SECTION 4
HIGH SPEED SAMPLING AND HIGH SPEED ADCs, *Walt Kester*

INTRODUCTION

High speed ADCs are used in a wide variety of real-time DSP signal-processing applications, replacing systems that used analog techniques alone. The major reason for using digital signal processing are (1) the cost of DSP processors has gone down, (2) their speed and computational power has increased, and (3) they are reprogrammable, thereby allowing for system performance upgrades without hardware changes. DSP offers solutions that cannot be achieved in the analog domain, i.e. V.32 and V.34 modems.

However, in order for digital signal processing techniques to be effective in solving an analog signal processing problem, appropriate cost effective high speed ADCs must be available. The ADCs must be tested and specified in such a way that the design engineer can relate the ADC performance to specific system requirements, which can be more demanding than if they were used in purely analog signal processing systems. In most high speed signal processing applications, AC performance and wide dynamic range are much more important than traditional DC performance. This requires that the ADC manufacturer not only design the right ADCs but specify them as completely as possible to cover a wide variety of applications.

Another important aspect of integrating ADCs into a high speed system is a complete understanding of the sampling process and the distortion mechanisms which ultimately limit system performance. High speed sampling ADCs first were used in instrumentation and signal processing applications, where much emphasis was placed on time-domain performance. While this is still important, applications of ADCs in communications also require comprehensive frequency-domain specifications.

Modern IC processes also allow the integration of more analog functionality into the ADC, such as on-board references, sample-and-hold amplifiers, PGAs, etc. This makes them easier to use in a system by minimizing the amount of support circuitry required.

Another driving force in high speed ADC development is the trend toward lower power and lower supply voltages. Most high speed sampling ADCs today operate on either dual or single 5V supplies, and there is increasing interest in single-supply converters which will operate on 3V or less for battery powered applications. Lower supply voltages tend to increase a circuit's sensitivity to power supply noise and ground noise, especially mixed-signal devices such as ADCs and DACs.

The trend toward lower cost and lower power has led to the development of a variety of high speed ADCs fabricated on standard 0.6 micron CMOS processes. Making a precision ADC on a digital process (no thin film resistors are available) is a real challenge to the IC circuit designer. ADCs which require the maximum in

performance still require a high speed complementary bipolar process (such as Analog Devices' XFCB) with thin film resistors.

The purpose of this section is to equip the engineer with the proper tools necessary to understand and select ADCs for high speed systems applications. Making intelligent tradeoffs in the system design requires a thorough understanding of the fundamental capabilities and limitations of state-of-the-art high speed sampling ADCs.

HIGH SPEED SAMPLING ADCs

- **Wide Acceptance in Signal Processing and Communications**
- **Emphasis on Dynamic Performance**
- **Trend to Low Power, Low Voltage, Single-Supply**
- **More On-Chip Functionality: PGAs, SHA, Digital Filters, etc.**
- **Process Technology:**
  - **Low Cost CMOS: Up to 12-bits @ 10MSPS**
  - **High Speed Complementary Bipolar: Up to 12-bits @ 70MSPS**
  - **Statistical Matching Techniques Rather than Thin Film Laser Trimming**

ANALOG DEVICES

4.1

FUNDAMENTALS OF HIGH SPEED SAMPLING

The sampling process can be discussed from either the frequency or time domain or both. Frequency-domain analysis is applicable to communications, so that's what we will consider.

First consider the case of a single frequency sinewave of frequency f_a sampled at a frequency f_S by an ideal impulse sampler (see top diagram in Figure 4.2). 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, ...

ANALOG SIGNAL f_a SAMPLED @ f_S USING IDEAL SAMPLER HAS IMAGES (ALIASES) AT $|\pm Kf_S \pm f_a|$, K = 1, 2, 3, . . .

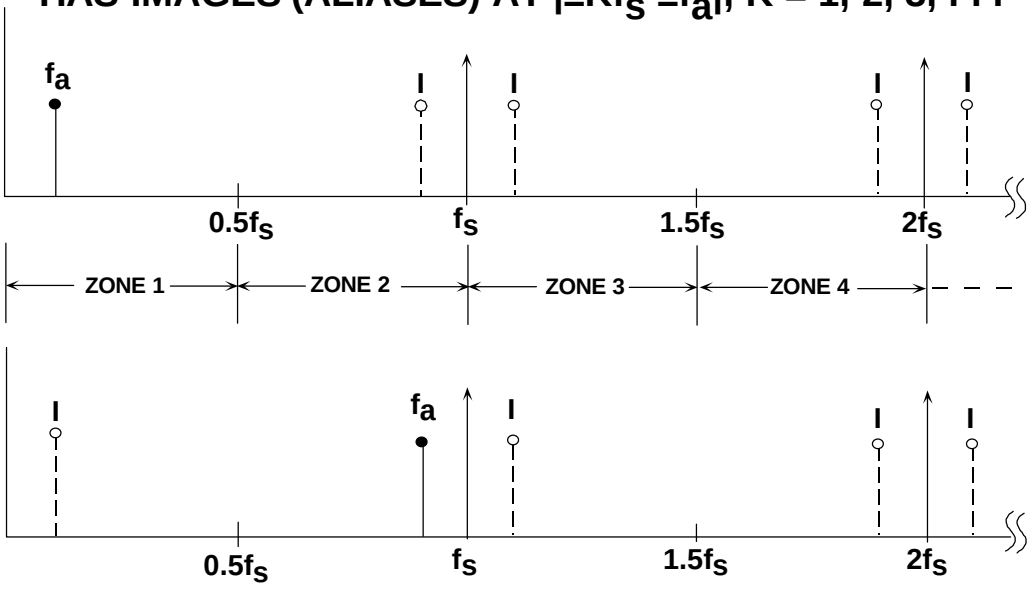

The *Nyquist* bandwidth is defined to be the frequency spectrum from DC to $f_S/2$. The frequency spectrum is divided into an infinite number of *Nyquist zones*, each having a width equal to $0.5f_S$ as shown. In practice, the ideal sampler is replaced by an ADC followed by an FFT processor. The FFT 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 4.2, bottom diagram) 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 4.2, top diagram, 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 4.3.

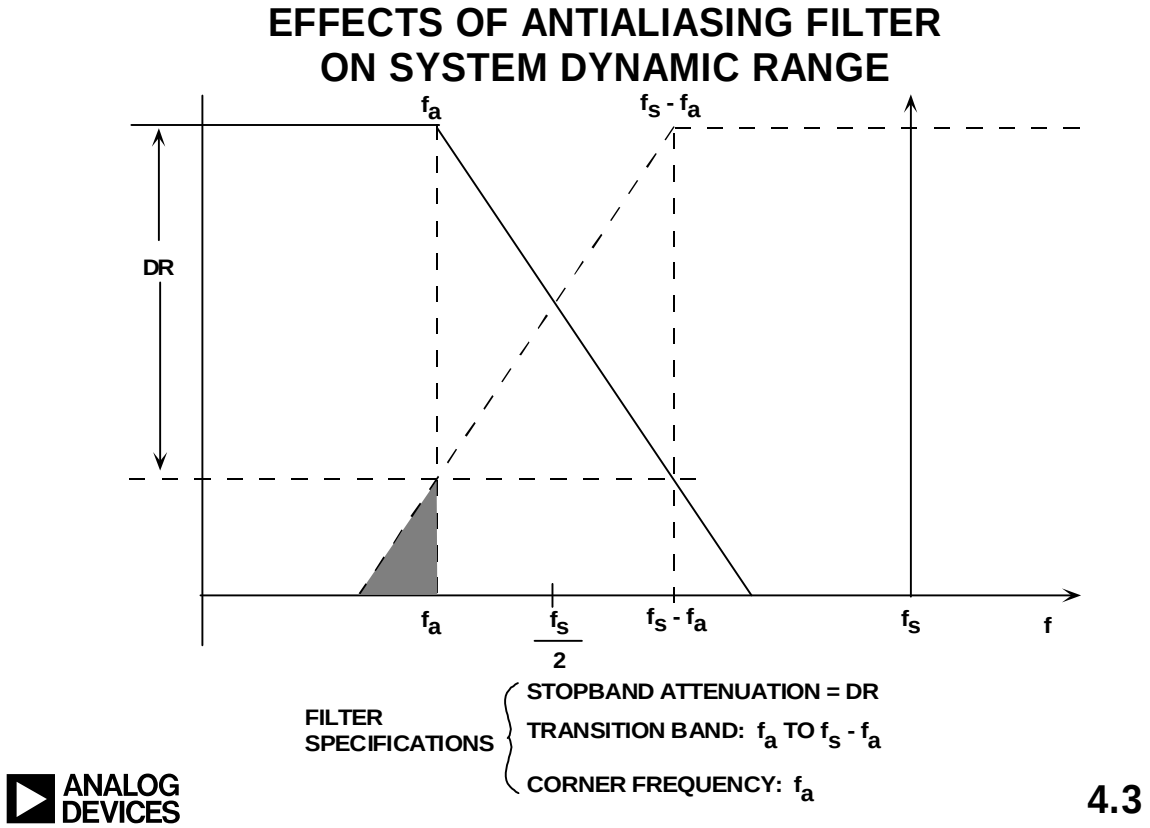

Assume that the input signal has fullscale components well above the maximum frequency of interest, f_a. The diagram shows how fullscale frequency components above $f_S - f_a$ are aliased back into the bandwidth DC to f_a. These aliased components are indistinguishable from actual signals and therefore 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 4.3, 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 stopband attenuation, DR. The required system dynamic range is chosen based on our requirement for signal fidelity.

Filters have to 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). As an example, the normalized response of the TTE, Inc., LE1182 11-pole elliptic antialiasing filter is shown in Figure 4.4. Notice that this filter is specified to achieve at least 80dB attenuation between f_C and $1.2f_C$. The corresponding passband ripple, return loss, delay, and phase response are also shown in Figure 4.4. This custom filter is available in corner frequencies up to 100MHz and in a choice of PC board, BNC, or SMA with compatible packages.

CHARACTERISTICS OF TTE, INC., LE1182-SERIES 11-POLE ELLIPTICAL FILTER

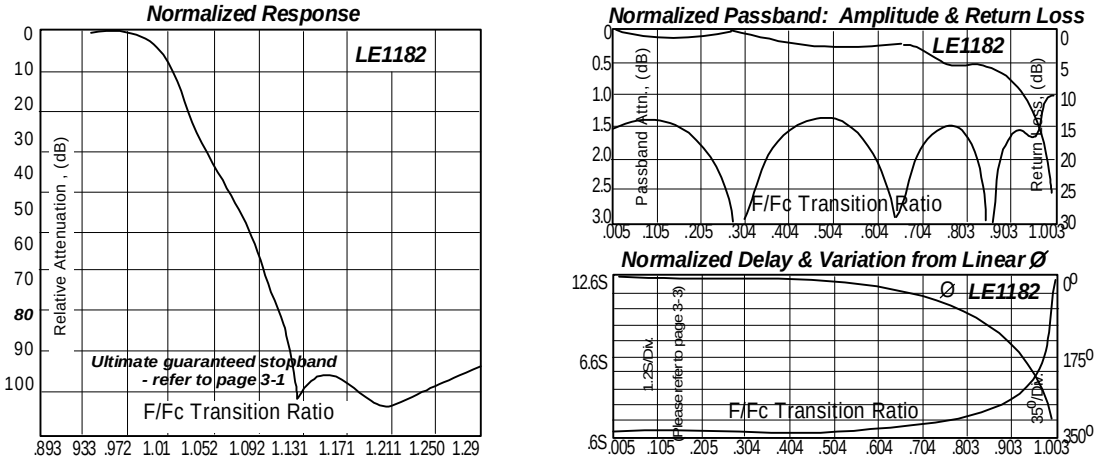

Reprinted with Permission of
TTE, Inc., 2251 Barry Ave., Los Angeles, CA 90064

ANALOG DEVICES

4.4

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 4.5 which shows the effects of

increasing the sampling frequency while maintaining the same analog corner frequency, f_a,and the same dynamic range, DR, requirement.

INCREASING SAMPLING FREQUENCY RELAXES REQUIREMENT ON ANTIALIASING FILTER

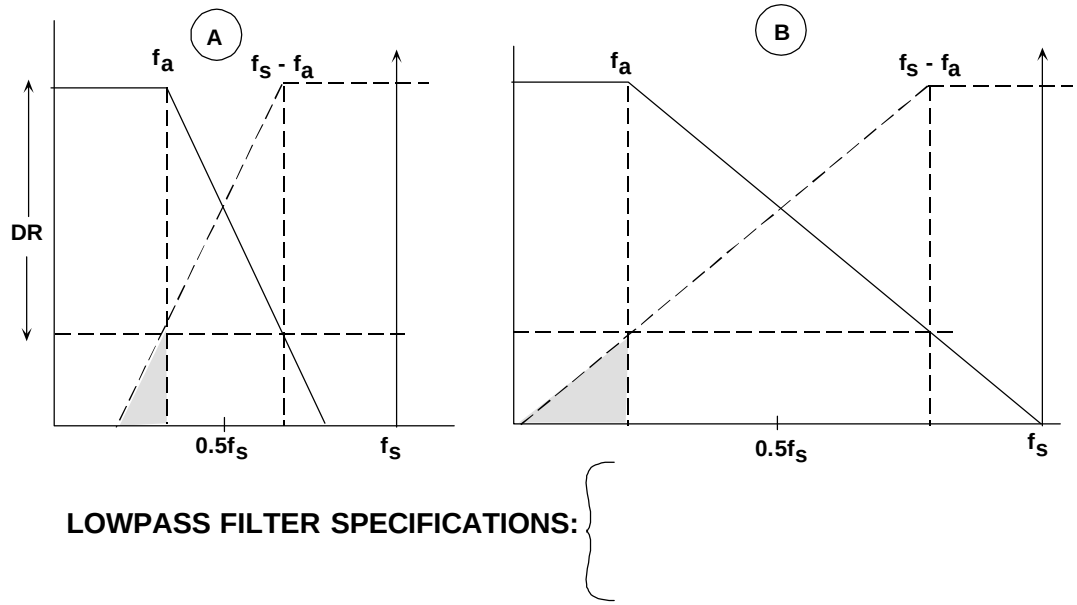

The above design process is started by choosing an initial sampling rate of 2 to 4 times f_a. Determine the filter specifications based on the required dynamic range and see if such a filter is realizable within the constraints of the system cost and performance. If not, consider a higher sampling rate which may require using a faster ADC.

