatively random noise spectrum, where the SFDR is now about 92dBc. In both cases, the rms value of all the noise components is $q/\sqrt{12}$, but in the first case, the noise is concentrated at harmonics of the fundamental.

DYNAMIC PERFORMANCE ANALYSIS OF AN IDEAL N-BIT ADC

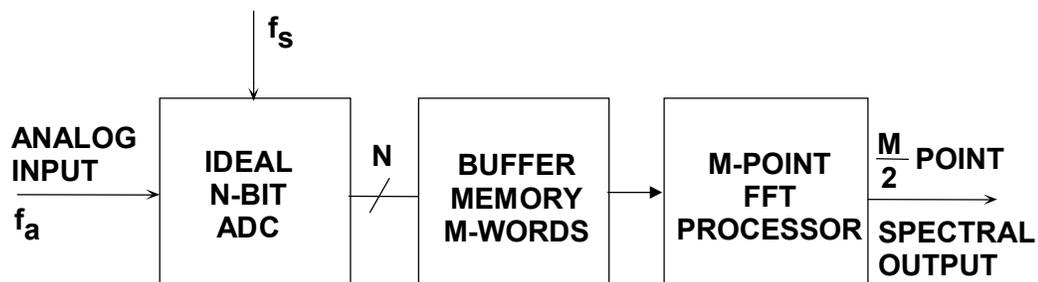


Figure 2.17

EFFECT OF RATIO OF SAMPLING CLOCK TO INPUT FREQUENCY ON SFDR FOR IDEAL 12-BIT ADC

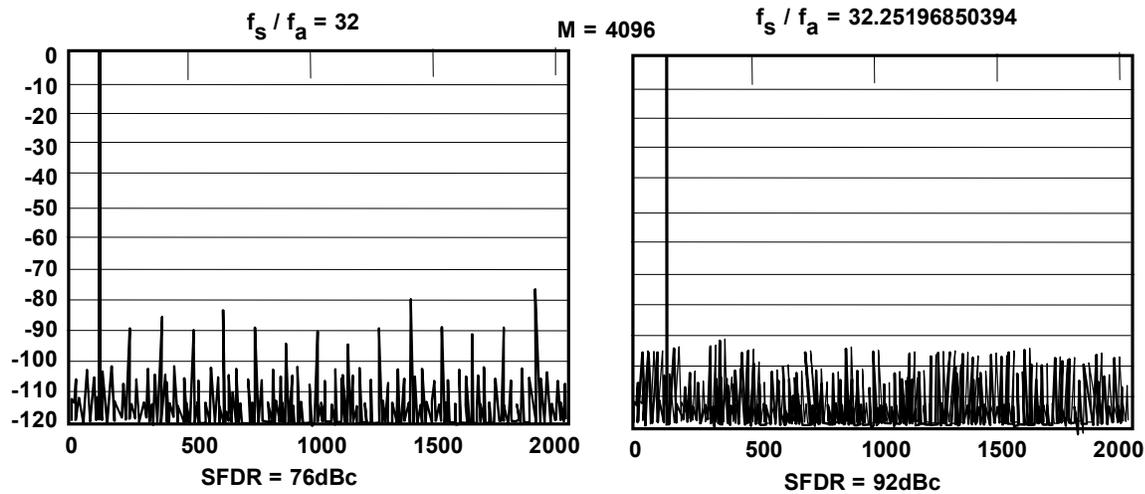


Figure 2.18

Note that this variation in the apparent harmonic distortion of the ADC is an artifact of the sampling process and the correlation of the quantization error with the input frequency. In a practical ADC application, the quantization error generally appears as random noise because of the random nature of the wideband input signal and the additional fact that there is a usually a small amount of system noise which acts as a *dither* signal to further randomize the quantization error spectrum.

It is important to understand the above point, because single-tone sinewave FFT testing of ADCs is a universally accepted method of performance evaluation. In order to accurately measure the harmonic distortion of an ADC, steps must be taken to ensure that the test setup truly measures the ADC distortion, not the artifacts due to quantization noise correlation. This is done by properly choosing the frequency ratio and sometimes by injecting a small amount of noise (dither) with the input signal.

Now, return to Figure 2.18, and note that the average value of the noise floor of the FFT is approximately 100dB below full scale, but the theoretical SNR of a 12-bit ADC is 74dB. The FFT noise floor is *not* the SNR of the ADC, because the FFT acts like an analog spectrum analyzer with a bandwidth of f_s/M , where M is the number of points in the FFT. The theoretical FFT noise floor is therefore $10\log_{10}(M/2)$ dB below the quantization noise floor due to the so-called *processing gain* of the FFT (see Figure 2.19). In the case of an ideal 12-bit ADC with an SNR of 74dB, a 4096-point FFT would result in a processing gain of $10\log_{10}(4096/2) = 33$ dB, thereby

resulting in an overall FFT noise floor of $74+33=107\text{dBc}$. In fact, the FFT noise floor can be reduced even further by going to larger and larger FFTs; just as an analog spectrum analyzer's noise floor can be reduced by narrowing the bandwidth. When testing ADCs using FFTs, it is important to ensure that the FFT size is large enough so that the distortion products can be distinguished from the FFT noise floor itself.

NOISE FLOOR FOR AN IDEAL 12-BIT ADC USING 4096-POINT FFT

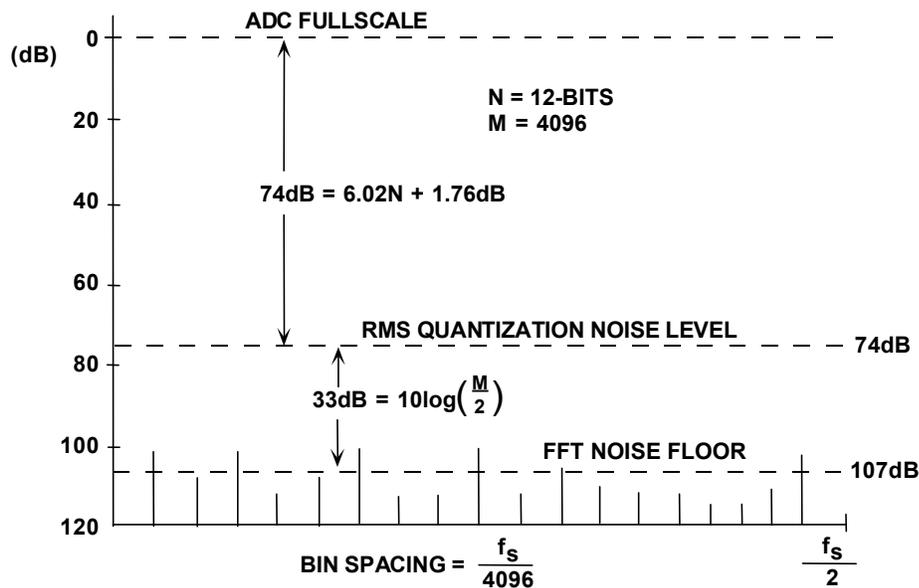


Figure 2.19

Distortion and Noise in Practical ADCs

A practical sampling ADC (one that has an integral sample-and-hold), regardless of architecture, has a number of noise and distortion sources as shown in Figure 2.20. The wideband analog front-end buffer has wideband noise, non-linearity, and also finite bandwidth. The SHA introduces further non-linearity, bandlimiting, and aperture jitter. The actual quantizer portion of the ADC introduces quantization noise, and both integral and differential non-linearity. In this discussion, assume that sequential outputs of the ADC are loaded into a buffer memory of length M and that the FFT processor provides the spectral output. Also assume that the FFT arithmetic operations themselves introduce no significant errors relative to the ADC. However, when examining the output noise floor, the FFT processing gain (dependent on M) must be considered.

ADC MODEL SHOWING NOISE AND DISTORTION SOURCES

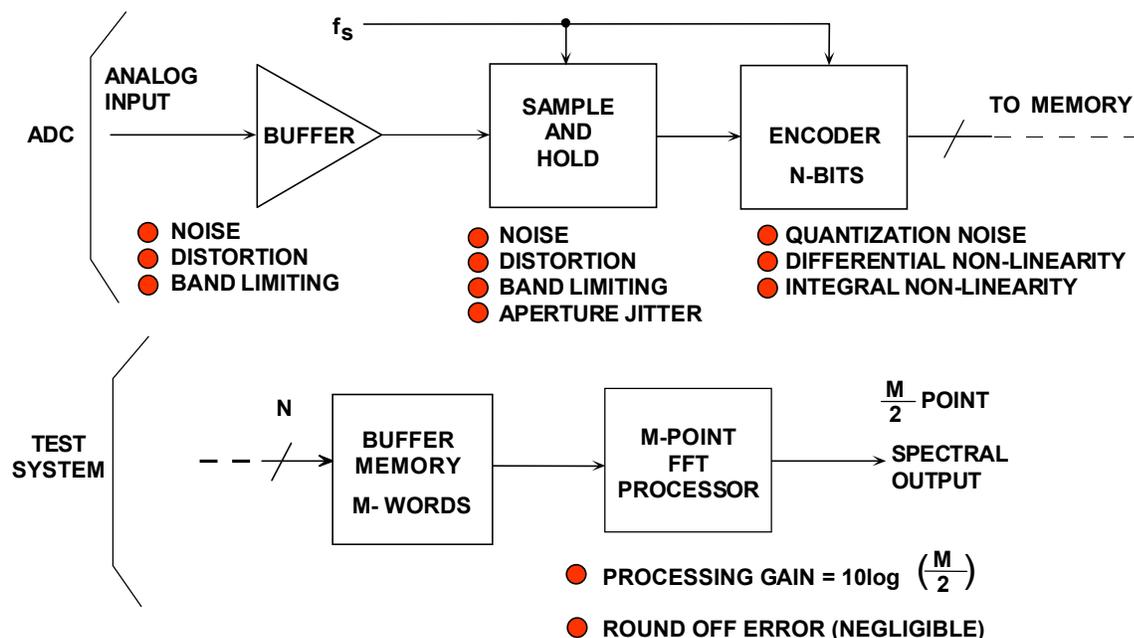


Figure 2.20

Equivalent Input Referred Noise (Thermal Noise)

The wideband ADC internal circuits produce a certain amount of wideband rms noise due to thermal and kT/C effects. This noise is present even for DC input signals, and accounts for the fact that the output of most wideband (or high resolution) ADCs is a distribution of codes, centered around the nominal value of a DC input (see Figure 2.21). To measure its value, the input of the ADC is grounded, and a large number of output samples are collected and plotted as a histogram (sometimes referred to as a *grounded-input* histogram). Since the noise is approximately Gaussian, the standard deviation of the histogram is easily calculated (see Reference 3), corresponding to the effective input rms noise. It is common practice to express this rms noise in terms of LSBs, although it can be expressed as an rms voltage.

