NTIA Report 96-328

RF and IF Digitization in Radio Receivers: Theory, Concepts, and Examples

J.A. Wepman
J.R. Hoffman

report series

U.S. DEPARTMENT OF COMMERCE • National Telecommunications and Information Administration

NTIA Report 96-328

RF and IF Digitization in Radio Receivers: Theory, Concepts, and Examples

J.A. Wepman
J.R. Hoffman

U.S. DEPARTMENT OF COMMERCE

March 1996
Reissued February 2014

DISCLAIMER

Certain commercial equipment, components, instruments, and materials are identified in this paper to provide some examples of current technology. In no case does such identification imply recommendation or endorsement by the National Telecommunications and Information Administration, nor does it imply that the material or equipment identified is necessarily the best available for the purpose. Furthermore, examples of technology identified in this paper are not intended to be all-inclusive; they represent only a sampling of what is available.

CONTENTS

RF AND IF DIGITIZATION IN RADIO RECEIVERS: THEORY, CONCEPTS, AND EXAMPLES

Jeffery A. Wepman and J. Randy Hoffman[1]

Hardware development of analog-to-digital converters (ADCs) and digital signal processors, including specialized integrated circuits, has advanced rapidly within the last few years. These advances have paved the way for development of radio receivers using digitization at the IF and in some cases at the RF. Applications for these receivers are expected to increase rapidly in areas such as cellular mobile, satellite, and personal communications services (PCS) systems. The constraints placed on these receivers due to hardware limitations of these devices are investigated in this paper. Some examples of state-of-the-art ADCs, signal processors, and specialized integrated circuits are listed. Various quantization techniques, nonlinear compression devices, postdigitization algorithms for improving dynamic range, sampling downconverters, and specialized integrated circuits are discussed as they are expected to be useful in the development of these types of receivers. Several examples of radio receivers employing digitization at the IF and RF are also presented.

Key words: analog-to-digital converters; automatic gain control devices; digital signal processors; digitization; logarithmic amplifiers; quantization; radio receivers; sampling downconverters; signal-to-noise ratio; spurious free dynamic range

1 INTRODUCTION

As advances in technology provide increasingly faster and less expensive digital hardware, more of the traditionally analog functions of a radio receiver will be replaced with software or digital hardware. The final goal for radio receiver design is to directly digitize the RF signal at the output of the receive antenna and therefore implement all receiver functions in either digital hardware or software. The trend in receiver design is evolving toward this goal by incorporating digitization closer and closer to the receive antenna for systems at increasingly higher frequencies and wider bandwidths. Analog RF front-ends with digitization at either baseband or IF are currently being implemented in many arenas.

There is keen interest in replacing analog hardware with digital signal processing in radio receivers for several reasons. One reason is the potential for the reduction in product development time since changes can be implemented in software instead of altering the hardware [1]. Digital technology can offer a more ideal performance for implementing signal-processing functions. The repeatability and temperature stability can be substantially better. Functions that

[1] Jeffrey A. Wepman is with and J. Randy Hoffman was previously with the Institute for Telecommunication Sciences, National Telecommunications and Information Administration, U.S. Department of Commerce, Boulder, CO 80305.

are not implementable in analog hardware can be implemented in software. An example is the design of finite impulse response (FIR) filters that simultaneously can achieve sharp rolloff and linear phase. Another advantage is that digitally implemented signal-processing functions do not require the tuning or "tweaking" typically required in an analog implementation to achieve the desired performance [2]. (Proper operation of digital processing circuitry does require some level of synchronization, however.) Cost-effective multipurpose radios can be designed to allow reception of different modulation types and bandwidths simply by changing the software that controls the radio. The final benefit is the cost savings in implementing the receiver.

As radio receiver design evolves so that direct digitization of the RF input signal becomes more commonplace, these systems will have to go through the process of spectrum certification before they can be implemented and used by Government agencies. The process of spectrum certification includes an electromagnetic compatibility (EMC) analysis. Development of EMC analysis methodologies and a spectrum certification process for radio receivers using digitization at the RF is required to help the National Telecommunications and Information Administration (NTIA) manage the Federal radio spectrum for Government agencies in the most efficient manner possible.

Methods for analyzing EMC in traditional receivers (such as the superheterodyne) are well established. EMC analysis of these new receivers that utilize digitization of the RF signal at the front-end may be different. Information currently requested by NTIA for receiver equipment characteristics that is used in the EMC analyses may no longer be relevant for these new types of receivers. Detailed knowledge of how these receivers operate is therefore required to help develop appropriate methods of EMC analysis. This report provides more information on these types of radio receivers. In Section 2, we discuss analog-to-digital converters (ADCs), one of the important components needed in radio receivers using digitization at the RF or IF. The requirements, practical limitations, and potential problems for ADCs are presented. Section 3 includes the signal-processing requirements and limitations for radio receivers that digitize at the RF or IF. Some devices and techniques that may be useful for receivers employing direct digitization of the RF are described in Section 4. These include 1) methods of nonuniform quantization, 2) nonlinear amplitude compression devices, 3) algorithms for improving dynamic range, 4) sampling downconverters, and 5) specialized integrated circuits. Section 5 presents some examples of radios that digitize at the RF or IF. Section 6 provides a brief summary of this investigation and some recommendations for further work in this area.

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 1(a) shows the spectrum of the original analog signal $F(f)$. Figure 1(b) 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 1(b) 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 1(c). 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 1(d) shows the spectrum overlap that occurs when sampling at rates less than $2f_{max}$.

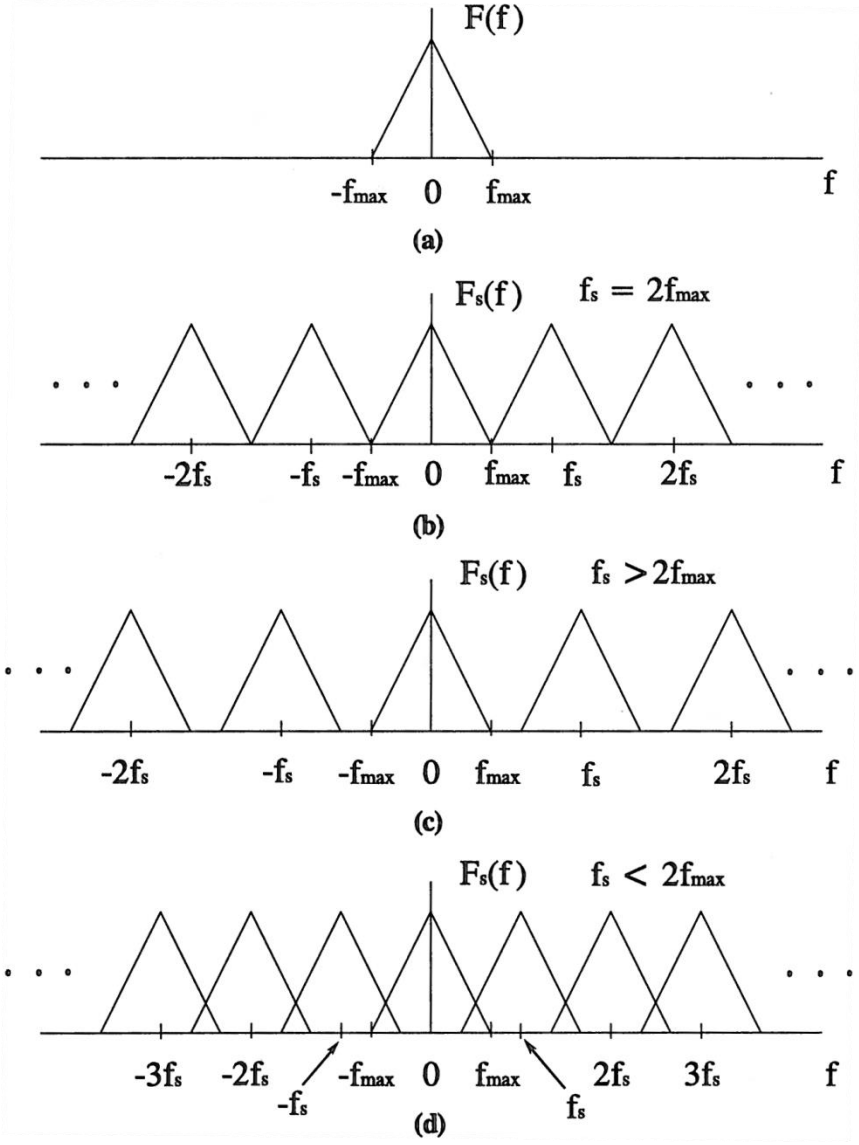

Figure 1. Spectrum of: (a) a bandlimited continuous-time analog signal; (b) the signal sampled at $f_s = 2f_{max}$; (c) the signal sampled at $f_s > 2f_{max}$; and (d) the signal sampled at $f_s < 2f_{max}$.