The antialiasing filter requirements can 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 XdB 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 – XdB. 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 4.6A 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.

Consider the case shown in Figure 4.6B, 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).

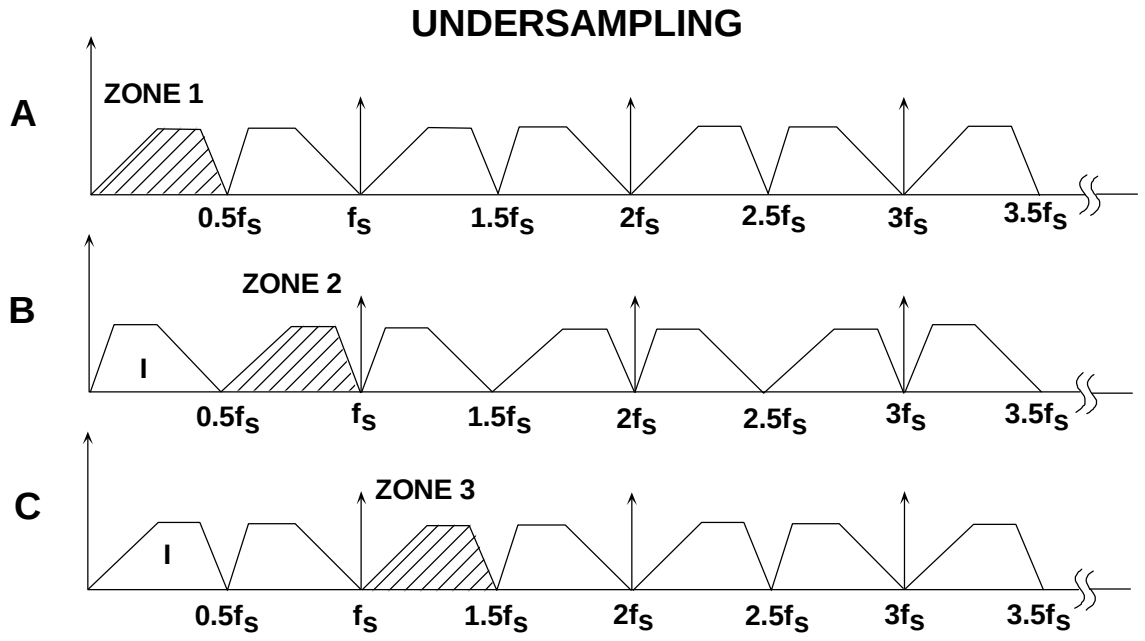

Figure 4.6C shows the sampled signal restricted to the third Nyquist zone. Note that the first Nyquist zone image has no frequency reversal. In fact, the sampled signal frequencies may lie in *any* unique Nyquist zone, and the first Nyquist zone image is still an accurate representation (with the exception of the frequency reversal which occurs when the signals are located in even Nyquist zones). At this point we can clearly state 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.

ANTIALIASING FILTERS IN UNDERSAMPLING APPLICATIONS

Figure 4.7 shows a signal in the second Nyquist zone centered around a carrier frequency, f_C, whose lower and upper frequencies are f_1 and f_2. The antialiasing filter is a bandpass filter. The desired dynamic range is DR, which defines the filter stopband attenuation. The upper transition band is f_2 to $2f_S–f_2$, and the lower is f_1 to $f_S – f_1$. As in the case of baseband sampling, the antialiasing filter requirements can be relaxed by proportionally increasing the sampling frequency, but f_C must also be increased so that it is always centered in the second Nyquist zone.

ANTIALIASING FILTER FOR UNDERSAMPLING

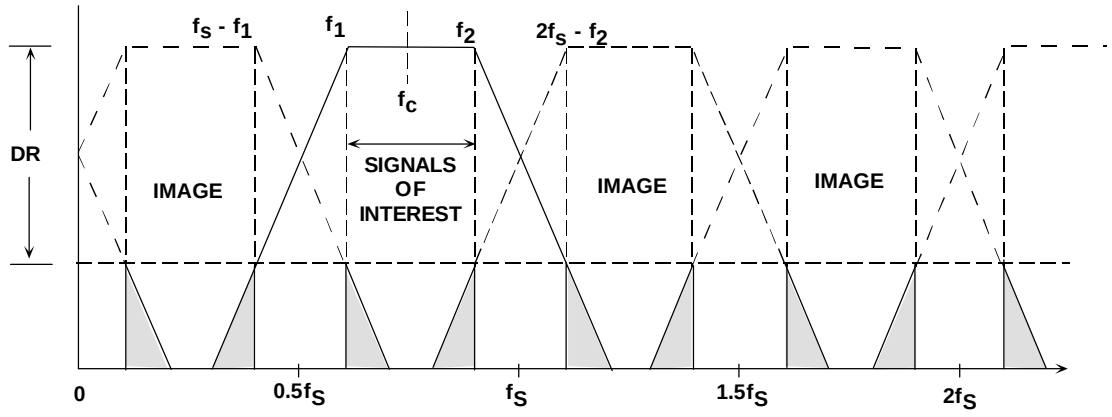

BANDPASS FILTER SPECIFICATIONS:
- **STOPBAND ATTENUATION = DR**
- **TRANSITION BAND: f_2 TO $2f_S - f_2$**
  f_1 TO $f_S - f_1$
- **CORNER FREQUENCIES: f_1, f_2**

Two key equations can be used to select the sampling frequency, f_S, given the carrier frequency, f_C, and the bandwidth of its signal, Δf. The first is the Nyquist criteria:

$$f_S > 2\Delta f \qquad \text{Eq. 1}$$

The second equation ensures that f_C is placed in the center of a Nyquist zone:

$$f_S = \frac{4f_C}{2NZ-1}, \quad \text{Eq. 2}$$

where NZ = 1, 2, 3, 4, ... and NZ corresponds to the Nyquist zone in which the carrier and its signal fall (see Figure 4.8).

NZ is normally chosen to be as large as possible while still maintaining $f_S > 2\Delta f$. This results in the minimum required sampling rate. If NZ is chosen to be odd, then f_C and it's signal will fall in an odd Nyquist zone, and the image frequencies in the first Nyquist zone will not be reversed. Tradeoffs can be made between the sampling frequency and the complexity of the antialiasing filter by choosing smaller values of NZ (hence a higher sampling frequency).

CENTERING AN UNDERSAMPLED SIGNAL WITHIN A NYQUIST ZONE

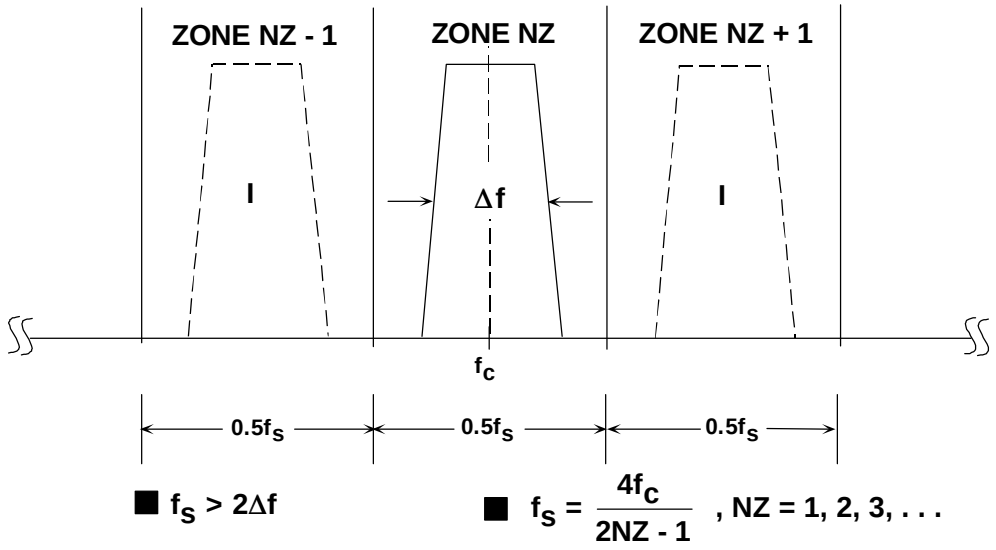

As an example, consider a 4MHz wide signal centered around a carrier frequency of 71MHz. The minimum required sampling frequency is therefore 8MSPS. Solving Eq. 2 for NZ using f_C = 71MHz and f_S = 8MSPS yields NZ = 18.25. However, NZ must be an integer, so we round 18.25 to the next lowest integer, 18. Solving Eq. 2 again for f_S yields f_S = 8.1143MSPS. The final values are therefore f_S = 8.1143MSPS, f_C = 71MHz, and NZ = 18.

Now assume that we desire more margin for the antialiasing filter, and we select f_S to be 10MSPS. Solving Eq. 2 for NZ, using f_C = 71MHz and f_S = 10MSPS yields NZ = 14.7. We round 14.7 to the next lowest integer, giving NZ = 14. Solving Eq. 2 again for f_S yields f_S = 10.519MSPS. The final values are therefore f_S = 10.519MSPS, f_C = 71MHz, and NZ = 14.

The above iterative process can also be carried out starting with f_S and adjusting the carrier frequency to yield an integer number for NZ.

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 digitizing a DC input signal is ±1/2LSB. 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 √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/√12 (see Figure 4.9).

IDEAL N-BIT ADC QUANTIZATION NOISE

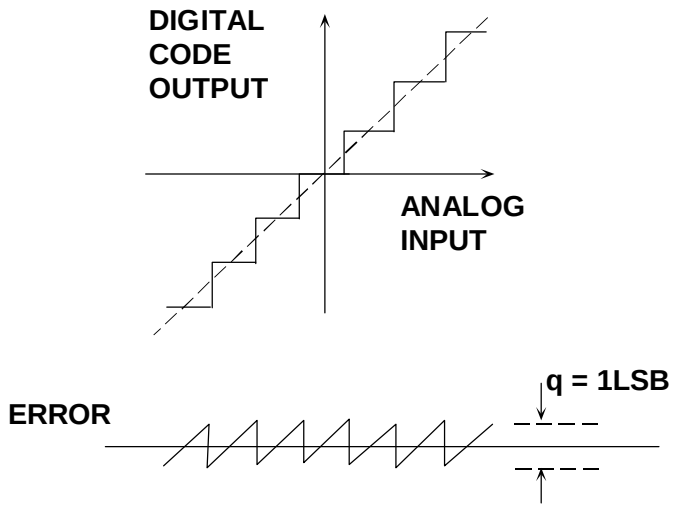

- **RMS ERROR = q/√12**
- **SNR = 6.02N + 1.76dB + 10log [f_S / 2•BW] FOR FS SINEWAVE**

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:

SNR = 6.02N + 1.76dB,

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 $f_S/2$.* 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:

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

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 $q/\sqrt{12}$, 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 4.10), 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 4.11, 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

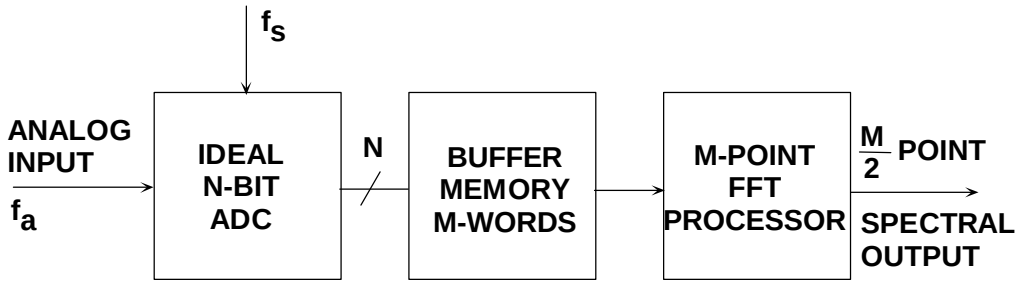

EFFECT OF RATIO OF SAMPLING CLOCK TO INPUT FREQUENCY ON SFDR FOR IDEAL 12-BIT ADC

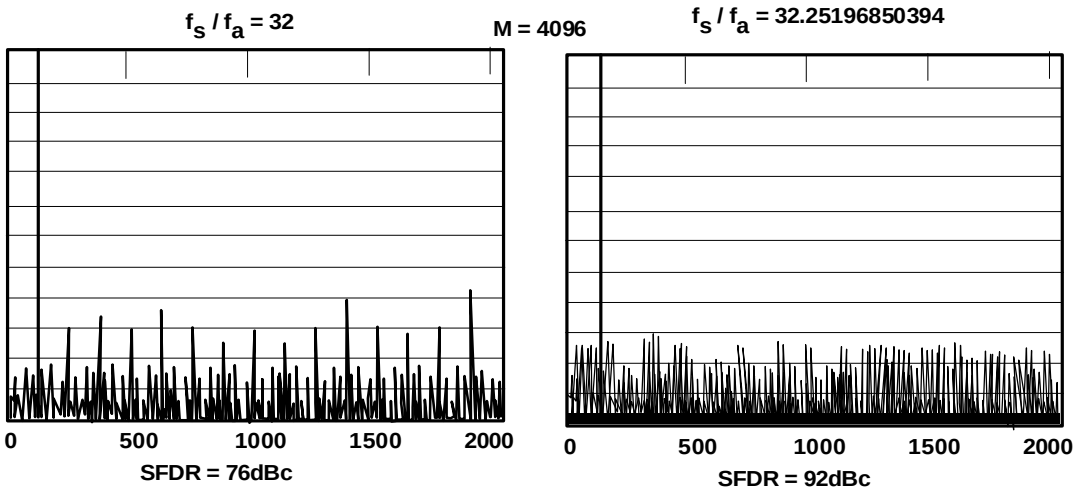

Note that this variation in the apparent harmonic distortion of the ADC is an artifact of the sampling process and the correlation of the quantization error with the input frequency. In a practical ADC application, the quantization error generally appears as random noise because of the random nature 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. (For further discussions on dither, see Section 5 of this book).

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 4.11, and note that the average value of the noise floor of the FFT is greater than 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, rather than $f_S/2$. 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 4.12). 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 $10\log_{10}(4096/2)$ = 33dB, thereby resulting in an overall FFT noise floor of 74+33=107dBc. 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.

