

2. ANALOG-TO-DIGITAL CONVERTERS

The ADC is a key component in any radio that uses direct digitization of the RF input signal or that uses digitization after an initial downconversion to an IF. The other key component is the digital signal processor (discussed in Section 3).

2.1 Sampling Methods and Analog Filtering

The sampling process is critical for radio receivers using digitization at the RF or IF. The content of the resulting sampled signal waveform is highly dependent on the relationship between the sampling rate employed and the minimum and maximum frequency components of the analog input signal. Some common sampling techniques that utilize a uniform spacing between the samples include sampling at twice the maximum frequency, oversampling, quadrature sampling, and bandpass sampling (also called downsampling or direct downconversion). Sampling techniques with nonuniform spacing between the samples do exist but they are not widely used and therefore are not considered in this report.

When a continuous-time analog signal is sampled uniformly, the spectrum of the original signal $F(f)$ is repeated at integer multiples of the sampling frequency (i.e., $F(f)$ becomes periodic). This is an inherent effect of sampling and cannot be avoided. This phenomenon is shown graphically in Figure 1. Figure 1a shows the spectrum of the original analog signal $F(f)$. Figure 1b shows the spectrum of the sampled signal $F_s(f)$ using a sampling rate of $f_s = 2f_{\max}$.

2.1.1 Sampling at Twice the Maximum Frequency

The general theorem for sampling a bandlimited analog signal (a signal having no frequency components above a certain frequency f_{\max}) requires that the sampling rate be at least two times the highest frequency component of the analog signal $2f_{\max}$. This ensures that the original signal can be reconstructed exactly from the samples. Figure 1b shows an example of sampling a bandlimited signal with a maximum frequency of f_{\max} at $f_s = 2f_{\max}$. Note that the copies of $F(f)$ that are present in $F_s(f)$ do not overlap. As the sampling rate is increased beyond $2f_{\max}$, the copies of $F(f)$ that are present in $F_s(f)$ are spread even farther apart. This is shown in Figure 1c. Sampling a bandlimited signal at rates equal to or greater than $2f_{\max}$ guarantees that spectrum overlap (often called aliasing) does not occur and that the original analog signal can be reconstructed exactly [3,4]. Figure 1d shows the spectrum overlap that occurs when sampling at rates less than $2f_{\max}$.

2.1.2 Out-of-Band Energy

Two practical problems arise when sampling at the $2f_{\max}$ rate: defining what a bandlimited signal is for real systems and analog filtering before the ADC stage. A theoretically defined bandlimited

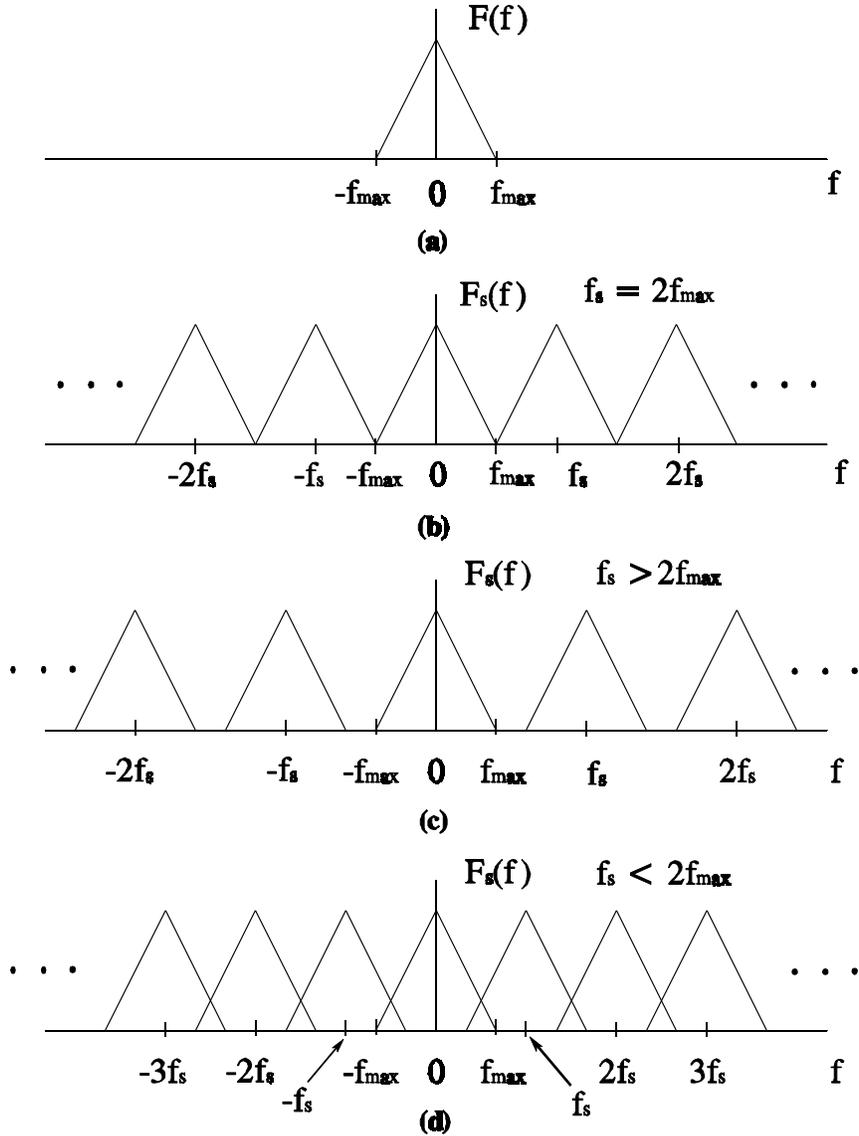


Figure 1. Spectrum of: (a) a bandlimited continuous-time analog signal; (b) the signal sampled at $f_s = 2f_{\max}$; (c) the signal sampled at $f_s > 2f_{\max}$; and (d) the signal sampled at $f_s < 2f_{\max}$.

signal is a signal with no frequency components above a certain frequency. When considering real signals such as an RF signal at the input of a radio receiver, however, signals of all frequencies are always present. While all frequencies are always present, it is the amplitude of these frequencies that is the important factor. In particular, the relative amplitude of the undesired signals to the desired signal is important. When digitizing an RF or IF signal at the $2f_{\max}$ rate in a radio receiver, undesired signals (above one-half the sampling rate) of a sufficient amplitude can create spectrum overlap and distort the desired signal. This phenomenon is illustrated in Figure 2. Figure 2a shows the spectrum of the analog input signal with its desired and undesired components. If this signal is sampled at two times the highest frequency in the desired signal f_d , the resulting spectrum of the sampled signal $F_s(f)$ is shown in Figure 2b. Note that spectrum overlap has occurred here (i.e., the spectrum of the undesired signal occurs within

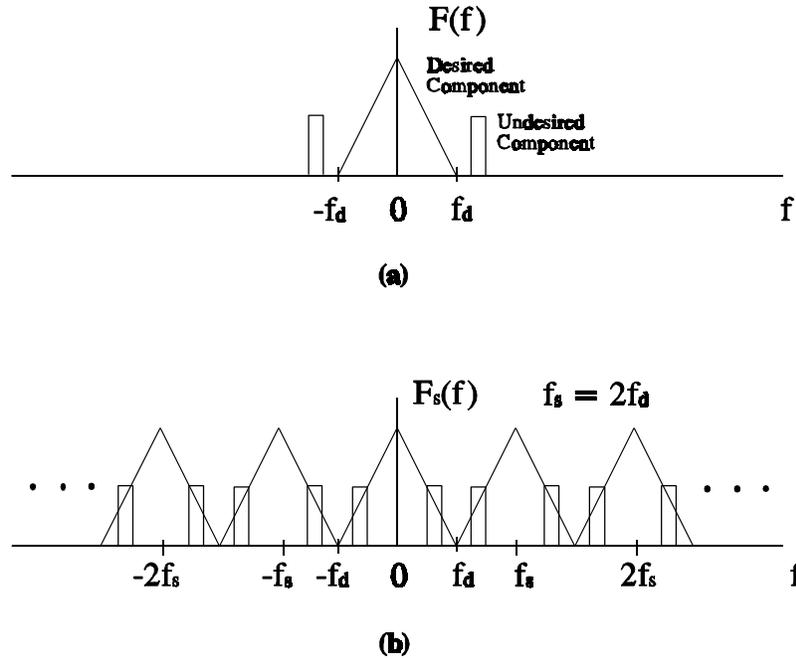


Figure 2. Spectrum of: (a) a continuous-time analog signal with a desired and undesired component and (b) the signal sampled at $f_s = 2f_d$.

the spectrum of the desired signal). This causes distortion in the reconstructed desired signal. This effect raises an important question: “How large do signals occurring above $f_s/2$ need to be for the distortion of the desired signal to be caused predominantly by spectrum overlap and not ADC nonlinearities?” Nonlinearities in the ADC cause spurious responses in the ADC output spectrum. Distortion due to spectrum overlap can be said to predominate distortion due to ADC nonlinearities when the undesired signals appearing in the frequency band from 0 to $f_s/2$ due to spectrum overlap exceed the largest spurious response of the ADC due to nonlinearities. Therefore, undesired signals appearing in the frequency band from 0 to $f_s/2$ due to spectrum overlap must be lower in power than the largest spurious response of the ADC. In other words, distortion of the desired signal is predominated by ADC nonlinearities (and not spectrum overlap) if signals higher in frequency than $f_s/2$ are lower in power than the largest spurious response of the ADC. This can be quite a stringent requirement. Depending upon the details of the radio system, this requirement may be eased.