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 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 2(a) 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 2(b). Note that spectrum overlap has occurred here (i.e., the spectrum of the undesired signal occurs within the spectrum of the desired signal). This causes distortion in the reconstructed desired signal.

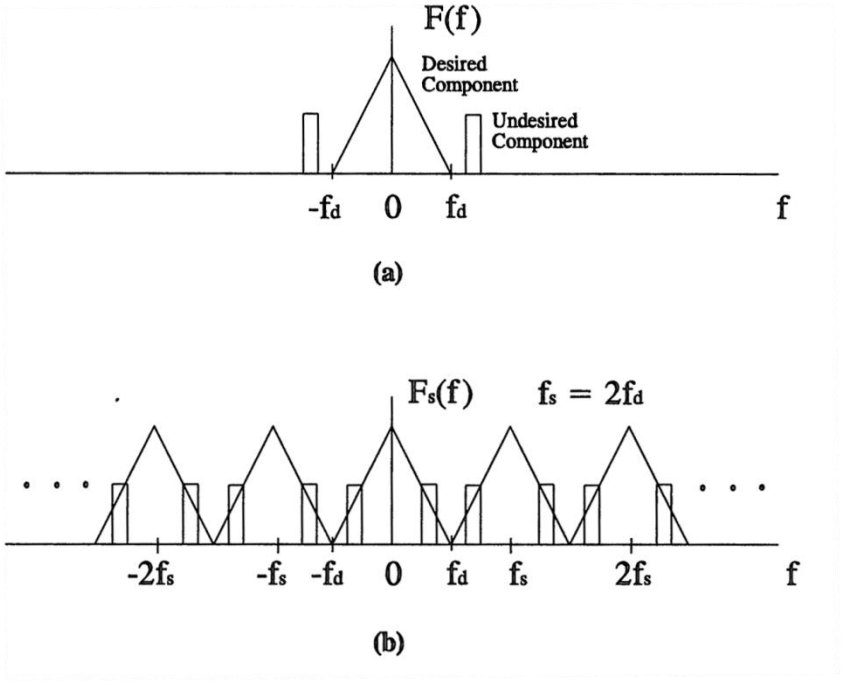

Figure 2. Spectrum of: (a) a continuous-time analog signal with a desired and undesired component and (b) the signal sampled at $f_s=2f_d$

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 ADCs 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 ADCs 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]

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

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

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 ADCs 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 ADCs are designed to operate on signals with maximum frequencies up to one-half the sampling rate. In other words, conventional ADCs typically are not suitable for bandpass sampling applications where the maximum input frequencies are greater than the sampling rate. Furthermore, for conventional ADCs, many manufacturers provide specifications only at frequencies well below one-half the maximum sampling rate. In general, performance of ADCs typically degrades with increased input frequency. Therefore, in using ADCs 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 ADCs 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.

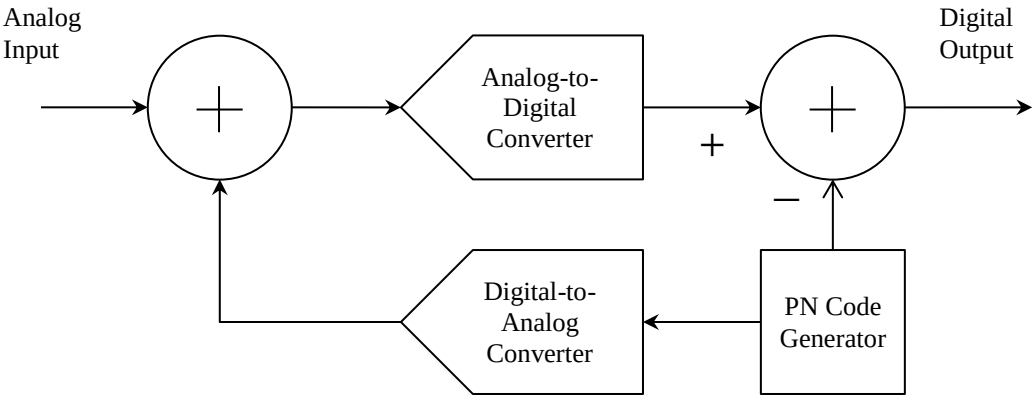

Figure 3. Block diagram of subtractive dithering.

Commercially available ADCs 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 ADCs 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 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_{rn} can be computed as the thermal noise power in the given receiver bandwidth (BW) plus the receiver noise figure (NF). This is given as

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

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 ADCs 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 ADCs 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 ADCs 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 ADCs 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

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

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 ($P_{qn} = q^2/12R$), 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.

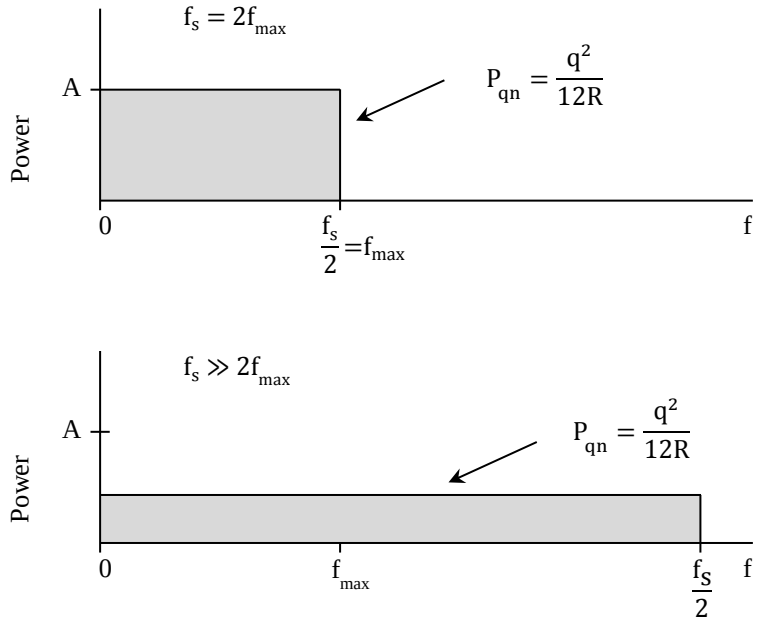

Figure 4. Frequency-spreading of quantization noise power due to oversampling.

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

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

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 ADCs

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.[2]

A specification sometimes used for real ADCs 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 ADCs. 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.

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 ADCs 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

[2] The 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).

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_{p-p}), the typical SFDR specification is 81 dB while the SNR specification is 66.5 dB (from −40 - +85 °C) [12].