There are various ways to characterize the AC performance of ADCs. In the early years of ADC technology (over 30 years ago) there was little standardization with respect to AC specifications, and measurement equipment and techniques were not well understood or available. Over nearly a 30 year period, manufacturers and customers have learned more about measuring the dynamic performance of converters, and the specifications shown in Figure 2.22 represent the most popular ones used today. Practically all the specifications represent the converter's performance in the frequency domain. The FFT is the heart of practically all these measurements and is discussed in detail in Section 5 of this book.

EFFECT OF INPUT-REFERRED NOISE ON ADC "GROUNDED INPUT" HISTOGRAM

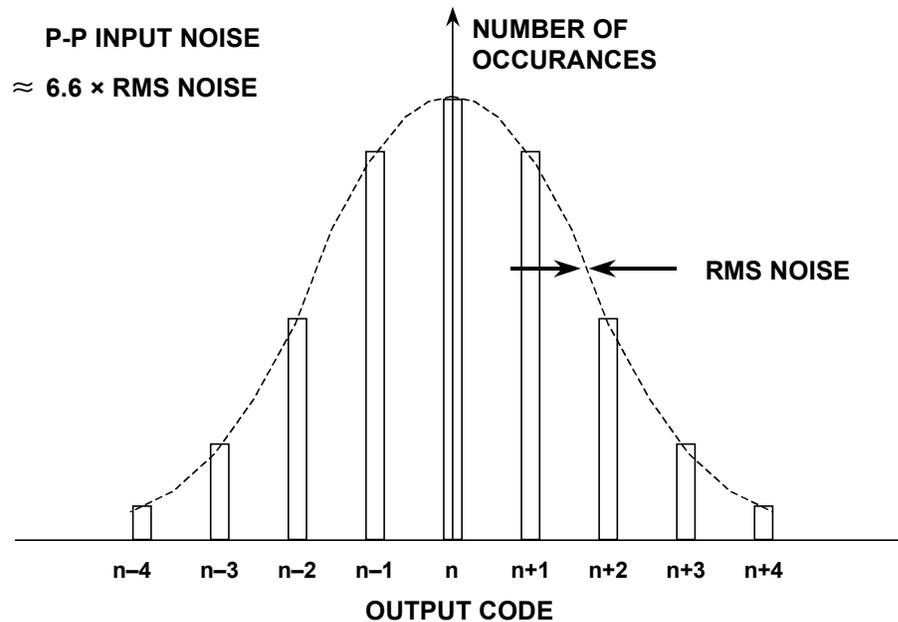


Figure 2.21

QUANTIFYING ADC DYNAMIC PERFORMANCE

- Harmonic Distortion
- Worst Harmonic
- Total Harmonic Distortion (THD)
- Total Harmonic Distortion Plus Noise (THD + N)
- Signal-to-Noise-and-Distortion Ratio (SINAD, or S/N +D)
- Effective Number of Bits (ENOB)
- Signal-to-Noise Ratio (SNR)
- Analog Bandwidth (Full-Power, Small-Signal)
- Spurious Free Dynamic Range (SFDR)
- Two-Tone Intermodulation Distortion
- Multi-tone Intermodulation Distortion

Figure 2.22

Integral and Differential Non-Linearity Distortion Effects

One of the first things to realize when examining the nonlinearities of data converters is that the transfer function of a data converter has artifacts which do not occur in conventional linear devices such as op amps or gain blocks. The overall integral non-linearity of an ADC is due to the integral non-linearity of the front-end and SHA as well as the overall integral non-linearity in the ADC transfer function. However, *differential non-linearity is due exclusively to the encoding process* and may vary considerably dependent on the ADC encoding architecture. Overall integral non-linearity produces distortion products whose amplitude varies as a function of the input signal amplitude. For instance, second-order intermodulation products increase 2dB for every 1dB increase in signal level, and third-order products increase 3dB for every 1dB increase in signal level.

The differential non-linearity in the ADC transfer function produces distortion products which not only depend on the amplitude of the signal but the positioning of the differential non-linearity along the ADC transfer function. Figure 2.23 shows two ADC transfer functions having differential non-linearity. The left-hand diagram shows an error which occurs at midscale. Therefore, for both large and small signals, the signal crosses through this point producing a distortion product which is relatively independent of the signal amplitude. The right-hand diagram shows another ADC transfer function which has differential non-linearity errors at 1/4 and 3/4 full scale. Signals which are above 1/2 scale peak-to-peak will exercise these codes and produce distortion, while those less than 1/2 scale peak-to-peak will not.

TYPICAL ADC / DAC DNL ERRORS

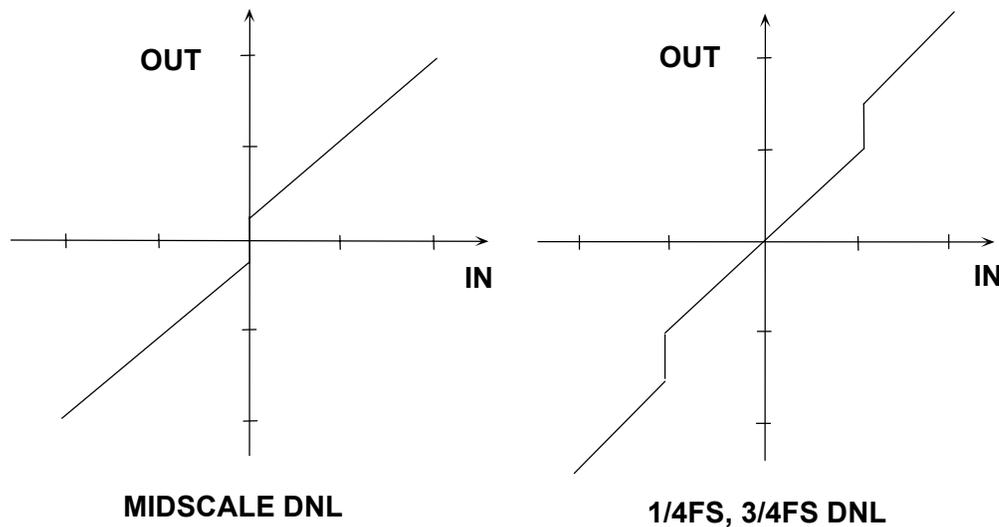


Figure 2.23

Most high-speed ADCs are designed so that differential non-linearity is spread across the entire ADC range. Therefore, for signals which are within a few dB of full scale, the overall integral non-linearity of the transfer function determines the distortion products. For lower level signals, however, the harmonic content becomes

dominated by the differential non-linearities and does not generally decrease proportionally with decreases in signal amplitude.

Harmonic Distortion, Worst Harmonic, Total Harmonic Distortion (THD), Total Harmonic Distortion Plus Noise (THD + N)

There are a number of ways to quantify the distortion of an ADC. An FFT analysis can be used to measure the amplitude of the various harmonics of a signal. The harmonics of the input signal can be distinguished from other distortion products by their location in the frequency spectrum. Figure 2.24 shows a 7MHz input signal sampled at 20MSPS and the location of the first 9 harmonics. Aliased harmonics of f_a fall at frequencies equal to $|\pm Kf_s \pm nf_a|$, where n is the order of the harmonic, and $K = 0, 1, 2, 3, \dots$. The second and third harmonics are generally the only ones specified on a data sheet because they tend to be the largest, although some data sheets may specify the value of the *worst* harmonic. *Harmonic distortion* is normally specified in dBc (decibels below *carrier*), although at audio frequencies it may be specified as a percentage. Harmonic distortion is generally specified with an input signal near full scale (generally 0.5 to 1dB below full scale to prevent clipping), but it can be specified at any level. For signals much lower than full scale, other distortion products due to the DNL of the converter (not direct harmonics) may limit performance.

**LOCATION OF HARMONIC DISTORTION PRODUCTS:
INPUT SIGNAL = 7MHz, SAMPLING RATE = 20MSPS**

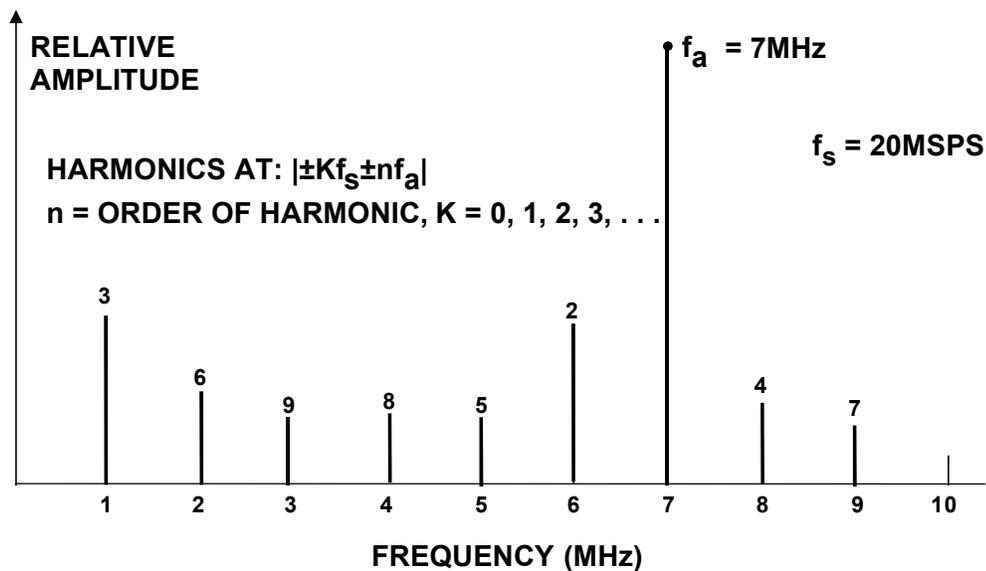


Figure 2.24

Total harmonic distortion (THD) is the ratio of the rms value of the fundamental signal to the mean value of the root-sum-square of its harmonics (generally, only the first 5 are significant). THD of an ADC is also generally specified with the input signal close to full scale, although it can be specified at any level.