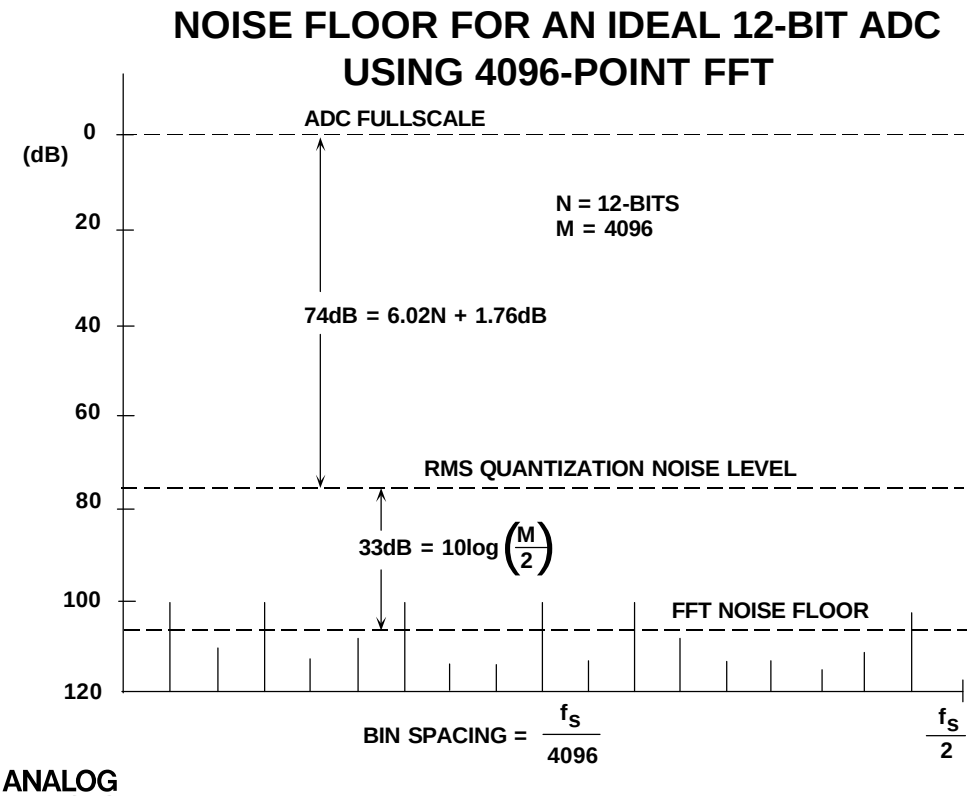

ANALOG DEVICES

4.12

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 4.13. 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

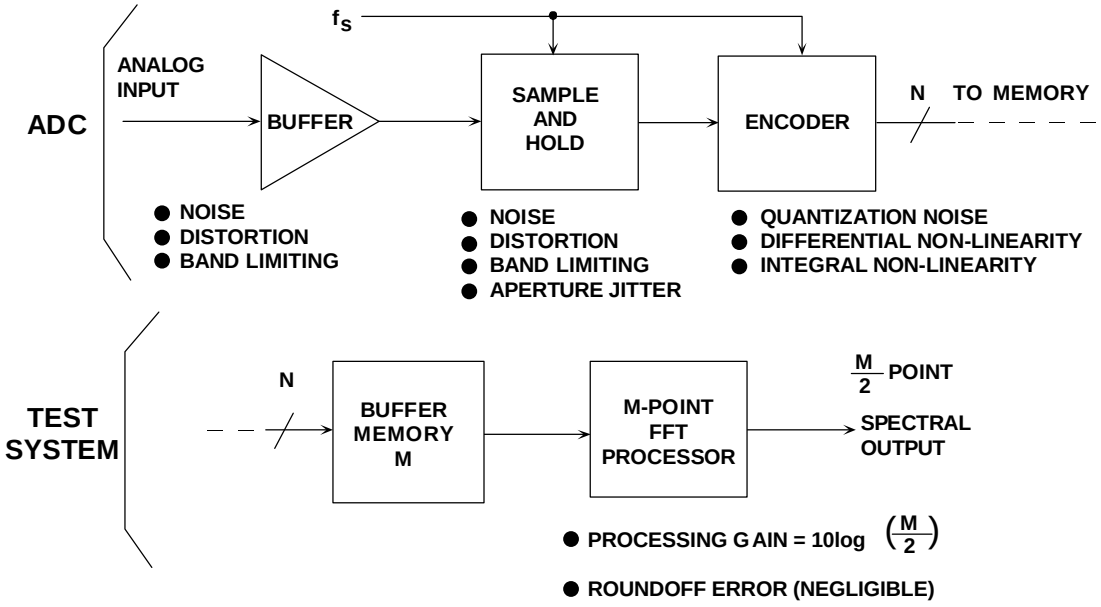

ANALOG DEVICES

4.13

Equivalent Input Referred Noise (Thermal Noise)

The wideband ADC internal circuits produce a certain amount of wideband rms noise due to thermal effects. This noise is present even for DC input signals, and accounts for the fact that the output of most wideband ADCs is a distribution of codes, centered around the nominal value of a DC input (see Figure 4.14). 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.

HISTOGRAM OF 5000 CONVERSIONS FOR A DC INPUT SHOWS 5 LSB p-p OR 0.8LSB RMS EQUIVALENT INPUT NOISE

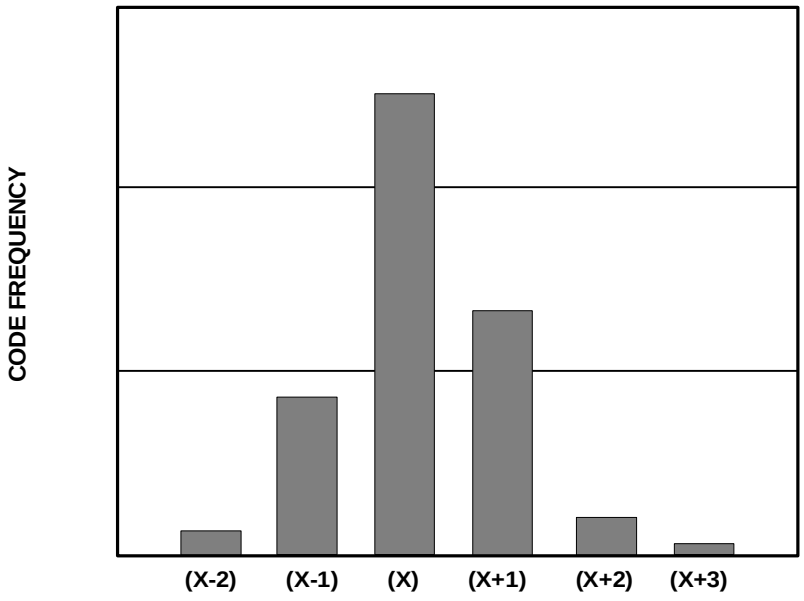

ANALOG DEVICES

4.14

Integral and Differential Non-Linearity

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.

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**
- **Noise Power Ratio (NPR)**

ANALOG DEVICES **4.15**

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 4.16 shows two ADC transfer functions containing 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, while those less and 1/2 scale peak-to-peak will not.

ADC DNL ERRORS

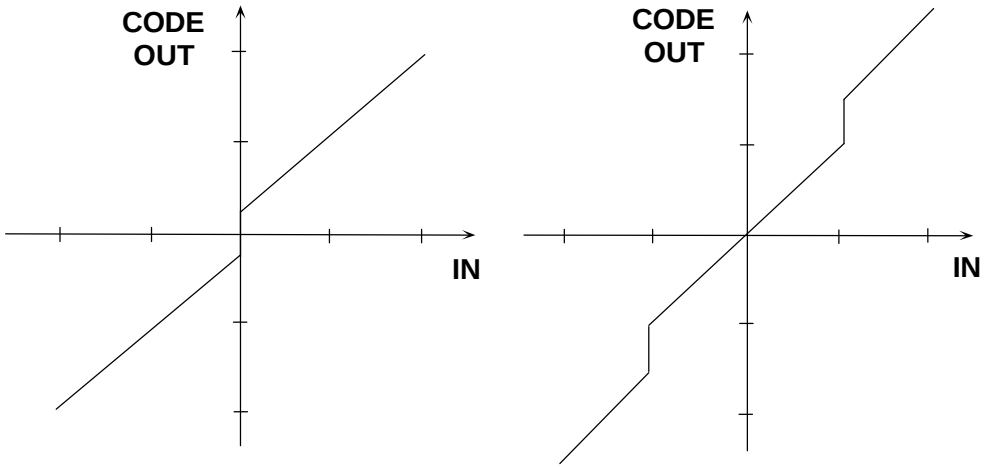

MIDSCALE DNL **1/4FS, 3/4FS DNL**

ANALOG DEVICES **4.16**

The design of most high-speed ADCs is such 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 as shown in Figure 4.17. The harmonics of the input signal can be distinguished from other distortion products by their location in the frequency spectrum. The figure 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,.... 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 specified with an input signal near full scale (generally 0.5 to 1dB below full scale to prevent clipping). For signals much lower than full scale, other distortion products (not direct harmonics) may limit performance.

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

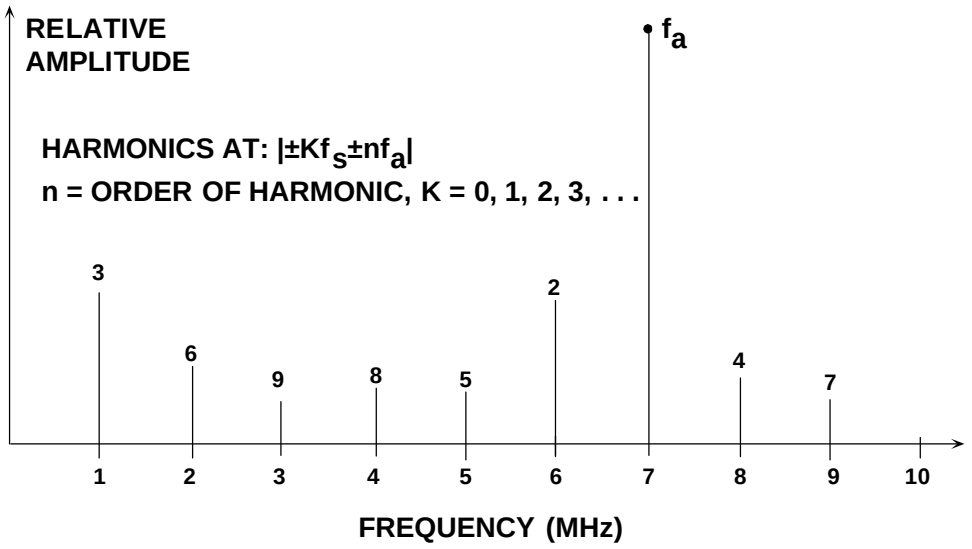

ANALOG DEVICES

4.17

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.

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 4.19.

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):**

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

- **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

AD9220 12-BIT, 10MSPS ADC SINAD AND ENOB VS. INPUT FREQUENCY FOR SAMPLING RATE = 10MSPS: SINGLE-ENDED DRIVE, V_{cm} = +2.5V, INPUT SPAN = 2V p-p

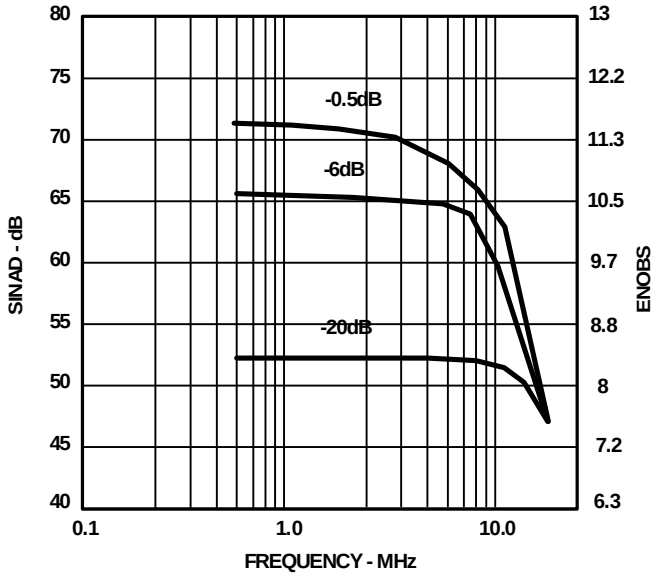

ANALOG DEVICES

4.19

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:

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

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 also, 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 4.20 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

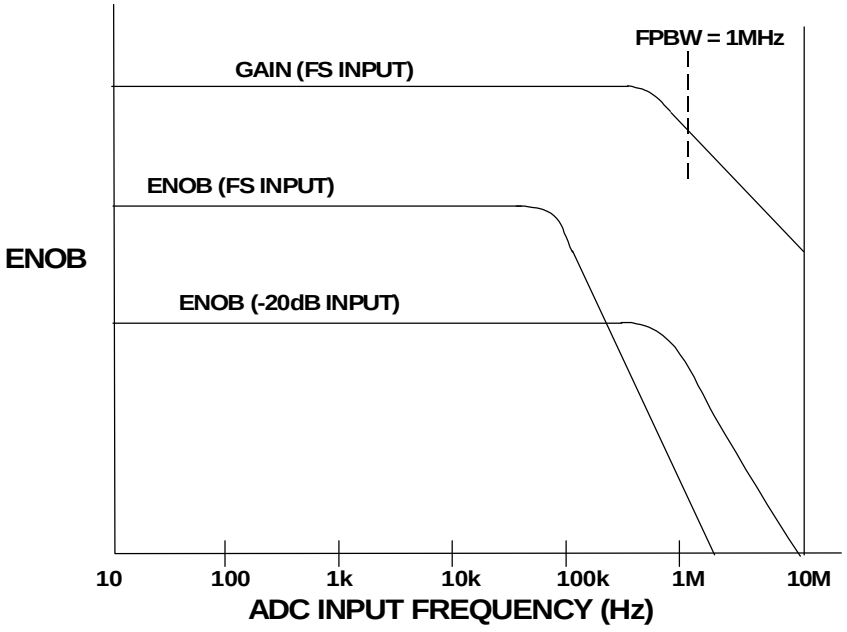

ANALOG DEVICES

4.20

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).

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 4.21. 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.

