SECTION 12

SELECTING THE RIGHT HIGH SPEED ADC

- DISCRETE TIME SAMPLING OF ANALOG SIGNALS
- PROPERLY SPECIFYING THE ADC SAMPLING RATE AND THE ANTIALIASING FILTER
- PROPERLY SPECIFYING THE ADC RESOLUTION AND DYNAMIC RANGE REQUIREMENTS
- HIGH SPEED FLASH CONVERTERS AND ERROR SOURCES
- HIGH SPEED SUBRANGING ADCS AND ERROR SOURCES
- APERTURE DELAY TIME (OR EFFECTIVE APERTURE DELAY TIME)
- ERRORS DUE TO APERTURE JITTER
- TRANSIENT RESPONSE OR SETTLING TIME
- OVERVOLTAGE RECOVERY TIME
- METASTABILITY AND BIT ERROR RATES IN ADCS

SECTION 12

SELECTING THE RIGHT HIGH SPEED ADC

Walt Kester

Determining the ADC requirement in a signal processing system is often a challenge which involves many tradeoffs. We have discussed high speed signal amplification and transmission in detail and explored many of the various options available. When selecting the ADC, you must continue to keep the goal of preserving signal fidelity and dynamic range in mind, while resisting the temptation to overspecify.

Overspecifying the ADC, its antialiasing filter, and other peripheral circuitry is especially dangerous, because it may result in high cost and even unrealizable component requirements. In order to intelligently specify the ADC portion of the system, you must first understand the fundamental concepts of sampling and quantization and their effects on the signal.

KEY ELEMENTS OF A BASEBAND SAMPLED DATA SYSTEM

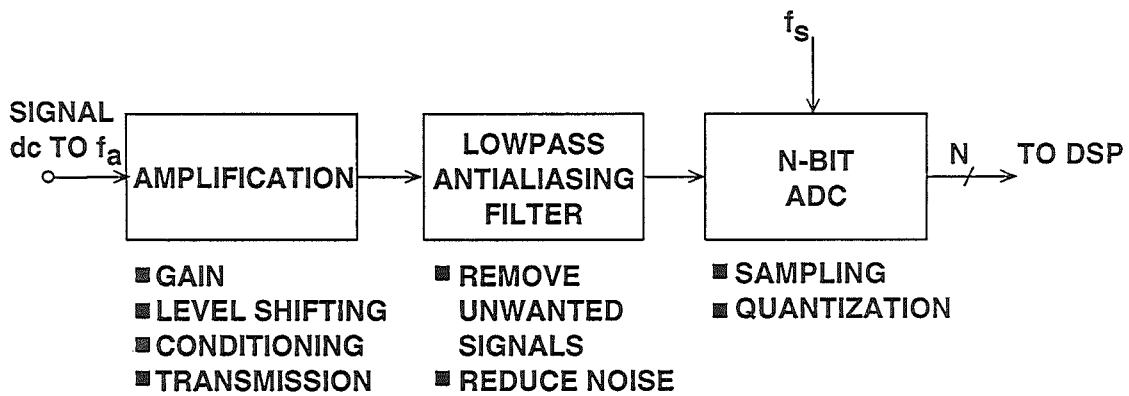

Figure 12.1

In this section, we will consider the traditional problem of sampling and quantizing a baseband signal whose bandwidth lies between dc and an upper frequency of interest, f_a. This is often referred to as *Nyquist, or Sub-Nyquist Sampling*. The topic of *Super-Nyquist* sampling (sometimes called *undersampling*) where the signal of interest falls outside of the Nyquist bandwidth (dc to $f_S/2$) is treated in Section 15.

DISCRETE TIME SAMPLING OF ANALOG SIGNALS

The concept of discrete time and amplitude sampling of an analog signal is shown in Figure 12.2. The continuous analog data must be sampled at discrete intervals, t_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 Shannon's Information Theorem and Nyquist's Criteria given in Figure 12.3.

DISCRETE SAMPLING OF AN ANALOG SIGNAL

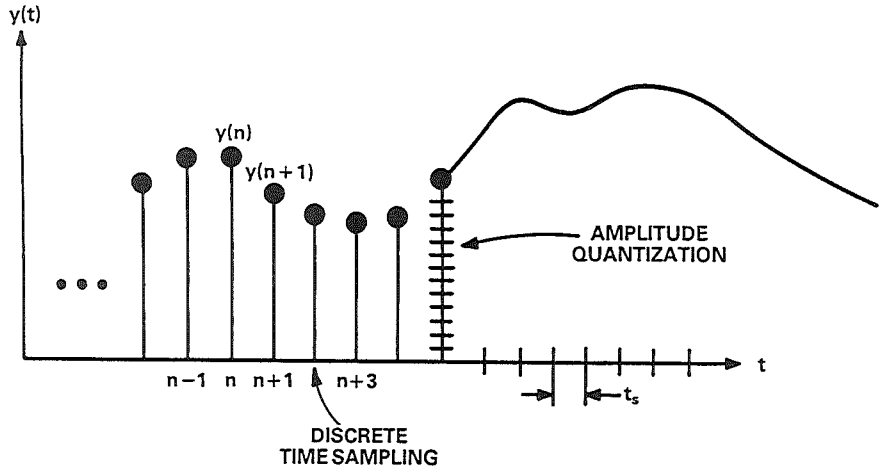

Figure 12.2

SHANNON'S INFORMATION THEOREM AND NYQUIST'S CRITERIA

- Shannon:

  An Analog Signal with a *Bandwidth* of f_a Must be Sampled at a Rate $f_S > 2f_a$ in Order to Avoid the Loss of Information.

- Nyquist:

  If $f_S < 2f_a$, then a Phenomena Called *Aliasing* Will Occur

Figure 12.3

In order to understand the implications of *aliasing* in both the time and frequency domain, first consider the four cases of a time domain representation of a sampled sinewave signal shown in Figure 12.4. In the Case 1, it is clear that an adequate number of samples have been taken to preserve the information about the sinewave. In Case 2 of the figure, only four samples per cycle are taken; still an adequate number to preserve the information. Case 3 represents the ambiguous limiting condition where $f_S = 2f_a$. If the relationship between the sampling points and the sinewave were such that the sinewave was being sampled at precisely the zero crossings (rather than at the peaks, as shown in the illustration), then all information regarding the sinewave would be lost. Case 4 of Figure 12.4 represents the situation where $f_S < 2f_a$, and the information obtained from the samples indicates a sinewave having a frequency which is lower than $f_S/2$, i.e. the out-of -band signal is *aliased* into the Nyquist bandwidth between dc and $f_S/2$. As the sampling rate is further decreased, and the analog input frequency f_a approaches the sampling frequency f_S, the aliased signal approaches dc in the frequency spectrum.

TIME DOMAIN EFFECTS OF ALIASING

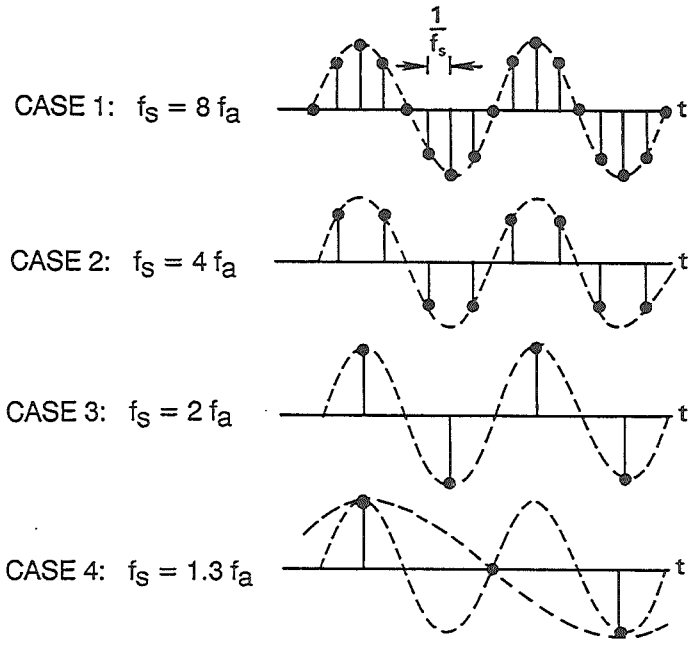

Figure 12.4

The corresponding frequency domain representation of the above scenario is shown in Figure 12.5. Note that sampling the analog signal f_a at a sampling rate f_S actually produces two alias frequency components, one at f_S+f_a, and the other at f_S-f_a. The upper alias, f_S+f_a, seldom presents a problem, since it lies outside the Nyquist bandwidth. It is the lower alias component, f_S-f_a, which causes problems when the input signal exceeds the Nyquist bandwidth, $f_S/2$.

From Figure 12.5, we make the important observation that *regardless* of where the analog signal being sampled happens to lie in the frequency spectrum, the effects of sampling will cause either the actual signal or an aliased component to fall within the Nyquist bandwidth between dc and $f_S/2$. Therefore, any signals which fall outside the bandwidth of interest, whether they be spurious tones or random noise, must be adequately filtered *before* sampling. If unfiltered, the sampling process will alias them back within the Nyquist bandwidth where they will corrupt the signals of interest.

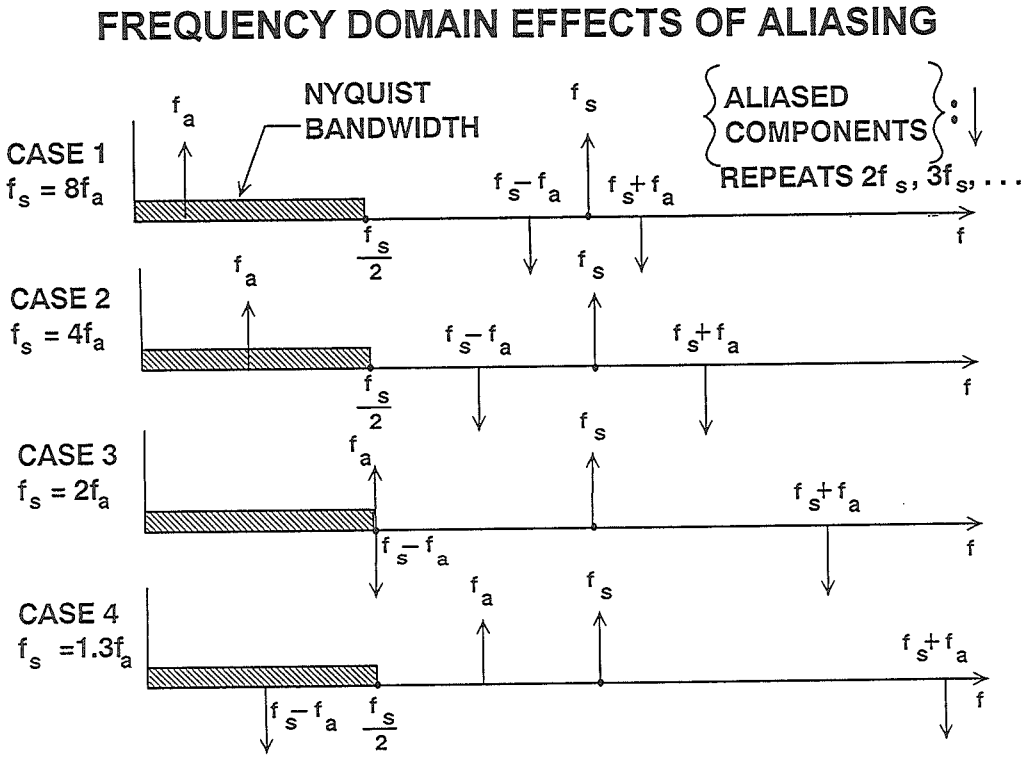

Figure 12.5

PROPERLY SPECIFYING THE ADC SAMPLING RATE AND THE ANTIALIASING FILTER

Properly specifying the ADC sampling rate basically involves trading off higher ADC sampling rates against increased antialiasing filter complexity. The first step is to know the characteristics of the signal being processed. 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. We have now reached the first decision point, since there is no such thing as a perfect analog lowpass filter.

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 12.6.