Total harmonic distortion plus noise (THD+ N) is the ratio of the rms value of the fundamental signal to the mean value of the root-sum-square of its harmonics plus all noise components (excluding DC). The bandwidth over which the noise is measured must be specified. In the case of an FFT, the bandwidth is DC to $f_s/2$. (If the bandwidth of the measurement is DC to $f_s/2$, THD+N is equal to SINAD - see below).

Signal-to-Noise-and-Distortion Ratio (SINAD), Signal-to-Noise Ratio (SNR), and Effective Number of Bits (ENOB)

SINAD and SNR deserve careful attention, because there is still some variation between ADC manufacturers as to their precise meaning. Signal-to-noise-and-Distortion (SINAD, or S/N+D) is the ratio of the rms signal amplitude to the mean value of the root-sum-square (RSS) of all other spectral components, *including harmonics*, but excluding DC. SINAD is a good indication of the overall dynamic performance of an ADC as a function of input frequency because it includes all components which make up noise (including thermal noise) and distortion. It is often plotted for various input amplitudes. SINAD is equal to THD+N if the bandwidth for the noise measurement is the same. A typical plot for the AD9220 12-bit, 10MSPS ADC is shown in Figure 2.26.

SINAD, ENOB, AND SNR

■ **SINAD (Signal-to-Noise-and-Distortion Ratio):**

- ◆ The ratio of the rms signal amplitude to the mean value of the root-sum-squares (RSS) of all other spectral components, including harmonics, but excluding DC.

■ **ENOB (Effective Number of Bits):**

$$\text{ENOB} = \frac{\text{SINAD} - 1.76\text{dB}}{6.02}$$

■ **SNR (Signal-to-Noise Ratio, or Signal-to-Noise Ratio Without Harmonics):**

- ◆ The ratio of the rms signal amplitude to the mean value of the root-sum-squares (RSS) of all other spectral components, excluding the first 5 harmonics and DC

Figure 2.25

AD9220 12-BIT, 10MSPS ADC SINAD AND ENOB FOR VARIOUS INPUT SIGNAL LEVELS

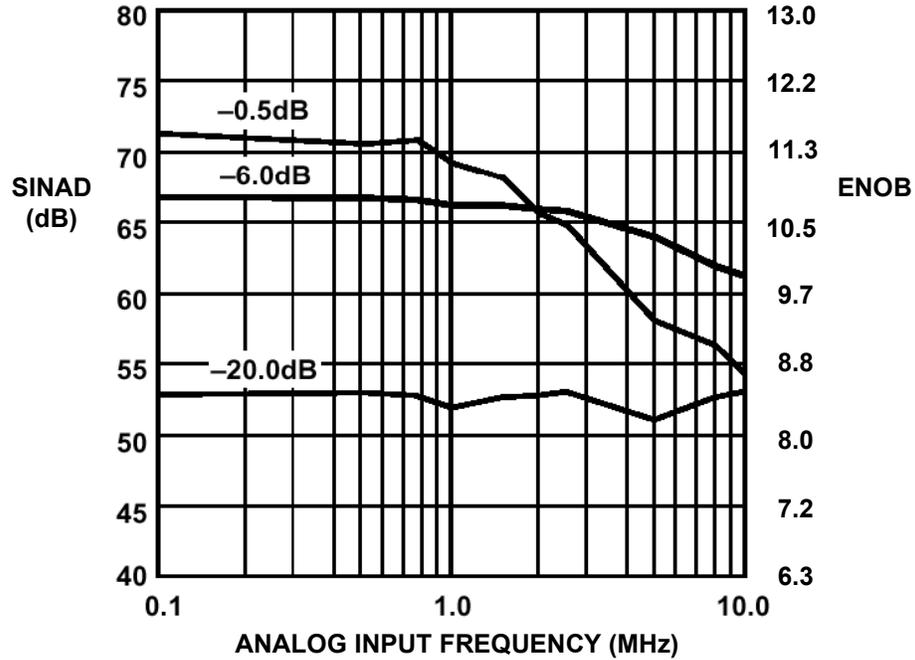


Figure 2.26

The SINAD plot shows where the AC performance of the ADC degrades due to high-frequency distortion and is usually plotted for frequencies well above the Nyquist frequency so that performance in undersampling applications can be evaluated. SINAD is often converted to *effective-number-of-bits* (ENOB) using the relationship for the theoretical SNR of an ideal N-bit ADC: $SNR = 6.02N + 1.76dB$. The equation is solved for N, and the value of SINAD is substituted for SNR:

$$ENOB = \frac{SINAD - 1.76dB}{6.02}$$

Signal-to-noise ratio (SNR, or *SNR-without-harmonics*) is calculated the same as SINAD except that the signal harmonics are excluded from the calculation, leaving only the noise terms. In practice, it is only necessary to exclude the first 5 harmonics since they dominate. The SNR plot will degrade at high frequencies, but not as rapidly as SINAD because of the exclusion of the harmonic terms.

Many current ADC data sheets somewhat loosely refer to SINAD as SNR, so the engineer must be careful when interpreting these specifications.

Analog Bandwidth

The analog bandwidth of an ADC is that frequency at which the spectral output of the *fundamental* swept frequency (as determined by the FFT analysis) is reduced by 3dB. It may be specified for either a small signal (SSBW- small signal bandwidth),

or a full scale signal (FPBW- full power bandwidth), so there can be a wide variation in specifications between manufacturers.

Like an amplifier, the analog bandwidth specification of a converter does not imply that the ADC maintains good distortion performance up to its bandwidth frequency. In fact, the SINAD (or ENOB) of most ADCs will begin to degrade considerably before the input frequency approaches the actual 3dB bandwidth frequency. Figure 2.27 shows ENOB and full scale frequency response of an ADC with a FPBW of 1MHz, however, the ENOB begins to drop rapidly above 100kHz.

ADC GAIN (BANDWIDTH) AND ENOB VERSUS FREQUENCY SHOWS IMPORTANCE OF ENOB SPECIFICATION

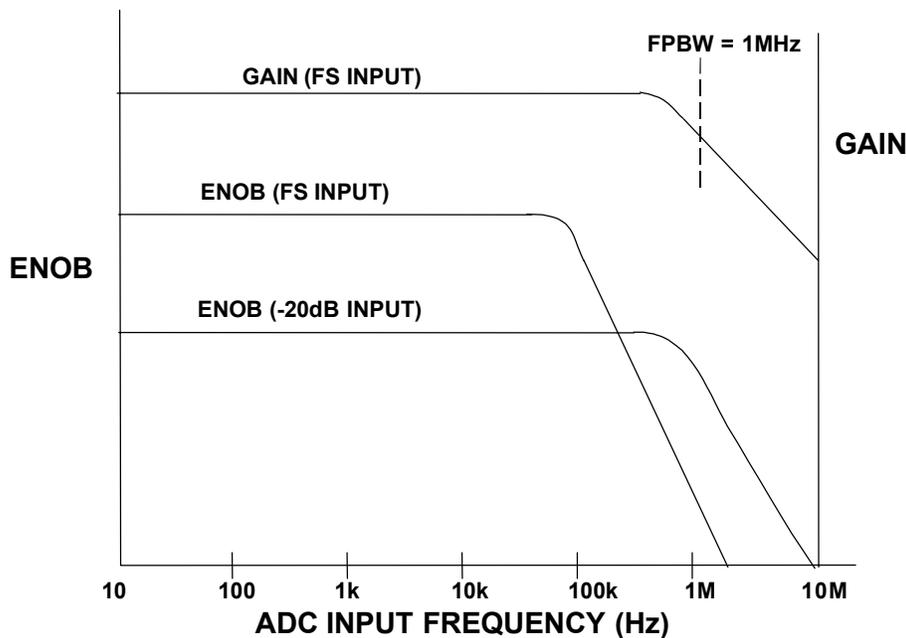


Figure 2.27

Spurious Free Dynamic Range (SFDR)

Probably the most significant specification for an ADC used in a communications application is its spurious free dynamic range (SFDR). The SFDR specification is to ADCs what the third order intercept specification is to mixers and LNAs. SFDR of an ADC is defined as the ratio of the rms signal amplitude to the rms value of the *peak spurious spectral content* (measured over the entire first Nyquist zone, DC to $f_s/2$). SFDR is generally plotted as a function of signal amplitude and may be expressed relative to the signal amplitude (dBc) or the ADC full scale (dBFS) as shown in Figure 2.28.

SPURIOUS FREE DYNAMIC RANGE (SFDR)

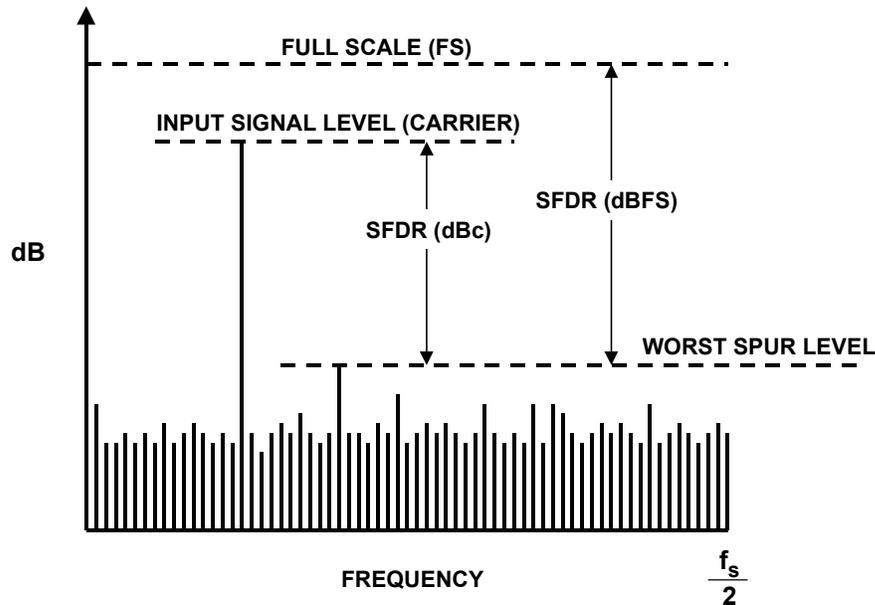


Figure 2.28

For a signal near full scale, the peak spectral spur is generally determined by one of the first few harmonics of the fundamental. However, as the signal falls several dB below full scale, other spurs generally occur which are not direct harmonics of the input signal. This is because of the differential non-linearity of the ADC transfer function as discussed earlier. Therefore, SFDR considers *all* sources of distortion, regardless of their origin.