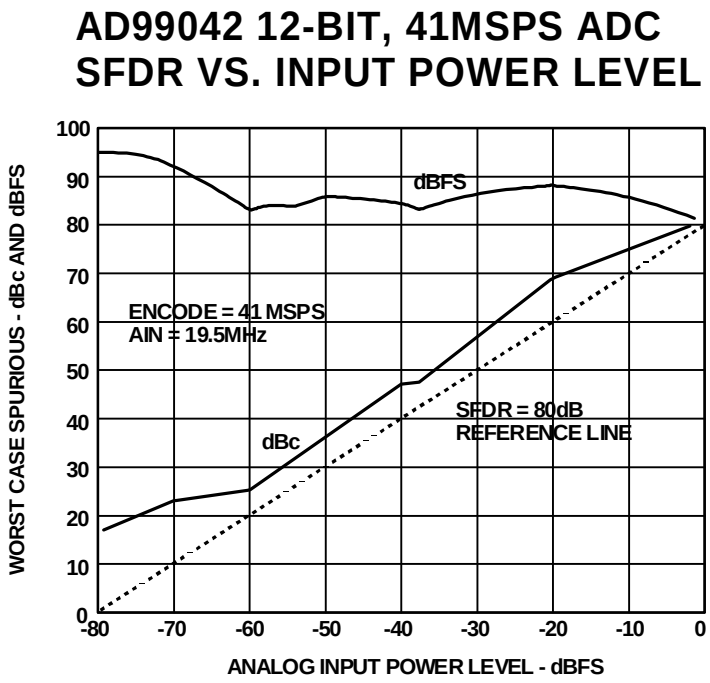

SFDR is generally much greater than the ADCs theoretical N-bit SNR (6.02N + 1.76dB). For example, the AD9042 is a 12-bit ADC with an SFDR of 80dBc and a typical SNR of 65dBc (theoretical SNR is 74dB). This is because there is a fundamental distinction between noise and distortion measurements. The process 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.

Two Tone Intermodulation Distortion

Two tone IMD is measured by applying two spectrally pure sinewaves to the ADC at frequencies f1 and f2, 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 4.22. Notice that the second-order products fall at frequencies which can be removed by digital filters. However, the third-order products 2f2–f1 and 2f1–f2 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.

SECOND AND THIRD-ORDER INTERMODULATION PRODUCTS FOR f_1 = 5MHz, f_2 = 6MHz

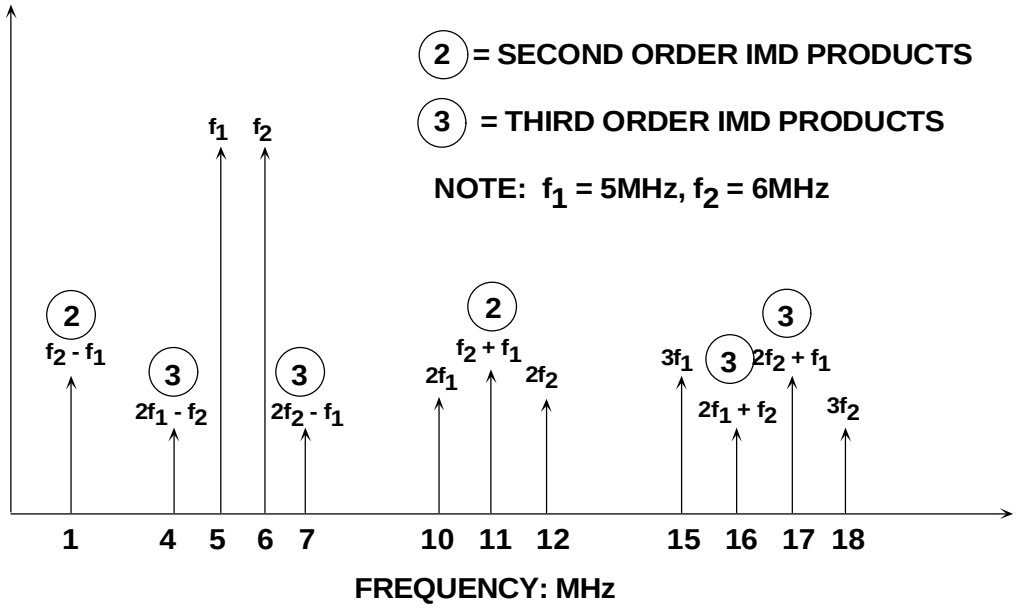

ANALOG DEVICES

4.22

Note, however, that if the two tones are close to $f_S/4$, then the aliased third harmonic of the fundamental can make the identification of the actual 2f2–f1 and 2f1–f2 products difficult. Similarly, if the two tones are close to $f_S/3$, the aliased second harmonic may interfere with the measurement.

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. This is illustrated in Figure 4.23 for the AD9042, where two-tone SFDR is plotted as a function of signal level. The plot indicates that the distortion floor ranges from 85 to 90dBFS regardless of the input signal amplitude.

AD9042 12-BIT, 41MSPS ADC TWO-TONE SFDR

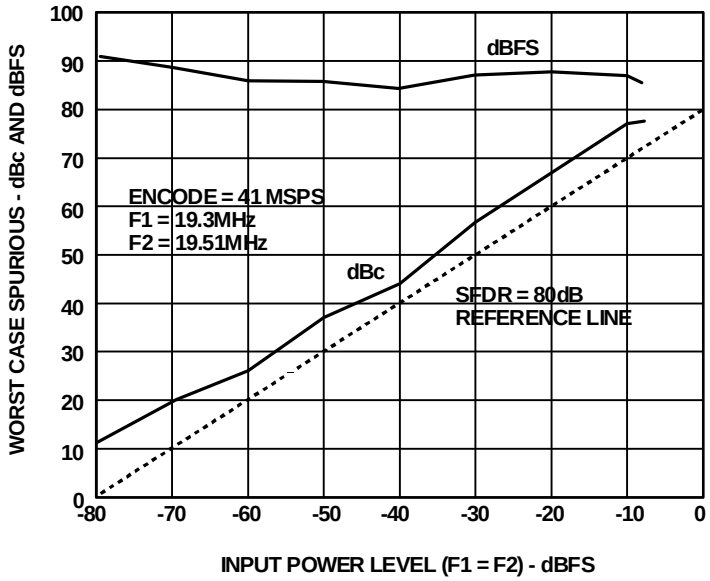

ANALOG DEVICES

4.23

Noise Power Ratio (NPR)

Noise power ratio testing has been used extensively to measure the transmission characteristics of Frequency Division Multiplexed (FDM) communications links (see Reference 4). In a typical FDM system, 4kHz wide voice channels are "stacked" in frequency bins for transmission over coaxial, microwave, or satellite equipment. At the receiving end, the FDM data is demultiplexed and returned to 4kHz individual baseband channels. In an FDM system having more than approximately 100 channels, the FDM 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 4.24).

NOISE POWER RATIO (NPR) MEASUREMENTS

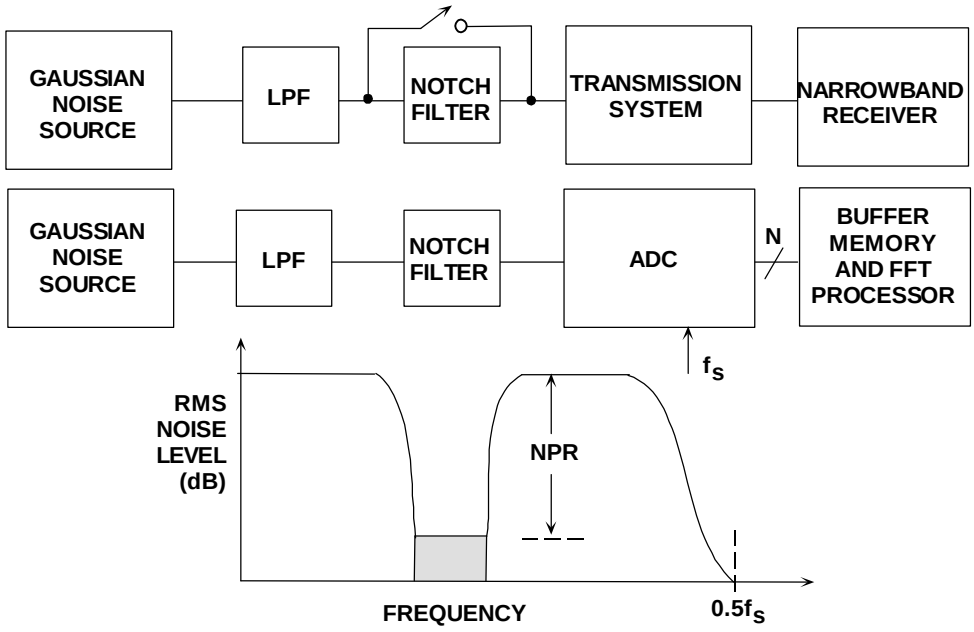

ANALOG DEVICES

4.24

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.

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. FDM 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 4.25 (see Reference 5). Peak NPR and corresponding loading levels are shown in Figure 4.26.

THEORETICAL NPR FOR 10, 11, 12-BIT ADCs

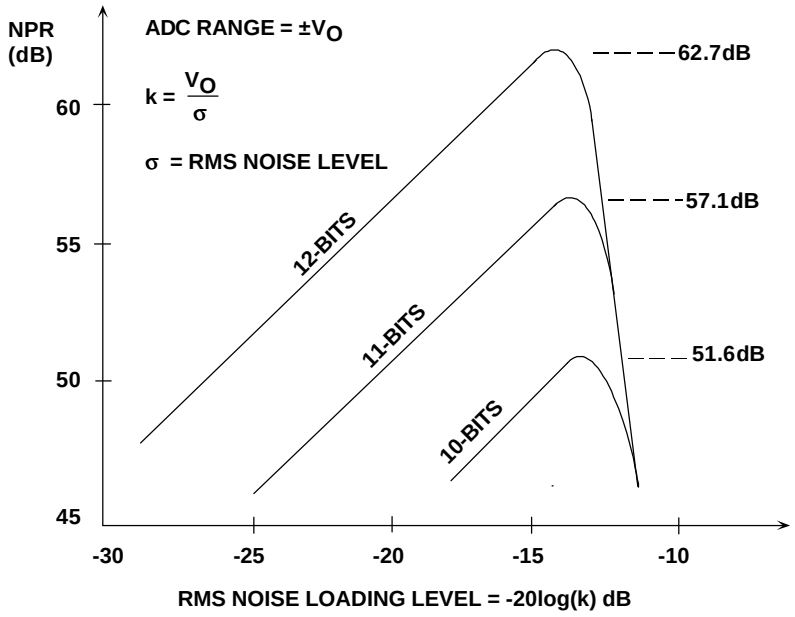

ANALOG DEVICES

4.25

THEORETICAL NPR SUMMARY

BITS	k OPTIMUM	k(dB)	MAX NPR (dB)
8	3.92	11.87	40.60
9	4.22	12.50	46.05
10	4.50	13.06	51.56
11	4.76	13.55	57.12
12	5.01	14.00	62.71
13	5.26	14.41	68.35
14	5.49	14.79	74.01
15	5.72	15.15	79.70
16	5.94	15.47	85.40

ADC Range = $\pm V_O$
$k = V_O / \sigma$
σ = RMS Noise Level

ANALOG DEVICES

4.26

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 FDM 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 4.27 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.

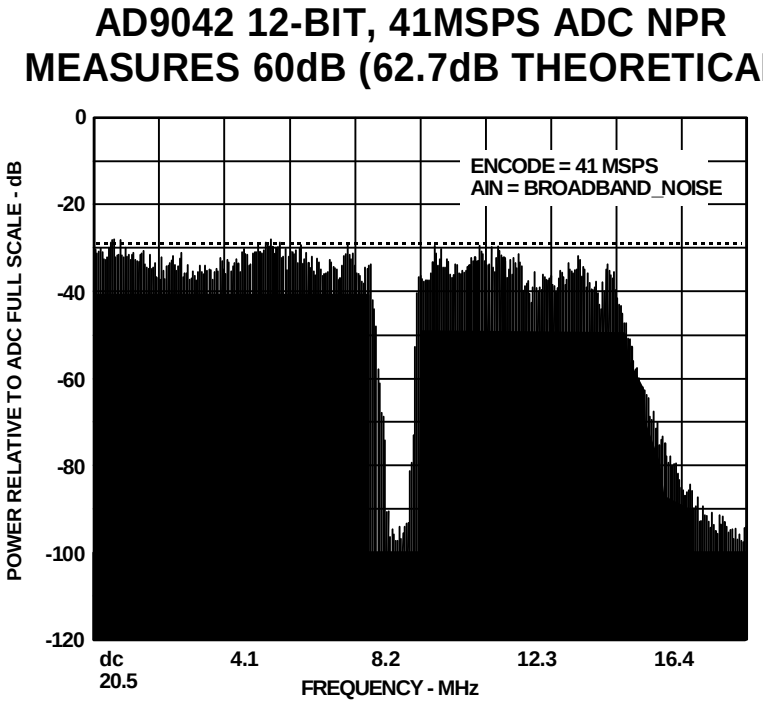

ANALOG DEVICES

4.27

Aperture Jitter and Aperture Delay

Another reason that the SNR of an ADC decreases with input frequency may be deduced from Figure 4.28, 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 4.29. 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