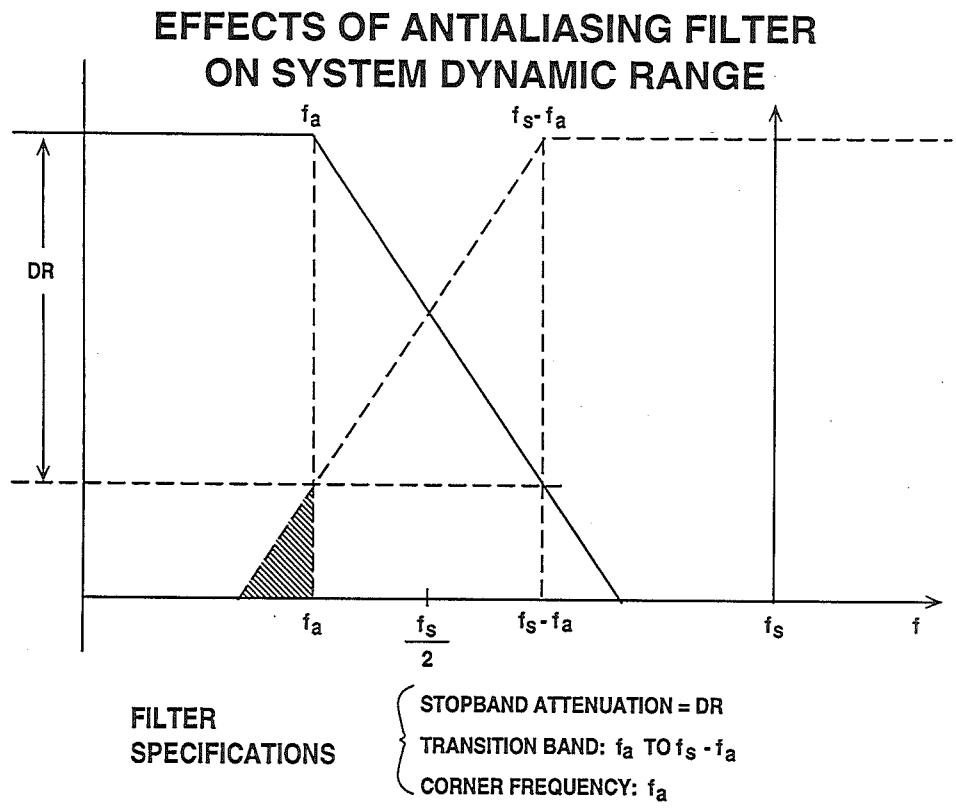

Figure 12.6

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 12.6, 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. We choose the required system dynamic range based on our requirement for signal fidelity.

Filters become more complex as the transition band becomes sharper, all other things being equal. For instance, a Butterworth filter design 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.

Other filter designs are generally more suited to high speed applications where the requirement for a sharp transition band is combined with requirements on in-band flatness and linear phase response. Elliptic filters are popular choices for high speed antialiasing filters.

There are a number of companies which specialize in designing custom analog filters. As an example, the normalized response of the TTE, Inc., LE1182 11-pole elliptic antialiasing filter is shown in Figure 12.7. Notice that this filter is specified to achieve at least 80dB attenuation between f_c and $1.2f_c$ (Reference 1). The corresponding passband ripple, return loss, delay, and phase response are also shown in Figure 12.7. This custom filter is available in corner frequencies up to 100MHz and in a choice of PC board, BNC, or SMA compatible packages.

CHARACTERISTICS OF TTE, INC., L31182-SERIES 11-POLE ELLIPITICAL FILTER

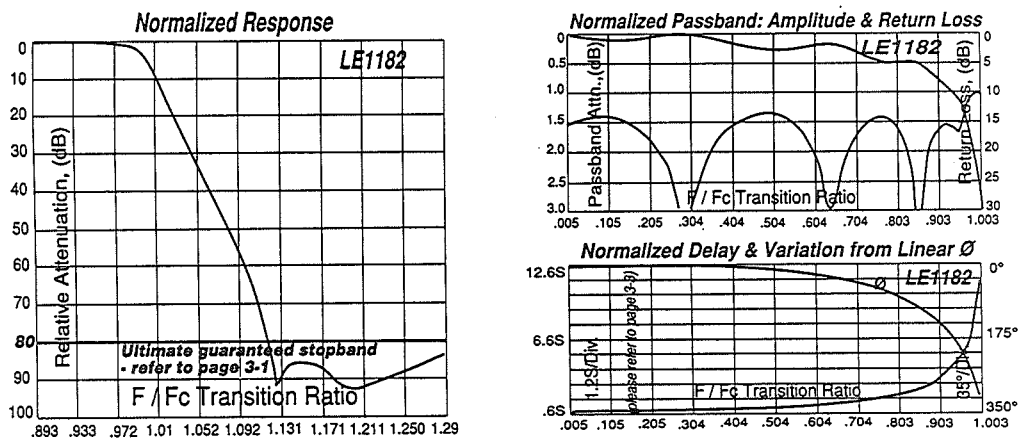

Reprinted with Permission of
TTE, Inc., 2251 Barry Ave., Los Angeles, CA 90064

12

Figure 12.7

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. Remember that high speed ADCs are not oversampling as in the case of Sigma-Delta converters.

You can begin the above design process 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 and a faster ADC.

You can also relax the antialiasing filter requirements somewhat if you know that you are never going to get a fullscale signal at the stopband frequency $f_s - f_a$. In many applications it is very rare that fullscale signals will occur at this frequency. If you know that your 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.

PROPERLY SPECIFYING THE ADC RESOLUTION AND DYNAMIC RANGE REQUIREMENTS

So far, we have discussed only the effects of ADC sampling and aliasing on the system dynamic range. The effects of dividing the signal amplitude into a finite number of discrete quantization levels must also be considered.

Figure 12.8 shows a table of relative bit sizes for various resolution ADCs. The fullscale input range is chosen to be approximately 2V which is popular for high speed ADCs. The bit size (or LSB weight, q) is determined by dividing the fullscale range of the converter by the number of possible quantization levels. Hence, a 10bit ADC having 1024 discrete levels has an LSB weight of 2.048V/1024, or 2mV.

BIT SIZES, THEORETICAL QUANTIZATION NOISE, AND SNR FOR 2.048V FULLSCALE CONVERTERS

Resolution (N Bits)	1 LSB = q	% FS	ppm FS	dB FS (6N)	RMS Quantization Noise, q/√12	Theoretical Fullscale SNR (dB)
6	32mV	1.56	15625	36	9.2mV	37.9
8	8mV	0.39	3906	48	2.3mV	50.0
10	2mV	0.098	977	60	580µV	62.0
12	500µV	0.024	244	72	144µV	74.0
14	125µV	0.0061	61	84	36µV	86.0
16	31µV	0.0015	15	96	13µV	98.1

Figure 12.8

The selection process for determining the ADC resolution should begin by determining the ratio between the largest signal (fullscale) and the smallest signal you wish the ADC to detect. Convert this ratio to dB, and divide by 6. This is your *minimum* ADC resolution requirement for dc signals. You will actually need more resolution to account for extra signal headroom, since ADCs act as hard limiters at both ends of their range. Remember that this computation is for dc or low frequency signals and that the ADC performance will degrade somewhat as the input signal slewrate increases. What will actually occur is that the final ADC resolution will be dictated by dynamic performance at high frequencies. This will lead to the selection of an ADC which probably has more resolution at dc than is actually required.

Also shown in the table of Figure 12.8 is the theoretical rms quantization noise produced by a perfect N-bit ADC. It has been demonstrated that this noise can usually be treated as random noise which is uncorrelated with the input signal. The quantization noise is spread uniformly over the entire Nyquist bandwidth dc to $f_S/2$. The theoretical fullscale rms sinewave signal-to-noise ratio (SNR) may then be calculated using the well known formula, SNR = 6.02N + 1.76dB.

In actual practice, sampling ADCs are evaluated for their dynamic performance by first applying a spectrally pure sinewave input and then performing an FFT on the ADC output data as shown in Figures 12.9 and 12.10. The FFT output can be used to calculate harmonic distortion, THD, and SNR. The actual SNR is then compared to the theoretical SNR. The measured SNR may be substituted in the SNR formula, and the equation solved for N. The resulting value for N is called the *effective number of bits, or ENOBs.*

12

ADC DYNAMIC TESTING

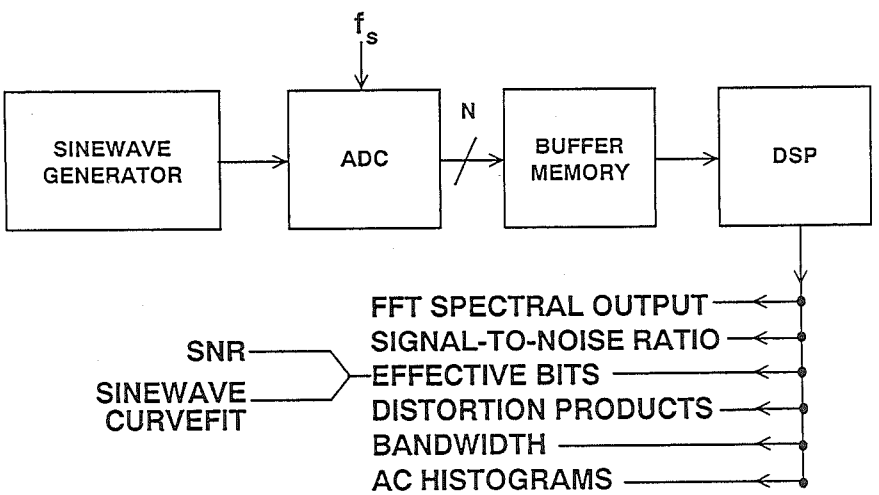

Figure 12.9

4096 POINT FFT OUTPUTS FOR AD9022 12-BIT, 20MSPS ADC

F_S = 20 MSPS

F_{in} = 0.54 MHz

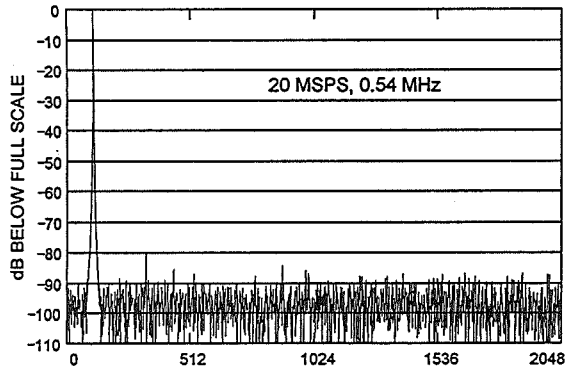

F_{in} = 9.3 MHz

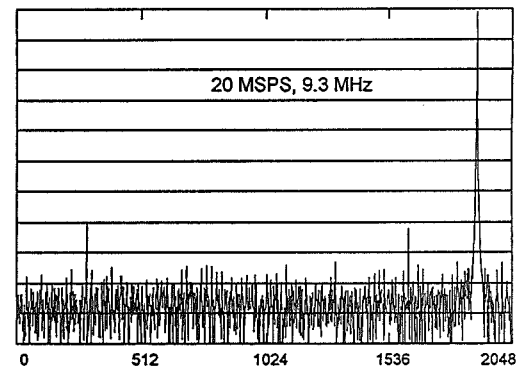

Figure 12.10

QUANTIZATION THEORY BASICS

- RMS Quantization Noise in Nyquist Bandwidth, $f_S/2$:
  $q/\sqrt{12}$, q = LSB Weight
- Fullscale Sinewave RMS Signal to RMS Noise Ratio in Nyquist Bandwidth:

  SNR = 6.02N + 1.76dB
- Effective Number of Bits (ENOB):

$$ENOB = \frac{SNR_{ACTUAL} - 1.76dB}{6.02}$$

Figure 12.11

The SNR and effective bits measurement is made for both low, intermediate, and high frequency input signals. Practically all ADCs exhibit some degradation in SNR and ENOBs at higher input signal frequencies due to various sources of ac nonlinearities. Figure 12.12 shows the S/(N + D) for the AD9022 12 bit, 20MSPS monolithic sampling ADC.