The AD9042 is a 12-bit, 41MSPS wideband ADC designed for communications applications where high SFDR is important. The SFDR for a 19.5MHz input and a sampling frequency of 41MSPS is shown in Figure 2.29. Note that a minimum of 80dBc SFDR is obtained over the entire first Nyquist zone (DC to 20MHz). The plot also shows SFDR expressed as dBFS.

SFDR is generally much greater than the ADCs theoretical N-bit SNR ($6.02N + 1.76$ dB). For example, the AD9042 is a 12-bit ADC with an SFDR of 80dBc and a typical SNR of 65dBc (theoretical SNR is 74dB). This is because there is a fundamental distinction between noise and distortion measurements. The process gain of the FFT (33dB for a 4096-point FFT) allows frequency spurs well below the noise floor to be observed. Adding extra resolution to an ADC may serve to increase its SNR but may or may not increase its SFDR.

AD9042 12-BIT, 41MSPS ADC SFDR VS. INPUT POWER LEVEL

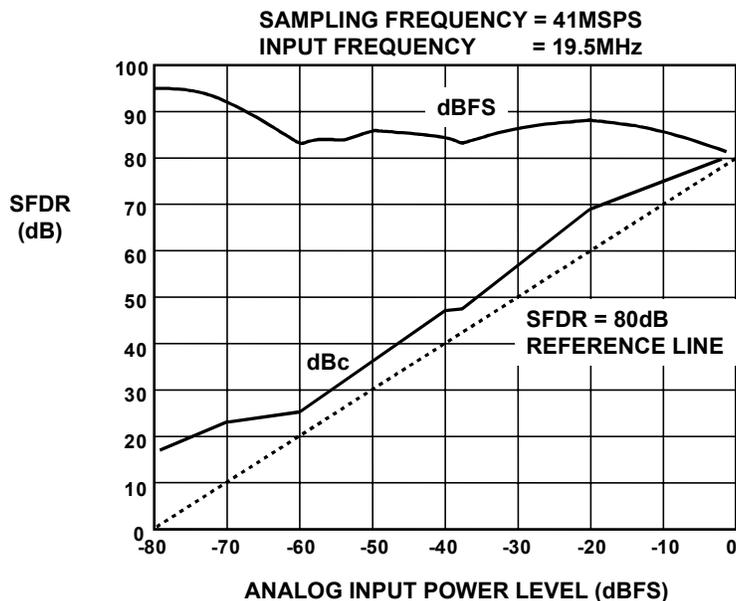


Figure 2.29

Two Tone Intermodulation Distortion (IMD)

Two tone IMD is measured by applying two spectrally pure sinewaves to the ADC at frequencies f_1 and f_2 , usually relatively close together. The amplitude of each tone is set slightly more than 6dB below full scale so that the ADC does not clip when the two tones add in-phase. The location of the second and third-order products are shown in Figure 2.30. Notice that the second-order products fall at frequencies which can be removed by digital filters. However, the third-order products $2f_2-f_1$ and $2f_1-f_2$ are close to the original signals and are more difficult to filter. Unless otherwise specified, two-tone IMD refers to these third-order products. The value of the IMD product is expressed in dBc relative to the value of *either* of the two original tones, and not to their sum.

Note, however, that if the two tones are close to $f_s/4$, then the aliased third harmonics of the fundamentals can make the identification of the actual $2f_2-f_1$ and $2f_1-f_2$ products difficult. This is because the third harmonic of $f_s/4$ is $3f_s/4$, and the alias occurs at $f_s - 3f_s/4 = f_s/4$. Similarly, if the two tones are close to $f_s/3$, the aliased second harmonics may interfere with the measurement. The same reasoning applies here; the second harmonic of $f_s/3$ is $2f_s/3$, and its alias occurs at $f_s - 2f_s/3 = f_s/3$.

SECOND AND THIRD-ORDER INTERMODULATION PRODUCTS FOR $f_1 = 5\text{MHz}$, $f_2 = 6\text{MHz}$

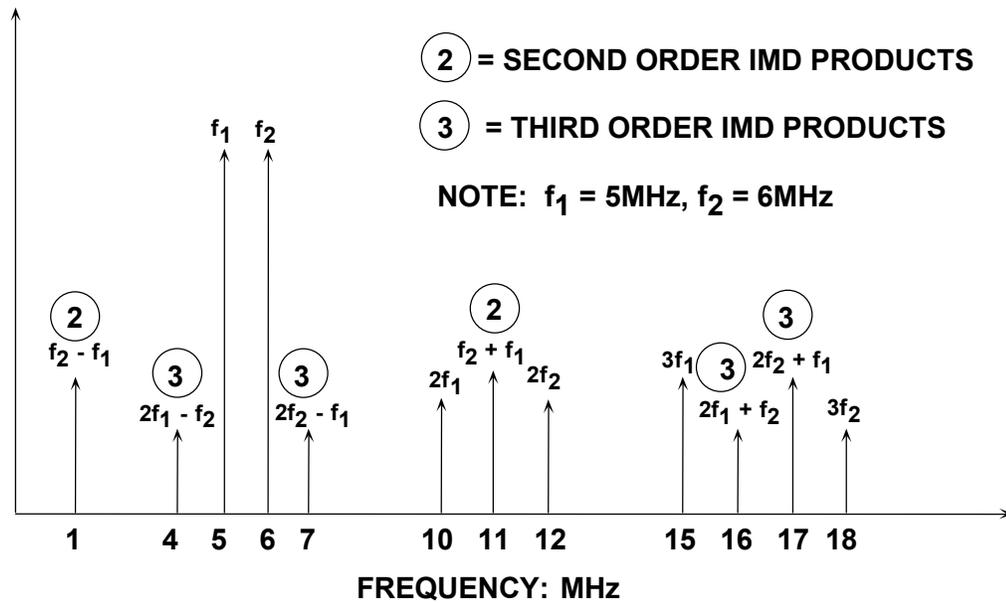


Figure 2.30

The concept of *second and third-order intercept points* is not valid for an ADC, because the distortion products do not vary in a predictable manner (as a function of signal amplitude). The ADC does not gradually begin to compress signals approaching full scale (there is no 1dB compression point); it acts as a *hard limiter* as soon as the signal exceeds the ADC input range, thereby suddenly producing extreme amounts of distortion because of clipping. On the other hand, for signals much below full scale, the distortion floor remains relatively constant and is independent of signal level.

Multi-tone SFDR is often measured in communications applications. The larger number of tones more closely simulates the wideband frequency spectrum of cellular telephone systems such as AMPS or GSM. Figure 2.31 shows the 4-tone intermodulation performance of the AD6640 12-bit, 65MSPS ADC. High SFDR increases the receiver's ability to capture small signals in the presence of large ones, and prevent the small signals from being masked by the intermodulation products of the larger ones.

MULTITONE TESTING: AD6640 12-BIT, 65MSPS ADC

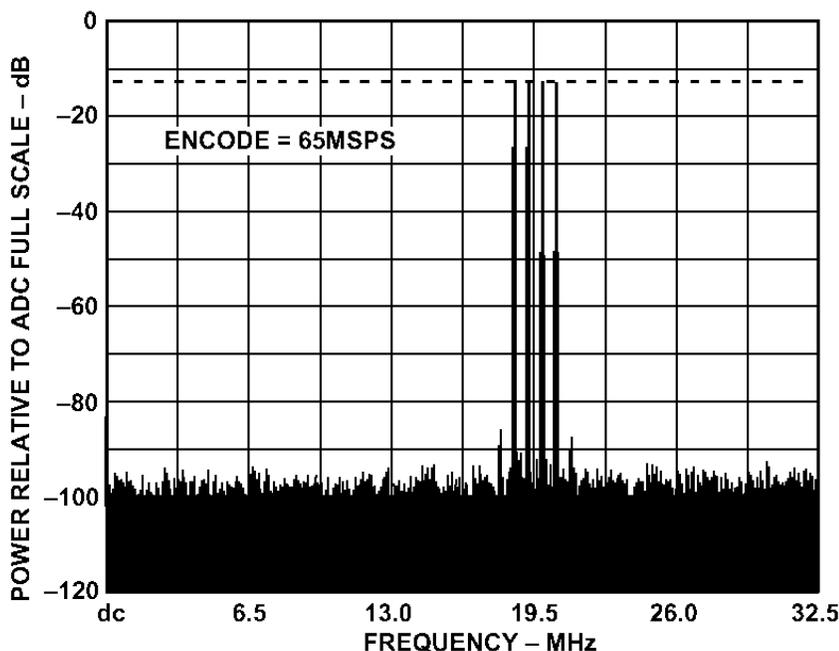


Figure 2.31

Noise Power Ratio (NPR)

Noise power ratio testing has been used extensively to measure the transmission characteristics of Frequency Division Multiple Access (FDMA) communications links (see Reference 4). In a typical FDMA system, 4kHz wide voice channels are "stacked" in frequency bins for transmission over coaxial, microwave, or satellite equipment. At the receiving end, the FDMA data is demultiplexed and returned to 4kHz individual baseband channels. In an FDMA system having more than approximately 100 channels, the FDMA signal can be approximated by Gaussian noise with the appropriate bandwidth. An individual 4kHz channel can be measured for "quietness" using a narrow-band notch (bandstop) filter and a specially tuned receiver which measures the noise power inside the 4kHz notch (see Figure 2.32).