To determine ways to “ease” this requirement, the following questions should be asked: “How much distortion of the desired signal is tolerable?,” “Do the bandwidth and frequency content of both the desired signal in the frequency band from 0 to $f_s/2$ and the undesired signals above the frequency band from 0 to $f_s/2$ effect the distortion of the desired signal?” These questions are best answered by considering the details of the specific radio system such as the type of source information (voice, data, video, etc.); the desired signal bandwidth; the modulation and coding techniques; the undesired signal characteristics (bandwidth, power, and type of signal); and the performance criterion used to evaluate the reception quality of the desired signal. System

simulation is a valuable tool for providing answers to these questions for specific radio systems and operating environments.

2.1.3 Realizable Anti-Aliasing Filters

Analog filtering before the ADC stage is intimately related to the definition of bandlimiting. Where the definition of bandlimiting involves the content of the signals that may be present, analog filtering before the ADC represents a signal-processing stage where certain frequencies can be attenuated. It is important to know both the signals that can be present before filtering and the amount of attenuation that the filter offers for different frequencies. With knowledge of both of these, the true spectrum of the signal to be digitized can be determined. Sampling at two times the maximum desired frequency presents a large and often impractical demand on the filter used before digitization (the anti-aliasing filter). Ideally, an anti-aliasing filter placed before an ADC would pass all of the desired frequencies up to some cutoff frequency and provide infinite attenuation for frequencies above the cutoff frequency. Then sampling at $f_s = 2f_{\max}$ would be two times the cutoff frequency and no spectrum overlap would occur. Unfortunately, practical, realizable filters cannot provide this type of “brickwall” response. The attenuation of real filters increases more gradually from the cutoff frequency to the stopband. Therefore, for a given cutoff frequency on a real filter, sampling at two times this cutoff frequency will produce some spectrum overlap. The steeper the transition from the passband to the stopband and the more attenuation in the stopband, the less the sampled signal will be distorted by spectrum overlap. In general, more complicated filters are required to achieve steeper transitions and higher attenuation in the stopband. Therefore, more complicated filters are required to reduce the distortion in the sampled signal due to spectrum overlap for a given sampling rate. Limitations on the practical implementation of analog filters make high-order, steep rolloff filters difficult to realize. Also, as the steepness of the rolloff is increased, the phase response tends to become more nonlinear. This can create distortion of the desired receive signal since different frequencies within a signal can be delayed in time by different amounts.

2.1.4 Oversampling

Sampling at rates greater than $2f_{\max}$ is called oversampling. One of the benefits of oversampling is that the copies of $F(f)$ that are present in $F_s(f)$ become increasingly separated as the sampling rate is increased beyond $2f_{\max}$. For an analog signal with a given frequency content and a given anti-aliasing filter with a cutoff frequency of f_c , sampling at two times the cutoff frequency produces a certain amount of distortion due to spectrum overlap. When sampling at a higher rate, a simpler anti-aliasing filter with a more gradual transition from passband to stopband and less stopband attenuation can be used without any increase in the distortion due to spectrum overlap. Therefore, oversampling can minimize the requirements of the anti-aliasing filter. The tradeoff, of course, is that faster ADC's are required to digitize relatively low frequency signals.

2.1.5 Quadrature Sampling

In quadrature sampling the signal to be digitized is split into two signals. One of these signals is multiplied by a sinusoid to downconvert it to a zero center frequency and form the in-phase component of the original signal. The other signal is multiplied by a 90° phase-shifted sinusoid to downconvert it to a zero center frequency and form the quadrature-phase component of the original signal. Each of these components occupies only one-half of the bandwidth of the original signal and can be sampled at one-half the sampling rate required for the original signal. Therefore, quadrature sampling reduces the required sampling rate by a factor of two at the expense of using two ADC's instead of one.

2.1.6 Bandpass Sampling for Direct Downconversion

Sampling at rates lower than $2f_{\max}$ still can allow for an exact reconstruction of the information content of the analog signal if the signal is a bandpass signal. An ideal bandpass signal has no frequency components below a certain frequency f_l and above a certain frequency f_h . Typically, bandpass signals have $f_l \gg f_h - f_l$. For a bandpass signal, the minimum requirement on the sampling rate to allow for exact reconstruction is that the sampling rate be at least two times the bandwidth $f_h - f_l$ of the signal.

Sampling at a rate two times the bandwidth of a signal is called the Nyquist sampling rate. When the signal is a baseband signal (a signal with frequency content from DC to f_{\max}) the Nyquist sampling rate is $2f_{\max}$. For bandpass signals, however, the Nyquist sampling rate is $2(f_h - f_l)$. To ensure that spectrum overlap does not occur when sampling rates are between two times the bandwidth of the bandpass signal and two times the highest frequency in the bandpass signal, the sampling frequency f_s must satisfy [4]

$$\frac{2f_h}{k} \leq f_s \leq \frac{2f_l}{(k-1)} \quad \text{where } k \text{ is restricted to integer values that satisfy } 2 \leq k \leq \frac{f_h}{(f_h - f_l)}$$

$$\text{and } (f_h - f_l) \leq f_l .$$

These equations show that only certain ranges of sampling rates can be used if spectrum overlap is to be prevented.

Bandpass sampling can be used to downconvert a signal from a bandpass signal at an RF or IF to a bandpass signal at a lower IF. Since the bandpass signal is repeated at integer multiples of the sampling frequency, selecting the appropriate spectral replica of the original bandpass signal provides the downconversion function.

Bandpass sampling holds promise for radio receivers that digitize directly at the RF or IF since the desired input signals to radio receivers are normally bandpass signals. Theoretically, bandpass sampling allows sampling rates to be much lower than those required by sampling at two or more times the highest frequency content of the bandpass signal. This implies that ADC's with slower sampling rates (and therefore potentially higher performance, lower power consumption, or lower cost) may be used. An important practical limitation, when using

bandpass sampling, is that the ADC must still be able to effectively operate on the highest frequency component in the signal. This specification is usually given as the analog input bandwidth for the ADC.

Conventional ADC's are designed to operate on signals with maximum frequencies up to one-half the sampling rate. In other words, conventional ADC's typically are not suitable for bandpass sampling applications where the maximum input frequencies are greater than the sampling rate. Furthermore, for conventional ADC's, many manufacturers provide specifications only at frequencies well below one-half the maximum sampling rate. In general, performance of ADC's typically degrades with increased input frequency. Therefore, in using ADC's for frequencies near one-half the maximum sampling rate or for bandpass sampling applications, the specifications of the converter must be determined and carefully examined at the desired input frequencies. In addition, when bandpass sampling, stringent requirements on analog bandpass filters (steep rolloffs) are needed to prevent distortion of the desired signal from strong adjacent channel signals.

2.2 Effects of Quantization Noise, Distortion, and Receiver Noise

This section addresses the relationships between quantization noise, harmonic distortion, and receiver noise. The ADC's best suited to RF and IF processing that have widespread availability use uniform quantization. In uniform quantization, the voltage difference between each quantization level is the same. Other methods of quantization include logarithmic (A-law and μ -law), adaptive, and differential quantization. These methods currently are used in source coding. A discussion of these quantization techniques is given in Section 4.

In uniform quantization, the analog signal cannot be represented exactly with only a finite number of discrete amplitude levels. Therefore, some error is introduced into the quantized signal. The error signal is the difference between the analog signal and the quantized signal. Statistically, the error signal is assumed to be uniformly distributed within a quantization level. Using this assumption, the mean squared quantization noise power P_{qn} is

$$P_{qn} = \frac{q^2}{12R}$$

where q is the quantization step size and R is the input resistance of the ADC [5]. In an ideal ADC, this representation of the quantization noise power is accurate to within a dB for input signals that are not correlated with the sampling clock.

If the analog input into an ADC is periodic, the error signal is also periodic. This periodic error signal includes harmonics of the analog input signal and results in harmonic distortion. Furthermore, harmonics that fall above $f_s/2$ appear in the frequency band from 0 to $f_s/2$ due to aliasing.

This harmonic distortion that occurs through the quantization process is quite undesirable in radio receiver applications; it becomes difficult if not impossible to distinguish the harmonics caused by quantization and the spurious and harmonic components of the actual input signal. Dithering is a technique that is commonly used to reduce this harmonic distortion.