S/(N+D) AND EFFECTIVE BIT PERFORMANCE OF AD9022 12-BIT, 20MSPS SAMPLING ADC

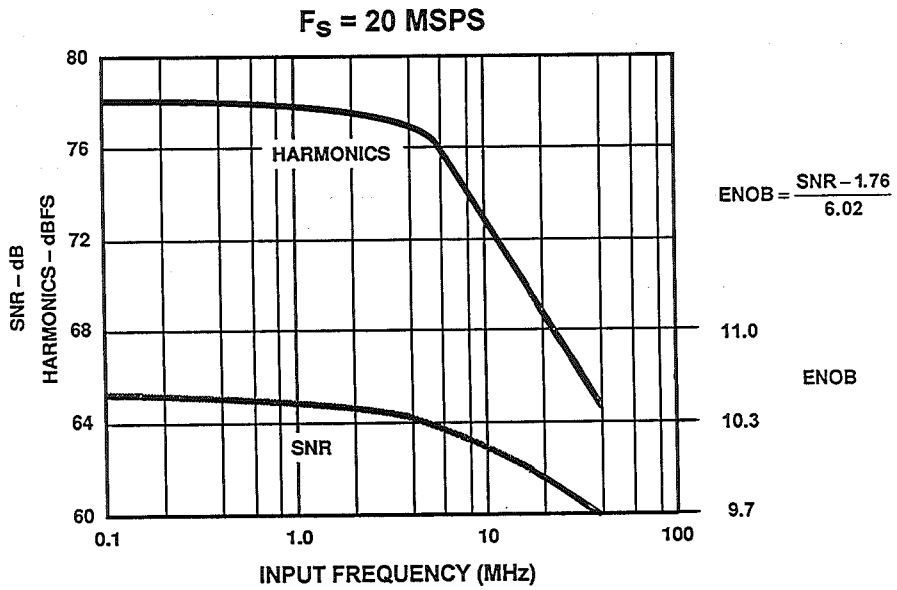

Figure 12.12

The AD9022 is representative of a class of high-performance *sampling* ADCs. Sampling ADCs have the sample-and-hold function on-chip and are characterized in terms of both dc and ac specifications.

MODERN HIGH SPEED ADCs

- Most are *Sampling* ADCs Containing on-chip SHA Function as Opposed to *Encoders*, which have no SHA
- Interface Between SHA and ADC Handled on-chip
- Complete DC and AC Specifications Usually Provided: SNR, THD, SFDR, ENOB, Bandwidth, etc.
- Input Bandwidth is Usually Much Greater than $f_S/2$

Figure 12.13

ADC FULL-POWER BANDWIDTH

- The Frequency at Which the Amplitude of the *Fundamental* Component in the FFT Output is Down 3dB
- FPBW Usually > f_S (Except for $\Sigma\Delta$ ADCs)
- Must Examine ENOB and THD at FPBW Frequency - Usually Reduced
- Example: AD9022 FPBW = 100MHz, 9.7 ENOB @ 40MHz INPUT
- Use FPBW or Small Signal BW (if Greater Than FPBW) for Noise Calculations

Figure 12.14

The input bandwidth of the ADC should be considerably greater than the highest-frequency baseband signal of interest. The full-power bandwidth (FPBW) of an ADC is that input frequency at which the amplitude of the reconstructed FFT *fundamental* is reduced by 3dB for a fullscale input. Generally, though, there is considerable loss of *resolution* at frequencies well below this. Full-power bandwidth must be examined in conjunction with SNR, ENOB, and THD in order to determine the actual dynamic performance of the ADC at the FPBW frequency. The small signal ADC bandwidth is approximately equal to the FPBW if there is no slewrate limiting.

The next step in the process of determining the ADC resolution is to determine the effective bit (ENOB) requirement or SNR at the highest input frequency of interest, f_a. Most modern sampling ADCs have these specifications and also curves similar to that of Figure 12.12. Remember that the S/(N+D) calculation includes all distortion products as well as those due to quantization.

The *peak spurious* or peak harmonic component is the largest spectral component excluding the input signal and dc (measured with FFT techniques). This value is expressed in dB relative to the rms value of the input signal. This specification is also referred to as *spurious free dynamic range*, or SFDR. In applications such as digital spectral analysis using FFT techniques, harmonic distortion, THD, or spurious free dynamic range (SFDR) may be of greater concern than the actual broadband rms noise level.

The FFT takes a discrete number of time samples, M, and converts them into M/2 discrete spectral components. The spacing between the spectral lines is $\Delta f = f_S/M$. If an FFT is performed on broadband quantization noise which has a bandwidth of $f_S/2$, the average value of the noise contained in each FFT frequency cell is $10\log_{10}(M/2)$ dB less than the rms value of the quantization noise. This is illustrated in Figure 12.15. This is equivalent to sweeping an analog spectrum analyzer from dc to $f_S/2$ with the bandwidth set to Δf. The average value of the noise components in each frequency bin can be reduced 3dB by doubling the record length M. Using deeper FFTs, averaging the results of a number of FFTs, or other filtering techniques may also be used to reduce the rms noise floor and allow greater dynamic range. (This topic will be addressed in greater detail in the section on wide dynamic range applications).

THE EFFECTIVE NOISE FLOOR OF AN M-POINT FFT (MEASUREMENT BANDWIDTH = f_S/M) IS MUCH LESS THAN THE RMS VALUE OF THE QUANTIZATION NOISE (MEASUREMENT BANDWIDTH = $f_S/2$)

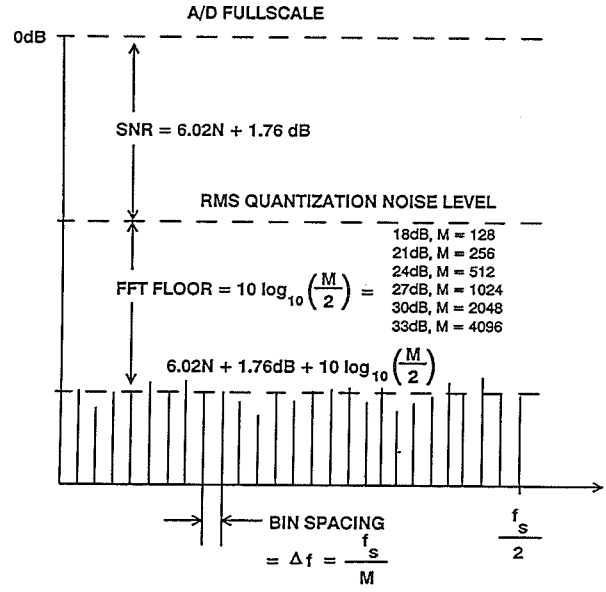

Figure 12.15

SPURIOUS FREE DYNAMIC RANGE AS A FUNCTION OF INPUT SIGNAL LEVEL FOR THE AD9014 14-BIT, 10MSPS ADC

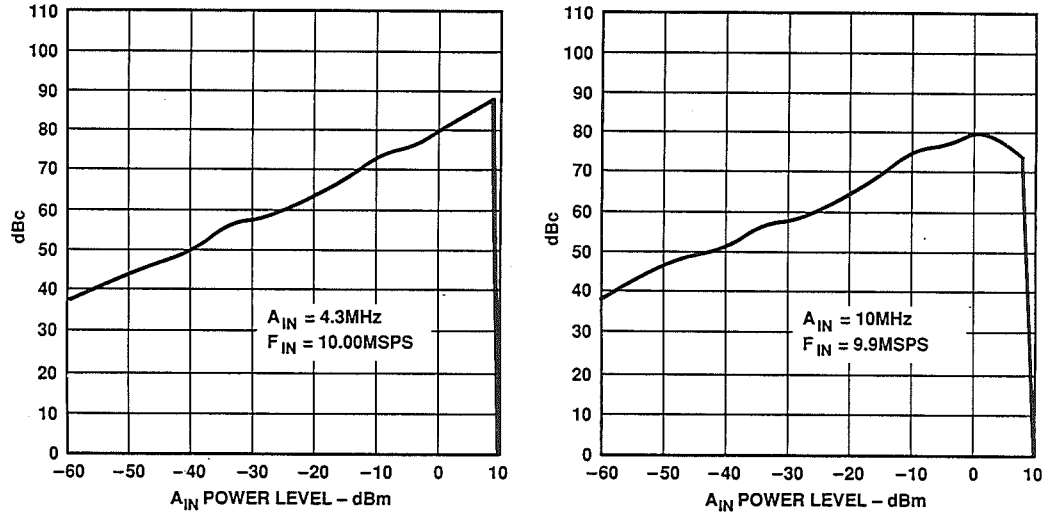

Figure 12.16

The SFDR of an ADC is generally a function of both input frequency and amplitude. Figure 12.16 shows the SFDR versus the input signal level for the AD9014 14 bit 10MSPS ADC. The data is given for two input frequencies: 4.3MHz and 9.9MHz. Notice that the SFDR at 4.3MHz reaches its maximum value at fullscale. For the 9.9MHz input, however, the point of maximum SFDR occurs several dB below fullscale.

As discussed above, SFDR should not be confused with SNR. SNR depends more on the rms value of the quantization noise and is therefore a function of the number of ADC bits. SFDR, on the other hand, depends more on the linearity of the ADC and is relatively independent of the number of actual bits. This is dramatically illustrated in Figure 12.17, where SNR and SFDR is shown for the AD9014 ADC. The top curve in the figure shows the SFDR of the AD9014 utilizing 14, 12, and 10 bits of the ADC. The bottom three curves show the SNR of the AD9014 operating with 14, 12, and 10 bits. The level of the internally generated spurs will not rise as bits are omitted, but the broadband rms noise floor rises for each bit that is dropped.

AD9014 SFDR AND SNR VERSUS FREQUENCY FOR 14, 12, AND 10 ACTIVE BITS

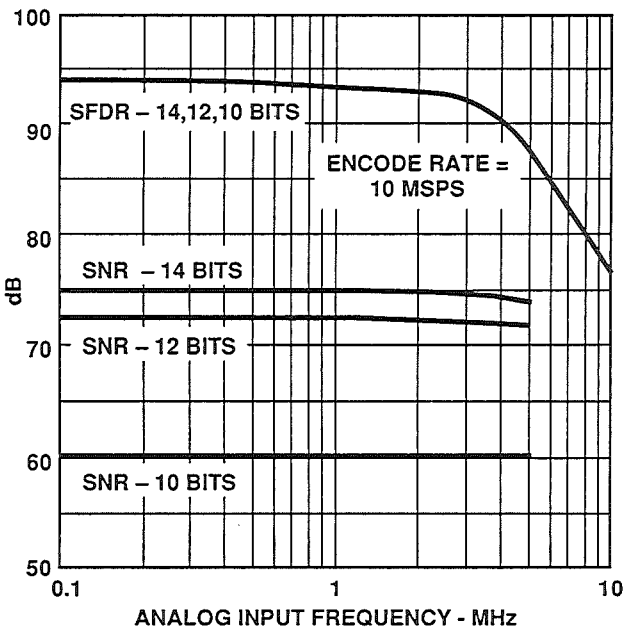

Figure 12.17

The distortion produced by an ADC cannot be analyzed in terms of second and third-order intercepts as in the case of an amplifier. This is because there are two components of distortion in a high performance ADC. One component is due to the non-linearity associated with the analog front end amplifier and the sample-and-hold. This non-linearity has the familiar "bow" or "s"-shaped curve shown in Figure 12.18. The distortion associated with this type of non-linearity is sometimes referred to as *soft* distortion and produces low-order distortion products. This component of distortion behaves in the traditional manner and is a function of signal level. In a practical ADC, however, the soft distortion is usually much less than the other component of distortion in an ADC which is due to the nonlinearity of the encoder transfer function itself. This function is more likely to have discrete points of discontinuity across the signal range as shown in Figure 12.18.