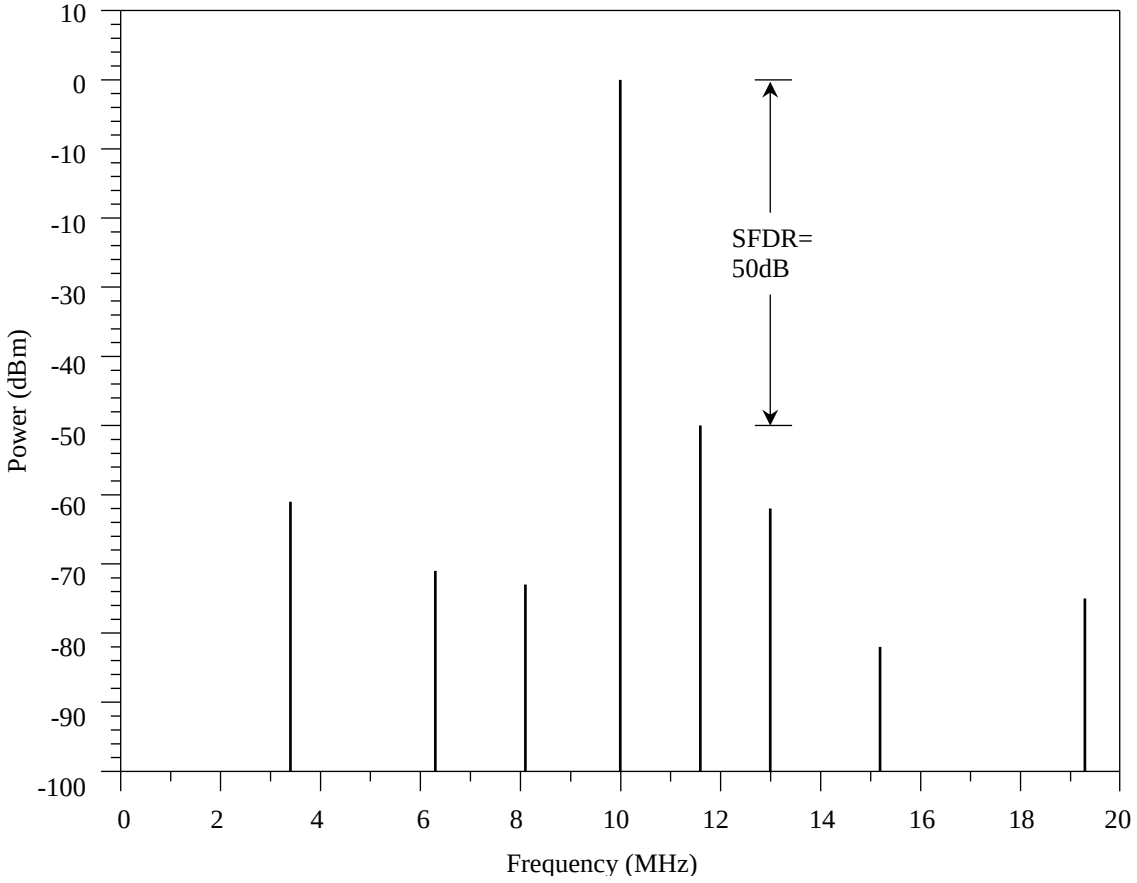

Figure 5. Example ADC output spectrum showing the spurious free dynamic range.

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 ADCs. 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 (ΣΔ) converters, the maximum SFDR occurs at a full-scale input level. In other types of practical ADCs, 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 ADCs. Prediction of the SFDR for practical ADCs 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 ADCs 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 ADCs 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	$\frac{\text{Power Spectral Density of Noise}^* \text{ Outside Freq. Band of Notch Filter}}{\text{Power Spectral Density of Noise Inside Freq. Band of Notch Filter}}$
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 ADCs in general and do not occur in ΣΔ converters.

2.4 ADC Conversion Methods

Many methods for implementing ADCs 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 ADCs belong to the feedback class of ADCs. 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 ADCs of more than 8 bits typically are not available commercially. Linearity is a problem in flash ADCs as observed in degraded SFDR performance.

One technique used to implement high-speed ADCs combines two separate B-bit ADCs (usually flash ADCs) 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 ADCs 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 ADCs 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 ADCs. Subranging ADCs 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 6(a) while the ramp input is shown in Figure 6(b). The output of the first ADC is a quantized version of the ramp as shown in Figure 6(c). 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 6(d)) 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.)

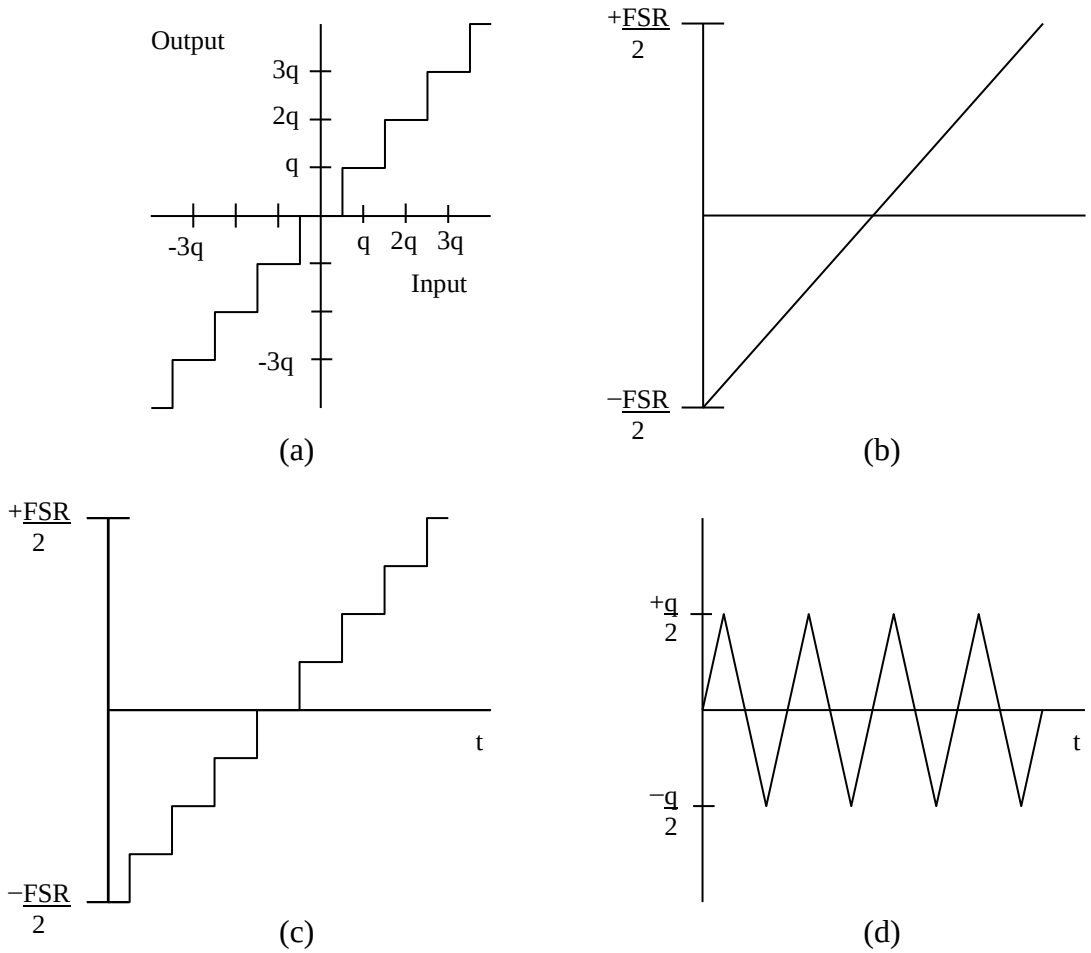

Figure 6. (a) Ideal ADC transfer function; (b) input ramp signal; (c) output (quantized) ramp signal; and (d) repetitive ramp difference signal.

Subranging ADCs 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 ADCs. While the internal ADCs within a subranging ADC have traditionally been flash ADCs, other types of ADCs 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 ADCs.

The cascaded magnitude amplifier (MA) ADC consists of $B-1$ MAs 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 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 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 MAs only. Very high-speed operation is achieved because the response time of the amplifiers is very fast.