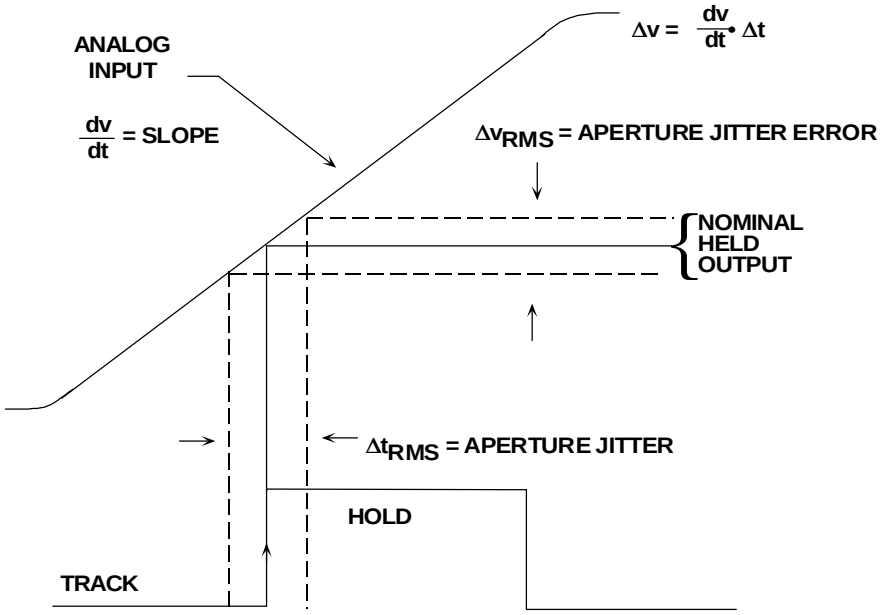

ANALOG DEVICES

4.28

SNR DUE TO APERTURE AND SAMPLING CLOCK JITTER

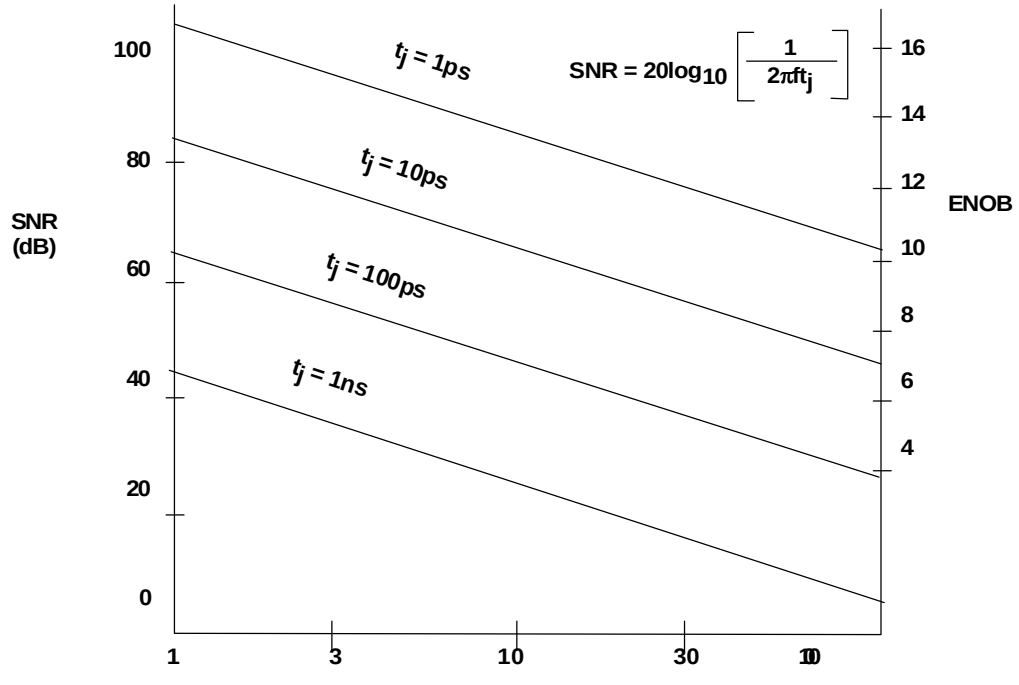

ANALOG DEVICES

4.29

A decade 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 4.30), 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.

EFFECTIVE APERTURE DELAY TIME

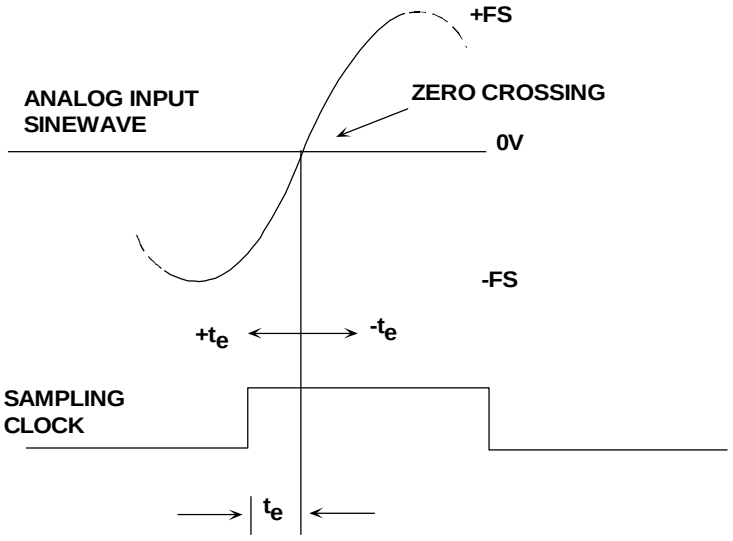

ANALOG DEVICES

4.30

HIGH SPEED ADC ARCHITECTURES

Successive Approximation ADCs

The successive approximation (SAR) ADC architecture has been used for decades and is still a popular and cost effective form of converter for sampling frequencies of 1MSPS or less. A simplified block diagram of a SAR ADC is shown in Figure 4.31. On the START CONVERT command, all the bits of the successive approximation register (SAR) are reset to "0" except the MSB which is set to "1". Bit 1 is then tested in the following manner: If the DAC output is greater than the analog input, the MSB is reset, otherwise it is left set. The next most significant bit is then tested by setting it to "1". If the DAC output is greater than the analog input, this bit is reset,

otherwise it is left set. The process is repeated with each bit in turn. When all the bits have been set, tested, and reset or not as appropriate, the contents of the SAR correspond to the digital value of the analog input, and the conversion is complete.

SUCCESSIVE APPROXIMATION ADC

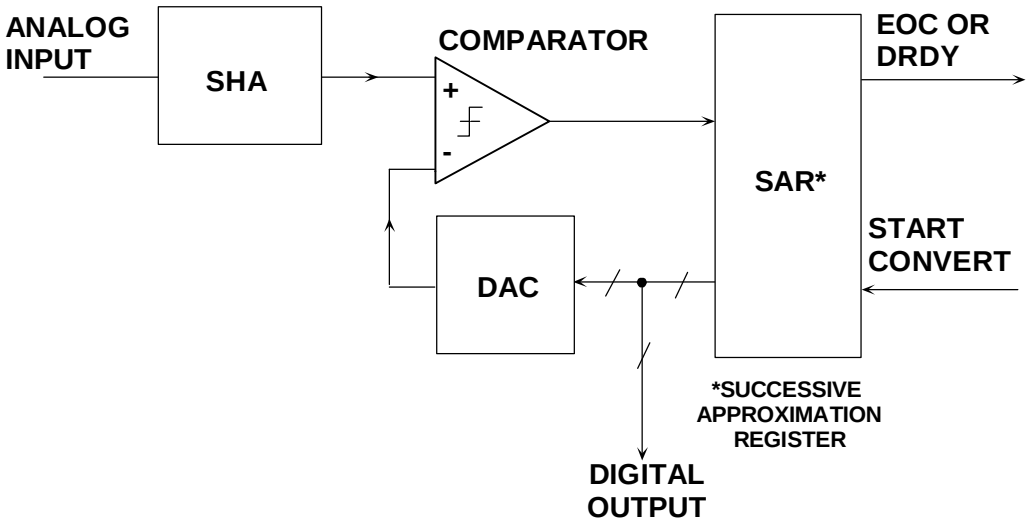

ANALOG DEVICES

4.31

An N-bit conversion takes N steps. It would seem on superficial examination that a 16-bit converter would have a conversion time that is twice as long as an 8-bit one, but this is not the case. In an 8-bit converter, the DAC must settle to 8-bit accuracy before the bit decision is made, whereas in a 16-bit converter, it must settle to 16-bit accuracy, which takes a lot longer. In practice, 8-bit successive approximation ADCs can convert in a few hundred nanoseconds, while 16-bit ones will generally take several microseconds.

The classic SAR ADC is only a quantizer: no sampling takes place, and for an accurate conversion, the input must remain constant for the entire conversion period. Most modern SAR ADCs are sampling types and have an internal sample-and-hold so that they can process AC signals. They are specified for both AC and DC applications. A SHA is required in a SAR ADC because the signal must remain constant during the entire N-bit conversion cycle.

The accuracy of a SAR ADC depends primarily on the accuracy (differential and integral linearity, gain, and offset) of the internal DAC. Until recently, this accuracy was achieved using laser trimmed thin film resistors. Modern SAR ADCs utilize CMOS switched capacitor charge redistribution DACs. This type of DAC depends on the accurate ratio matching and stability of on-chip capacitors rather than thin film resistors. For resolutions greater than 12-bits, on-chip autocalibration techniques using an additional *calibration DAC* and the accompanying logic can accomplish the same thing as thin film laser trimmed resistors, at much less cost. Therefore, the entire ADC can be made on a standard sub-micron CMOS process.

The successive approximation ADC has a very simple structure, is low power, and has reasonably fast conversion times (<1MSPS). It is probably the most widely used ADC architecture, and will continue to be used for medium speed and medium resolution applications.

Current 12-bit SAR ADCs achieve sampling rates up to about 1MSPS, and 16-bit ones up to about 300kSPS. Examples of typical state-of-the-art SAR ADCs are the AD7892 (12-bits at 600kSPS), the AD976/977 (16-bits at 100kSPS), and the AD7882 (16-bits at 300kSPS).

Flash Converters

Flash ADCs (sometimes called *parallel* ADCs) are the fastest type of ADC and use large numbers of comparators. An N-bit flash ADC consists of 2^N resistors and 2^N-1 comparators arranged as in Figure 4.32. Each comparator has a reference voltage which is 1 LSB higher than that of the one below it in the chain. For a given input voltage, all the comparators below a certain point will have their input voltage larger than their reference voltage and a "1" logic output, and all the comparators above that point will have a reference voltage larger than the input voltage and a "0" logic output. The 2^N-1 comparator outputs therefore behave in a way analogous to a mercury thermometer, and the output code at this point is sometimes called a *thermometer* code. Since 2^N-1 data outputs are not really practical, they are processed by a decoder to an N-bit binary output.

FLASH OR PARALLEL ADC

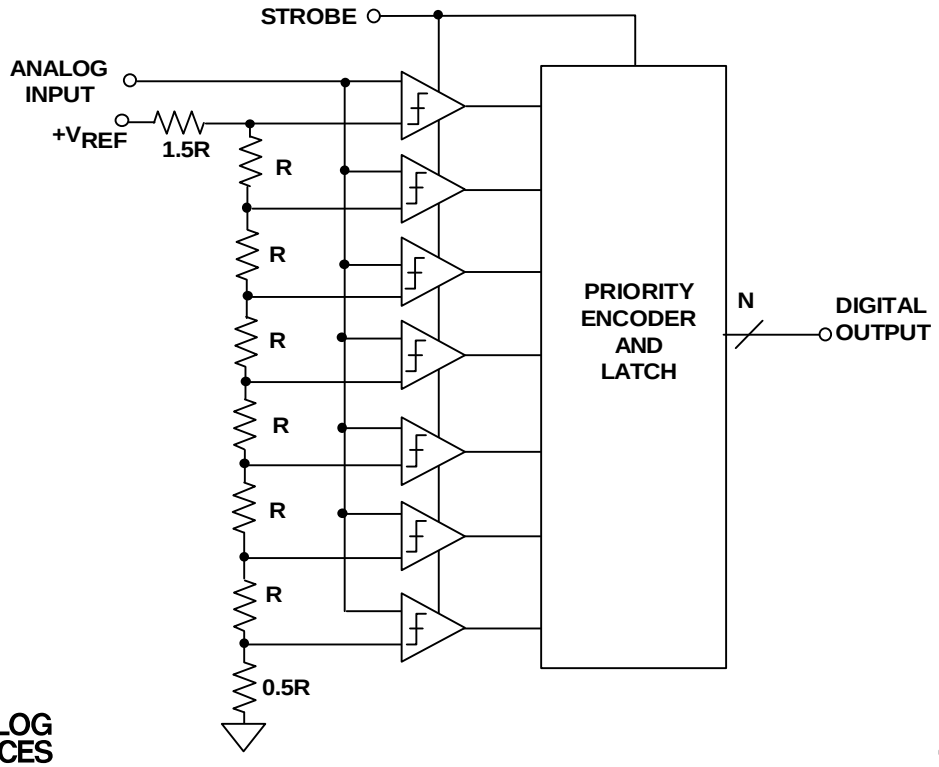

ANALOG DEVICES

4.32

The input signal is applied to all the comparators at once, so the thermometer output is delayed by only one comparator delay from the input, and the encoder N-bit output by only a few gate delays on top of that, so the process is very fast. However, the architecture uses large numbers of resistors and comparators and it limited to low resolutions, and if it is to be fast, each comparator must run at relatively high power levels. Hence, the problems of flash ADCs include limited resolution, high power dissipation because of the large number of high speed comparators (especially at sampling rates greater than 50MSPS), and relatively large (and therefore expensive) chip sizes. In addition, the resistance of the reference resistor chain must be kept low to supply adequate bias current to the fast comparators, so the voltage reference has to source quite large currents (>10 mA).

In practice, flash converters are available up to 10-bits, but more commonly they have 8-bits of resolution. Their maximum sampling rate can be as high as 500 MSPS, and input full-power bandwidths in excess of 300 MHz.

But as mentioned earlier, full-power bandwidths are not necessarily full-resolution bandwidths. Ideally, the comparators in a flash converter are well matched both for DC and AC characteristics. Because the strobe is applied to all the comparators simultaneously, the flash converter is inherently a sampling converter. In practice, there are delay variations between the comparators and other AC mismatches which cause a degradation in ENOB at high input frequencies. This is because the inputs are slewing at a rate comparable to the comparator conversion time.

The input to a flash ADC is applied in parallel to a large number of comparators. Each has a voltage-variable junction capacitance, and this signal-dependent capacitance results in all flash ADCs having reduced ENOB and higher distortion at high input frequencies. A model is shown in Figure 4.33, where the input capacitance is modeled as a fixed 10pF capacitor in parallel with a variable capacitor (modeled as a diode with a zero-bias junction capacitance of 6pF). As the input changes from –FS to +FS, the total input capacitance changes from about 12.5 to 16pF. The wideband external drive amplifier is isolated from the flash converter by a 50Ω series resistor. The distortion of this circuit degrades from about 70dBc at 1MHz to 35dBc at 100MHz.