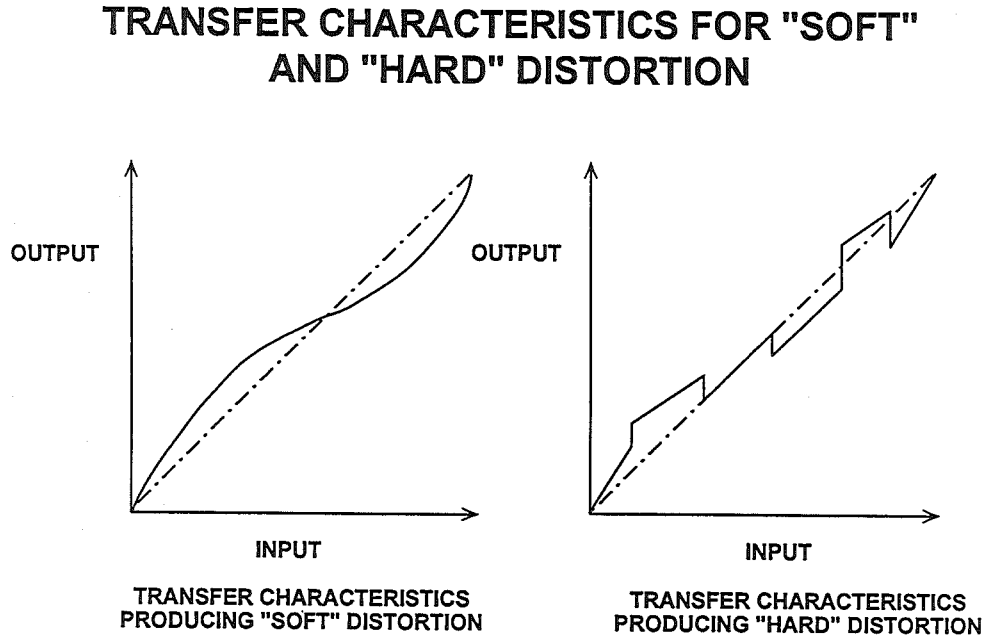

Figure 12.18

The actual location of the points of discontinuity depends on the particular ADC architecture, but nevertheless such discontinuities occur in practically all high speed ADCs. Non-linearity of this type produces high-order distortion products which are relatively unpredictable with respect to input signal level. For lower-amplitude signals, this constant level *hard* distortion causes the SFDR of the ADC to decrease as input amplitude decreases. The soft distortion in a well-designed ADC generally only comes into play for high frequency large-amplitude input signals where it may rise above the hard distortion floor. This can be observed in Figure 12.16, where the 4.3MHz input data indicates a relatively constant hard distortion floor as the signal amplitude is increased to fullscale. The 9.9MHz data, however, indicates an increase in soft distortion as the input signal approaches fullscale.

HIGH SPEED FLASH CONVERTERS AND ERROR SOURCES

The ADC architecture which offers the highest possible sampling rate is the flash converter. Recent advances in VLSI process technology and circuit design techniques have made flash ADCs with up to 10 bits of resolution practical. As well as offering high sampling rates for digitizing high frequency signals, flash converters are often used as building blocks for higher resolution ADCs.

A block diagram of a typical flash converter is shown in Figure 12.19. The analog input signal to be digitized is applied simultaneously to $2^N - 1$ latched comparators, where N is the number of bits. The reference voltage input for each comparator is derived from a resistive voltage divider string. The reference voltage for each comparator is one LSB higher than the one immediately below it.

When an analog signal is present at the input of the comparator bank, all comparators which have a reference voltage below the level of the input signal will assume a logic "1" output. The comparators which have their reference voltage above the input signal will assume a logic "0" output. The result is often referred to as a *thermometer code* (shown in Figure 12.20), and is applied to the decoding logic stage. This decoding can be accomplished in a variety of ways (such as a simple priority encoder) and ultimately results in the digital output word.

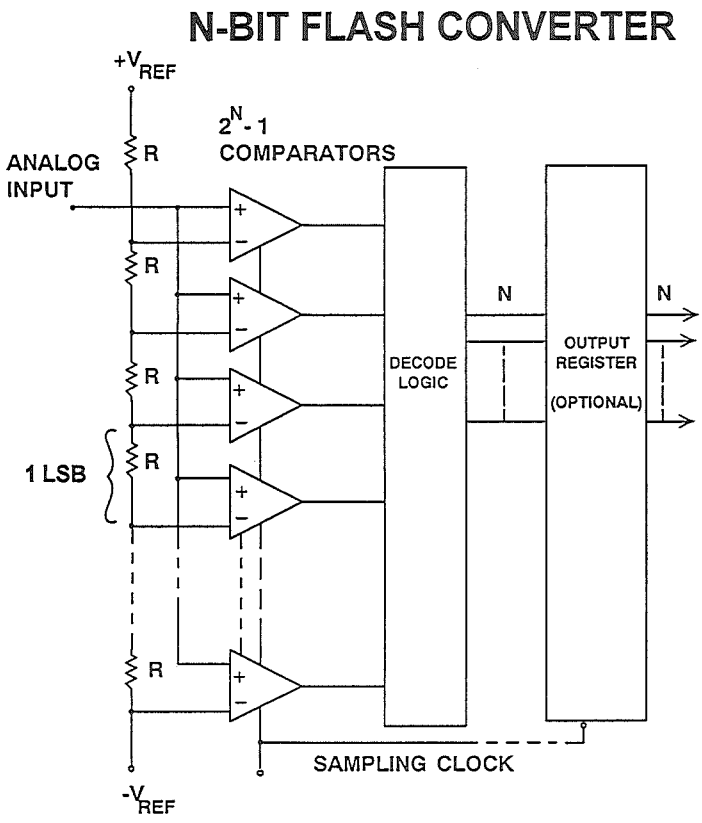

Figure 12.19

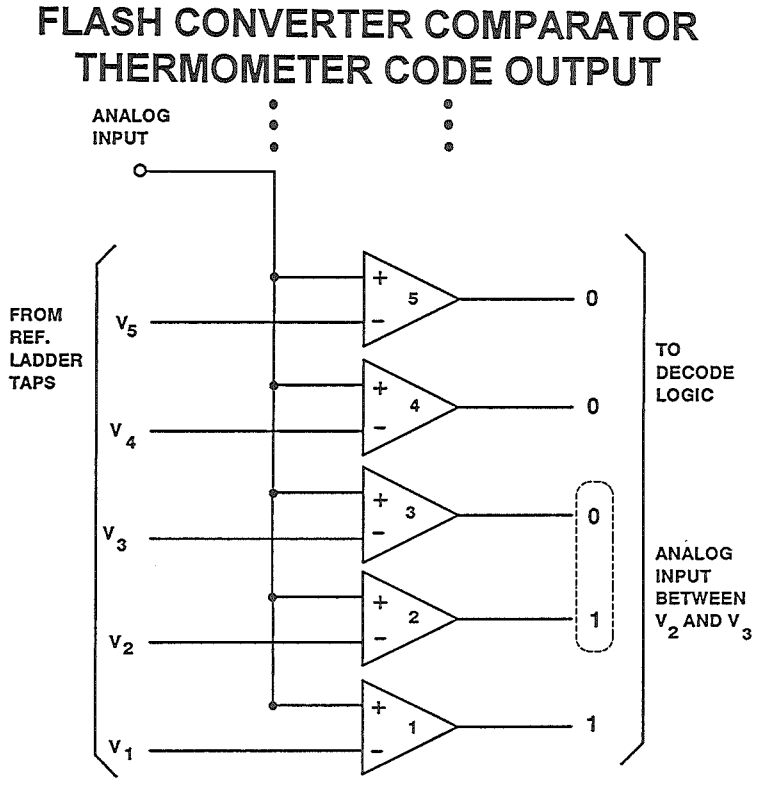

Figure 12.20

The primary source of dc nonlinearity in a flash converter is the comparator input offset voltage variation. In a good flash converter this nonlinearity should be less than ±½LSB. The reference voltage resistor ladder string rarely contributes significant error as it has inherently good ratio matching.

Dynamic errors in a flash converter come from several sources, the chief source being the ac mismatch between comparators. In an ideal flash converter with perfectly matched (ac and dc) comparators, each comparator latches simultaneously when the sampling clock is asserted. This effectively provides an internal SHA function, and the overall ADC nonlinearity is independent of the input signal slewrate. In actual flash converters, however, the delay and bandwidth matching of the comparators is not perfect. Although this mismatch does not affect dc performance, it will introduce slewrate-dependent nonlinearities. The effect is to reduce the SNR and ENOB performance of the flash converter at high frequencies. Well designed flash converters, however, are capable of maintaining good performance at fairly high input frequencies without the necessity for an external SHA. The SNR and ENOB performance of the AD9058 dual 8 bit 50MSPS flash converter is shown in Figure 12.21. Notice that 40dB SNR (6.4 ENOBs) is achieved at an input frequency of 20MHz.

SNR, EFFECTIVE BIT, AND HARMONIC DISTORTION OF THE AD9058 DUAL 8-BIT, 50MSPS FLASH CONVERTER

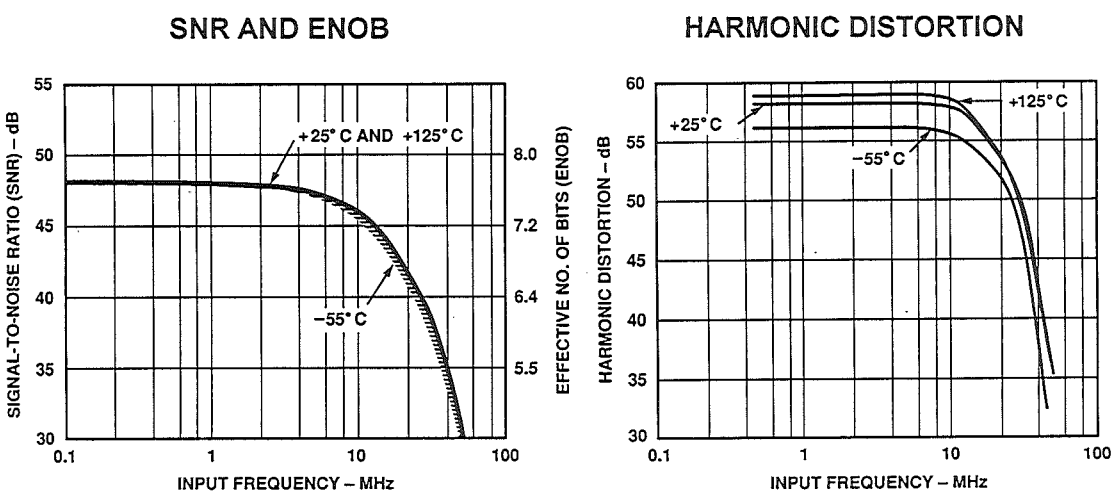

Figure 12.21

Another source of dynamic errors in a flash converter is the signal-dependent input capacitance. The input stage of each comparator in a flash converter is a long tail, common emitter differential pair. The analog input is applied to one base, and the reference ladder voltage to the other base. As the input signal changes, the corresponding input capacitance will also change. The problem is compounded because of the large number of comparators in parallel.

The input circuit of the AD9028/ AD9038 8 bit, 300MSPS flash converter is modeled in Figure 12.22. The signal-dependent capacitance is modeled as a reverse-biased diode. The total input capacitance as a function of signal level is also shown in Figure 12.22.

Wideband, low-distortion current feedback amplifiers such as the AD9617 are ideal for driving this type of flash converter. However, a series resistor of approximately 50Ω is required to isolate the amplifier from the converter input capacitance in order to prevent peaking and to maintain stability. Because of the series resistor and the signal-dependent capacitance, harmonic distortion will result as shown in Figure 12.23. The series isolation resistor should therefore be no larger than is necessary to maintain op amp stability. Large resistor values will increase the distortion and limit the input bandwidth. Datasheets for flash converters should provide you with recommended drive circuits for optimum performance.