Dithering is a method of randomizing the quantization noise by adding an additional noise signal to the input of the ADC [6]. Several types of techniques are used for dithering. Perhaps the most basic technique adds wideband thermal noise to the input of the ADC. This can be accomplished by summing the output of a noise diode with the input signal before digitization by the ADC. This also can be achieved by simply placing an amplifier before the ADC and providing enough gain to boost the receiver noise to a level that minimizes the spurious responses of the ADC. These techniques reduce the levels of the spurious responses by randomizing the quantization noise. In other words, for periodic input signals that would normally produce harmonics in the ADC output, the addition of a dithering signal spreads the energy in these harmonic components into random noise, thus reducing the amplitude of the spurious components.

A disadvantage of adding wideband noise to the input of the ADC is that the SNR is degraded. The amount of degradation depends upon the amount of noise power added to the input of the ADC. Adding a noise power equal to the quantization noise power degrades the SNR by 3 dB [5].

Two techniques commonly are used to prevent degradation in the SNR while dithering. The first technique filters noise from a wideband noise source before the noise is added to the ADC input. Filtering limits the noise power to a frequency range that is outside of the receiver's bandwidth. Hence, over the receiver's bandwidth, the SNR is not degraded.

The other technique used to prevent degradation of the SNR is called subtractive dithering (Figure 3). A pseudorandom noise (PN) code generator is used to generate the dithering signal. The digital output of the PN code generator is converted into an analog noise signal using a digital-to-analog converter. This noise signal is added to the input signal of the ADC. The digital output of the PN code generator is then subtracted from the output of the ADC, again preserving the SNR of the ADC [7]. An example of dithering, achieved by boosting the receiver system noise with an amplifier, is given in the following paragraphs.

Commercially available ADC's typically have a full scale range (FSR) of 1-20 V. The FSR of the ADC is the difference between the maximum and the minimum analog input voltages to the ADC. Dividing the FSR by the number of quantization levels 2^B , where B is the number of bits of the ADC, provides the quantization step size q . For an 8-bit ADC with an FSR of 2.5 V, the quantization step size is 9.77 mV. To compute the quantization noise power, the effective input resistance of the ADC must be known.

Components in radio receivers typically have a 50-ohm input and output impedance. The input impedance of ADC's is usually higher than this and is not well specified. Therefore, when interfacing an RF component with an ADC, as is necessary for digitization at the RF or IF, this

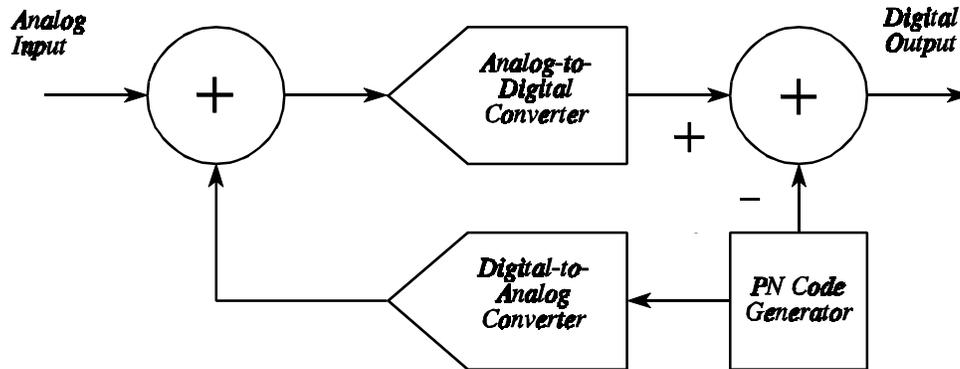


Figure 3. Block diagram of subtractive dithering.

impedance mismatch must be considered. A simple method of impedance matching is to place a 50-ohm resistive load at the input of the ADC. This forces the effective input resistance of the ADC to be close to 50 ohms. The quantization noise power then can be computed. Assuming a 50-ohm effective input resistance R to the ADC in this example, the quantization noise power equals -38 dBm. For a noise-limited receiver, the receiver noise power P_m can be computed as the thermal noise power in the given receiver bandwidth (BW) plus the receiver noise figure (NF). This is given as

$$P_{rn} = -174 \text{ dBm} + 10 \log_{10} \text{BW (Hz)} + \text{NF (dB)} .$$

For a receiver with a 10-MHz BW and a 6-dB NF, the receiver noise power is -98 dBm. Therefore, a gain of 60 dB is required to boost the receiver noise to the quantization noise power level. For an ADC of higher resolution, less gain would be needed since the quantization noise power would be smaller. Also, wider receiver bandwidths and higher receiver noise figures would require less gain since the receiver noise power would be larger. Nevertheless, for most practical receiver and ADC combinations, an amplifier with automatic gain control is necessary before the ADC. The automatic gain control is designed so that the receiver noise roughly equals the quantization noise power level for low-level signals and the input signal power does not exceed the ADC's FSR for high-level signals.

The 2- to 30-MHz single-sideband (SSB) receiver presented in [8] shows an example of a receiver using this dithering technique. Digitization occurs at the 456-kHz IF after dual downconversion. The Collins Radio Division of Rockwell International uses this type of scheme in many of their receivers.

2.3 Important Specifications

In this section, theoretical signal-to-noise ratio (SNR) due to quantization noise and aperture jitter is discussed. Practical specifications for real ADC's are then presented.

2.3.1 Theoretical Signal-to-Noise Ratio Specifications

For radio receiver applications where the amplitude of the desired signal falls within the ADC's FSR, and the bandwidth of the desired signal is equal to $f_s/2$, the SNR of an ADC is a useful specification. The theoretical maximum SNR of ADC's generally is assumed to be $6B$ (dB), where B is the number of bits of resolution of the ADC. A more precise expression providing the maximum possible theoretical SNR can be derived based on some assumptions about the noise and the input signal. First, it is assumed that the noise present is due to quantization error only. The amplitude of this quantization noise is assumed to be a random variable uniformly distributed over one quantization step. Assuming a sinusoidal input with an amplitude equal to the FSR of the ADC, the maximum possible theoretical SNR is given as

$$\text{SNR} = 6.02B + 1.76 + 10 \log_{10} \left(\frac{f_s}{2f_{\max}} \right) \quad (\text{dB}) \quad (1)$$

where f_s is the sampling frequency and f_{\max} is the maximum frequency of the input analog signal [2,9]. The commonly stated theoretical SNR of $6B$ (dB) is an approximation to this equation when $f_s = 2f_{\max}$ and the 1.76 dB is neglected. From this equation, note that as the sampling frequency is increased beyond $2f_{\max}$, the SNR increases. This occurs because the quantization noise power, which is a fixed amount and independent of bandwidth ($P_{\text{qn}} = q^2/12R$), is spread out over an increasingly wider band as the sampling frequency is increased. This lessens the amount of the quantization noise that falls within the 0 to $f_s/2$ band. Figure 4 shows this phenomenon. Consequently, oversampling increases the maximum possible SNR. Such oversampling is sometimes used to realize a greater maximum SNR than at first appears possible. An 8-bit ADC, with a sampling rate of 20 Msamples/s, for example, can provide 68 dB rather than 48 dB of maximum SNR for 100-kHz signals in the passband if appropriate digital filtering is used to recover the 100-kHz signal.

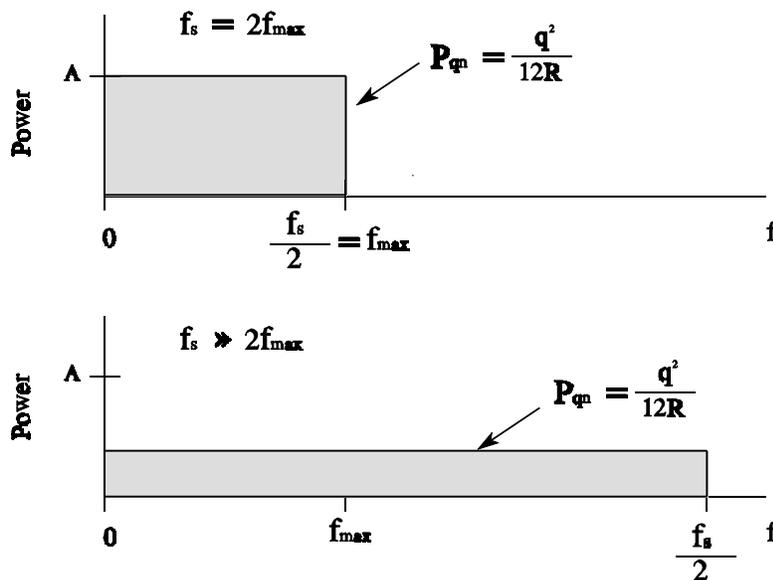


Figure 4. Frequency-spreading of quantization noise power due to oversampling.

Besides being limited by the quantization step size (resolution), the SNR of the ADC also is limited by aperture jitter. Aperture jitter is the variation in time of the exact sampling instant. Aperture jitter can be caused externally by jitter in the sampling clock, or internally since the sampling switch does not open at precise times. Aperture jitter causes a phase modulation of the sampled signal and thus results in an additional noise component in the sampled signal [10]. The maximum analog input frequency of the ADC is limited by this aperture jitter since the SNR due to aperture jitter (SNR_{aj}) degrades as the input frequency increases. The SNR_{aj} is given as

$$\text{SNR}_{aj} = 20 \log_{10} \left(\frac{1}{2 \pi f_{\max} t_a} \right)$$

where t_a is the aperture jitter of the ADC [2]. For sampling at $f_s = 2f_{\max}$, both the SNR due to quantization noise and the SNR due to aperture jitter can be combined to give the overall SNR [11].