Noise Power Ratio (NPR) measurements are straightforward. With the notch filter out, the rms noise power of the signal inside the notch is measured by the narrowband receiver. The notch filter is then switched in, and the residual noise inside the slot is measured. The ratio of these two readings expressed in dB is the NPR. Several slot frequencies across the noise bandwidth (low, midband, and high) are tested to characterize the system adequately. NPR measurements on ADCs are made in a similar manner except the analog receiver is replaced by a buffer memory and an FFT processor.

NOISE POWER RATIO (NPR) MEASUREMENTS

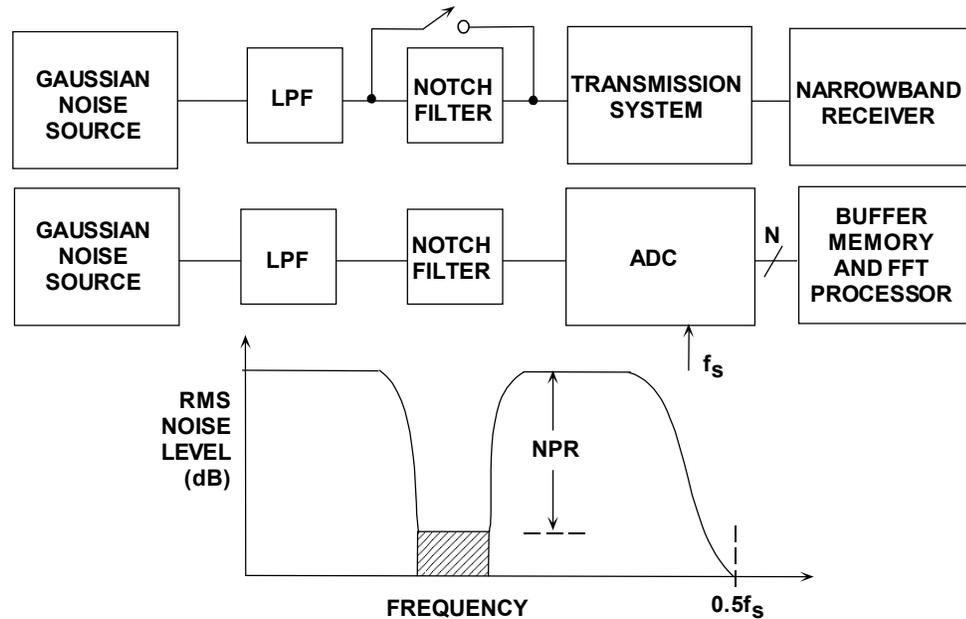


Figure 2.32

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

In a digital system containing an ADC, the noise within the slot is primarily quantization noise when low levels of noise input are applied. The NPR curve is linear in this region. As the noise level increases, there is a one-for-one correspondence between the noise level and the NPR. At some level, however, "clipping" noise caused by the hard-limiting action of the ADC begins to dominate. A theoretical curve for 10, 11, and 12-bit ADCs is shown in Figure 2.33 (see Reference 5).

In multi-channel high frequency communication systems, NPR can also be used to simulate the distortion caused by a large number of individual channels, similar to an FDMA system. A notch filter is placed between the noise source and the ADC, and an FFT output is used in place of the analog receiver. The width of the notch filter is set for several MHz as shown in Figure 2.34 for the AD9042. NPR is the "depth" of the notch. An ideal ADC will only generate quantization noise inside the notch, however a practical one has additional noise components due to

intermodulation distortion caused by ADC non-linearity. Notice that the NPR is about 60dB compared to 62.7dB theoretical.

THEORETICAL NPR FOR 10, 11, 12-BIT ADCs

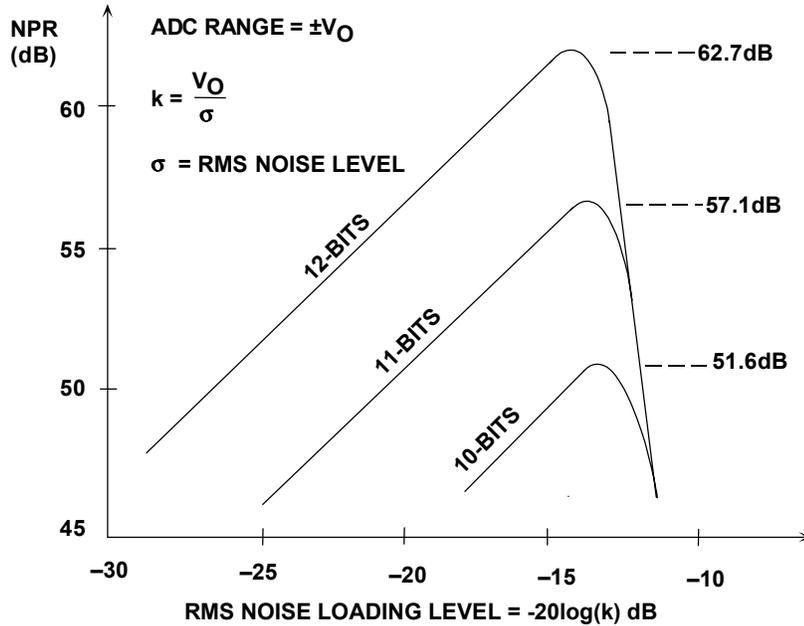


Figure 2.33

AD9042 12-BIT, 41MSPS ADC NPR MEASURES 60dB (62.7dB THEORETICAL)

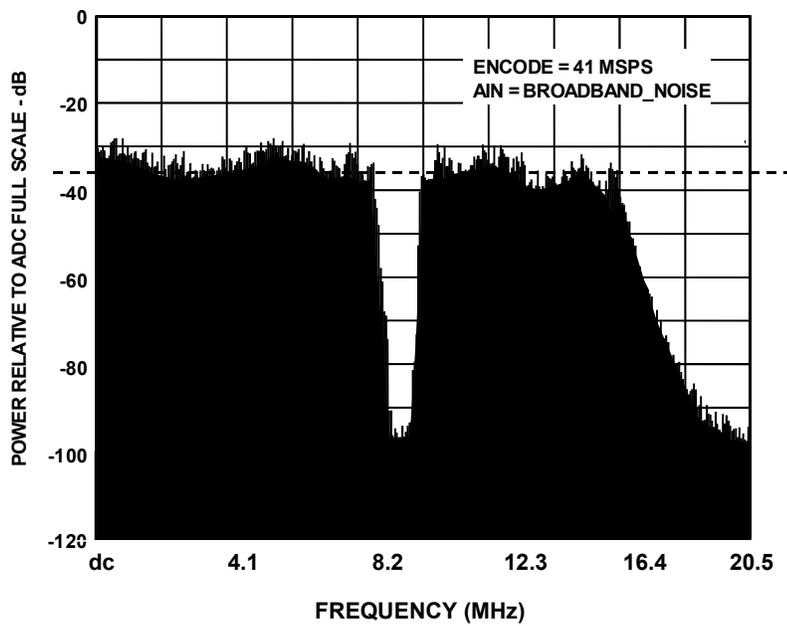


Figure 2.34

Aperture Jitter and Aperture Delay

Another reason that the SNR of an ADC decreases with input frequency may be deduced from Figure 2.35, which shows the effects of phase jitter (or aperture time jitter) on the sampling clock of an ADC (or internal in the sample-and-hold). The phase jitter causes a voltage error which is a function of slew rate and results in an overall degradation in SNR as shown in Figure 2.36. This is quite serious, especially at higher input/output frequencies. Therefore, extreme care must be taken to minimize phase noise in the sampling/reconstruction clock of any sampled data system. This care must extend to all aspects of the clock signal: the oscillator itself (for example, a 555 timer is absolutely inadequate, but even a quartz crystal oscillator can give problems if it uses an active device which shares a chip with noisy logic); the transmission path (these clocks are very vulnerable to interference of all sorts), and phase noise introduced in the ADC or DAC. A very common source of phase noise in converter circuitry is aperture jitter in the integral sample-and-hold (SHA) circuitry.

EFFECTS OF APERTURE AND SAMPLING CLOCK JITTER

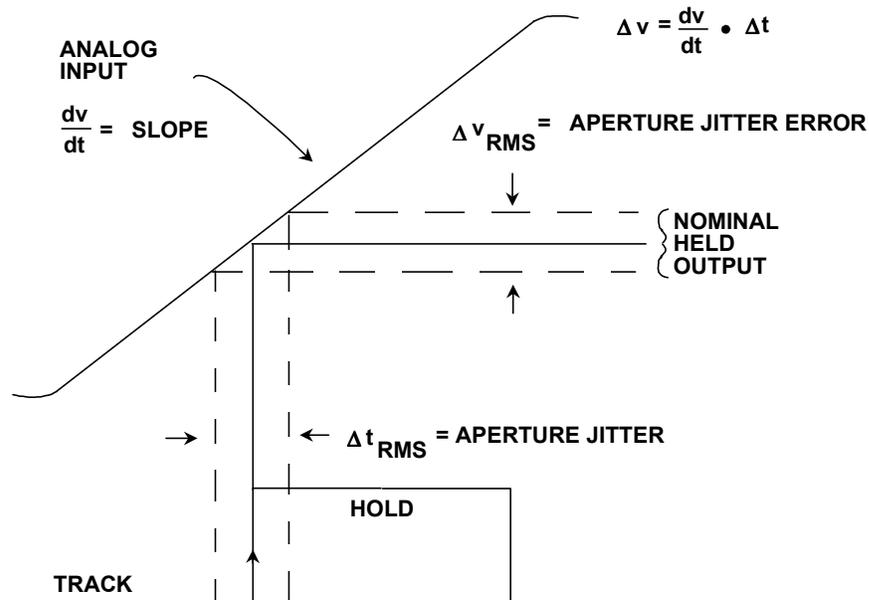


Figure 2.35

Two decades or so ago, sampling ADCs were built up from a separate SHA and ADC. Interface design was difficult, and a key parameter was aperture jitter in the SHA. Today, most sampled data systems use *sampling* ADCs which contain an integral SHA. The aperture jitter of the SHA may not be specified as such, but this

is not a cause of concern if the SNR or ENOB is clearly specified, since a guarantee of a specific SNR is an implicit guarantee of an adequate aperture jitter specification. However, the use of an additional high-performance SHA will sometimes improve the high-frequency ENOB of a even the best sampling ADC by presenting "DC" to the ADC, and may be more cost-effective than replacing the ADC with a more expensive one.