SIGNAL-DEPENDENT INPUT CAPACITANCE CAUSES DISTORTION AT HIGH FREQUENCIES

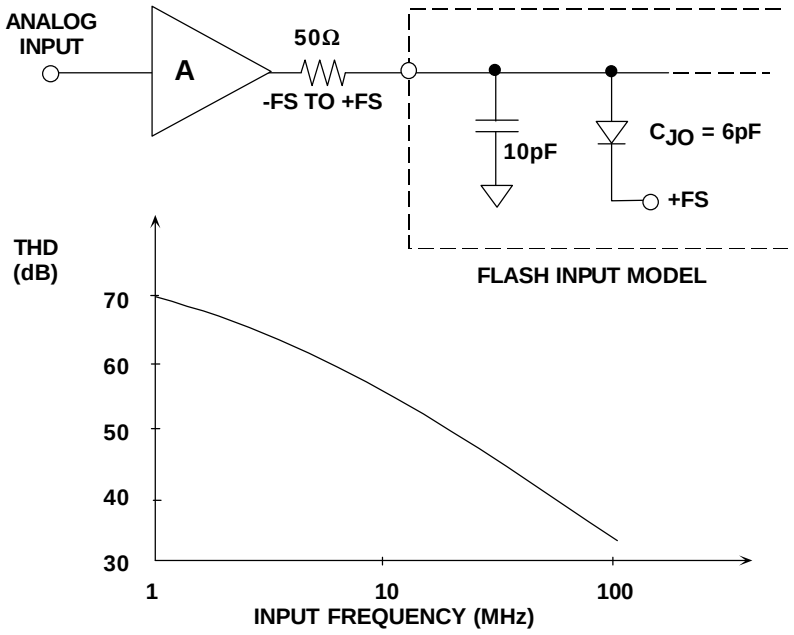

ANALOG DEVICES

4.33

High data rate digital communications applications such as set-top boxes for direct broadcast satellites (DBS) require dual 6 or 8-bit high speed ADCs to perform quadrature demodulation. A dual flash converter ensures good matching between the two ADCs. The AD9066 (dual 6-bit, 60MSPS) flash converter is representative of this type of converter. The AD9066 is fabricated on a BiCMOS process, operates on a single +5V supply, and dissipates 400mW. The effective bit performance of the device is shown in Figure 4.34. Note that the device maintains greater than 5 ENOBs up to 60MSPS analog input.

AD9066 DUAL 6-BIT, 60MSPS ADC ENOB VS. ANALOG INPUT FREQUENCY

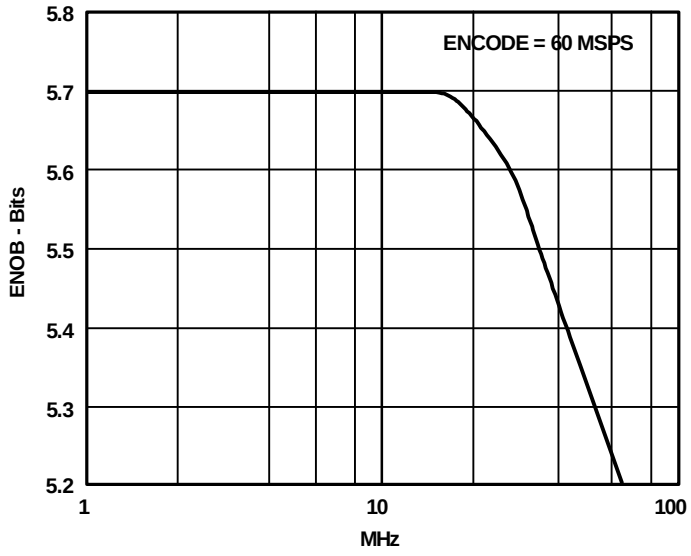

ANALOG DEVICES

4.34

Part of the reason for the excellent performance of the AD9066 is the use of an interpolation scheme that reduces the number of differential amplifiers required by a factor of two (see Reference 6). The architecture enables 64 possible quantization levels to be determined with only 32 preamplifiers which drive 63 latches. This keeps the input capacitance to a minimum (10pF) and reduces total power dissipation of the device. The basic interpolation circuit is shown in Figure 4.35.

"INTERPOLATING" FLASH REDUCES THE NUMBER OF PREAMPLIFIERS BY FACTOR OF TWO

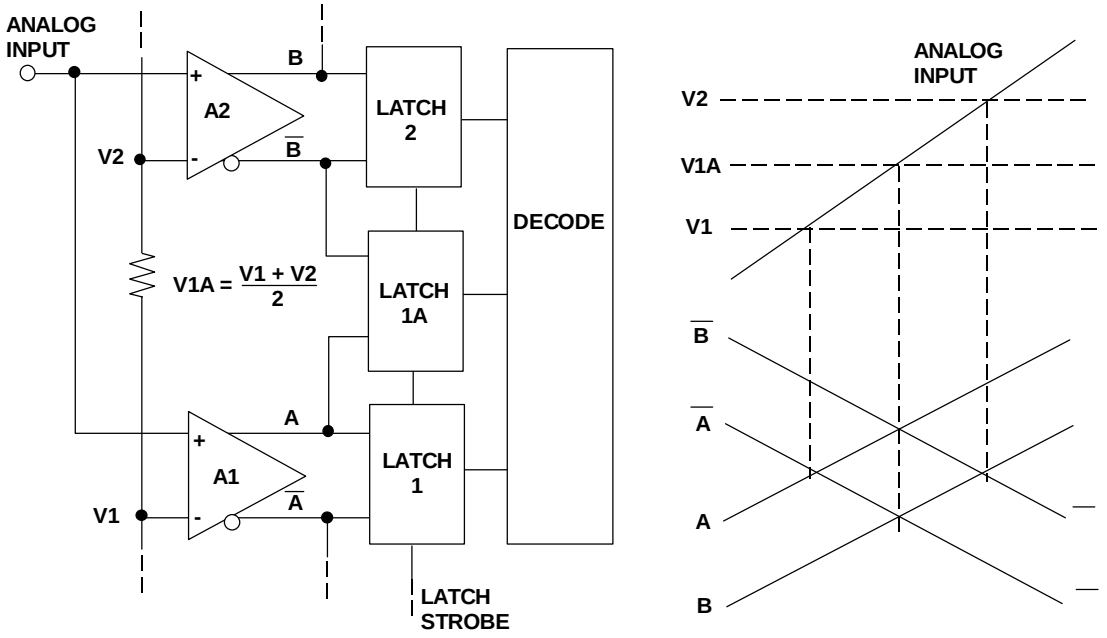

ANALOG DEVICES

4.35

The preamplifiers are low-gain g_m stages whose bandwidth is proportional to the tail currents of the differential pairs. Consider the case for a positive-going ramp input which is initially below the reference to AMP A1, V1. As the input signal approaches V1, the differential output of A1 approaches zero (i.e., A = $\overline{A}$), and the decision point is reached. The output of A1 drives the differential input of LATCH 1. As the input signals continues to go positive, A continues to go positive, and $\overline{B}$ begins to go negative. The interpolated decision point is determined when A = $\overline{B}$. As the input continues positive, the third decision point is reached when B = $\overline{B}$. This novel architecture reduces the ADC input capacitance and thereby minimizes its change with signal level and the associated distortion. The input capacitance of the AD9066 is only about 10pF. Key specifications for the device are summarized in Figure 4.36.

AD9066 DUAL 6-BIT, 60MSPS FLASH ADC KEY SPECIFICATIONS

- **Input Range: 500mV p-p**
- **Input Impedance: 50kΩ || 10pF**
- **ENOB: 5.7bits @ 15.5MHz Input**
- **On-Chip Reference**
- **Power Supply: Single +5V**
- **Power Dissipation: 400mW**

- **Package: 28-pin SOIC**
- **Ideal for Quadrature Demodulation**

ANALOG DEVICES

4.36

Subranging (Pipelined) ADCs

Although it is not practical to make flash ADCs with high resolution, flash ADCs are often used as subsystems in "subranging" ADCs (sometimes known as "half-flash ADCs"), which are capable of much higher resolutions (up to 16-bits).

A block diagram of an 8-bit subranging ADC based upon two 4-bit flash converters is shown in Figure 4.37. Although 8-bit flash converters are readily available at high sampling rates, this example will be used to illustrate the theory. The conversion process is done in two steps. The first four significant bits (MSBs) are digitized by the first flash (to better than 8-bits accuracy), and the 4-bit binary output is applied to a 4-bit DAC (again, better than 8-bit accurate). The DAC output is subtracted from the held analog input, and the resulting residue signal is amplified and applied to the second 4-bit flash. The outputs of the two 4-bit flash converters are then combined into a single 8-bit binary output word. If the residue signal range does not exactly fill the range of the second flash converter, non-linearities and perhaps missing codes will result.

8-BIT SUBRANGING ADC

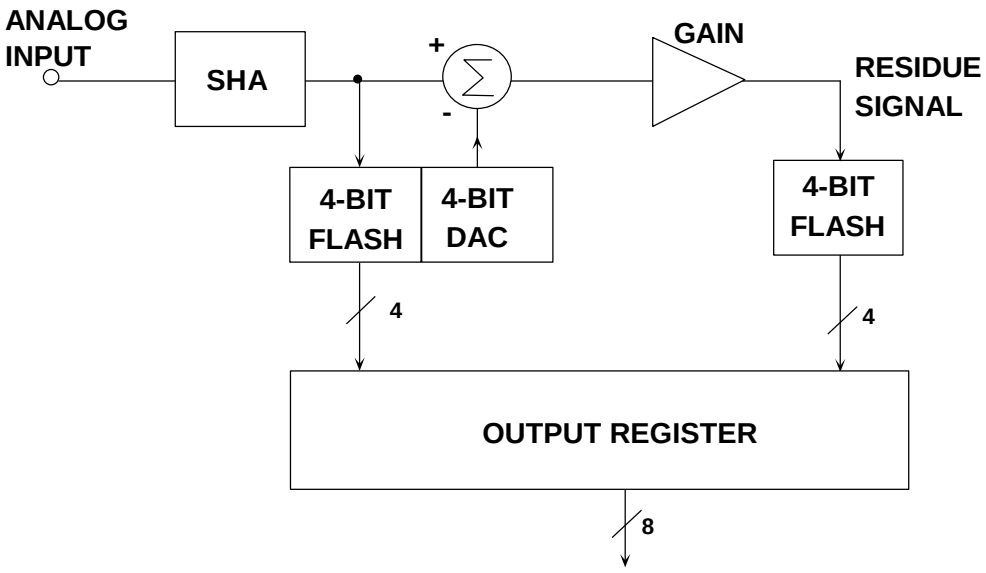

ANALOG DEVICES

4.37

Modern subranging ADCs use a technique called *digital correction* to eliminate problems associated with the architecture of Figure 4.37. A simplified block diagram

of a 12-bit digitally corrected subranging (DCS) ADC is shown in Figure 4.38. The architecture is similar to that used in the AD9042 12-bit, 41MSPS ADC. Note that a 6-bit and an 7-bit ADC have been used to achieve an overall 12-bit output. These are not flash ADCs, but utilize a *magnitude-amplifier (MagAmp™)* architecture which will be described shortly.

AD9042 12-BIT, 41MSPS PIPELINED SUBRANGING ADC WITH DIGITAL ERROR CORRECTION

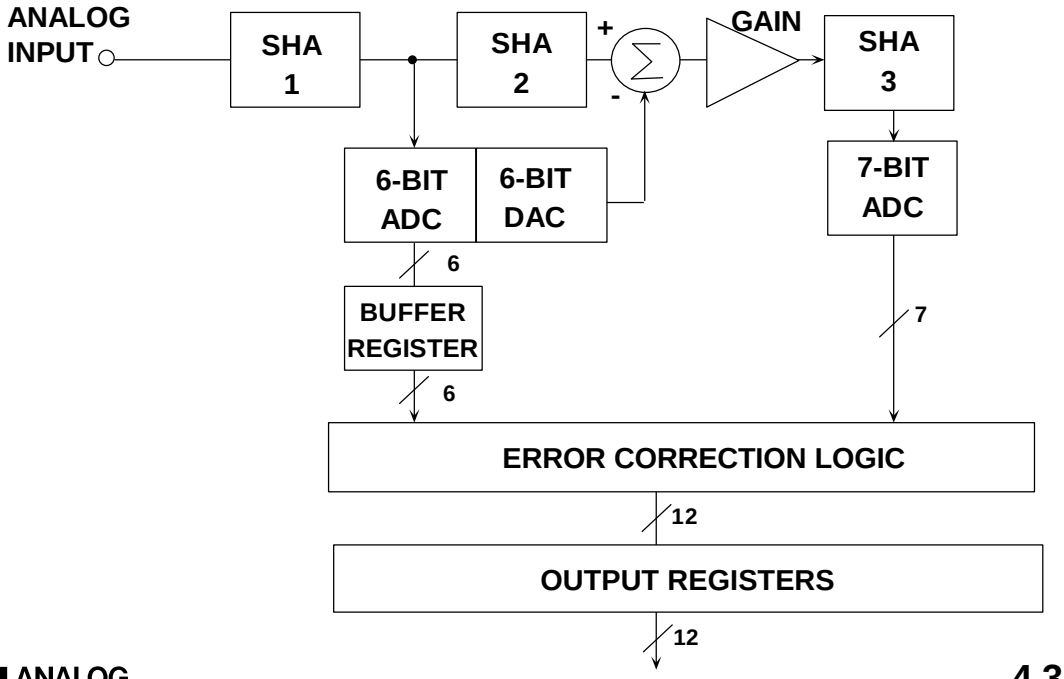

ANALOG DEVICES

4.38

If there were no errors in the first-stage conversion, the 6-bit "residue" signal applied to the 7-bit ADC by the summing amplifier would never exceed one-half of the range of the 7-bit ADC. The extra range in the second ADC is used in conjunction with the error correction logic (usually just a full adder) to correct the output data for most of the errors inherent in the traditional uncorrected subranging converter architecture. It is important to note that the 6-bit DAC must be better than 12-bit accurate, because the digital error correction does not correct for DAC errors. In practice, "thermometer" or "fully-decoded" DACs using one current switch per level (63 switches in the case of a 6-bit DAC) are often used instead of a "binary" DAC to ensure excellent differential and integral linearity and minimum switching transients.

The second SHA delays the held output of the first SHA while the first-stage conversion occurs, thereby maximizing throughput. The third SHA serves to *deglitch* the residue output signal, thereby allowing a full conversion cycle for the 7-bit ADC to make its decision (the 6 and 7-bit ADCs in the AD9042 are bit-serial *MagAmp* ADCs which require more settling time than a flash converter).

This multi-stage conversion technique is sometimes referred to as "pipelining." Additional shift registers in series with the digital outputs of the first-stage ADC ensure that its output is ultimately time-aligned with the last 7 bits from the second ADC when their outputs are combined in the error correction logic. A

pipelined ADC therefore has a specified number of clock cycles of *latency*, or *pipeline delay* associated with the output data. The leading edge of the sampling clock (for sample N) is used to clock the output register, but the data which appears as a result of that clock edge corresponds to sample N – L, where L is the number of clock cycles of latency. In the case of the AD9042, there are two clock cycles of latency. Key specifications for the AD9042 are shown in Figure 4.39.