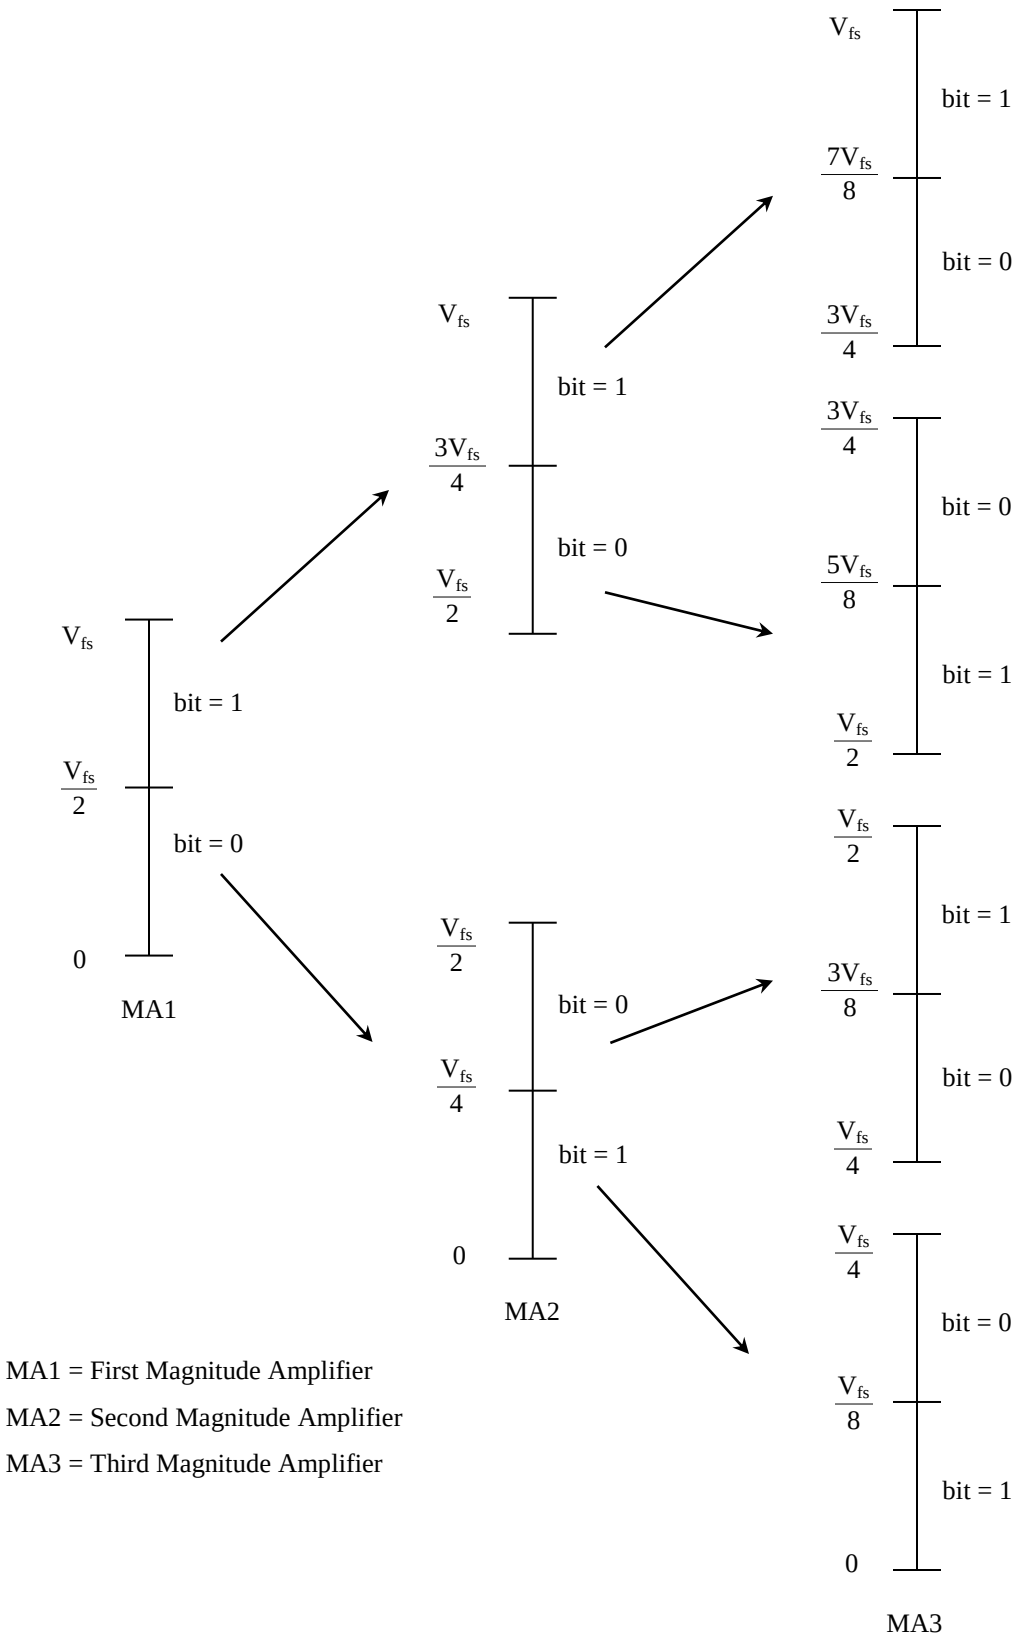

Figure 7. Operation of the cascaded magnitude amplifier ADC.

Integrating ADCs 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 ADCs are the dual slope and charge balancing methods. While these types of ADCs are highly linear and are good at rejecting input noise, they are quite slow.

A relatively new type of ADC is the ΣΔ converter. The first-order ΣΔ converter is the most basic ΣΔ converter (Figure 8). It consists of a ΣΔ 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.

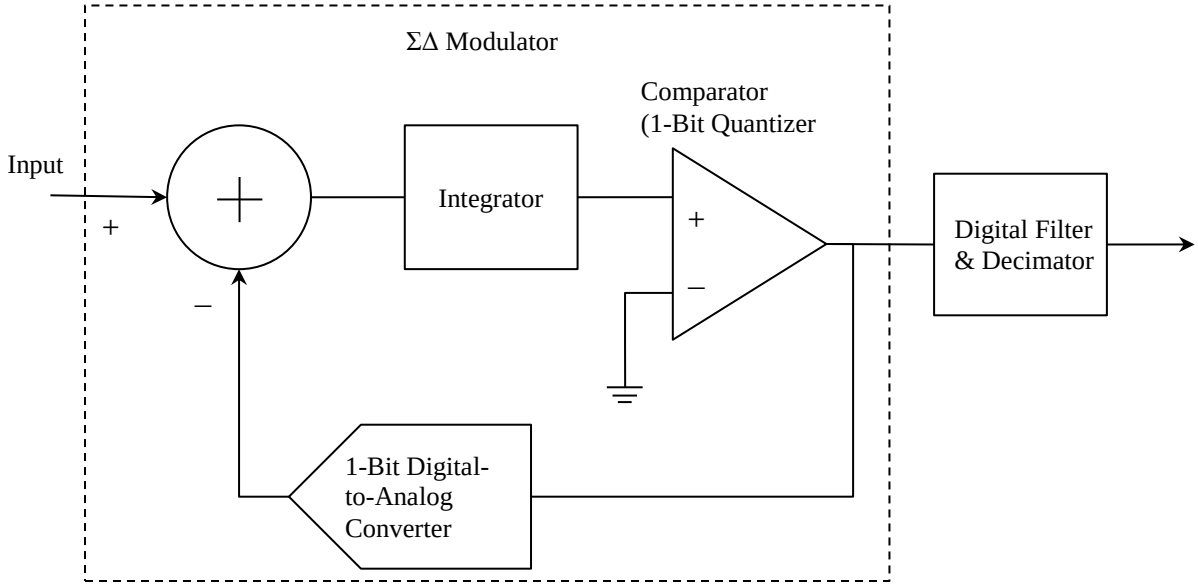

Figure 8. First-order ΣΔ ADC.

The operation of the ΣΔ converter relies upon the effects of oversampling. ΣΔ 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

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

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 ΣΔ converters must use other techniques in addition to oversampling.

The other key component in ΣΔ 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).