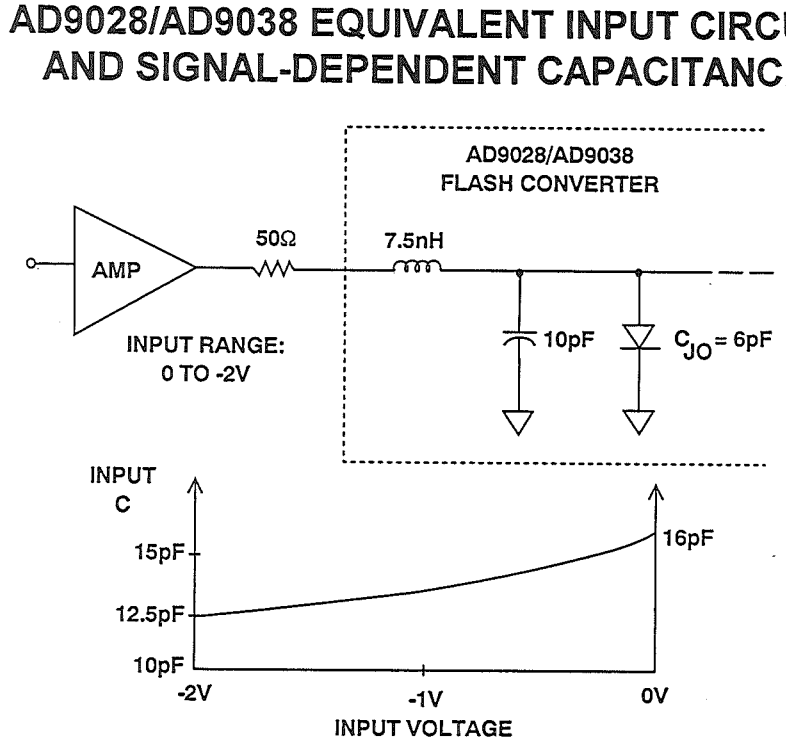

Figure 12.22

SIMULATED THD DUE TO SIGNAL-DEPENDENT INPUT CAPACITANCE

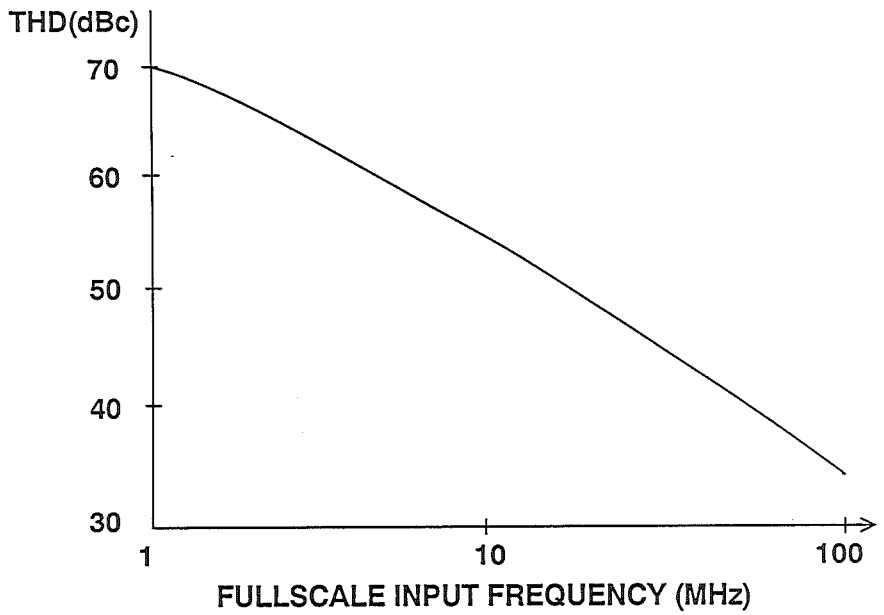

Figure 12.23

FLASH CONVERTER PRIMARY ERROR SOURCES

- Static Errors:
  - Comparator Input Offset Voltage Mismatch
- Dynamic Errors:
  - Comparator Bandwidth and Delay Mismatch
  - Signal-Dependent Input Capacitance
  - Layout Parasitics

Figure 12.24

HIGH SPEED SUBRANGING ADCS AND ERROR SOURCES

A block diagram of an 8-bit subranging ADC based upon two 4-bit flash converters is shown in Figure 12.25. 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 converter by the summing amplifier. 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.

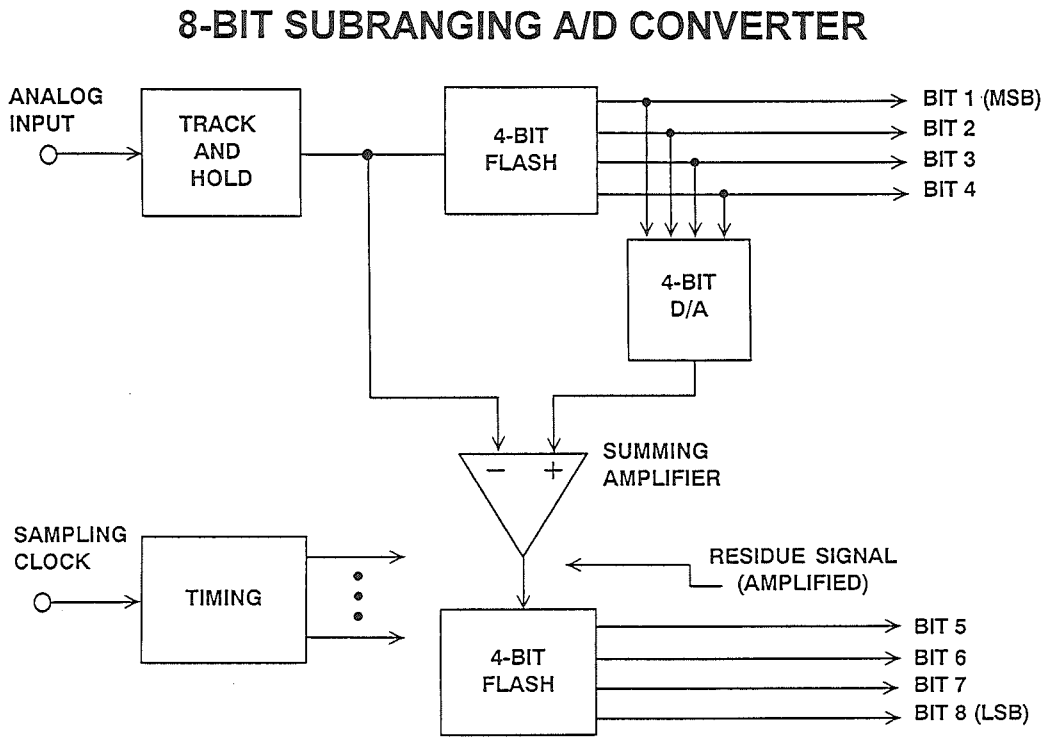

Figure 12.25

Modern subranging ADCs use a technique called *digital correction* to eliminate problems associated with the architecture of Figure 12.25. A block diagram of a 12-bit digitally corrected subranging (DCS) ADC is shown in Figure 12.26. Note that a 5-bit and an 8-bit flash converter have been used to achieve an overall 12-bit output. If there were no errors, the 5-bit "residue" signal applied to the 8-bit flash converter by the summing amplifier would never exceed one-half of the range of the 8-bit flash. The extra range in the second flash converter is used in conjunction with the error correction logic (usually just an adder) to correct the output data for most of the errors inherent in the traditional uncorrected subranging converter architecture.

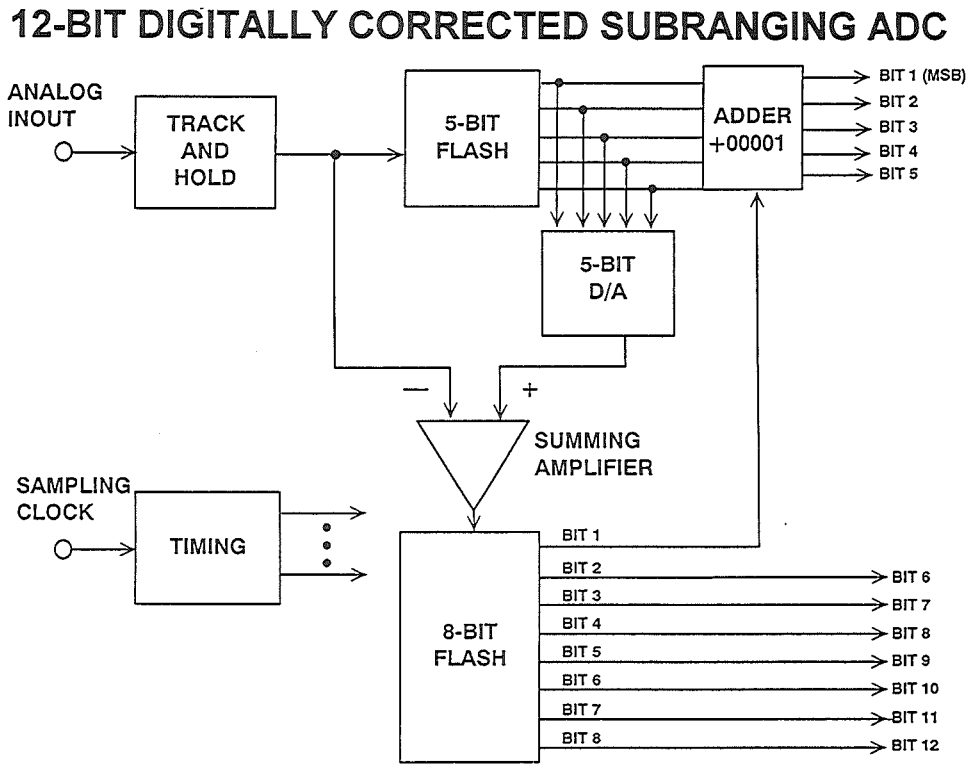

Figure 12.26

A block diagram of the AD9040 10 bit, 40MSPS ADC is shown in Figure 12.27. This ADC makes use of two stages. The held analog value of the first track-and-hold is applied to a 5 bit flash converter and a pair of internal T/Hs (shown on the diagram as a single unit). The T/Hs pipeline the analog signal to the amplifier array through a residue ladder and switching circuit while the 5 bit flash converter resolves the most significant bits (MSBs) of the held analog voltage. When the 5 bit flash converter has completed its cycle, its output activates 1-of-32 ladder switches; these, in turn, cause the correct residue signal to be applied to the error amplifier array. The output of the error amplifier is applied to a 6 bit flash converter whose output supplies the five least significant bits (LSBs) of the digital output along with one bit of error correction for the main range converter. Decoding logic combines the bits from the two converters and presents the result as a 10 bit parallel word.

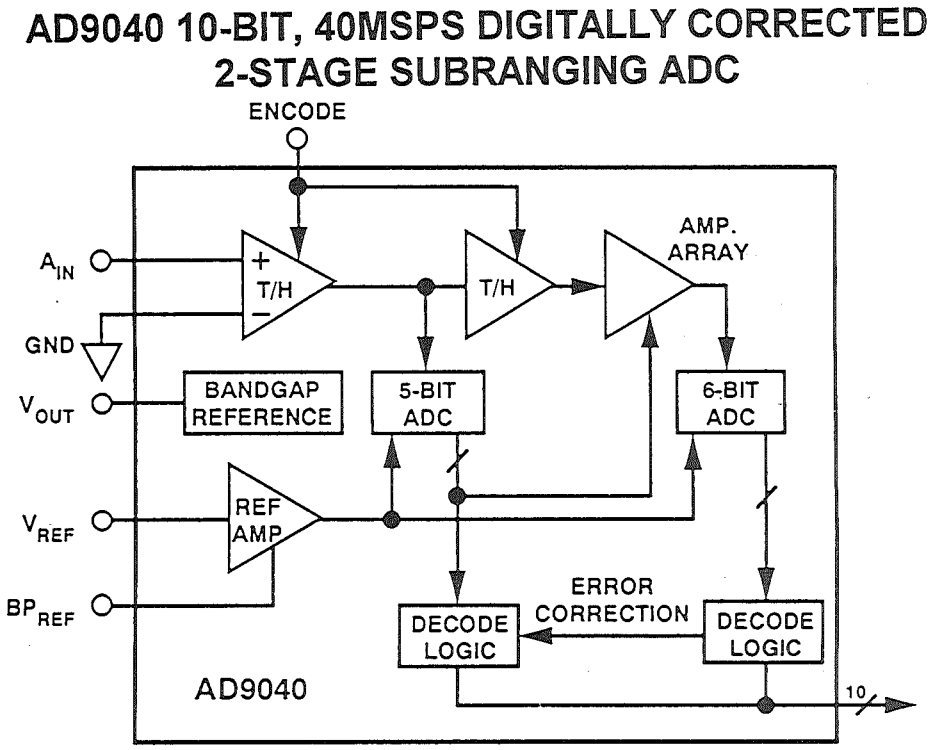

Figure 12.27

The AD9022 12 bit, 20MSPS monolithic ADC operates on the same basic principles described above, except that three internal stages are used to perform the total conversion (see Figure 12.28). The held value of the analog input is first applied to a 5 bit flash converter which performs the first coarse conversion. The first 5 bit flash converter output also drives a 5 bit DAC whose output is subtracted from the held analog signal to form the first residue. The second T/H pipelines the residue and applies it to a gain-of-16 amplifier which drives a second 5 bit flash converter. The second 5 bit flash output is combined with the first flash converter output with 1 bit of error correction. The second 5 bit flash converter also drives a second 5 bit DAC whose output is subtracted from the first residue to form a second residue . The second residue is applied to a gain-of-8 amplifier which drives a 4 bit flash converter. The final 4 bit flash converter output is combined with the outputs of the two 5 bit converters in a logic block which performs the digital error correction and generates the final ADC output word.