2.3.2 Practical Specifications for Real ADC's

The SNR in a real ADC can be determined by measuring the residual error. Residual error is the combination of quantization noise, random noise, and nonlinear distortion (i.e., all of the undesired components of the output signal from the ADC). The residual error for an ADC is found by using a sinusoidal input into the ADC. An estimate of the input signal is subtracted from the output of the ADC; the remaining signal is the residual error. The mean squared (MS) power of the residual error then is computed. The SNR then is found by dividing the mean squared power of the input signal by the mean squared power of the residual error.¹

A specification sometimes used for real ADC's instead of the SNR is the effective number of bits (ENOB). This specification is defined as the number of bits required in an ideal ADC so that the mean squared noise power in the ideal ADC equals the mean squared power of the residual error in the real ADC.

The spurious free dynamic range (SFDR) is another useful specification for ADC's. One definition of the SFDR assumes a single tone sinusoidal input into the ADC. Measurement of this SFDR is made by taking the Fast Fourier Transform (FFT) of the output of the ADC. This provides the frequency spectrum of the output of the ADC and is plotted as the ADC output power in dB vs. frequency. The SFDR is then the difference between the power in the sinusoidal input signal and the peak power of the largest spurious signal in the ADC output spectrum. An example of determining the SFDR from the ADC output spectrum is shown in Figure 5. In this idealized ADC output spectrum, the input signal is a 10-MHz sinusoid. Various spurious responses are shown. The SFDR is 50 dB.

¹The SNR is often (and more accurately) called the signal-to-noise plus distortion ratio (SINAD) when distortion is included with the noise (as in this case).

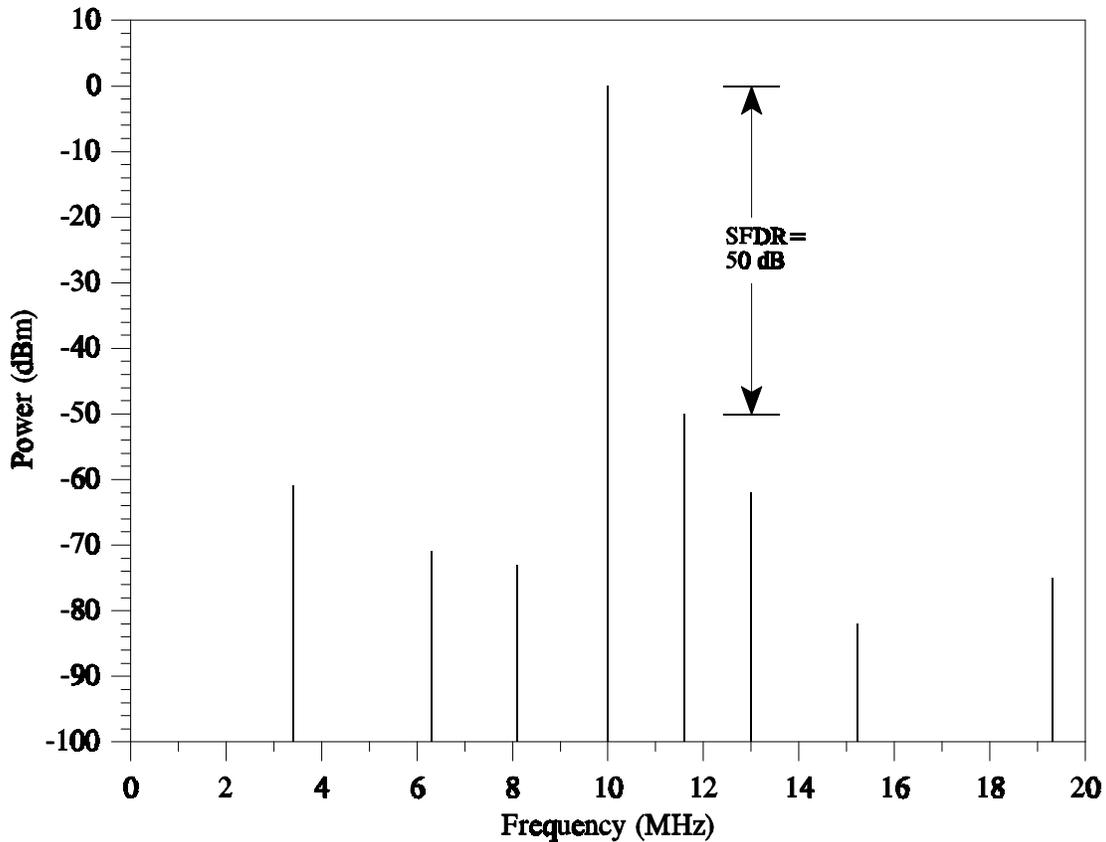


Figure 5. Example ADC output spectrum showing the spurious free dynamic range.

SFDR allows one to assess how well an ADC can detect simultaneously a very small signal in the presence of a very large signal. Hence, it is an important specification for ADC's used in radio receiver applications. A common misconception is that the SFDR of the ADC is equivalent to the SNR of the ADC. In fact, there is typically a large difference between the SFDR and the SNR of an ADC. The SNR is the ratio between the signal power and the power of the residual error. The SFDR, however, is the ratio between the signal power and the peak power of only the largest spurious product that falls within the band of interest. Therefore, the SFDR is not a direct function of bandwidth; it does not necessarily change with a change in bandwidth, but it may. Since the power of the residual error includes quantization noise, random noise, and nonlinear distortion within the entire 0 to $f_s/2$ band, the power of the residual error can be much higher than the peak power of the largest spurious product. Hence, the SFDR can be much larger than the SNR [5]. A practical example of this can be seen from the specifications for the Analog Devices AD9042 monolithic ADC. With a 19.5-MHz analog input signal, 1 dB below full scale (the full-scale input is $1 V_{pp}$), the typical SFDR specification is 81 dB while the SNR specification is 66.5 dB (from -40 - $+85$ °C) [12].

The SFDR specification is useful for applications when the desired signal bandwidth is smaller than $f_s/2$. In this case, a wide band of frequencies is digitized and results in a given SNR. The desired signal then is obtained by using a narrowband digital bandpass filter on this entire band

of frequencies. The SNR is improved by this digital-filtering process since the power of the residual error is decreased by filtering. The SFDR specification for the ADC is important because a spurious component still may fall within the bandwidth of the digital filter; hence, the SFDR, unlike the SNR, does not necessarily improve by the digital-filtering process. However, several techniques are available to improve the SFDR. Dithering (discussed in Section 2.2) improves the SFDR of ADC's. Additionally, postdigitization-processing techniques such as state variable compensation [13], phase-plane compensation [14], and projection filtering [15] have been used to improve SFDR.

For an ideal ADC, and in practical sigma-delta ($\Sigma \Delta$) converters, the maximum SFDR occurs at a full-scale input level. In other types of practical ADC's, however, the maximum SFDR occurs at input levels at least several dB below the full-scale input level. This occurs because as the input levels approach full-scale (within several dB), the response of the ADC becomes more nonlinear and more distortion is exhibited. Additionally, due to random fluctuations in the amplitude of real input signals, as the input signal level approaches the FSR of the ADC, the probability of the signal amplitude exceeding the FSR increases. This causes additional distortion from clipping. Therefore, it is extremely important to avoid input signal levels that closely approach the full-scale level in ADC's. Prediction of the SFDR for practical ADC's is difficult, therefore measurements are usually required to characterize the SFDR.

In the preceding discussion on SFDR, a sinusoidal ADC input signal was assumed. However, intermodulation distortion (IMD) due to multitone inputs is important in ADC's used for wideband radio receiver applications. To characterize this IMD due to multitone inputs, another definition of the SFDR could be used. In this case, the SFDR is the ratio of the combined signal power of all of the multitone inputs to the peak power of the largest spurious signal in the ADC output spectrum. A current example of test equipment to generate multitone inputs produces up to 48 tones.

The noise power ratio (NPR) specification is useful in applications such as mobile cellular radio, where the spectrum of a signal to be digitized consists of many narrowband channels and where adjacent channel interference can degrade system performance. Particularly, the NPR provides information on the effectiveness of an ADC in limiting crosstalk between channels [13].

The NPR is measured by using a noise input signal into the ADC. This noise signal has a flat spectrum that is bandlimited to a frequency that is less than one-half the sampling frequency. Additionally, a narrow band of frequencies is removed from the noise signal using a notch filter. This noise spectrum is used as the input signal to the ADC. The frequency spectrum of the output of the ADC then is determined. The NPR then is computed by dividing the power spectral density of the noise outside the frequency band of the notch filter by the power spectral density of the noise inside the frequency band of the notch filter [5].