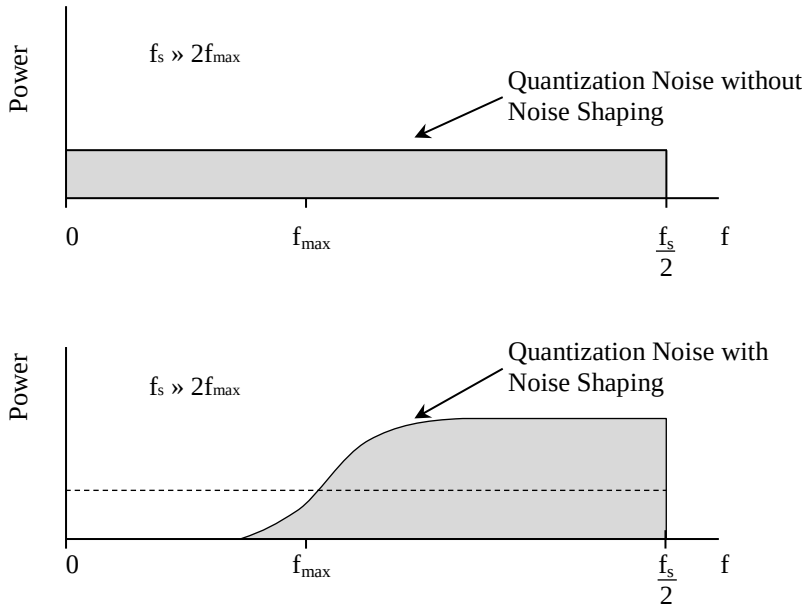

Figure 9. Noise shaping in ΣΔ ADCs.

The effects of the integrator on the quantization noise of the ΣΔ converter can be seen mathematically by considering a linearized model of the ΣΔ 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 ΣΔ 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

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

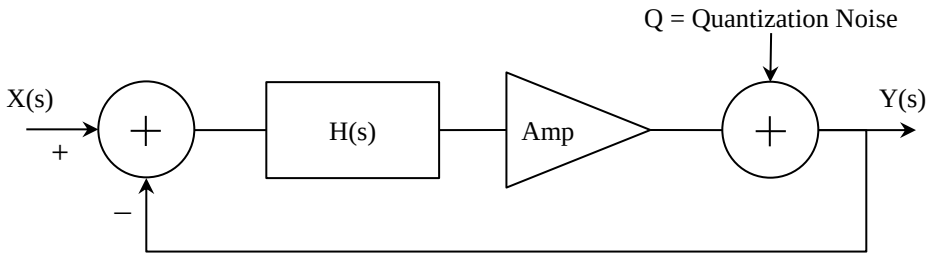

Figure 10. Linearized model of the ΣΔ modulator

This shows that at low frequencies (s « 1) the output is primarily a function of the input signal *X(s)* and not the quantization noise. For high frequencies (s » 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 ΣΔ 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 ΣΔ converters can provide the same apparent resolution with less oversampling than lower-order ΣΔ converters [19]. ΣΔ modulators higher than second-order provide some difficult design challenges. Instability becomes possible and must be considered carefully in the design.

After the ΣΔ 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 ΣΔ 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 ΣΔ converters described are designed to operate on baseband signals. A new type of converter, the bandpass ΣΔ converter, shows great potential for radio receiver applications for digitization at the RF or IF. This converter architecture is identical to the traditional ΣΔ converter except that the integrators are replaced by bandpass filters and a bandpass digital filter after the ΣΔ 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 ΣΔ converters are currently a very promising research and development topic.

The ΣΔ converter has a couple of advantages over the more traditional types of ADCs. 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 ΣΔ 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 ADCs usually incorporates both a sample-and-hold amplifier (SHA) and a quantizer. Many ADCs provide a SHA as an integral part of the ADC. External SHAs also can be used with ADCs 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 SHAs 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 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].

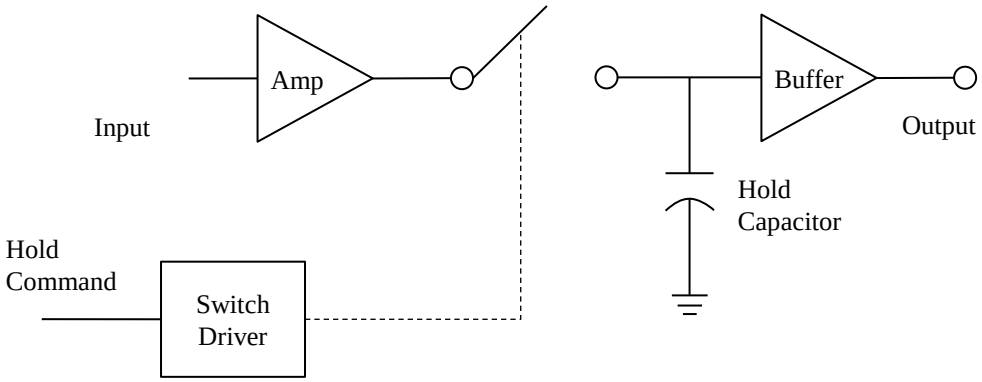

Figure 11. Example showing the basic components of a sample-and-hold circuit.

Sometimes SHAs are called track-and-hold amplifiers (THAs). 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 ADCs are required, especially for direct digitization of the RF or digitization of wideband IF. Because of this, successive-approximation, subranging, flash, and bandpass ΣΔ ADCs are the most likely types of ADCs to be used for these applications.

2.5 ADC Performance vs. Sampling Rate

The performance of ADCs continues to improve at a rapid rate. For radio receiver applications using digitization at the RF or IF, ADCs 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 ADCs. Digital sampling oscilloscopes are another example of applications that encourage the development of higher-performance, faster ADCs.

Interleaving is a common technique used to increase the sampling rate beyond the capability of a single ADC. In this technique, multiple ADCs 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 ADCs 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 ADCs are typically implemented as flash ADCs 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 ADCs used to digitize narrower bandwidths are often practical ADC architectures. In this case, ADCs 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

Resolution (Number of Bits)	Sampling Rate (Msamples/s)	Manufacturer
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.

3 DIGITAL SIGNAL-PROCESSING REQUIREMENTS AND LIMITATIONS

Besides ADCs, digital signal processing is another key element in radios using digitization of the RF or IF. The amount of time required for signal processing is of critical importance in radio receiver applications. The required processing time is a function of the received signal bandwidth, the speed of the processor, and the number and complexity of the algorithms required to perform the needed radio receiver operations. These operations are application-specific and may include some or all of the following: downconversion, filtering, multiple access processing, demultiplexing, frequency despreading, demodulation, synchronization, channel decoding, decryption, and source decoding [22], [23]. Because of the wide variety of algorithms possible in radio receiver applications, digital signal-processing limitations are more difficult to discuss than ADCs.

The intent of this section is to discuss the requirements and limitations of digital signal processing. It is not intended to provide a detailed presentation of the wide variety of digital signal-processing techniques and algorithms that are available. Many books and papers are available to provide detailed information on digital signal-processing techniques and algorithms. One fundamental book on digital signal processing is [24].

3.1 Processors

Many different processors are available to provide digital signal processing. These processors vary substantially in speed of operation, physical size, and cost. Speed is usually a critical requirement in selecting a processor. Other factors including dynamic range, arithmetic precision, cost, and size are also important considerations when choosing a processor.

A common method used to increase total processing speed beyond that of a single processor is to employ multiple processors operating in parallel. Assuming a given processor, by putting more and more of these processors together and operating them in parallel, higher and higher processing speeds can be achieved. This, of course, also increases power consumption, size, and cost.

Many radio receiver applications require processors with small physical size and relatively low cost. For these cases, single chip processors are the preferred choice. Single chip processors can be general purpose microprocessors (such as the Intel 80486), digital signal processors (such as the Texas Instruments TMS320C40), or specialized integrated circuits for dedicated processing tasks (such as the Harris HSP50016 Digital Downconverter). Some specialized radio receiver applications may not be bound by stringent physical size and cost limitations. Therefore, processors of all types are considered in this report, ranging from single chip general purpose microprocessors to supercomputers.

Computations in digital signal processing can be performed using fixed-point arithmetic or floating-point arithmetic, although many of the computations require floating-point arithmetic. The advantage of floating-point arithmetic over fixed-point arithmetic is that it permits the use of numbers with a much greater dynamic range. This is important in many digital signal-processing operations.

In fixed-point arithmetic, the position of the decimal point in the register where each operand is stored always is assumed to be the same. In floating-point arithmetic, each operand is represented by a number stored in a register representing a fraction or integer. A number stored in a second register specifies the position of the decimal point of the number stored in the first register. Some processors do not have floating-point hardware and require floating-point operation to be implemented in software. Software implementation of floating-point arithmetic is typically much slower than hardware implementation.