AD9022 12-BIT, 20MSPS DIGITALLY CORRECTED 3-STAGE SUBRANGING ADC

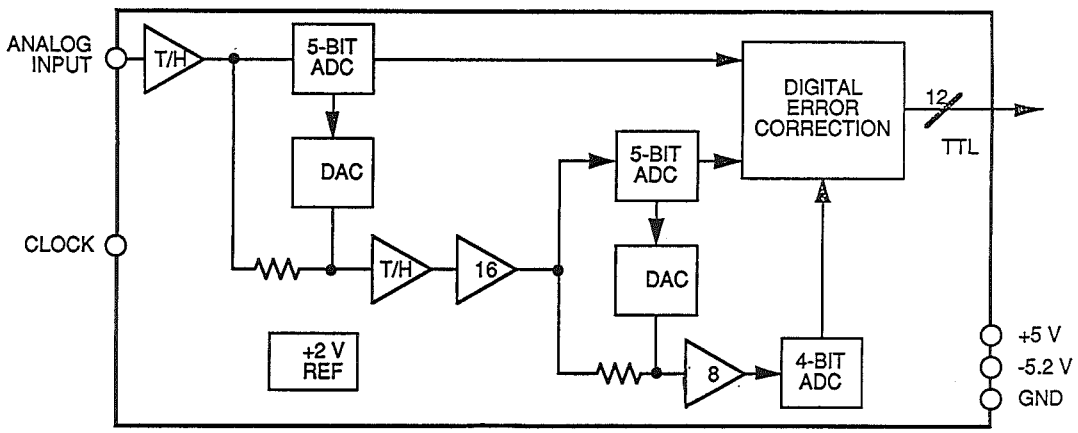

Figure 12.28

At slower sampling rates it is possible to use a single flash converter in multiple passes to perform the conversions as shown in Figure 12.29. The AD679 is a 14 bit 100kSPS ADC which uses this *recursive subranging* approach. The 4 bit flash first produces a 4 bit representation of the analog input. This value is reconstructed by the DAC, and the difference between the DAC output is amplified (to take full advantage of the dynamic range of the 4 bit flash), and the whole cycle then repeats itself. After 5 cycles, the result is presented to the final output. A 1 bit overlap between cycles serves as error correction.

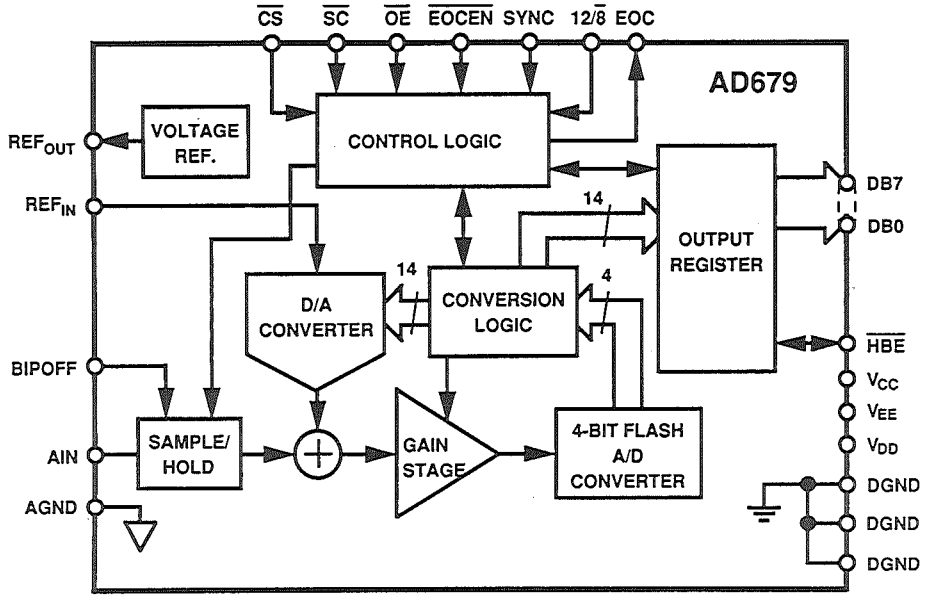

Figure 12.29

The error correction schemes described above are designed to correct for errors made in the early pipelined sub-conversions. Flash converter gain, offset, and linearity errors are corrected as long as the residue signals fall within the range of the next-stage flash conversion. These errors will not affect the linearity of the overall ADC transfer characteristic. Errors made in the final flash conversion, however, do translate directly as errors in the overall transfer function. Also, linearity errors or gain errors either in the DACs or the residue amplifiers will not be corrected and will show up as nonlinearities or non-monotonic behavior in the overall ADC transfer function. If the converter is not properly trimmed or drifts over temperature, errors such as those shown in Figure 12.31 may develop at the various DAC switching points along the ADC transfer function. A properly designed and trimmed ADC, however, should not exhibit errors such as these over the operating temperature range of the device.

SOME ERROR SOURCES IN DIGITALLY CORRECTED SUBRANGING ADCs

- Nonlinearities in Final Flash Conversion
- Internal DAC Linearity or Gain Errors
- Residue Amplifier Gain Errors
- Improper Laser Trimming of Thin Film Resistors
- Internal Timing Errors

Figure 12.30

IMPROPERLY TRIMMED SUBRANGING ADC LINEARITY ERRORS

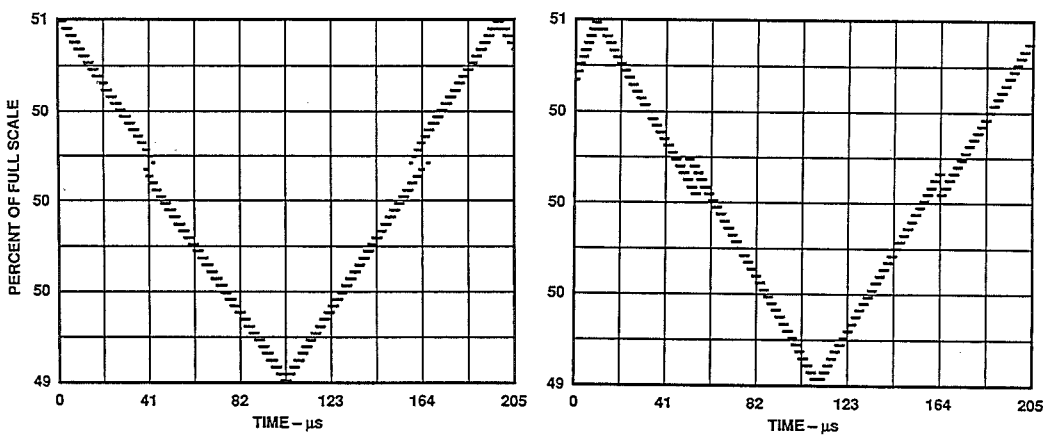

Figure 12.31

APERTURE DELAY TIME (OR EFFECTIVE APERTURE DELAY TIME)

Aperture delay time (sometimes called aperture time) is the amount of time from the leading edge of the sampling clock until the ADC actually takes the sample (see Figure 12.32). This specification is important because it helps the user to know when to apply the sampling clock with respect to the input signal timing. 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 the ADCs are required to track each other when processing dynamic signals.

MEASUREMENT OF EFFECTIVE APERTURE DELAY TIME

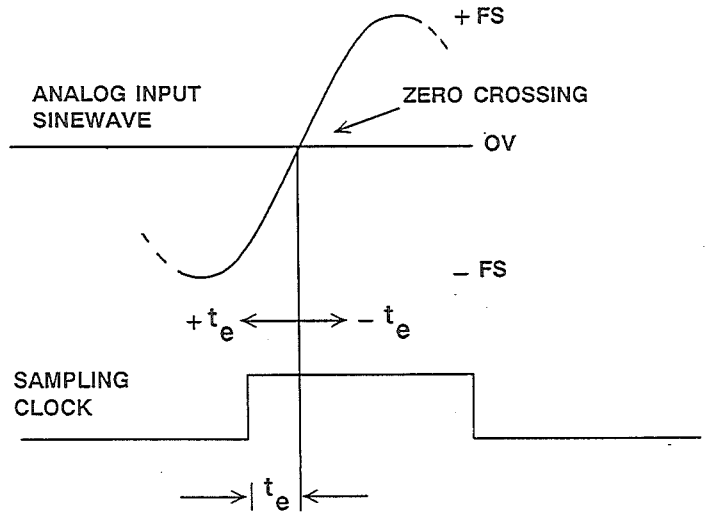

Figure 12.32

ERRORS DUE TO APERTURE JITTER

Aperture jitter is the sample-to-sample variation in the effective point in time at which the actual sample is taken (Figure 12.33). These errors generally emanate from several sources. In practical systems, the sampling clock is often phase-modulated by an external noise source; the source can be wideband random noise, power line noise, or digital noise due to poor layout, bypassing, and grounding techniques. While it is true that a portion of the total aperture jitter may be generated internal to the ADC, this component is rarely the dominant source of SNR degradation at high input frequencies. Nonlinearity in the ac transfer function and increased distortion are usually the major problems.

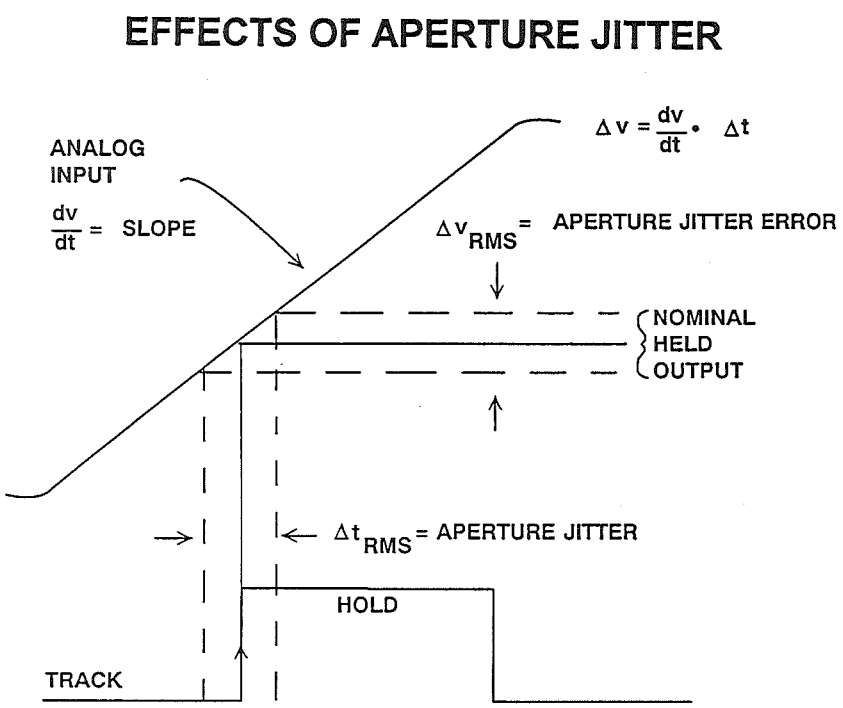

Figure 12.33

The resulting error can be expressed in terms of an rms time jitter. The corresponding rms voltage error caused by rms aperture jitter decreases the overall ADC signal-to-noise ratio. Phase jitter on the input sinewave can produce the same effect as jitter on the sampling clock. The SNR due exclusively to aperture jitter is plotted in Figure 12.34 as a function of fullscale sinewave input frequency for various values of aperture jitter. The equation for SNR due to aperture jitter is derived in Reference 2. Close examination of these curves indicates the importance of maintaining a low-jitter sampling clock if high values of SNR are required.