When using an ADC in a bandpass-sampling application where the maximum input frequency into the ADC is actually higher than one-half the sampling frequency, the full-power analog input bandwidth is an important specification. A common definition (although not universal) of full-power analog input bandwidth is the range from DC to the frequency where the amplitude of the output of the ADC falls to 3 dB below the maximum output level. This assumes a full-scale input

signal to the ADC. Typically, the ADC is operated at input frequencies below this bandwidth. Aside from full-power analog input bandwidth, it is important to examine the behavior of the other specifications such as SNR, SFDR, and NPR at the desired operating frequencies since these specifications typically vary with frequency. In addition to the SNR, SFDR, and NPR of real ADC's being a function of frequency, they are also a function of input signal amplitude. Table 1 provides a summary of the important ADC specifications for radio receiver applications.

Table 1. Summary of ADC Specifications for Radio Receiver Applications

Specification	Application	Definition
Signal-to-Noise Ratio (SNR)	Desired Signal BW Equal to $f_s/2$	$\frac{\text{MS Signal Power}}{\text{MS Power of Residual Error}}$
Spurious Free Dynamic Range (SFDR)	Desired Signal BW Less Than $f_s/2$	$\frac{\text{MS Signal Power}}{\text{Peak Power of the Largest Spurious Product}}$
Noise Power Ratio (NPR)	Desired Signal Spectrum Contains Many Narrowband Channels	$\frac{\text{Power Spectral Density of Noise * Outside Freq. Band of Notch Filter}}{\text{Power Spectral Density of Noise Inside Freq. Band of Notch Filter}}$
Full-Power Analog Input BW	Bandpass Sampling	Range from DC to Frequency Where Output Amplitude Falls to 3 dB Less Than Maximum **

*With an input signal having a bandlimited, flat noise spectrum and a narrow band of frequencies removed by a notch filter.

**For a full-scale input signal.

When testing an ADC, it is important to ensure that all quantization levels are tested. For single tone inputs, the relationship between the input signal frequency and the sampling rate must be chosen so that the same small set of quantization levels is not tested repeatedly. In other words, the samples should not always occur at the same amplitude levels of the input signal. For example, using an input frequency of $f_s/8$ is a poor choice since the same eight amplitude levels are sampled every period of the input signal (assuming that the input signal and the sampling clock are phase coherent) [16]. The histogram test can be used to ensure that all quantization levels are tested. In the histogram test, an input signal is applied to the ADC and the number of samples that are taken at each of the 2^B quantization levels are recorded. In an ideal ADC this histogram is identical to the probability density function of the amplitude values of the input signal. Comparing the histogram to the probability density function of the input signal gives an indication of the nonlinearity of the ADC. An examination of the histogram reveals whether all

of the different quantization levels are being tested. When no samples are recorded for a given quantization level, this level is either not being tested by the testing procedure (input signal and sampling rate) or the ADC is exhibiting a missing code. A missing code is a quantization level that is not present in the output of a real ADC that is present in the output of an ideal ADC. Missing codes are fairly rare in currently available ADC's in general and do not occur in $\Sigma\Delta$ converters.

2.4 ADC Conversion Methods

Many methods for implementing ADC's currently exist. Several of the most common techniques are presented below. The counter ADC uses a digital-to-analog converter (DAC) and increases the output of this DAC one quantization level at a time using a counter circuit until the output of the DAC equals the amplitude of the analog signal at a given time. The output of the counter then provides the digital representation of the analog input voltage. A major drawback to this type of converter is that it is fairly slow. An improvement to the counter type ADC is the tracking ADC. This type of converter is similar to the counter ADC except that an up-down counter is used in place of the ordinary counter. In this ADC, the output of the internal DAC is compared to the analog input signal. If the amplitude of the analog input signal is greater than the output of the DAC, the counter counts up; if it is less than the output of the DAC, the counter counts down. The tracking ADC is much faster than the counter ADC when there are only small changes in the amplitude of the input signal. For large changes in the input signal amplitude this type of ADC is still fairly slow.

The counter and tracking ADC's belong to the feedback class of ADC's. The successive-approximation ADC also belongs to this class. This type of ADC again uses a DAC in a feedback loop. For a conversion with this ADC, a register is used to set the most significant bit (MSB) in the DAC to 1. The output of the DAC is compared to the amplitude of the analog input. If the DAC output is greater than the analog input, the MSB of the DAC is cleared, otherwise it is kept set to 1. The next significant bit of the DAC is then set to 1 and the output of the DAC again is compared to the amplitude of the analog input. If the DAC output is greater than the analog input, this bit is cleared. This process continues for all B bits of the DAC. The input of the DAC provides the output of the ADC. The conversion is made in B steps making this technique quite efficient and hence reasonably fast. Successive-approximation is one of the most popular ADC techniques.

The parallel or flash ADC is used for applications that require the fastest possible digitization. In the current state-of-the-art technology, sampling rates on the order of 500-1000 Msamples/s for an 8-bit ADC most likely imply that a flash ADC is being used. This type of converter uses a bank of $2^B - 1$ voltage comparators in parallel where B is the number of bits of the ADC. The analog input signal is applied to one input on all of the voltage comparators while the other input to each comparator is a reference voltage corresponding to each of the $2^B - 1$ quantization levels. The reference voltages typically are generated by a voltage divider network. All comparators with reference voltages below the analog input signal produce a logical 1 output. The remaining comparators, with reference voltages equal to or above the input signal, produce a logical 0 output. The outputs of the comparators are then combined in a fast decoder circuit

to generate the output digital word of the ADC. Therefore, conversion takes place in only two steps (voltage comparison and decoding), making this technique the fastest of the commonly available techniques. A major limitation of this type of ADC is the large number of comparators required in the implementation. For a B-bit flash ADC, $2^B - 1$ comparators are needed. Since an 8-bit ADC requires 255 comparators and a 9-bit ADC requires 511 comparators, flash ADC's of more than 8 bits typically are not available commercially. Linearity is a problem in flash ADC's as observed in degraded SFDR performance.

One technique used to implement high-speed ADC's combines two separate B-bit ADC's (usually flash ADC's) to produce a single ADC with a resolution of 2B bits. For example, two 4-bit converters can be combined to provide an 8-bit converter. In this technique, the first 4-bit ADC digitizes the analog input. The output of the ADC is then converted back into an analog signal using a DAC. This signal then is subtracted from the original input analog signal producing a difference signal. This difference signal is then amplified and digitized using the second 4-bit ADC. The amplifier gain is set to provide a full-scale input signal into the second ADC. The outputs of both 4-bit ADC's then are combined using digital error correction logic to produce an 8-bit output representing the analog input signal [17]. This type of ADC is called a two-stage subranging ADC. Subranging ADC's with up to five stages are available. Signal delays (sometimes called pipeline delays) increase with each additional stage and must be considered in the design of subranging ADC's. Subranging ADC's exhibit a repetitive nonlinearity due to the nature of digitizing difference signals. This is visualized best by considering a ramp input into an ideal ADC representing the first stage of a subranging ADC. The transfer function of the ideal ADC is shown in Figure 6a while the ramp input is shown in Figure 6b. The output of the first ADC is a quantized version of the ramp as shown in Figure 6c. Subtracting a reconstructed version of the input ramp from the quantized version, produces a repetitive ramp difference waveform that repeats 2^B times where B is the number of bits in the first ADC. This difference signal (shown in Figure 6d) is amplified to produce a full-scale input to the second ADC. Therefore, the differential nonlinearity in the second ADC is exercised 2^B times. (Differential nonlinearity is defined as the deviation of any quantization step in the ADC from q , the theoretical quantization step size of the ADC.)

Subranging ADC's are becoming very popular since they can achieve high-speed operation with high resolution. They require far fewer comparators for a given resolution than flash ADC's. While the internal ADC's within a subranging ADC have traditionally been flash ADC's, other types of ADC's may be used. For example, a new architecture, the cascaded magnitude amplifier, has been used in the Analog Devices AD9042 (a 12-bit, 41-Msamples/s subranging ADC). This architecture provides a very high-speed conversion and greatly reduces the number of comparators required in the internal ADC's.