It should be noted that there is also a fixed component which makes up the ADC aperture time. This component, usually called *effective aperture delay time*, does not produce an error. It simply results in a time offset between the time the ADC is asked to sample and when the actual sample takes place (see Figure 2.37), and may be positive or negative. The variation or tolerance placed on this parameter from part to part is important in simultaneous sampling applications or other applications such as I and Q demodulation where two ADCs are required to track each other.

SNR DUE TO APERTURE AND SAMPLING CLOCK JITTER

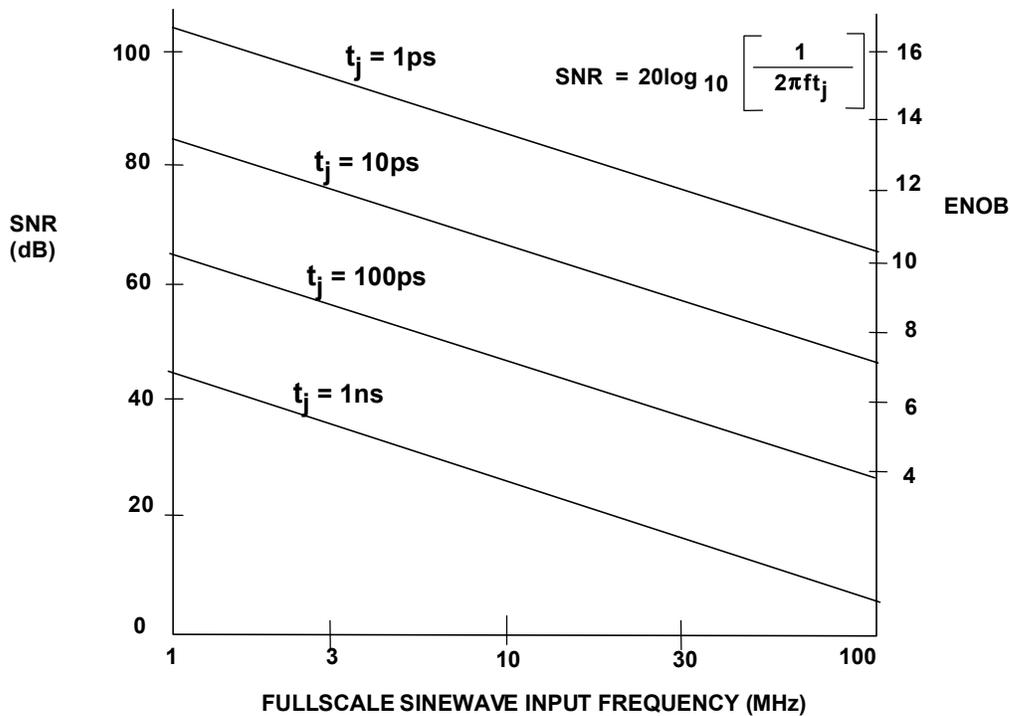


Figure 2.36

EFFECTIVE APERTURE DELAY TIME

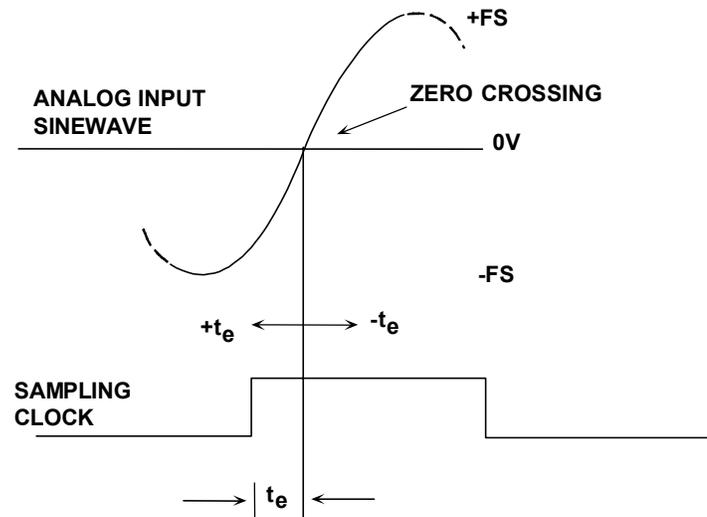


Figure 2.37

DAC DYNAMIC PERFORMANCE

The AC specifications which are most likely to be important with DACs are *settling time*, *glitch*, *distortion*, and *Spurious Free Dynamic Range (SFDR)*.

The settling time of a DAC is the time from a change of digital code to when the output comes within *and remains within* some error band as shown in Figure 2.38. With amplifiers, it is hard to make comparisons of settling time, since their error bands may differ from amplifier to amplifier, but with DACs the error band will almost invariably be ± 1 or $\pm \frac{1}{2}$ LSB.

The settling time of a DAC is made up of four different periods: the *switching time* or *dead time* (during which the digital switching, but not the output, is changing), the *slewing time* (during which the rate of change of output is limited by the slew rate of the DAC output), the *recovery time* (when the DAC is recovering from its fast slew and may overshoot), and the *linear settling time* (when the DAC output approaches its final value in an exponential or near-exponential manner). If the slew time is short compared to the other three (as is usually the case with current output DACs), then the settling time will be largely independent of the output step size. On the other hand, if the slew time is a significant part of the total, then the larger the step, the longer the settling time.

DAC SETTLING TIME

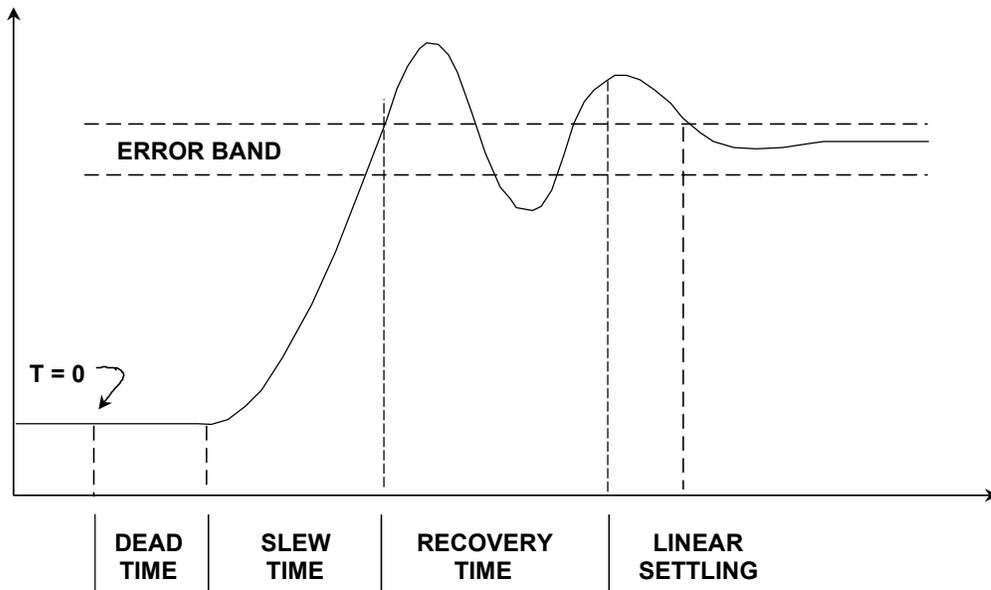


Figure 2.38

Ideally, when a DAC output changes it should move from one value to its new one monotonically. In practice, the output is likely to overshoot, undershoot, or both (see Figure 2.39). This uncontrolled movement of the DAC output during a transition is known as *glitch*. It can arise from two mechanisms: capacitive coupling of digital transitions to the analog output, and the effects of some switches in the DAC operating more quickly than others and producing temporary spurious outputs.

DAC TRANSITIONS (SHOWING GLITCH)

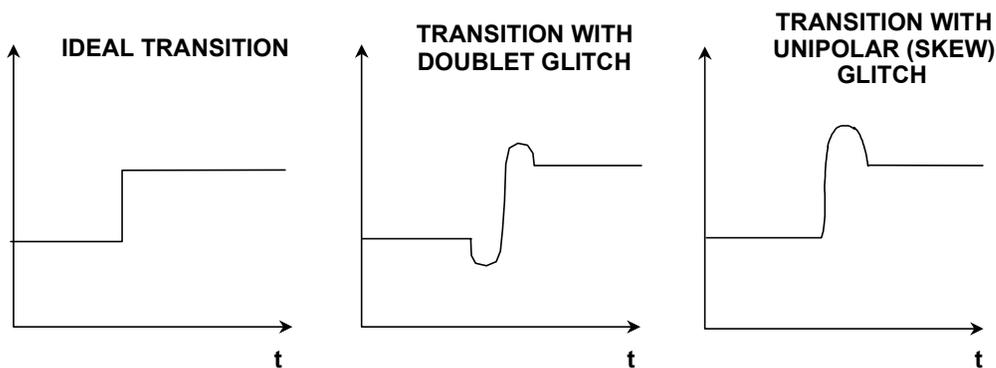


Figure 2.39

Capacitive coupling frequently produces roughly equal positive and negative spikes (sometimes called a *doublet* glitch) which more or less cancel in the longer term. The glitch produced by switch timing differences is generally unipolar, much larger and of greater concern.