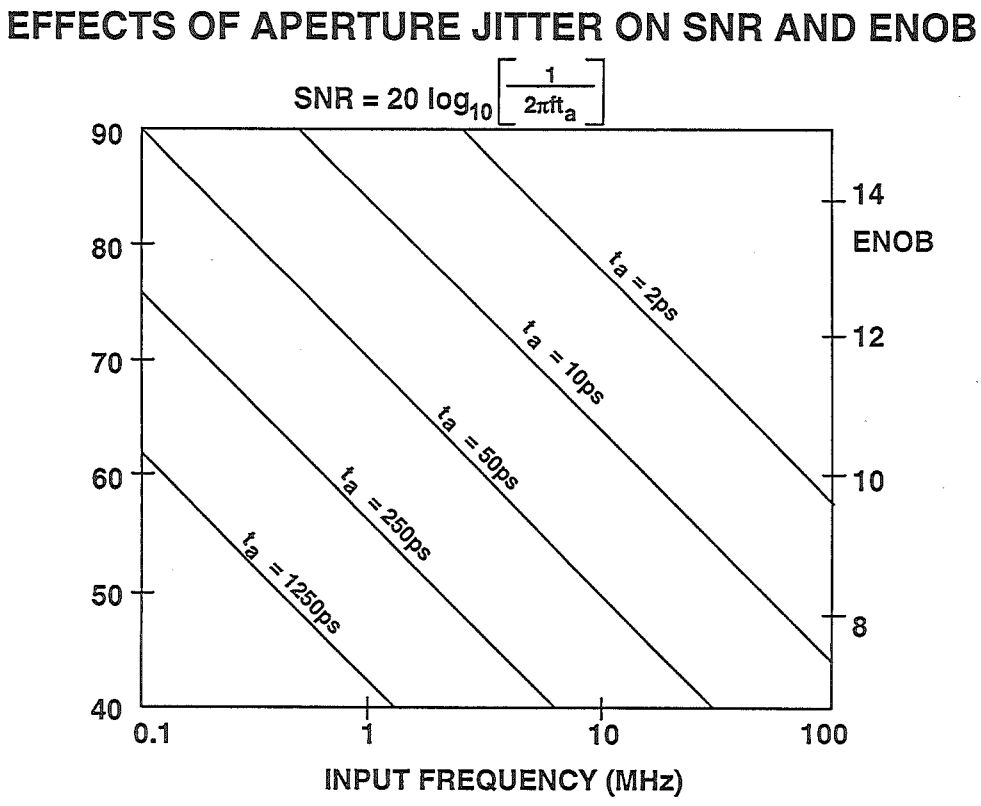

Figure 12.34

TRANSIENT RESPONSE OR SETTLING TIME

The transient response (or settling time) of an ADC is the time required for the ADC to settle to rated accuracy after the application of a fullscale step input. The typical response of the AD872 12 bit, 10MSPS ADC is shown in Figure 12.35. The AD872 can typically settle to 12 bit accuracy from a fullscale step input in less than 40ns.

This specification is critical in applications where the ADC is being driven by an analog multiplexer as shown in Figure 12.36. The multiplexer output can deliver a fullscale sample-to-sample change to the ADC input. If both the multiplexer and the ADC have not both settled to the required accuracy, channel-to-channel crosstalk will result.

AD872 12-BIT, 10MSPS ADC TRANSIENT RESPONSE

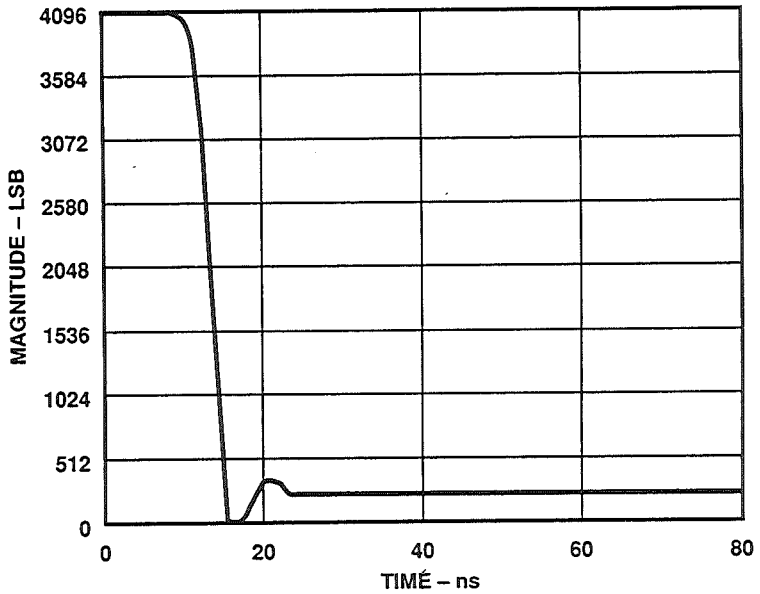

Figure 12.35

DATA ACQUISITION SYSTEM USING AN ANALOG MULTIPLEXER

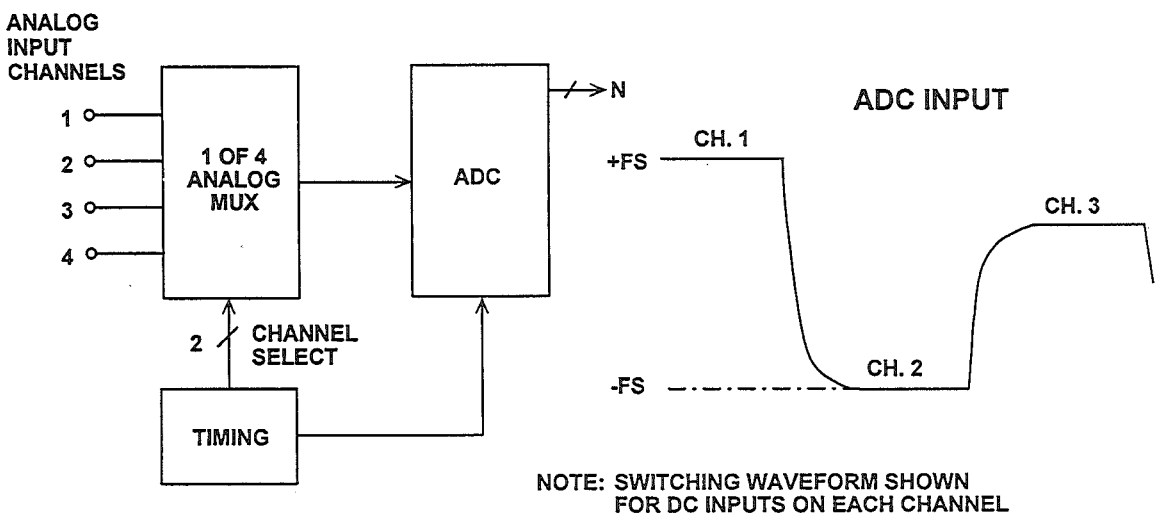

Figure 12.36

OVERVOLTAGE RECOVERY TIME

Overvoltage recovery time is defined as that amount of time required for an ADC to achieve a specified accuracy, measured from the time the overvoltage signal re-enters the converter's range, as shown in Figure 12.37. This specification is usually given for a signal which is 50% outside the ADC's input range. Needless to say, the ADC should act as an ideal limiter for out-of-range signals and should produce either the positive fullscale code or the negative fullscale code during the overvoltage condition. Some converters provide over- and under-range flags to allow gain-adjustment circuits to be activated. Care should always be taken to avoid overvoltage signals which will damage an ADC input.

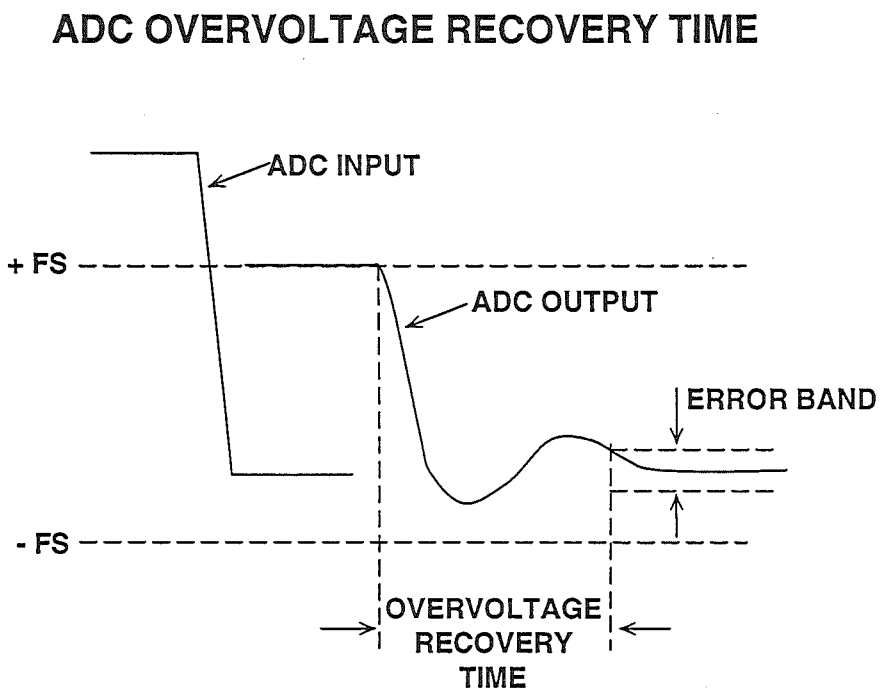

Figure 12.37

METASTABILITY AND BIT ERROR RATES IN ADCS

A primary concern in the design of many digital communications systems using ADCs is the bit error rate (BER). Unfortunately, ADCs contribute to the BER in ways that are not predictable by simple analysis. This section describes the mechanisms within the ADCs that can contribute to the error rate, ways to minimize the problem, and methods for measuring the BER.

Random noise, regardless of the source, creates a finite probability of errors (deviations from the expected output). Before describing the error code sources, however, it is important to define what constitutes an ADC error code. Noise generated prior to, or inside the ADC can be analyzed in the traditional manner. Therefore, an ADC error code is any deviation from the expected output that is not attributable to the equivalent input noise of the ADC. Figure 12.38 illustrates an exaggerated output of a pure sinewave applied to an ADC. Note that the SNR of the ADC creates some uncertainty in the output. These anomalies are not considered error codes, but are simply the result of ordinary noise and quantization. The large errors are more significant and are not expected. These errors are random and so infrequent that an SNR test of the ADC will rarely detect them. These types of errors plagued a few of the early ADCs for video applications, and were given the name *sparkle codes* because of their appearance on a TV screen as small white dots under certain test conditions. The large errors have also been called *rabbits* or *flyers*. In digital communications applications, this type of error increases the overall system bit error rate (BER).

EXAGGERATED OUTPUT OF ADC SHOWING ERROR CODES

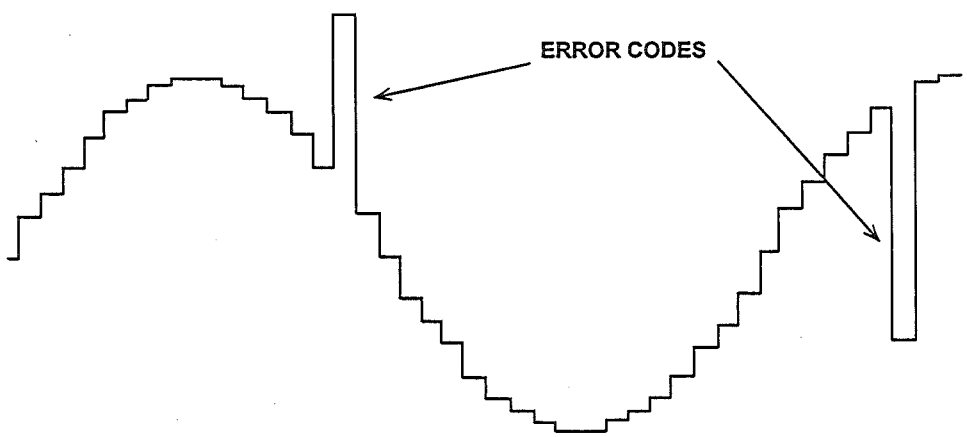

Figure 12.38

In order to understand the causes of the error codes, we will first consider the case of a simple flash converter. The comparators in a flash converter are latched comparators usually arranged in a master-slave configuration. If the input signal is in the center of the threshold of a particular comparator, that comparator will balance, and its output will take a longer period of time to reach a valid logic level after the application of the latch strobe than the outputs of its neighboring comparators which are being overdriven. This phenomenon is known as *metastability* and occurs when a balanced comparator cannot reach a valid logic level in the time allowed for decoding. If simple binary decoding logic is used to decode the thermometer code, a metastable comparator output may result in a large output code error. Consider the case of a simple 3 bit flash converter shown in Figure 12.39. Assume that the input signal is exactly at the threshold of Comparator 4 and random noise is causing the comparator to toggle between a "1" and a "0" output each time a latch strobe is applied. The corresponding binary output should be interpreted as either 011 or 100. If, however, the comparator output is in a metastable state, the simple binary decoding logic shown may produce binary codes 000, 011, 100, or 111. The codes 000 and 111 represent a one-half scale departure from the expected codes.