The cascaded magnitude amplifier (MA) ADC consists of B-1 MA's in series and a single comparator placed in series after the last MA. A diagram showing the operation of the cascaded MA ADC is given in Figure 7. Referring to this figure, the first MA compares the input signal to a voltage level $V_{fs}/2$ where V_{fs} is the full-scale input voltage of the ADC. If the input signal

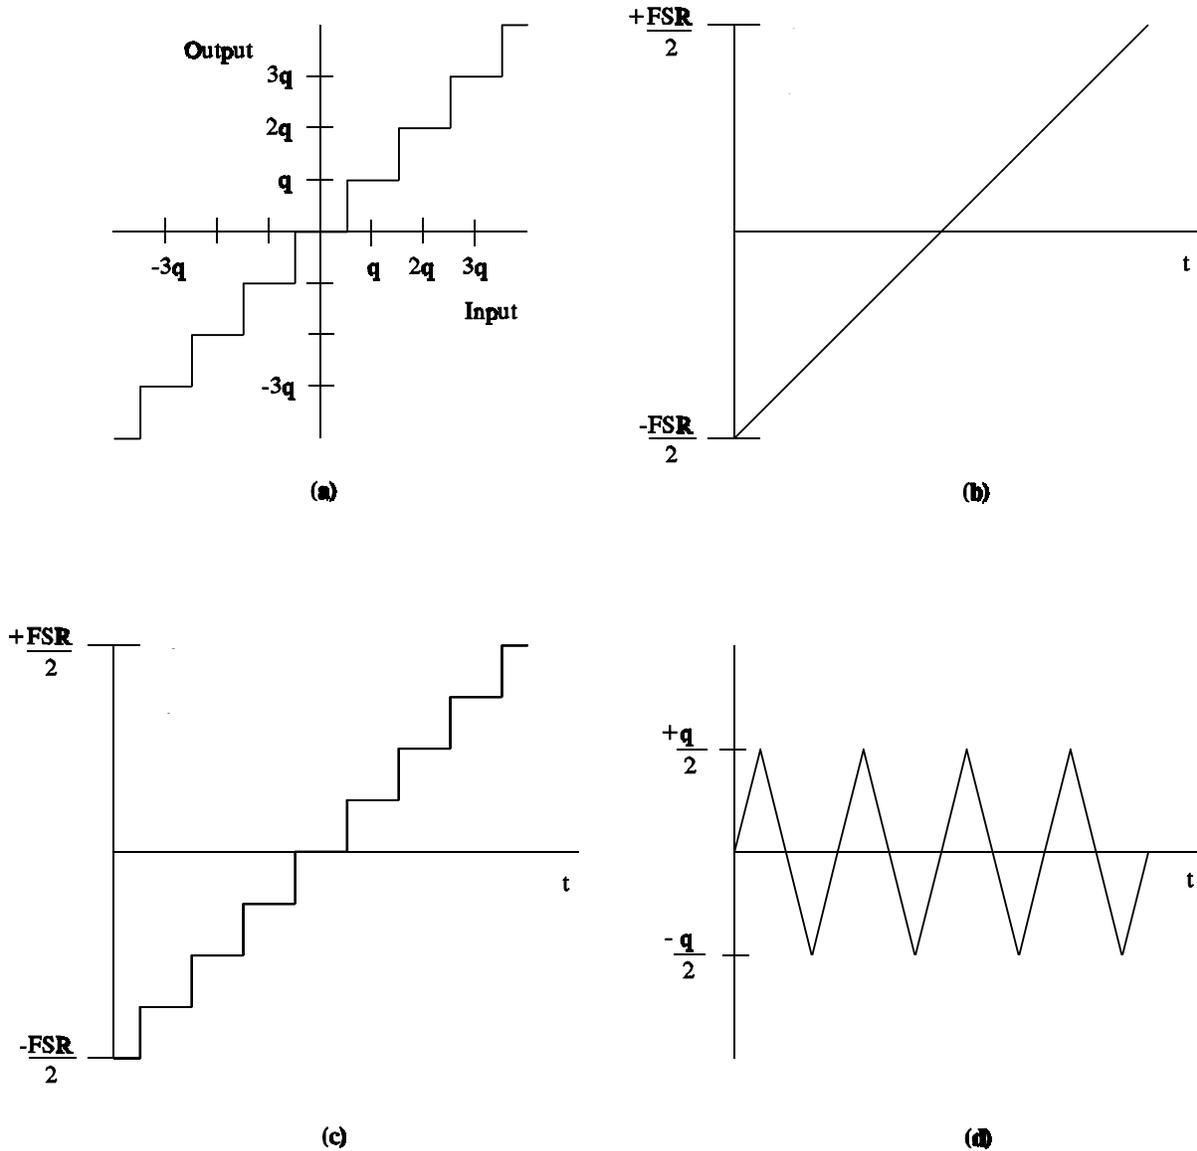


Figure 6. (a) Ideal ADC transfer function; (b) input ramp signal; (c) output (quantized) ramp signal; and (d) repetitive ramp difference signal.

is greater than $V_{fs}/2$, the bit representing this MA is set to 1. If the input signal is less than $V_{fs}/2$, the bit representing this MA is set to 0. Therefore, this first MA divides the full-scale voltage into two regions and the bit representing this MA is set according to the region in which the input voltage falls.

The next (second) MA uses the output of the first MA as its input. As shown in Figure 7, the second MA divides each of the two regions defined by the first MA into two additional regions. If the input voltage to the first MA is between V_{fs} and $V_{fs}/2$, the second MA determines if the input signal is between V_{fs} and $3V_{fs}/4$ or between $3V_{fs}/4$ and $V_{fs}/2$. The bit representing this MA

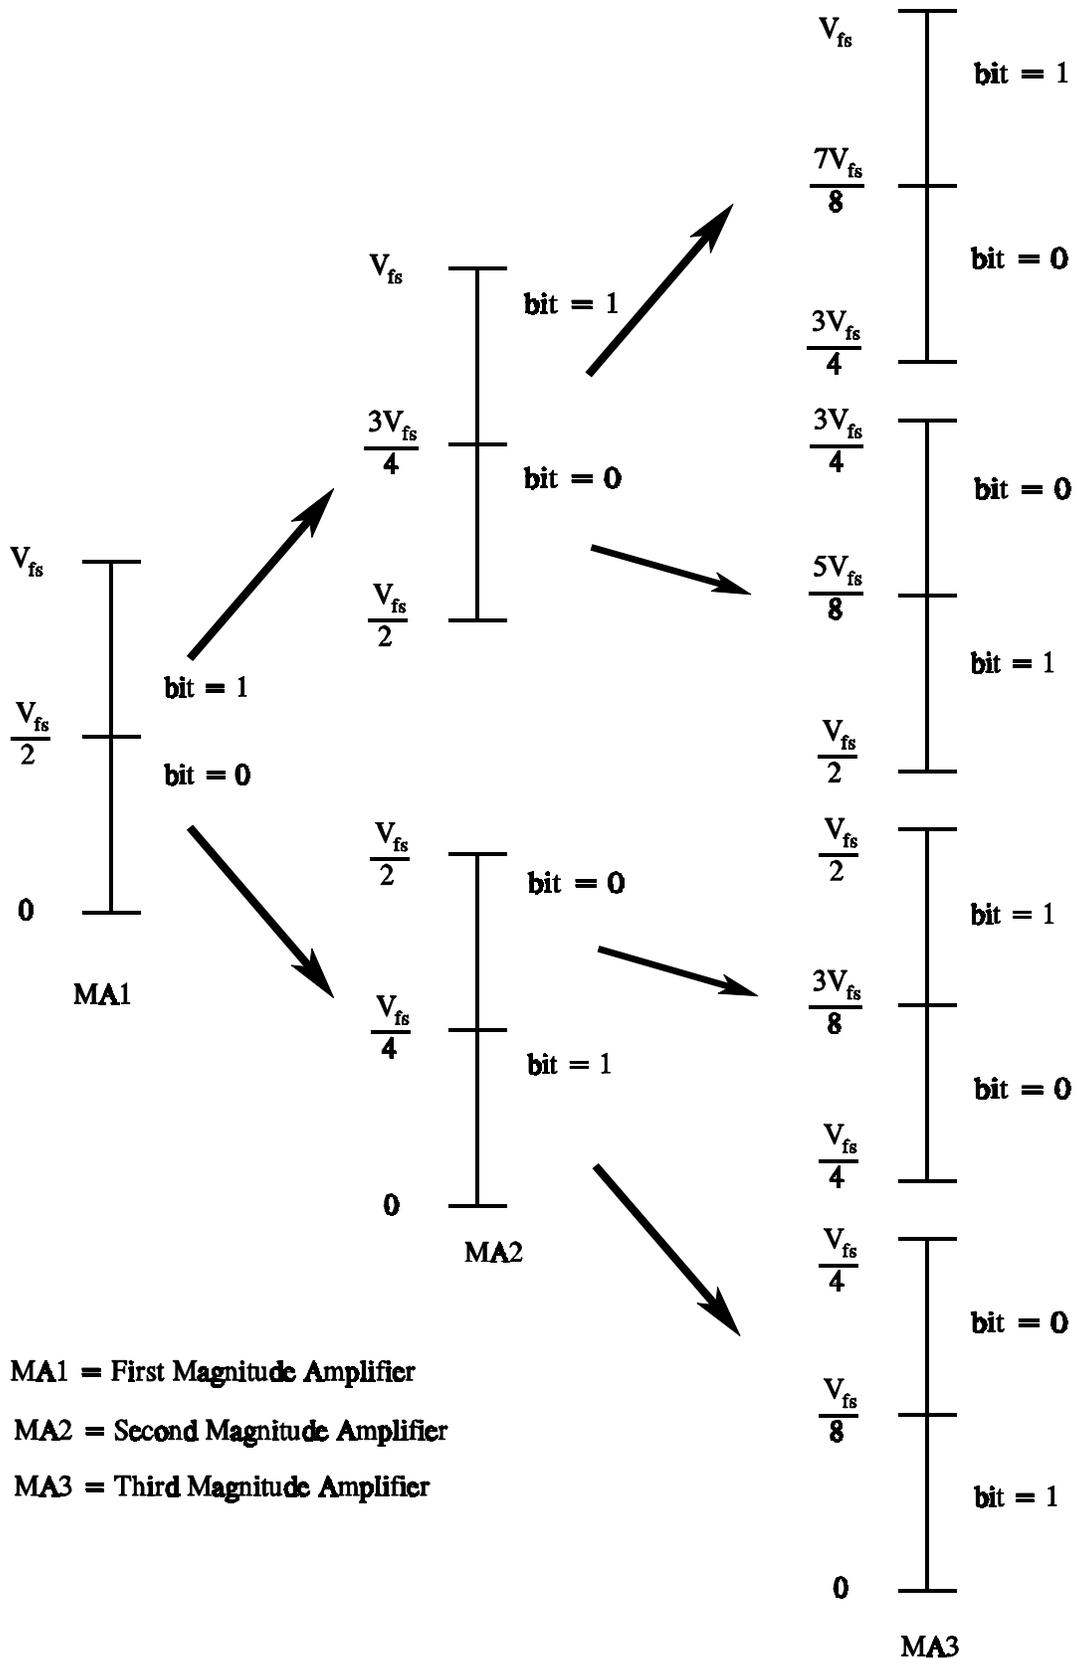


Figure 7. Operation of the cascaded magnitude amplifier ADC.