Because floating-point operations are so important in digital signal processing, the speed of processors is often specified in terms of millions of floating-point operations per second (MFLOPS). This parameter allows comparison of the processing speed of different processors and also allows determination of the time required to execute certain algorithms.

Many different benchmarks (such as the SPEC benchmarks, Whetstone, Dhrystone, and Linpack) are used to compare speeds between processors. Each benchmark provides a number indicating the relative speed of processing based on testing varying tasks. Results from the application of a benchmark to different processors can be compared. However, results between different benchmarks, in general, should not be compared. While these benchmarks are useful for comparing processor performance, the parameter chosen to compare processing speeds between processors in this report is the theoretical peak MFLOPS. This parameter was chosen due to its ease of availability for virtually all floating-point processors, its lack of dependence on specific benchmarking algorithms, and its relevancy to dedicated applications used in implementation of a radio receiver. The theoretical peak MFLOPS parameter gives the maximum possible speed of performance for the processor. It is found by computing the number of floating-point additions and multiplications (using the processor's full precision) that can be performed during a given time interval [25].

Some examples of the processing speed of various types of processors ranging from single chip processors to supercomputers are presented in Table 3. This table only gives a sampling of the range of capabilities that exist in digital signal processing. Many other processors, with varying capabilities, either exist or have been proposed. In addition, new developments with increasing capabilities are announced all the time. An extensive listing of processing speeds for many different computers is found in [25].

In certain situations, especially in the high-throughput case, overall processing performance is not limited by the processor speed but by the maximum data transfer rates of the peripheral components such as memory or I/O (input/output) ports. The inclusion of these factors in platform evaluation should not be ignored when choosing a processor.

Table 3. Examples of Processing Technology

Processing Speed*	Number of Processors	Platform	Manufacturer and Model
50 MFLOPS	1	DSP Chip	Texas Instruments TMS320C40

Processing Speed*	Number of Processors	Platform	Manufacturer and Model
120 MFLOPS	1	DSP Chip	Analog Devices ADSP-21060/62
400 MFLOPS	8	VME Board	Pentek 4285
800 MFLOPS	16	Computer Workstation	SUN Sparc 2000
6.48 GFLOPS	4	Supercomputer	Convex C4/XA-4
32 GFLOPS	4	Supercomputer	Hitachi S-3800/480
184 GFLOPS	3680	Massively Parallel Computer	Intel Paragon XPS140
236 GFLOPS	140	Massively Parallel Computer	National Aerospace Laboratory Numerical Wind Tunnel (Japan)

* Theoretical peak processing speed.

3.2 Real-Time Operation

For most radio receiver applications, real-time operation is important. In many types of processing, such as computing Fast Fourier Transforms (FFTs), the data is partitioned into blocks of a finite length. Processing is performed on the entire block of data. In this block type processing, assuming a single processor, real-time operation essentially means that all processing on a given block of data (including any required data transfers) is completed before all of the next block of data to be processed is captured. This concept is illustrated in Figure 12 [26].

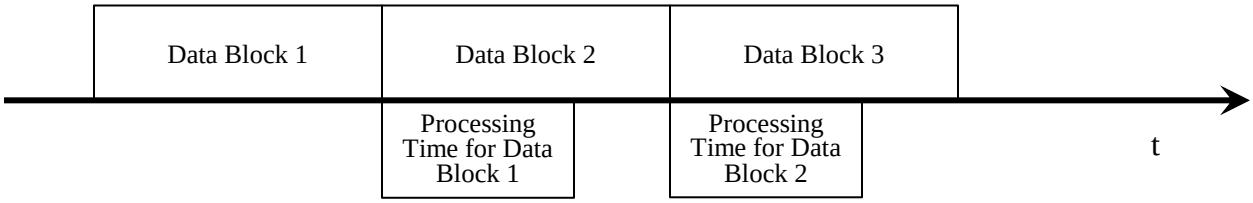

Figure 12. Real-time processing for block data using a single processor.

If the processing time (including any required data transfers) is longer than the time required to capture all of the next block of data (again assuming a single processor), data collection must be stopped until the processing is completed. At that time data collection may resume. Under these conditions some of the input data is missed and is not processed. This is an example of processing that does not take place in real time. Depending upon the application and the amount of data lost, this may not be acceptable.

This problem can be alleviated by using two or more processors operating cooperatively. This general technique is called multiprocessing and is used frequently. To illustrate how multiprocessing can be used to speed up overall data throughput, consider the previous example of data partitioned into data blocks but with two processors available instead of one. As shown in Figure 13, processor 1 operates on one block of data and processor 2 operates on the next block of data. The processors continue to operate on alternating data blocks. The processed data output is obtained by switching back and forth between the outputs of the two processors. This is sometimes called a "ping-pong" technique. Using this technique, the processing time of each processor can take longer than the time to capture the next block of data and still provide real-time output. The processing time cannot exceed the time to capture the next two blocks of data, however. This technique can be extended to more than two processors to achieve even faster overall data throughput.

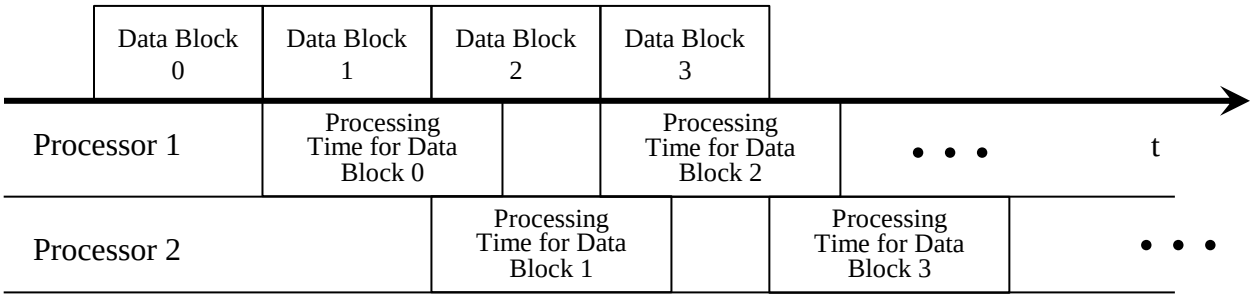

Figure 13. Real-time processing for block data using two processors.

3.3 Algorithms

Providing a general discussion on algorithms used for implementing radio receiver functions is difficult. This is due to the wide variety of types of receivers as well as the various ways of implementing the required receiver operations for each type of receiver. In short, algorithms are highly application-specific. The details of the many potential algorithms used in radio receivers are beyond the scope of this report. It is beneficial, however, to look at an example algorithm to observe the methodology in determining algorithm complexity vs. the potential for real-time operation. This type of assessment is crucial for radio receivers that use digitization at the RF or IF.

The FFT is an example of an algorithm frequently used in digital signal-processing and radio receiver applications. The FFT transforms time-domain samples of received signals into a set of frequency-domain samples, allowing operations on the received signals to be performed directly in the frequency domain. These received signals are typically bandpass signals when digitization occurs at the RF or IF. These bandpass signals may be digitized using either sampling at twice the maximum frequency, bandpass sampling, or oversampling. The FFT also can be applied to signals that have been downconverted to baseband but the signal must be split into co-phase and quadrature-phase components before digitization unless coherent downconversion has been used.