METASTABLE COMPARATOR OUTPUT STATES MAY PRODUCE ERROR CODES IN FLASH CONVERTERS

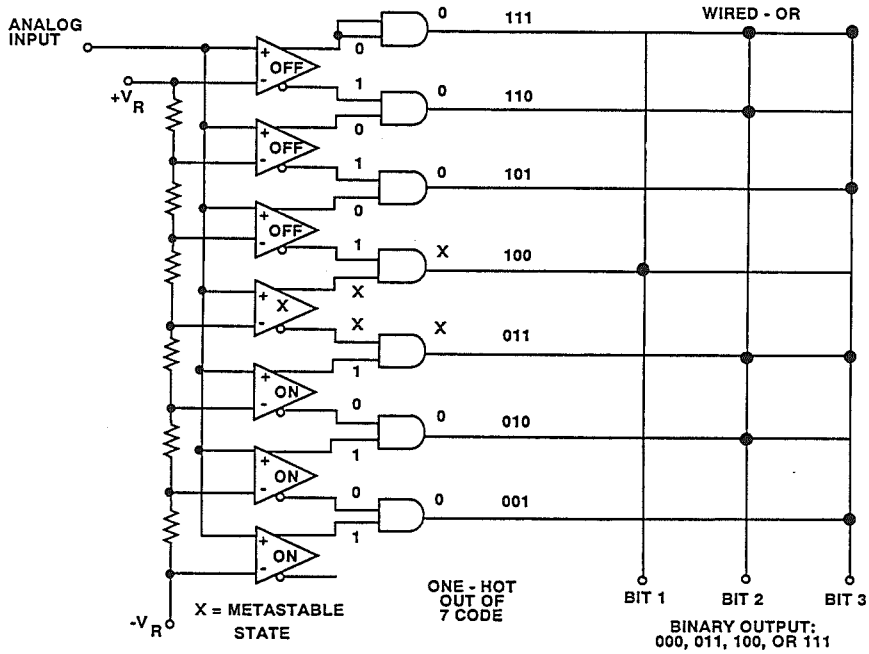

Figure 12.39

The probability of errors due to metastability increases as the sampling rate increases because less time is available for a metastable comparator to settle.

Various measures have been taken in flash converter designs to minimize the metastable state problem. Decoding schemes described in References 3 to 6 minimize the magnitude of these errors. Optimizing comparator designs for regenerative gain and small time constants is another way to reduce these problems.

Metastable state errors may also appear in subranging ADCs which make use of flash converters as building blocks. The same concepts apply, although the magnitudes and locations of the errors may be different.

The test system shown in Figure 12.40 may be used to test for BER in an ADC.

The analog input to the ADC is provided by a high stability low noise sinewave generator. The analog input level is set slightly greater than fullscale, and the frequency such that there is always slightly less than 1 LSB change between samples as shown in Figure 12.41. The test set uses series latches to acquire successive codes A and B. A logic circuit determines the absolute difference between A and B. This difference is then compared to the error limit, chosen to allow for expected random noise spikes and ADC quantization errors. Errors which cause the difference to be larger than the limit will increment the counters. The number of errors, E, are counted over a period of time, T. The error rate is then calculated as BER = $E/2Tf_S$.

ADC BIT ERROR RATE TEST SETUP

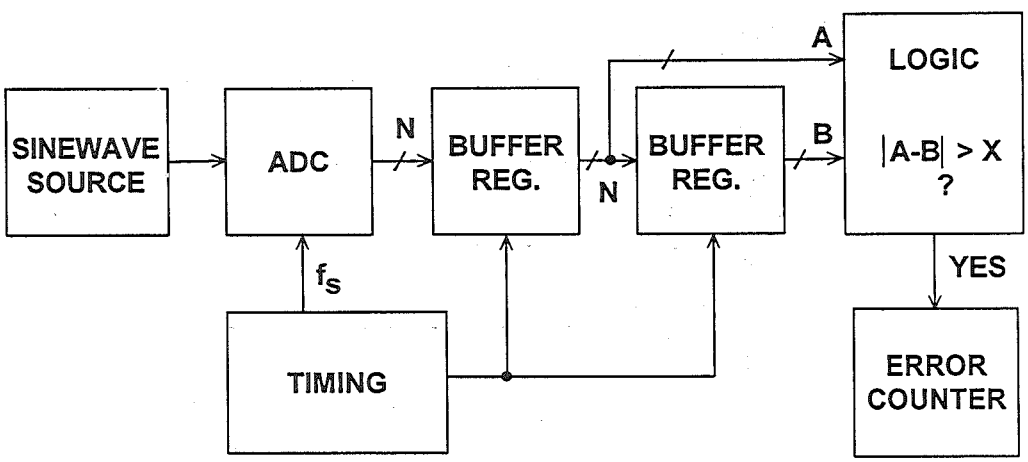

Figure 12.40

ADC ANALOG SIGNAL FOR LOW FREQUENCY BER TEST

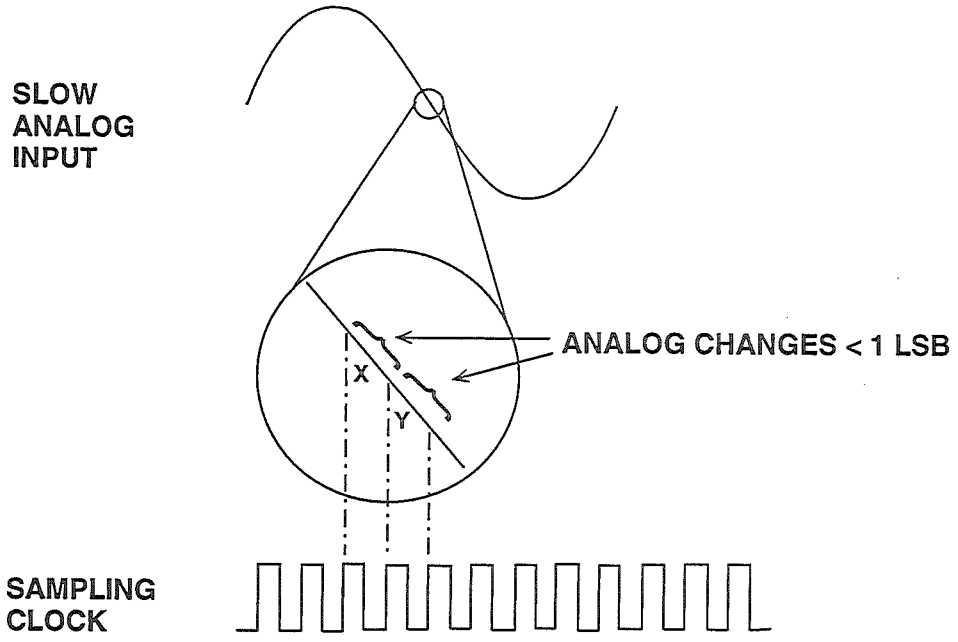

Figure 12.41

ADC ANALOG INPUT FOR HIGH FREQUENCY BER TEST

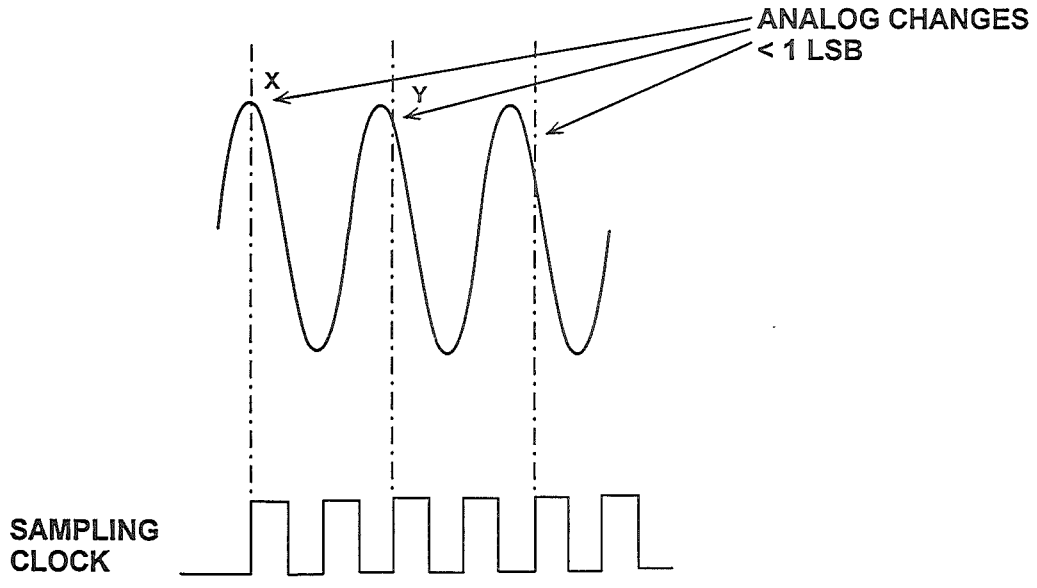

Figure 12.42

SELECTION OF INPUT FREQUENCIES FOR BER TESTING

- For Low Frequency Metastable State Errors:

$$f_{in} = \frac{f_S}{2^N \bullet 2\pi}$$

- For High Frequency Errors:

$$f_{in} = \frac{f_S}{2}\left[1 - \frac{1}{2^N \bullet 4\pi}\right]$$

- N = Number of ADC Bits, f_S = Sampling Rate

Figure 12.43

The same test can be conducted at high frequencies by applying an input frequency slightly offset from $f_S/2$ as shown in Figure 12.42. This causes the ADC to slew fullscale between conversions. Every other conversion is compared, and the "beat" frequency is chosen such that there is slightly less than 1 LSB change between alternate samples. The equations for calculating the proper frequencies for the low and high frequency BER tests are given in Figure 12.43.

Establishing the BER of a well-behaved ADC is a difficult, time-consuming task; a single unit can sometimes be tested for days without an error. For example, tests on the AD9002 8 bit flash converter operating at a sampling rate of 75MSPS yield a BER of approximately 3.7×10^{-12} (1 error per hour) with an error limit of 4 LSBs. Meaningful tests for longer periods of time require special attention to EMI/RFI effects (possibly requiring a shielded screen room), isolated power supplies, etc.

BIT ERROR RATE (BER) FOR 75MSPS SAMPLING RATE

Bit Error Rate (BER)	Average Time Between Errors
1×10^{-8}	1.3 seconds
1×10^{-9}	13.3 seconds
1×10^{-10}	2.2 minutes
1×10^{-11}	22 minutes
1×10^{-12}	3.7 hours
1×10^{-13}	1.5 days
1×10^{-14}	15 days

Figure 12.44

REFERENCES

1. **Active and Passive Electrical Wave Filter Catalog**, Vol. 34, TTE, Incorporated, 2251 Barry Avenue, Los Angeles, CA 90064.

2. **High Speed Design Seminar 1990**, Analog Devices, pp. I-86, I-87.

3. Christopher W. Mangelsdorf, *A 400-MHz Input Flash Converter with Error Correction*, **IEEE Journal of Solid-State Circuits**, Vol. 25, No. 1, February 1990, pp. 184-191.

4. Charles E. Woodward, *A Monolithic Voltage-Comparator Array for A/D Converters*, **IEEE Journal of Solid State Circuits**, Vol. SC-10, No. 6, December 1975, pp. 392-399.

5. Yukio Akazawa et. al., *A 400MSPS 8 Bit Flash A/D Converter*, **1987 ISSCC Digest of Technical Papers**, pp. 98-99.

6. A. Matsuzawa *et al.*, *An 8b 600MHz Flash A/D Converter with Multi-stage Duplex-gray Coding*, **Symposium VLSI Circuits, Digest of Technical Papers**, May 1991, pp. 113-114.

7. Ron Waltman and David Duff, *Reducing Error Rates in Systems Using ADCs*, **Electronics Engineer**, April 1993, pp. 98-104.