then is set to 1 or 0, respectively. Conversely, if the input voltage to the first MA is between 0 and $V_{fs}/2$, the second MA determines if the input signal is between $V_{fs}/2$ and $V_{fs}/4$ or between $V_{fs}/4$ and 0. The bit representing this MA then is set to 0 or 1, respectively.

Each subsequent MA (and the final comparator) further subdivides the regions in a similar manner, providing all of the necessary quantization levels (Figure 7 shows operation up to the third MA). Because of the way that the bit representing an MA is set (as seen in Figure 7), the output bits from each MA form a Gray code that represents the input signal voltage. In the Gray code, only one bit changes in the code word from one quantization level to another. Determination of the output bits is dependent upon the input signal propagating through the cascaded MA's only. Very high-speed operation is achieved because the response time of the amplifiers is very fast.

Integrating ADC's are another category of converters. They convert the analog input signal amplitude into a time interval that is measured subsequently. The most popular methods within this category of ADC's are the dual slope and charge balancing methods. While these types of ADC's are highly linear and are good at rejecting input noise, they are quite slow.

A relatively new type of ADC is the $\Sigma\Delta$ converter. The first-order $\Sigma\Delta$ converter is the most basic $\Sigma\Delta$ converter (Figure 8). It consists of a $\Sigma\Delta$ modulator, a digital filter, and a decimator. To understand how this converter works, an understanding of oversampling, noise shaping, digital filtering, and decimation is required.

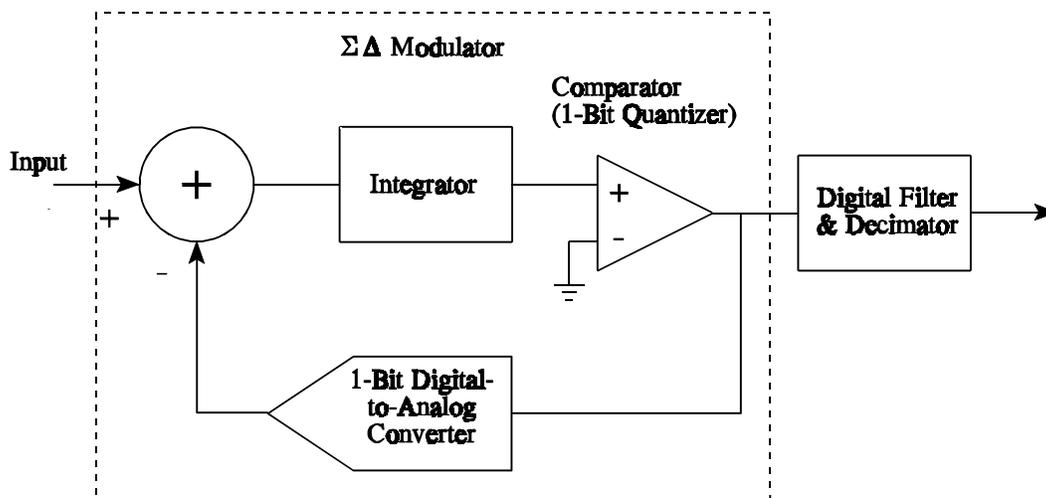


Figure 8. First-order $\Sigma\Delta$ ADC.

The operation of the $\Sigma\Delta$ converter relies upon the effects of oversampling. $\Sigma\Delta$ converters use a very low-resolution quantizer (typically a 1-bit quantizer) and sample at a rate much greater than $2f_{max}$. As discussed previously, sampling at rates faster than $2f_{max}$ provides an improvement in the SNR of the ADC. This occurs because the quantization noise, which is a fixed amount, is spread

out over a greater bandwidth as f_s increases beyond $2f_{\max}$. This improvement in SNR due to oversampling causes the low-resolution quantizer to appear to have a much higher resolution. This apparent higher resolution can be quantified by the ENOB and is found from

$$\text{ENOB} = \frac{\text{SNR} - 1.76\text{dB}}{6.02\text{dB}} . \quad (2)$$

This equation shows that the SNR must increase by approximately 6 dB in order for the ENOB to increase by 1 bit. As shown in (1), the sample rate f_s must be increased to four times greater than $2f_{\max}$ in order for the SNR to increase by approximately 6 dB. Each subsequent increase of 6 dB in the SNR requires a further increase in sampling rate of four times.

As seen from (1) and (2), to achieve an ENOB of 12 bits using a 1-bit quantizer, a sampling rate over 4 million times faster than $2f_{\max}$ is required. This obviously is not practical and shows that $\Sigma\Delta$ converters must use other techniques in addition to oversampling.

The other key component in $\Sigma\Delta$ converters is the integrator that is placed before the 1-bit quantizer. This integrator functions as a low-pass filter for the desired signals occurring at or below f_{\max} and as a high-pass filter for the quantization noise in the ADC. This shapes the quantization noise (which is normally flat across the band from 0 to $f_s/2$) so that very little of this noise occurs in the desired signal's band (0 to f_{\max}). Most of the quantization noise is shifted to frequencies above f_{\max} . This process is called noise shaping and is shown in Figure 9. The results of this noise shaping are that the desired apparent resolution (ENOB) can be achieved with much less oversampling than is predicted by (1) and (2).

The effects of the integrator on the quantization noise of the $\Sigma\Delta$ converter can be seen mathematically by considering a linearized model of the $\Sigma\Delta$ modulator portion of the converter. The block diagram of this model is shown in Figure 10. The quantizer is modeled as a unity gain amplifier with quantization noise added. Looking at this model in the frequency domain, the output of the $\Sigma\Delta$ modulator $Y(s)$ is given as

$$Y(s) = [X(s) - Y(s)] \left(\frac{1}{s} \right) + Q$$

where $X(s)$ is the input signal, $H(s)=1/s$ is the transfer function of the integrator, and Q is the quantization noise. This expression can be rewritten as

$$Y(s) = \frac{X(s)}{s+1} + \frac{Q \cdot s}{s+1} .$$

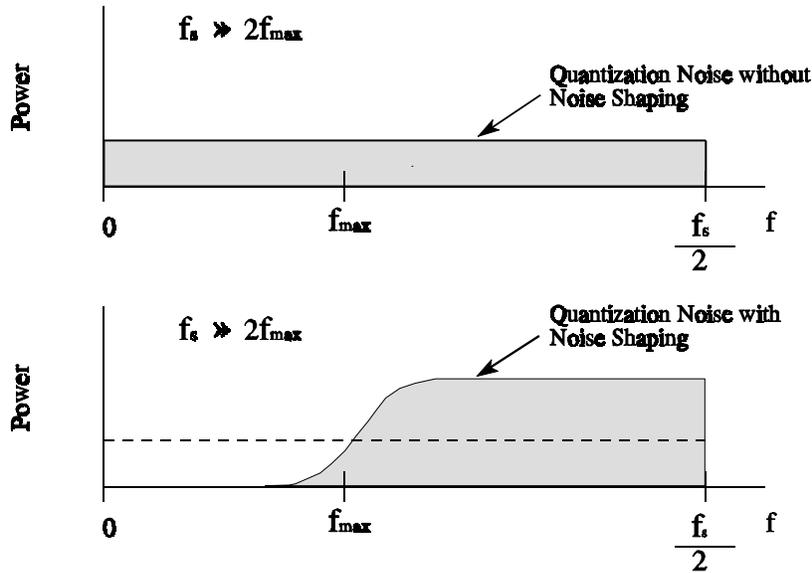


Figure 9. Noise shaping in $\Sigma\Delta$ ADC's.