Regardless of the sampling method employed, the resolution of the transformed signal in the frequency domain is a function of both the time spacing between the samples of the signal Δt in

the time domain and the number of samples N used in the computation of the FFT. The frequency spacing between samples in the frequency domain is then given as

$$\Delta f = \frac{1}{N\Delta t}(Hz) .$$

The maximum frequency of the spectrum is then

$$f_{\max} = \frac{N}{2}\Delta f = \frac{1}{2\Delta t}$$

since there are $N/2$ samples in the FFT computed from N real-valued time-domain samples. (Actually, there are $(N/2)+1$ samples in the FFT ranging from DC to f_{max} if both DC and f_{max} are included.) For a fixed N-point FFT (i.e., an FFT computed from N real-valued time-domain samples), the frequency spacing between the frequency domain samples Δf must be changed in order to change the maximum frequency of the spectrum. That requires changing the time spacing Δt between the time-domain samples. By decreasing Δt and holding N constant, the time duration of a block of N samples is reduced and the maximum frequency of the spectrum is thereby increased. Therefore, for fixed values of N, the higher the maximum frequency desired, the shorter the duration of the block of N samples must be. With a fixed number of samples N, a given processor, and a given FFT algorithm,[3] the processing time to compute an N-point FFT is fixed. For real-time operation, this computation of the FFT (including any other required processing such as windowing and any required data transfers) must be performed within the time taken to capture all N samples of the current data block assuming that a single processor is used. A parameter called real-time bandwidth then can be defined as the maximum frequency that can be processed in real time.

To achieve real-time processing, careful consideration of processor speed, the signal bandwidth (data rate), the number of computations required in implementing the signal-processing algorithms, and the speed of any necessary data transfers is required. An example showing how to estimate the amount of processing power required for real-time analysis is given below. For this example, the simplified case of looking at the time required to compute an FFT is investigated. While this example shows the methodology used to determine a required processing speed, the processing required for radio receiver applications normally would involve much more than computing a single FFT. In this simplified example, it is assumed that an input signal is sampled at a fixed rate and that the only processing performed on the sampled input signal is the FFT. No other processing (such as windowing or averaging) is performed. Data transfer time is also neglected.

Assuming a bandlimited input signal with a maximum frequency of 5 MHz, the $2f_{max}$ sampling rate would be 10 Msamples/s. For this sampling rate, the time between samples Δt is 100 ns. Assume that FFTs are computed from blocks of N = 1024 samples. Therefore, it takes $N\Delta t$ = 102.4 μs to capture a block of data. The number of floating-point operations (actually multiplications) required[3] to compute an N-point FFT is estimated as $N \log_2 N$. Therefore,

[3] There are many different FFT algorithms available requiring different numbers of floating-point operations. $N \log_2 N$ commonly is used as an approximation for the number of floating-point operations required in computing an FFT.

roughly 10,240 floating-point operations are required for the 1024-point FFT. In order to achieve real-time processing in the single processor case, the FFT must be computed within the time period required to capture a block of data (102.4 μs). The minimum required processing speed is then found by

$$\frac{\text{Number of floating point operations required}}{\text{Time to capture a block of data}} = \text{Minimum processor speed (FLOPS)}.$$

In this simplified case, the minimum processing speed is 100 MFLOPS. One can then compare the required processing speed to the processing speeds available for different types of processors such as those listed in Table 3.

4 POTENTIAL DEVICES AND METHODS USEFUL IN RADIOS EMPLOYING RF AND IF DIGITIZATION

This section describes some methods and devices that may be useful in implementing radio receivers where digitization occurs at the RF or IF. These methods and devices probably will be considered in the design of future radio receivers as the digitization of received signals progresses toward the receive antenna. Various quantization techniques are discussed since the authors believe that this area has great potential for use in radio receiver applications. Nonlinear compression devices are discussed, including log amplifiers and automatic gain control devices. Postdigitization algorithms for improving the SFDR provide extended dynamic range capability for presently available ADCs. Sampling downconverter technology is introduced. This technology is based on the theory of bandpass sampling and may be coupled with ADC technology to provide improved ADC performance for bandpass sampling applications. Finally, specialized integrated circuits for digital signal-processing tasks required in radio receiver applications are mentioned. An example of this type of technology, the Harris HSP50016 digital downconverter, is presented.

4.1 Quantization Techniques[4]

Signals of widely varying amplitudes are present in radio receiver applications. Digitization of these signals requires ADCs with large dynamic range. Most commercially available ADCs use uniform quantization (for which quantization levels are evenly spaced) and therefore require high resolution (a large number of bits) to achieve a large dynamic range. The problem for many RF applications is that ADCs using uniform quantization cannot easily meet both the large dynamic range and high sampling rates needed to digitize wide bandwidths at the RF or IF. In addition, another problem is that uniform quantizers exhibit an SNR that varies with desired signal amplitude; the SNR decreases as the input signal power to the quantizer decreases. (Ideally, a large SNR would be maintained over a large variation in input signal power.) In an effort to solve these problems, several quantization schemes other than uniform quantization have been implemented.

These different quantization techniques include μ-law (and A-law) quantization, adaptive quantization, and differential quantization [28]. While these techniques are typically associated with speech signal processing applications, they may be useful for radio receiver applications [27].

[4] When this report was originally published, in this section the authors reproduced with permission some material from L.R. Rabiner and R.W. Schafer, *Digital Processing of Speech Signals*, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1978, cited in this report as [27]. This textbook served as a fundamental reference text for both teaching the fundamentals of digital speech processing and designing practical systems for over 30 years. In 2010, Rabiner and Shafer published a new textbook on the same topic that incorporates advances in the technology of digital speech processing that occurred in the interval and supersedes the 1978 textbook. The authors have therefore revised this section to provide just an overview of the material. Readers who want more current in-depth information on this topic should consult L.R. Rabiner and R.W. Schafer, *Theory and Applications of Digital Speech Processing*, Englewood Cliffs, NJ: Prentice-Hall, Inc., 2010.

In μ-law quantization, a nonlinear compression (an approximation to logarithmic compression) is used on the input signal before uniform quantization[28]. This increases the dynamic range of the overall quantization. In addition, while for uniform quantizers the SNR varies considerably with small changes in input signal amplitude, the SNR remains fairly constant over a wide range of input signal amplitudes for μ-law quantizers. As an example of this, the SNR changes by 25 dB for input signal amplitude changes of 25 dB in a uniform quantizer. With appropriate scaling of the input signal amplitude, the SNR changes by less than 2 dB in a μ-law quantizer ($\mu = 500$) for input signal amplitude changes of 25 dB.

Adaptive quantization is an approach by which the properties of the quantizer are varied depending on the amplitude of the input signal. One example of adaptive quantization is when the step size of the quantizer is adjusted based on the input signal amplitude. Note that while the overall dynamic range is increased using adaptive quantization, the instantaneous dynamic range is not increased.[5]. The SNR can be kept relatively constant for input signal amplitude changes of 40 dB, for example, by selecting the ratio of maximum to minimum step size of 100 [29].

Finally, the technique of differential quantization involves the quantization of the difference between the actual input signal and an estimate of this input signal. Because this difference signal is much smaller than the actual input signal if a good estimate is obtained, the dynamic range is increased. The SNR can be maximized by optimizing the estimate of the input signal.

The potentially large processing load (number of computations within a given time interval) required to implement adaptive or differential quantization in radio receiver applications must be considered carefully. The speed of the processor must be able to handle the required processing load adequately.

4.2 Nonlinear Devices for Amplitude Compression

Logarithmic compression at the input to the quantizer can improve the SNR by making it independent of the input signal variance. In actual implementation, where system noise defines a minimum signal input amplitude, this form of compression can be achieved through the use of log amplifiers. Likewise, by using variable gain, it is possible to adapt the signal amplitude to the characteristics of the quantizer. This, too, can be implemented by using automatic gain control (AGC) amplifiers. Both log amplifiers and AGC amplifiers are discussed in greater detail below [30]–[33].

4.2.1 Log Amplifiers

There are several different types of log amplifiers, each suited to different applications. Common to all is some form of logarithmic compression of signal parameters. A simple operational amplifier circuit that uses the nonlinear (logarithmic) characteristics of a p-n junction can be used

[5] Dynamic range is the range of amplitude levels (minimum to maximum) that can be detected. Instantaneous dynamic range is the difference between the maximum and minimum signal amplitude levels for signals that can be detected simultaneously. Overall dynamic range is the difference between the maximum signal amplitude level that can be detected at any time and the minimum signal level that can be detected at any time.

[34]; however, these circuits suffer from temperature variations, limited dynamic range, and slow rise times. Depending on the application, these circuits may or may not be adequate. A more powerful technique is the approximation of the logarithmic function using the summation of linear (or curved) lines. Typically these are implemented with differential amplifiers using integrated circuit technology and may be purchased as discrete components or built into the front-end receiver circuitry. There are three basic types of these log amplifiers: detector log video amplifiers, successive detection log amplifiers, and true log IF amplifiers.