Glitches can be characterized by measuring the *glitch impulse area*, sometimes inaccurately called glitch energy. The term *glitch energy* is a misnomer, since the unit for glitch impulse area is Volt-seconds (or more probably $\mu\text{V}\cdot\text{sec}$ or $\text{pV}\cdot\text{sec}$. The *peak glitch area* is the area of the largest of the positive or negative glitch areas. The glitch impulse area is the net area under the voltage-versus-time curve and can be estimated by approximating the waveforms by triangles, computing the areas, and subtracting the negative area from the positive area. The midscale glitch produced by the transition between the codes 0111...111 and 1000...000 is usually the worst glitch. Glitches at other code transition points (such as 1/4 and 3/4 full scale) are generally less. Figure 2.40 shows the midscale glitch for a fast low-glitch DAC. The peak and net glitch areas are estimated using triangles as described above. Settling time is measured from the time the waveform leaves the initial 1LSB error band until it enters and remains within the final 1LSB error band. The step size between the transition regions is also 1LSB.

**DAC MIDSCALE GLITCH SHOWS 1.34pV-s
NET IMPULSE AREA AND SETTLING TIME OF 4.5ns**

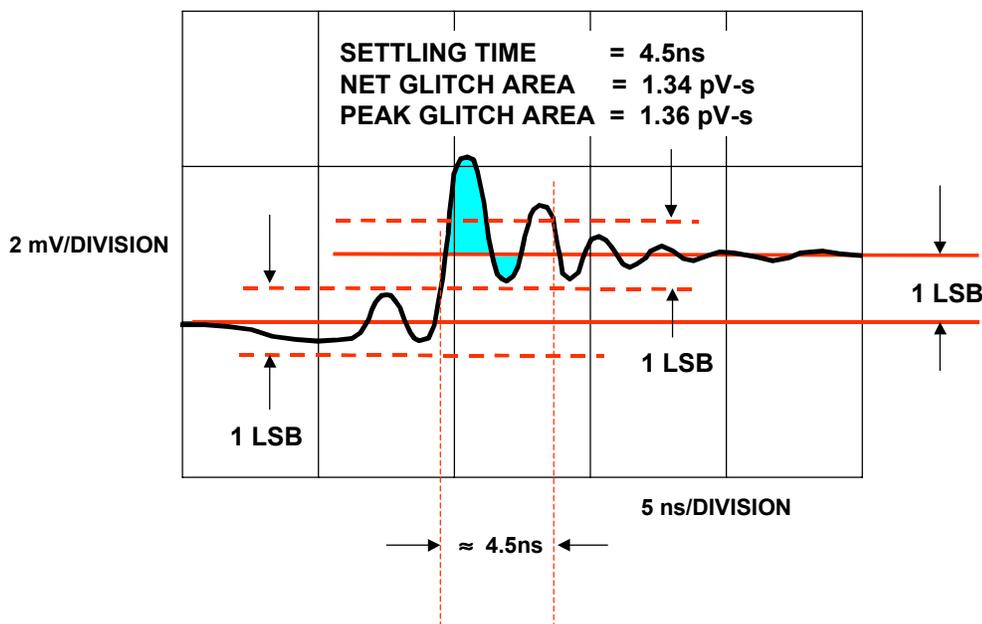


Figure 2.40

DAC settling time is important in applications such as RGB raster scan video display drivers, but frequency-domain specifications such as SFDR are generally more important in communications.

If we consider the spectrum of a waveform reconstructed by a DAC from digital data, we find that in addition to the expected spectrum (which will contain one or more frequencies, depending on the nature of the reconstructed waveform), there will also be noise and distortion products. Distortion may be specified in terms of harmonic distortion, Spurious Free Dynamic Range (SFDR), intermodulation distortion, or all of the above. Harmonic distortion is defined as the ratio of harmonics to fundamental when a (theoretically) pure sine wave is reconstructed, and is the most common specification. Spurious free dynamic range is the ratio of the worst spur (usually, but not necessarily always a harmonic of the fundamental) to the fundamental.

Code-dependent glitches will produce both out-of-band and in-band harmonics when the DAC is reconstructing a digitally generated sinewave as in a Direct Digital Synthesis (DDS) system. The midscale glitch occurs twice during a single cycle of a reconstructed sinewave (at each midscale crossing), and will therefore produce a second harmonic of the sinewave, as shown in Figure 2.41. Note that the higher order harmonics of the sinewave, which alias back into the Nyquist bandwidth (DC to $f_s/2$), cannot be filtered.

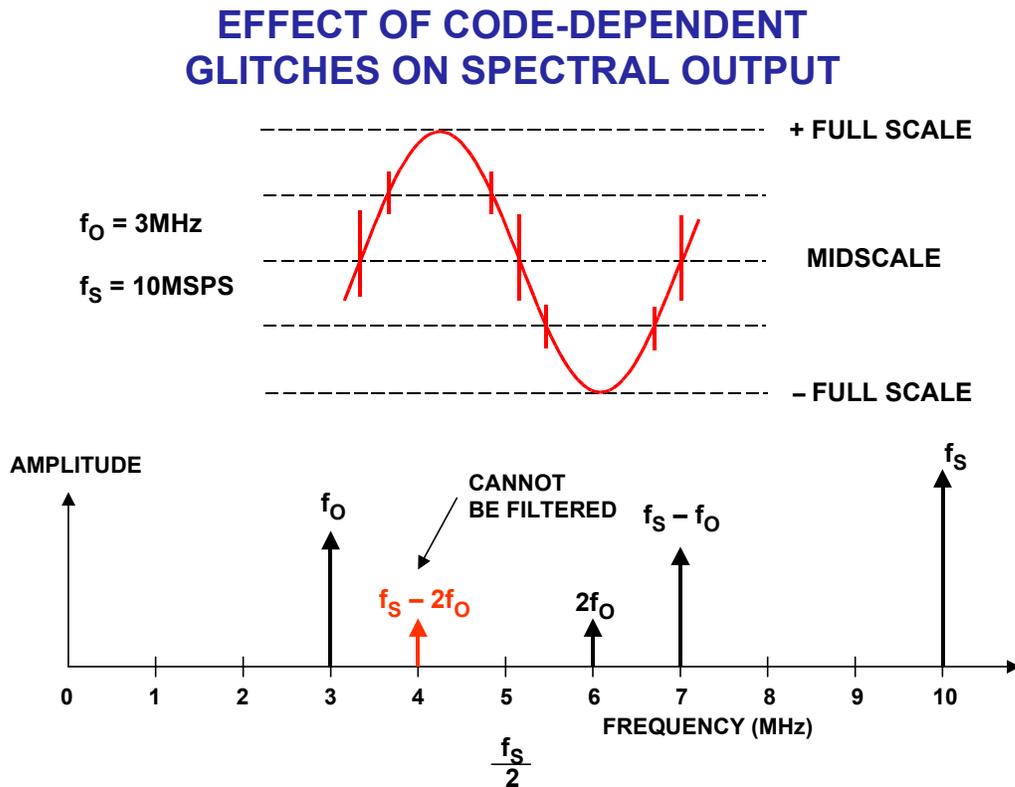


Figure 2.41

It is difficult to predict the harmonic distortion or SFDR from the glitch area specification alone. Other factors, such as the overall linearity of the DAC, also contribute to distortion. It is therefore customary to test reconstruction DACs in the frequency domain (using a spectrum analyzer) at various clock rates and output frequencies as shown in Figure 2.43. Typical SFDR for the 14-bit AD9772 Transmit DAC is shown in Figure 2.44. The clock rate is 65MSPS, and the output frequency is swept to 25MHz. As in the case of ADCs, quantization noise will appear as increased harmonic distortion if the ratio between the clock frequency and the DAC output frequency is an integer number. These ratios should be avoided when making the SFDR measurements.

CONTRIBUTORS TO DDS DAC DISTORTION

- Resolution
- Integral Non-Linearity
- Differential Non-Linearity
- Code-Dependent Glitches
- Ratio of Clock Frequency to Output Frequency (Even in an Ideal DAC)
- Mathematical Analysis is Difficult !