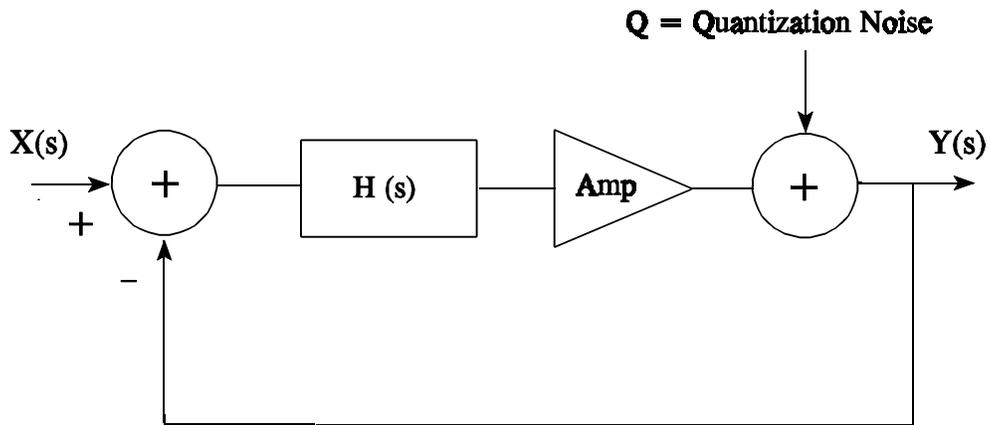


Figure 10. Linearized model of the $\Sigma\Delta$ modulator.

This shows that at low frequencies ($s \ll 1$) the output is primarily a function of the input signal $X(s)$ and not the quantization noise. For high frequencies ($s \gg 1$), $Y(s)$ is primarily a function of the quantization noise [18].

More than one integration and summing stage can be used in the modulator to provide even more noise shaping. Third and even higher-order $\Sigma\Delta$ converters have been designed. (The number of integrators determines the order of the modulator.) Higher-order modulators further decrease the amount of quantization noise in the desired signal's band by placing more of the quantization noise above f_{max} . Therefore, higher-order $\Sigma\Delta$ converters can provide the same apparent

resolution with less oversampling than lower-order $\Sigma\Delta$ converters [19]. $\Sigma\Delta$ modulators higher than second-order provide some difficult design challenges. Instability becomes possible and must be considered carefully in the design.

After the $\Sigma\Delta$ modulator, a digital filter is used. This digital filter is used to 1) filter the quantization noise above f_{\max} and 2) prevent aliasing when the signal is decimated. Decimation is a process of reducing the data rate by resampling a discrete-time signal at a lower rate. Decimation is useful in $\Sigma\Delta$ converters because the oversampling creates a data rate that is much higher than $2f_{\max}$. After filtering the quantization noise, the highest frequency component of the desired signal is only f_{\max} . Therefore, the required sampling rate only needs to be $2f_{\max}$ to fully reconstruct the desired input signal. Decimation is performed by saving only one out of every M samples to reduce the data rate to (or a little higher than) $2f_{\max}$.

Decimation may be combined with digital filtering for more efficient processing. FIR filters can be used to provide both filtering and decimation at the same time. This is true because the FIR filter output needs to be computed for only one out of every M input samples. Conversely, infinite impulse response (IIR) filters cannot be used for decimation because they rely on all of the input samples to produce the proper output. IIR filtering, if desired, can be performed after decimation.

The $\Sigma\Delta$ converters described are designed to operate on baseband signals. A new type of converter, the bandpass $\Sigma\Delta$ converter, shows great potential for radio receiver applications for digitization at the RF or IF. This converter architecture is identical to the traditional $\Sigma\Delta$ converter except that the integrators are replaced by bandpass filters and a bandpass digital filter after the $\Sigma\Delta$ modulator is used. Use of bandpass filters instead of integrators shapes the quantization noise such that it is moved both below and above the desired band of frequencies. This provides a bandpass region of low quantization noise. Bandpass $\Sigma\Delta$ converters are currently a very promising research and development topic.

The $\Sigma\Delta$ converter has a couple of advantages over the more traditional types of ADC's. Because of the high sampling rate, the attenuation requirements on the anti-aliasing filter can be lessened. Additionally, an improvement in the linearity of the ADC results in an improved SFDR. This advantage results from using a 1-bit or other low-resolution quantizer. One disadvantage of $\Sigma\Delta$ converters that are currently available is that they typically are limited to signal bandwidths below 150 kHz (for a 12-bit ENOB).

The design of high-speed ADC's usually incorporates both a sample-and-hold amplifier (SHA) and a quantizer. Many ADC's provide a SHA as an integral part of the ADC. External SHA's also can be used with ADC's but this requires careful design considerations of the SHA and ADC specifications, in addition to timing and interface issues between the external SHA and the ADC.

The purpose of the SHA in ADC applications is to keep the input signal constant during the ADC conversion. While there are many different SHA implementations, all SHA's consist of four basic components: an input amplifier, capacitor, output buffer, and switching circuit. An example SHA showing the basic components is given in Figure 11. The input amplifier provides a high input impedance to the input signal and supplies the necessary current to charge the hold

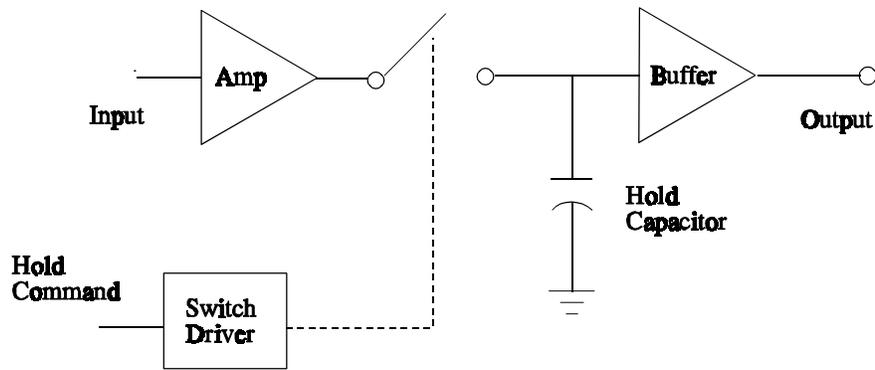


Figure 11. Example showing the basic components of a sample-and-hold circuit.

capacitor. When the switch closes, the SHA operates in the track mode and the voltage on the hold capacitor follows the input signal. When the switch opens, the SHA operates in the hold mode. Ideally, the voltage on the hold capacitor remains at its value before the switch opened. Since it has a high input impedance, the output buffer prevents the hold capacitor from discharging significantly. The hold command controls the operation of the switch and determines when the SHA is in the track or hold mode [20].

Sometimes SHA's are called track-and-hold amplifiers (THA's). The name used depends on how the device is used. When the device spends most of its time in the hold mode and just a short time in track mode (enough time to take a sample of the input), the device is called a SHA. When the device spends only a short time in the hold mode and most of its time in the track mode, it is called a THA [21].

For radio receiver applications, typically high-speed ADC's are required, especially for direct digitization of the RF or digitization of wideband IF. Because of this, successive-approximation, subranging, flash, and bandpass $\Sigma \Delta$ ADC's are the most likely types of ADC's to be used for these applications.

2.5 ADC Performance vs. Sampling Rate

The performance of ADC's continues to improve at a rapid rate. For radio receiver applications using digitization at the RF or IF, ADC's with both high sampling rates and high performance are desired. Unfortunately, there is a tradeoff between these two requirements. As a general trend, although not always true, the higher the performance of the ADC, the lower its maximum sampling rate will be. The goal of direct digitization at the RF in radio receivers at increasingly higher frequencies and wider bandwidths is one of the forces driving the development of higher-performance, faster ADC's. Digital sampling oscilloscopes are another example of applications that encourage the development of higher-performance, faster ADC's.

Interleaving is a common technique used to increase the sampling rate beyond the capability of a single ADC. In this technique, multiple ADC's of the same type are staggered in time to

achieve higher sampling rates. Each ADC is offset in time from the preceding ADC. For uniform sample spacing, this offset is determined by dividing the time interval between samples of a single ADC by the total number of ADC's to be used. The interleaving technique is used extensively in digital sampling oscilloscopes.

Examples of current high-speed ADC technology showing maximum sampling rates for various ADC resolutions are given in Table 2. The low-resolution (6- or 8-bit), high sampling rate ADC's are typically implemented as flash ADC's and therefore are limited in SFDR.

When selecting an ADC for a specific radio receiver application, in addition to the sampling rate, one must consider critical specifications that characterize the ADC performance such as the SNR, SFDR, NPR, and full-power analog input bandwidth. In certain applications such as channelized PCS and mobile cellular systems, instead of digitizing the entire band with a single high-speed ADC, parallel ADC's used to digitize narrower bandwidths are often practical ADC architectures. In this case, ADC's with better performance can be used since the demands of a high sampling rate are relieved.

Table 2. Examples of Current High-Speed ADC Technology

Resolution (Number of Bits)	Sampling Rate (Msamples/s)	Manufacturer
6	4000	Rockwell International
8	1000	Signal Processing Technology
8	2000*	Hewlett-Packard
8	3000	**
10	70	Pentek
12	50	Hughes Aircraft
12	100	**
14	24	Hughes Aircraft
18	10	Hewlett-Packard

*8000 Msamples/s with interleaving.

**Device in development; work is being sponsored by the Advanced Research Projects Agency (ARPA) of the U.S. Department of Defense.