AD9042 12-BIT, 41MSPS ADC KEY SPECIFICATIONS

- **Input Range: 1V peak-to-peak, V_{cm} = +2.4V**
- **Input Impedance: 250Ω to V_{cm}**
- **Effective Input Noise: 0.33LSBs rms**
- **SFDR at 20MHz Input: 80dB minimum**
- **SINAD (S/N+D) at 20MHz Input = 67dB**
- **Digital Outputs: TTL Compatible**
- **Power Supply: Single +5V**
- **Power Dissipation: 595mW**
- **Fabricated on High Speed Dielectrically Isolated Complementary Bipolar Process**

ANALOG DEVICES

4.39

The error correction scheme described above is designed to correct for errors made in the first conversion. Internal ADC gain, offset, and linearity errors are corrected as long as the residue signal fall within the range of the second-stage ADC. These errors will not affect the linearity of the overall ADC transfer characteristic. Errors made in the final conversion, however, do translate directly as errors in the overall transfer function. Also, linearity errors or gain errors either in the DAC or the residue amplifier will not be corrected and will show up as nonlinearities or non-monotonic behavior in the overall ADC transfer function.

So far, we have considered only two-stage subranging ADCs, as these are easiest to analyze. There is no reason to stop at two stages, however. Three-pass and four-pass subranging pipelined ADCs are quite common, and can be made in many different ways, usually with digital error correction.

A simplified block diagram of the AD9220 12-bit, 10MSPS single-supply, 250mW CMOS ADC is shown in Figure 4.40. The AD9221 (1.25MSPS, 60mW) and the AD9223 (3MSPS, 100mW) ADCs use the identical architecture but operate at lower power and lower sampling rates. This is a four-stage pipelined architecture with an additional bit in the second, third, and fourth stage for error correction. Because of

the pipelined architecture, these ADCs have a 3 clock-cycle latency (see Figure 4.41). Key specifications for the AD9220/9221/9223 are given in Figure 4.42.

AD9220/9221/9223 12-BIT PIPELINED CMOS ADC

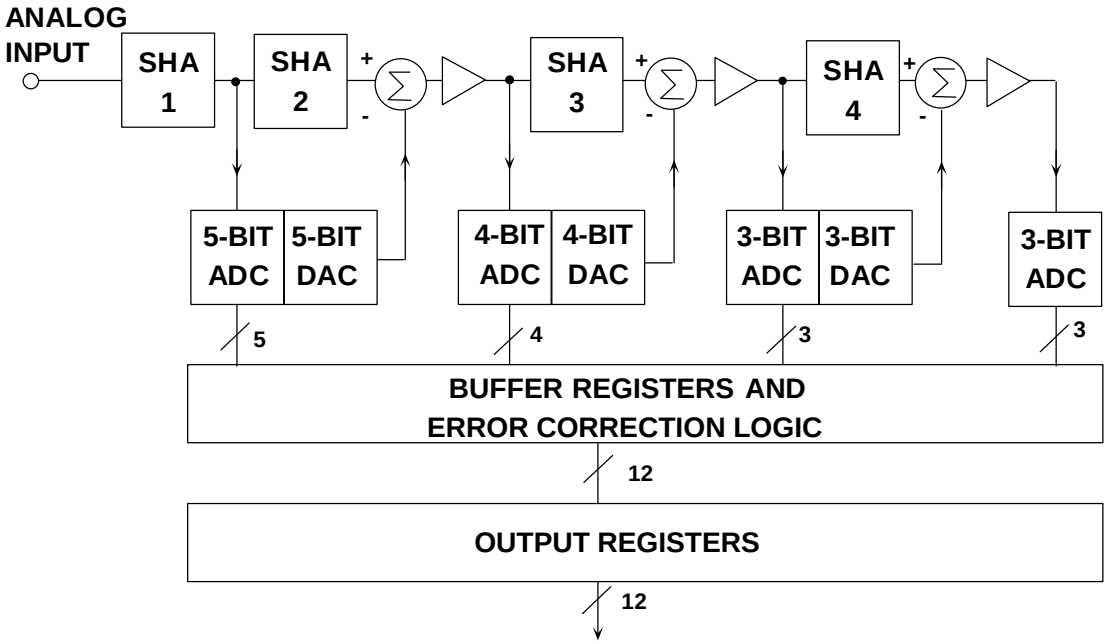

ANALOG DEVICES

4.40

LATENCY (PIPELINE DELAY) OF AD9220/9221/9223 ADC

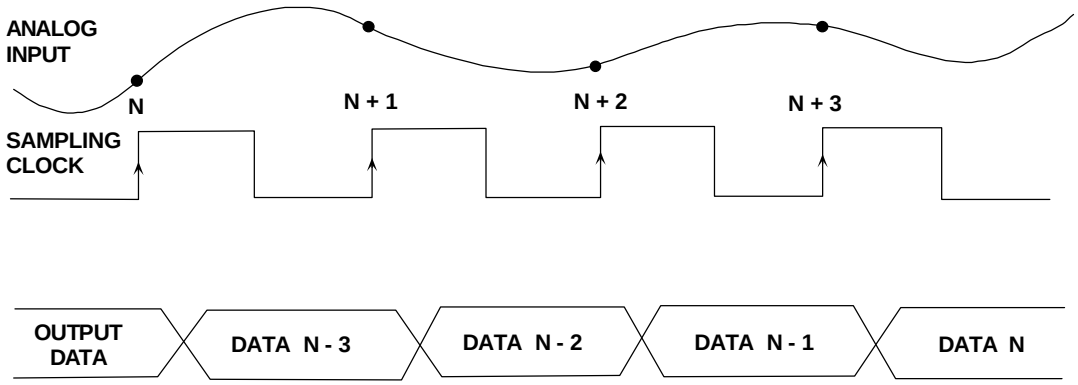

ANALOG DEVICES

4.41

AD9220, AD9221, AD9223
CMOS 12-BIT ADCs KEY SPECIFICATIONS

- **Family Members:**
  AD9221 (1.25MSPS), AD9223 (3MSPS), AD9220 (10MSPS)
- **Power Dissipation: 60, 100, 250mW, Respectively**
- **FPBW: 25, 40, 60MHz, Respectively**
- **Effective Input Noise: 0.1LSB rms (Span = 5V)**
- **SINAD: 71dB**
- **SFDR: 88dBc**
- **On-Chip Reference**
- **Differential Non-Linearity: 0.3LSB**
- **Single +5V Supply**
- **28-Pin SOIC Package**

ANALOG DEVICES

4.42

Bit-Per-Stage (Serial, or Ripple) ADCs

Various architectures exist for performing A/D conversion using one stage per bit. In fact, a multistage subranging ADC with one bit per stage and no error correction is one form. Figure 4.43 shows the overall concept. The SHA holds the input signal constant during the conversion cycle. There are N stages, each of which have a bit output and a residue output. The residue output of one stage is the input to the next. The last bit is detected with a single comparator as shown.

BIT-PER-STAGE, SERIAL, OR RIPPLE ADC

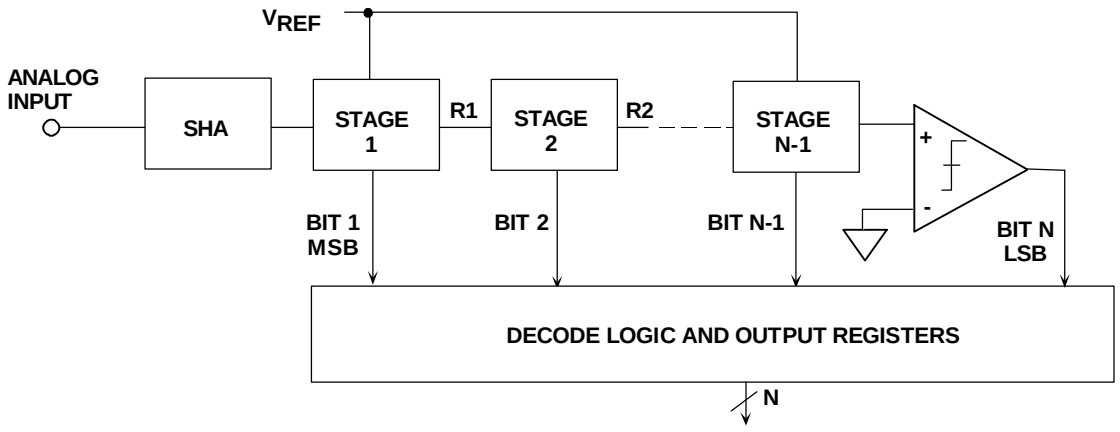

ANALOG DEVICES

4.43

The basic stage for performing a single binary bit conversion is shown in Figure 4.44. It consists of a gain-of-two amplifier, a comparator, and a 1-bit DAC. The comparator detects the zero-crossing of the input and is the binary bit output for that stage. The comparator also switches a 1-bit DAC whose output is summed with the output of the gain-of-two amplifier. The resulting residue output is then applied to the next stage.

SINGLE-STAGE OF BINARY ADC

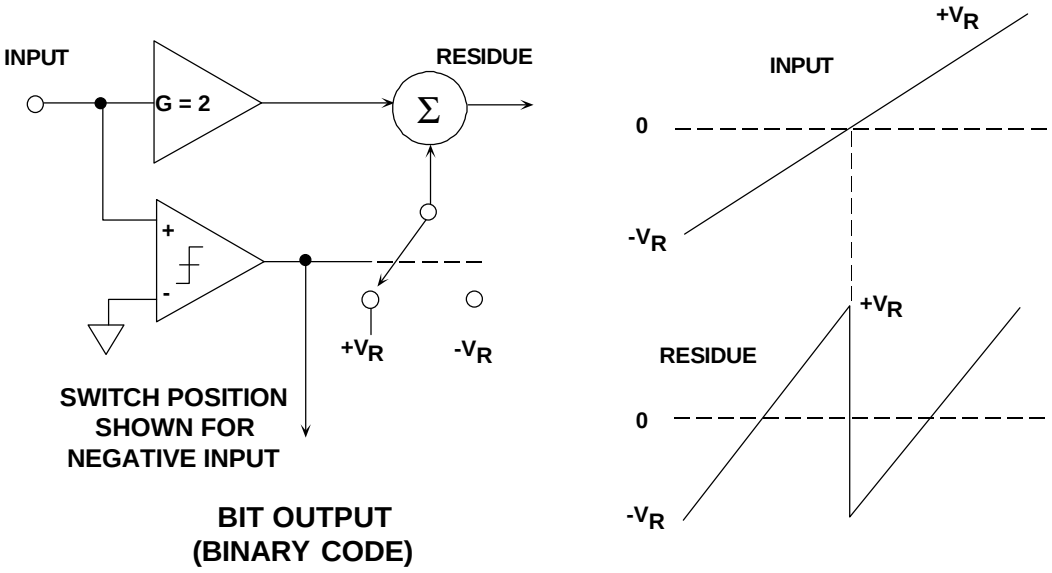

ANALOG DEVICES

4.44

A simplified 3-bit serial-binary ADC is shown in Figure 4.45, and the residue outputs are shown in Figure 4.46. Each residue output signal has discontinuities which correspond to the point where the comparator changes state and causes the DAC to switch. The fundamental problem with this architecture is the discontinuity in the residue output waveforms. Adequate settling time must be allowed for these transients to propagate through all the stages and settle at the final comparator input. The prospects of making this architecture operate at high speed are therefore dismal.

3-BIT SERIAL ADC WITH BINARY OUTPUT

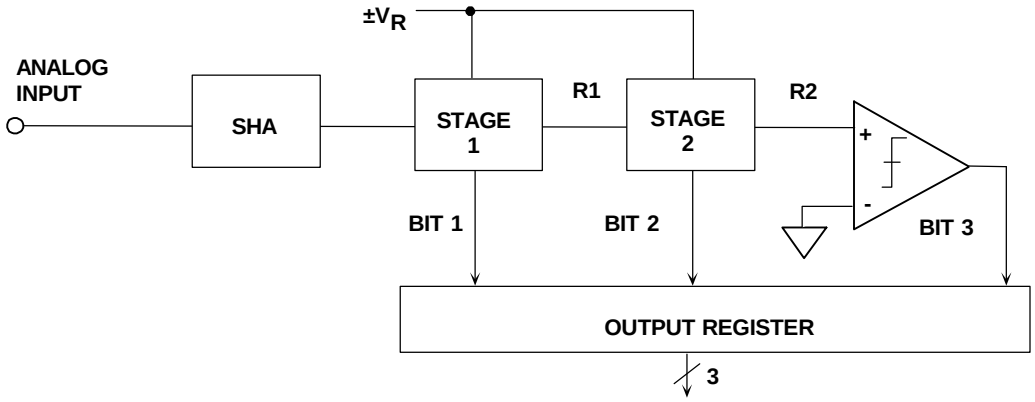

ANALOG DEVICES

4.45

INPUT AND RESIDUE WAVEFORMS OF 3-BIT BINARY RIPPLE ADC

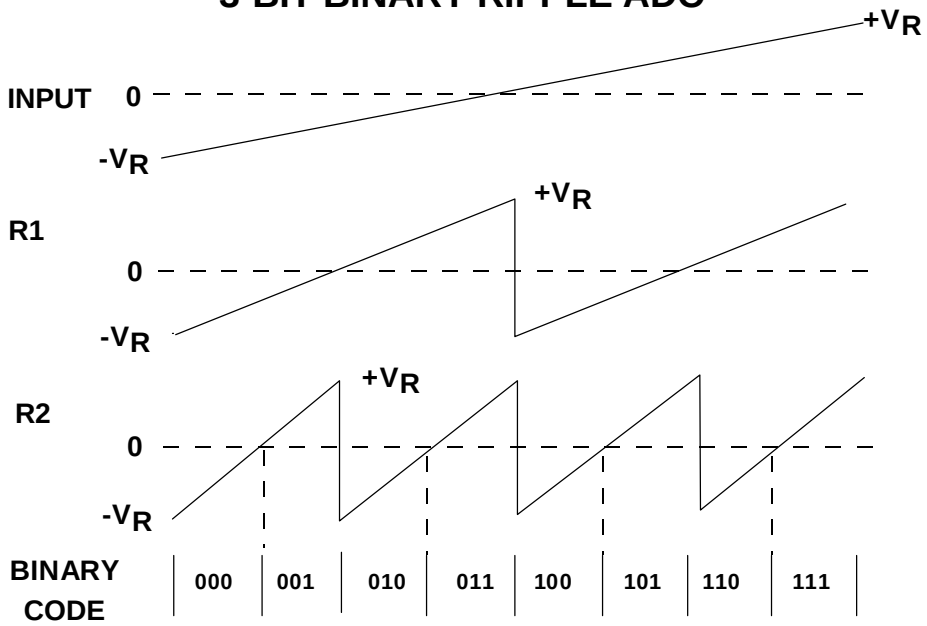

ANALOG DEVICES

4.46

A much better bit-per-stage architecture was developed by F.D. Waldhauer (Reference 7) based on absolute value amplifiers (magnitude amplifiers, or simply *MagAmps*™). This scheme has often been referred to as *serial-Gray* (since the output coding is in Gray code), or *folding* converter (References 8, 9, 10). The basic stage is

shown functionally in Figure 4.47. The comparator detects the polarity of the input signal and provides the Gray bit output for the stage. It also determines whether the overall stage gain is +2 or –2. The reference voltage V_R is summed with the switch output to generate the residue signal which is applied to the next stage. The transfer function for the folding stage is also shown in Figure 4.47.