Figure 2.42

TEST SETUP FOR MEASURING DAC SFDR

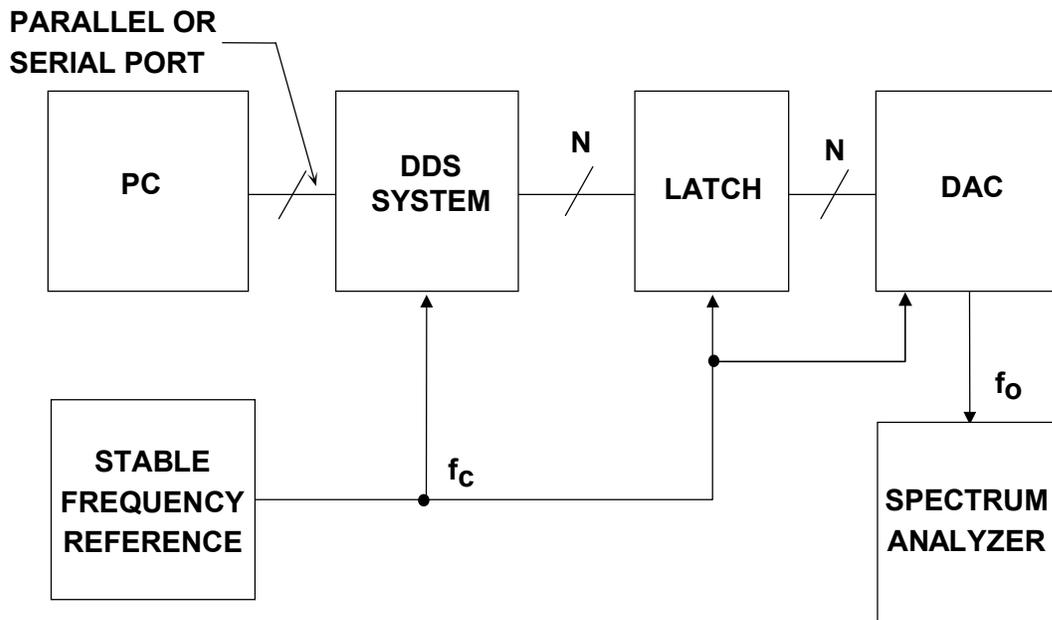


Figure 2.43

AD9772 14-BIT TxDAC™ SFDR, DATA UPDATE RATE = 65MSPS

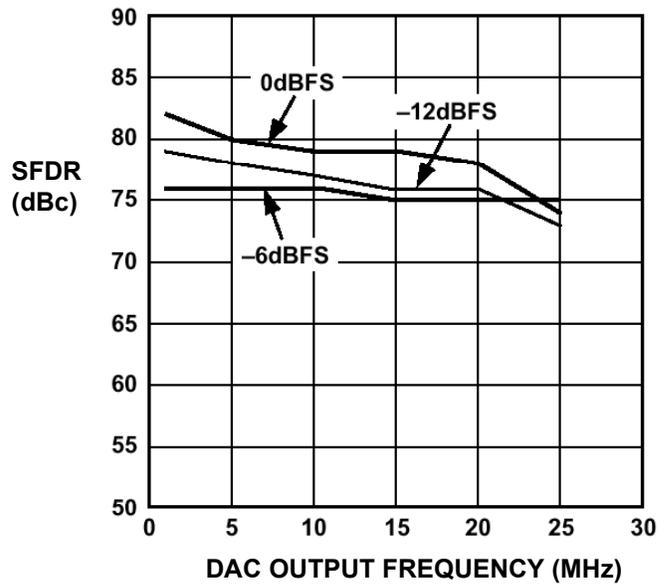


Figure 2.44

DAC $\sin(x)/x$ Frequency Rolloff

The output of a reconstruction DAC can be visualized as a series of rectangular pulses whose width is equal to the reciprocal of the clock rate as shown in Figure 2.45. Note that the reconstructed signal amplitude is down 3.92dB at the Nyquist frequency, $f_c/2$. An inverse $\sin(x)/x$ filter can be used to compensate for this effect in most cases. The images of the fundamental signal are also attenuated by the $\sin(x)/x$ function.

DAC SIN X/X ROLL OFF (AMPLITUDE NORMALIZED)

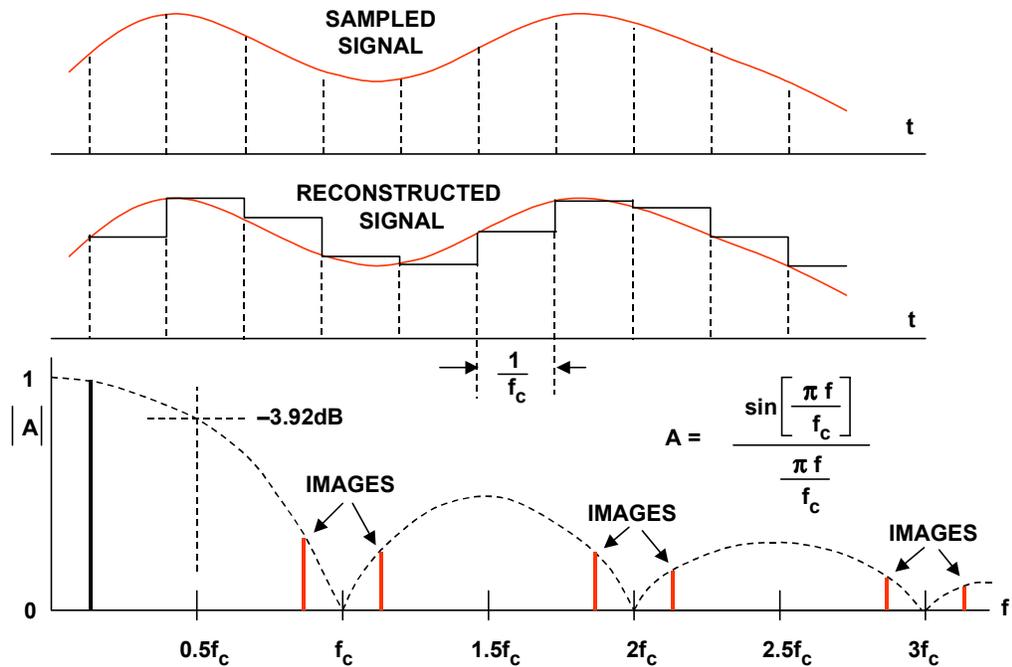


Figure 2.45

REFERENCES

1. **Active and Passive Electrical Wave Filter Catalog**, Vol. 34, TTE, Incorporated, 2251 Barry Avenue, Los Angeles, CA 90064.
2. W. R. Bennett, "Spectra of Quantized Signals", **Bell System Technical Journal**, No. 27, July 1948, pp. 446-472.
3. Steve Ruscak and Larry Singer, *Using Histogram Techniques to Measure A/D Converter Noise*, **Analog Dialogue**, Vol. 29-2, 1995.
4. M.J. Tant, **The White Noise Book**, Marconi Instruments, July 1974.
5. G.A. Gray and G.W. Zeoli, *Quantization and Saturation Noise due to A/D Conversion*, **IEEE Trans. Aerospace and Electronic Systems**, Jan. 1971, pp. 222-223.
6. Chuck Lane, *A 10-bit 60MSPS Flash ADC*, **Proceedings of the 1989 Bipolar Circuits and Technology Meeting**, IEEE Catalog No. 89CH2771-4, September 1989, pp. 44-47.
7. F.D. Waldhauer, *Analog to Digital Converter*, **U.S. Patent 3-187-325**, 1965.
8. J.O. Edson and H.H. Henning, *Broadband Codecs for an Experimental 224Mb/s PCM Terminal*, **Bell System Technical Journal**, 44, November 1965, pp. 1887-1940.
9. J.S. Mayo, *Experimental 224Mb/s PCM Terminals*, **Bell System Technical Journal**, 44, November 1965, pp. 1813-1941.
10. Hermann Schmid, **Electronic Analog/Digital Conversions**, Van Nostrand Reinhold Company, New York, 1970.
11. Carl Moreland, *An 8-bit 150MSPS Serial ADC*, **1995 ISSCC Digest of Technical Papers**, Vol. 38, p. 272.
12. Roy Gosser and Frank Murden, *A 12-bit 50MSPS Two-Stage A/D Converter*, **1995 ISSCC Digest of Technical Papers**, p. 278.
13. Carl Moreland, **An Analog-to-Digital Converter Using Serial-Ripple Architecture**, Masters' Thesis, Florida State University College of Engineering, Department of Electrical Engineering, 1995.
14. **Practical Analog Design Techniques**, Analog Devices, 1995, Chapter 4, 5, and 8.
15. **Linear Design Seminar**, Analog Devices, 1995, Chapter 4, 5.

16. **System Applications Guide**, Analog Devices, 1993, Chapter 12, 13, 15,16.
17. **Amplifier Applications Guide**, Analog Devices, 1992, Chapter 7.
18. Walt Kester, *Drive Circuitry is Critical to High-Speed Sampling ADCs*, **Electronic Design Special Analog Issue**, Nov. 7, 1994, pp. 43-50.
19. Walt Kester, *Basic Characteristics Distinguish Sampling A/D Converters*, **EDN**, Sept. 3, 1992, pp. 135-144.
20. Walt Kester, *Peripheral Circuits Can Make or Break Sampling ADC Systems*, **EDN**, Oct. 1, 1992, pp. 97-105.
21. Walt Kester, *Layout, Grounding, and Filtering Complete Sampling ADC System*, **EDN**, Oct. 15, 1992, pp. 127-134.
22. Robert A. Witte, *Distortion Measurements Using a Spectrum Analyzer*, **RF Design**, September, 1992, pp. 75-84.
23. Walt Kester, *Confused About Amplifier Distortion Specs?*, **Analog Dialogue**, 27-1, 1993, pp. 27-29.
24. **System Applications Guide**, Analog Devices, 1993, Chapter 16.
25. Frederick J. Harris, *On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform*, **IEEE Proceedings**, Vol. 66, No. 1, Jan. 1978, pp. 51-83.
26. Joey Doernberg, Hae-Seung Lee, David A. Hodges, *Full Speed Testing of A/D Converters*, **IEEE Journal of Solid State Circuits**, Vol. SC-19, No. 6, Dec. 1984, pp. 820-827.
27. Brendan Coleman, Pat Meehan, John Reidy and Pat Weeks, *Coherent Sampling Helps When Specifying DSP A/D Converters*, **EDN**, October 15, 1987, pp. 145-152.
28. Robert W. Ramierez, **The FFT: Fundamentals and Concepts**, Prentice-Hall, 1985.
29. R. B. Blackman and J. W. Tukey, **The Measurement of Power Spectra**, Dover Publications, New York, 1958.
30. James J. Colotti, *Digital Dynamic Analysis of A/D Conversion Systems Through Evaluation Software Based on FFT/DFT Analysis*, **RF Expo East 1987 Proceedings**, Cardiff Publishing Co., pp. 245-272.
31. **HP Journal**, Nov. 1982, Vol. 33, No. 11.
32. **HP Product Note 5180A-2**.

SAMPLED DATA SYSTEMS

33. **HP Journal**, April 1988, Vol. 39, No. 2.
34. **HP Journal**, June 1988, Vol. 39, No. 3.
35. Dan Sheingold, Editor, **Analog-to-Digital Conversion Handbook, Third Edition**, Prentice-Hall, 1986.
36. Lawrence Rabiner and Bernard Gold, **Theory and Application of Digital Signal Processing**, Prentice-Hall, 1975.
37. Matthew Mahoney, **DSP-Based Testing of Analog and Mixed-Signal Circuits**, IEEE Computer Society Press, Washington, D.C., 1987.
38. **IEEE Trial-Use Standard for Digitizing Waveform Recorders**, No. 1057-1988.
39. Richard J. Higgins, **Digital Signal Processing in VLSI**, Prentice-Hall, 1990.
40. M. S. Ghausi and K. R. Laker, **Modern Filter Design: Active RC and Switched Capacitors**, Prentice Hall, 1981.
41. Mathcad™ 4.0 software package available from MathSoft, Inc., 201 Broadway, Cambridge MA, 02139.
42. Howard E. Hilton, *A 10MHz Analog-to-Digital Converter with 110dB Linearity*, **H.P. Journal**, October 1993, pp. 105-112.