MagAmp STAGE FUNCTIONAL EQUIVALENT CIRCUIT

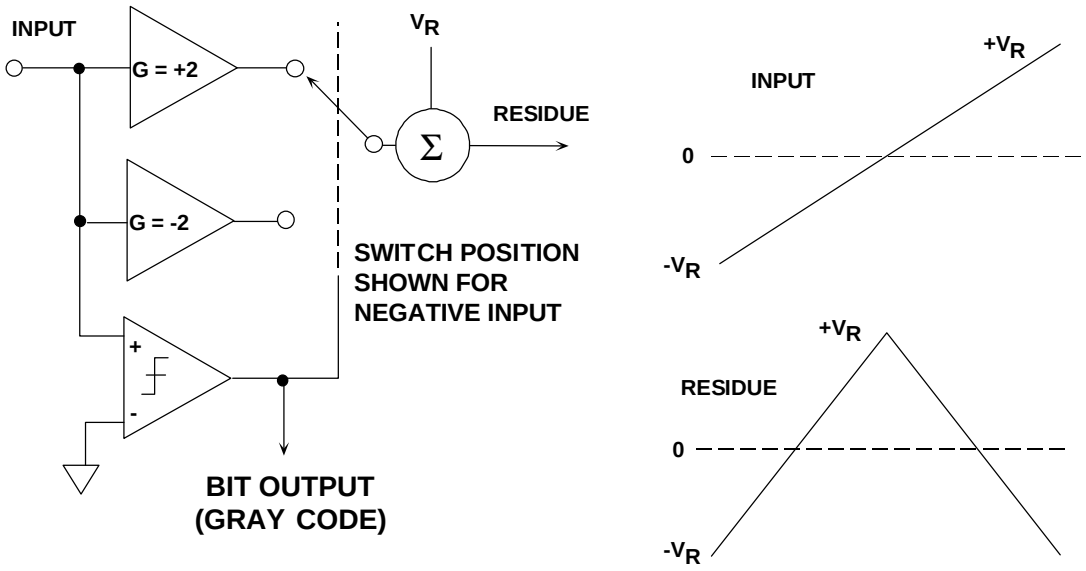

ANALOG DEVICES

4.47

A 3-bit MagAmp folding ADC is shown in Figure 4.48, and the corresponding residue waveforms in Figure 4.49. Notice that there is no abrupt transition in any of the folding stage output waveforms.

3-BIT MagAmp™ (FOLDING) ADC BLOCK DIAGRAM

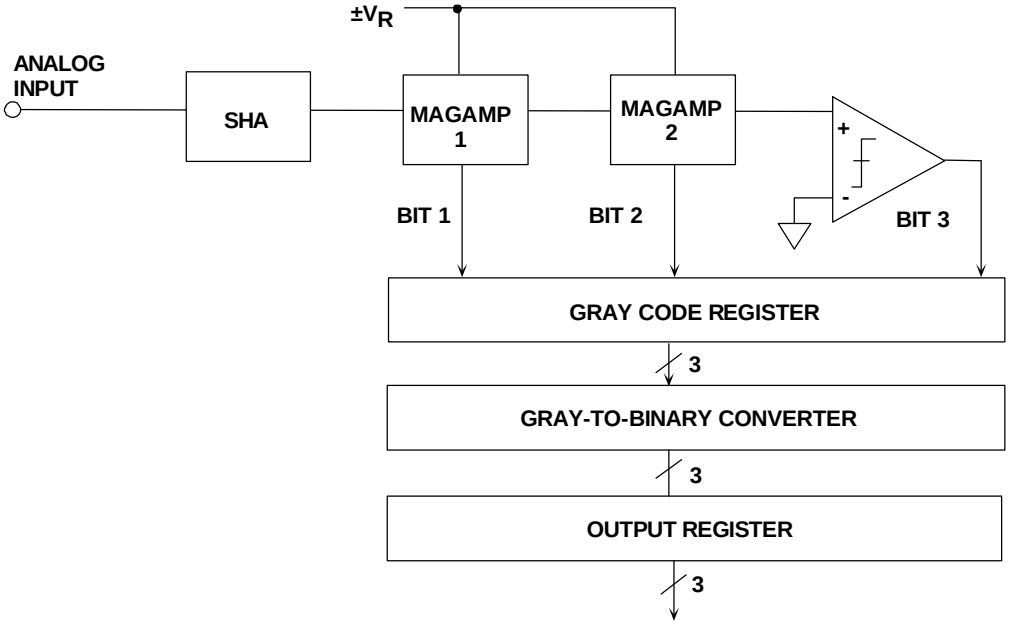

ANALOG DEVICES

4.48

INPUT AND RESIDUE WAVEFORMS FOR 3-BIT MagAmp ADC

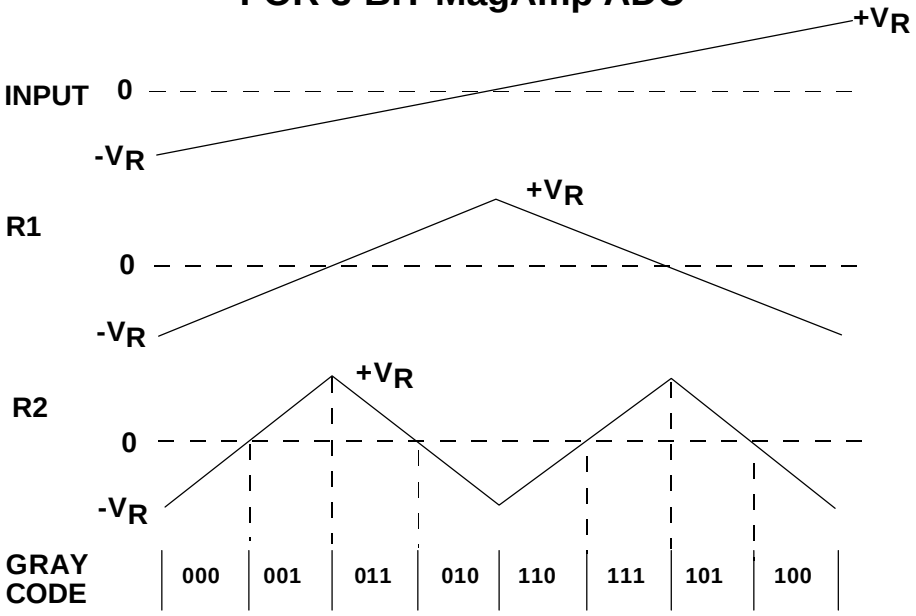

ANALOG DEVICES

4.49

The key to operating this architecture at high speeds is the folding stage. Early designs (see References 7, 8, 9) used discrete op amps with diodes inside the feedback loop to generate the folding transfer function. Modern IC circuit designs implement the transfer function using current-steering open-loop gain techniques

which can be made to operate much faster. Fully differential stages (including the SHA) also provide speed, lower distortion, and yield 8-bit accurate folding stages with no requirement for thin film resistor laser trimming.

An example of a fully differential gain-of-two MagAmp folding stage is shown in Figure 4.50 (see References 11, 12, 13). The differential input signal is applied to the degenerated-emitter differential pair Q1,Q2 and the comparator. The differential input voltage is converted into a differential current which flows in the collectors of Q1, Q2. If +IN is greater than –IN, cascode-connected transistors Q3, Q6 are on, and Q4, Q6 are off. The differential signal currents therefore flow through the collectors of Q3, Q6 into level-shifting transistors Q7, Q8 and into the output load resistors, developing the differential output voltage between +OUT and –OUT. The overall differential voltage gain of the circuit is two.

If +IN is less than –IN (negative differential input voltage), the comparator changes stage and turns Q4, Q5 on and Q3, Q6 off. The differential signal currents flow from Q5 to Q7 and from Q4 to Q8, thereby maintaining the same relative polarity at the differential output as for a positive differential input voltage. The required offset voltage is developed by adding a current I_{OFF} to the emitter current of Q7 and subtracting it from the emitter current of Q8.

The differential residue output voltage of the stage drives the next stage input, and the comparator output represents the Gray code output for the stage.

CIRCUIT DETAILS OF MagAmp STAGE

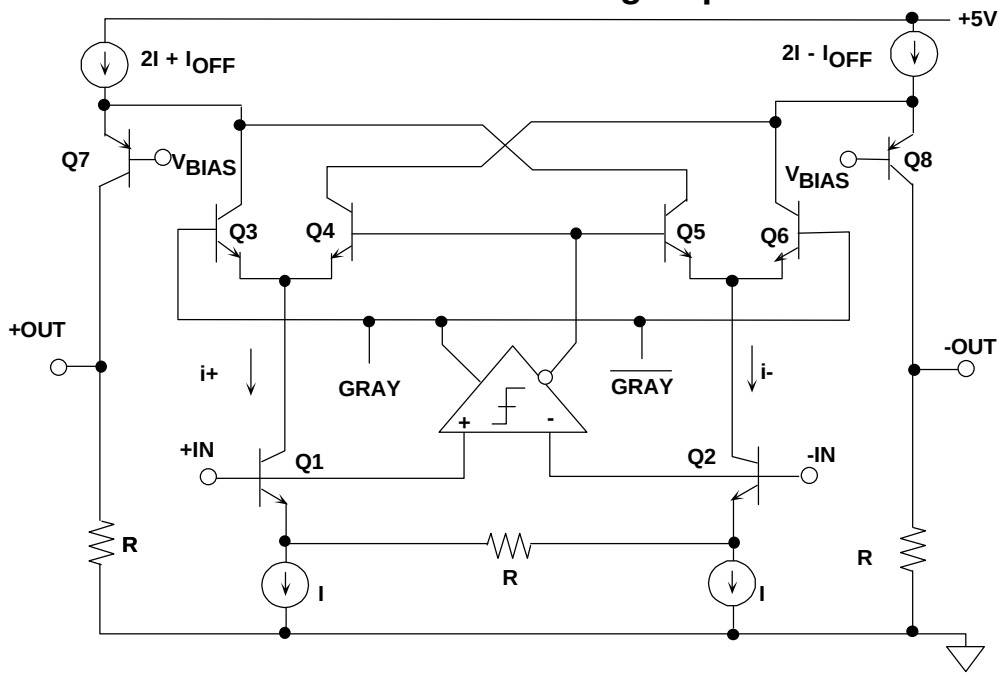

ANALOG DEVICES

4.50

The MagAmp architecture can be extended to sampling rates previously dominated by flash converters. The AD9059 8-bit, 60MSPS dual ADC is shown in Figure 4.51. The first five bits (Gray code) are derived from five differential MagAmp stages. The differential residue output of the fifth MagAmp stage drives a 3-bit flash converter, rather than a single comparator. The Gray-code output of the five MagAmps and the

binary-code output of the 3-bit flash are latched, all converted into binary, and latched again in the output data register. Key specifications for the AD9059 are shown in Figure 4.52.

AD9059 DUAL 8-BIT, 60MSPS ADC FUNCTIONAL DIAGRAM

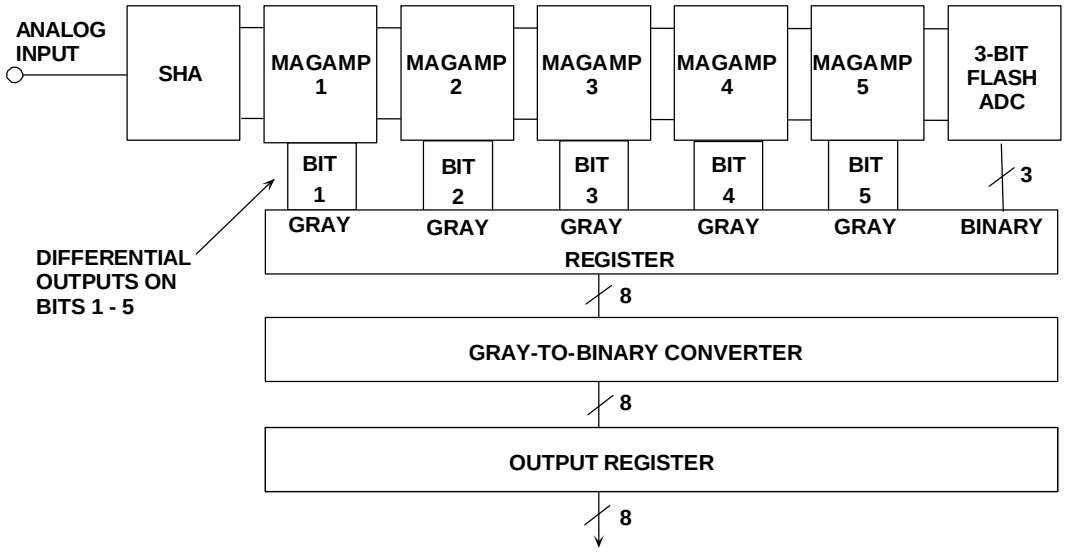

ANALOG DEVICES

4.51

AD9059 DUAL 8-BIT, 60MSPS ADC KEY SPECIFICATIONS

- **Input Range: 1V p-p, Vcm = +2.5V**
- **Input Impedance: 200kΩ || 5pF**
- **ENOB: 7.3 @ 10.3MHz Input**
- **On-Chip Reference**
- **Power Supply: Single +5V Supply (+5 or +3V Digital)**
- **Power Dissipation: 375mW (Power Down: 10mW)**
- **Package: 28-lead SSOP**
- **Ideal for Quadrature Demodulation in DBS Set-Top Boxes**

ANALOG DEVICES

4.52

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*, 1**995 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.

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 VSLI**, 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.