Detector log video amplifiers are suited to applications where phase and frequency information is not necessary. The envelope of the input signal simply is converted to a log-compressed video signal at the output. Total input dynamic range of these amplifiers is generally 50 dB. A typical application might be the demodulation and logarithmic compression of an amplitude-modulated signal.

The successive detection log amplifier provides two outputs. One is the same as the detector log video amplifier output, described above. The other is a limited IF signal. The former provides amplitude information expressed in logarithmic form and the latter provides phase and frequency information. The limited IF signal is a copy of the input signal except that its amplitude variation is compressed and limited by a transfer function similar to that in Figure 14. This figure shows the limited IF output power as a function of the input signal power. Typically these amplifiers have a total input dynamic range of 80 dB. A successive detection log amplifier recently has been announced by Microphase, Inc. that provides a 100-dB input dynamic range.

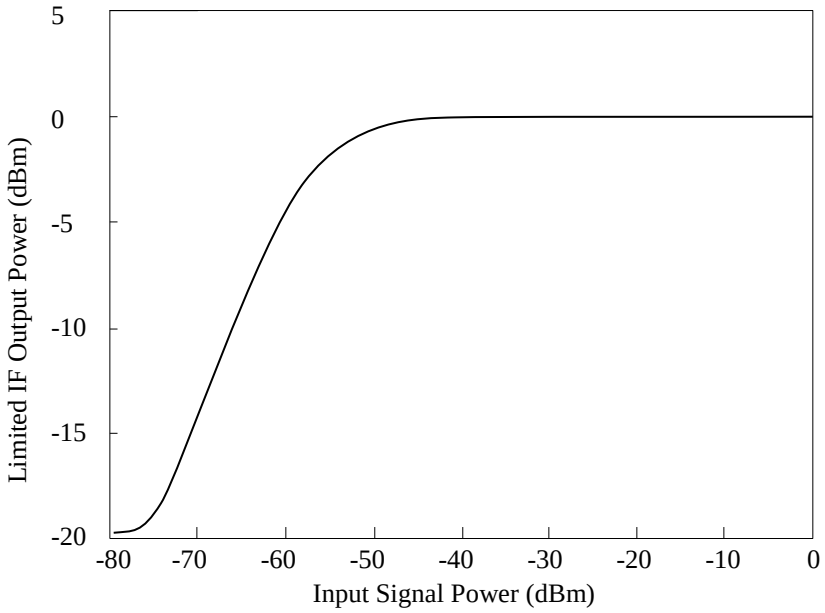

Figure 14. Limited IF output power as a function of the input signal power.

The “true” logarithmic IF amplifier is called so because the IF output is a bipolar logarithmic function of the IF input signal (without any limiting). This amplifier may also have an output that is the same as the detector log video amplifier. Due to the dual-polarity of the output signal, these amplifiers function well in logarithmic IF applications. As might be expected, there is some

deviation from the true logarithmic curve (toward a more linear function) as signals approach zero, but this is not generally a problem if the system noise power is set to the minimum signal input amplitude of the log amplifier (typically -80 dBm). These amplifiers also generally have a total input dynamic range of 80 dB. In addition, they inherently have low phase shifts over wide variations in input signal power.

Logarithmic amplifiers have wide instantaneous dynamic range. Instantaneous dynamic range means that at any point in time, all signals within the dynamic range of the amplifier appear at the output. Take for an example, an input containing two signals, one at -5 dBm and one at -75 dBm. For a logarithmic amplifier with an 80-dB dynamic range, it is possible to maintain the integrity of both signals at the output, but in a compressed form. Instead of the log amplifier, consider a variable attenuator placed before a fixed-gain amplifier. With the same input consisting of both the -5-dBm and -75-dBm signals, assume that the fixed-gain amplifier would saturate if no attenuation was used. Therefore, the attenuation of the variable attenuator must be increased to prevent amplifier saturation. In doing so, the noise figure of the combination of the variable attenuator and the fixed-gain amplifier will increase. This may cause the smaller signal (-75 dBm) to be overpowered by the noise and go undetected. Logarithmic amplifiers are, therefore, good for processing multiple signals where small signals are present simultaneously with large signals. Logarithmic amplifiers must be used carefully in receiver systems, however. Being nonlinear devices, they may cause distortion of the input signals. A careful analysis of the effects of distortion on the desired received signals must be performed.

4.2.2 Automatic Gain Control

AGC is used to normalize the output signal level despite variations in input signal power. A typical AGC device operates using a feedback loop whereby the power of the output of a variable gain device (amplifier or attenuator) is sampled and used to adjust the gain of the variable gain device itself. In this way, the output power is maintained at a relatively constant level. The response time of the AGC loop is perhaps the most critical design parameter and is dependent upon the application. Generally the AGC device is adjusted to react fast enough to normalize the overall power but slow enough to maintain the desired information content of the signal. Depending upon the application, the control voltage for the variable gain device may be used to decompress the output signal after digitization.

Typically AGC devices can have a total dynamic range of 80 dB. But because they operate by adjusting the gain of the system, the instantaneous dynamic range is much lower than that for logarithmic amplifiers. Therefore, these systems generally are better suited for narrower bandwidth systems where individual signals are isolated by appropriate filtering **before** the AGC. AGCs, unlike logarithmic amplifiers, are not effective in receivers where weak desired signals need to be detected in the presence of strong undesired signals.

AGC devices can be purchased as discrete components or built into the front-end receiver circuitry. As with any feedback system, stability is an important consideration. In radio receiver applications that use an AGC device before an ADC, there should be enough gain in the system so that the system noise at the output of the AGC device prior to digitization is slightly greater than the quantization noise power of the digitizer. To maximize the SNR, the normalized output

of the AGC device should be set as large as possible without causing the ADC to overload or generate excessive spurious responses. A more detailed discussion of AGC devices can be found in [35].

4.3 Postdigitization Algorithms for Improving Spurious Free Dynamic Range

As previously mentioned, there are numerous postdigitization techniques for optimizing the quantization process. Many of these techniques are used to increase the SNR of the quantizer by improving the predictor characteristics of differential quantization schemes. Another area of postdigitization processing provides compensation for the nonlinearities that occur in practical implementations of ADCs. As discussed in Section 2 of this report, these nonlinearities produce spurious signals that can reduce significantly the SFDR of the ADC. The purpose of this compensation is to suppress the spurious responses below the noise in the frequency band from 0 to $f_s/2$. Two of these techniques, phase-plane and state variable compensation, are discussed below [13], [14].

Both techniques are used to identify a set of correction factors that can be used to compensate for any nonlinearity throughout the full amplitude range of the ADC. In phase-plane compensation, the procedure for correcting the digitized signal is as follows: The input signal is split into two separate signals. One signal is fed into the ADC and the other is sent through an analog differentiator and then digitized by a second ADC. The differentiated signal is used to determine the instantaneous slope of the signal. The output of both ADCs then is used to determine the correction factor to be applied to the ADC output representing the digitized input signal. A table consisting of correction factors for each possible combination of quantization level and instantaneous slope is developed for an individual ADC based on measurements of that particular ADC. This table then is stored in RAM and is used to provide the correct ADC output for any given input signal amplitude. Studies using this technique show as much as a 15- to 16-dB (about 2.5 bits) improvement in the SFDR over uncompensated ADCs [14]. This improvement, however, is restricted to a narrow frequency band well below $f_s/2$.

In an effort to improve the SFDR for all frequencies in the 0 to $f_s/2$ frequency band, a state variable compensation technique was also proposed. This type of compensation is implemented by applying the input signal to an ADC and splitting the output of the ADC into two signals. One of these signals is used without modification while the other is delayed by a single clock cycle (one sample of the input signal). The two outputs, representing the quantization levels for the present and previous ADC outputs, then are used to determine the correction factor to be applied to the present ADC output. A table of correction factors for each possible combination of present and previous quantization levels is developed for an individual ADC based on measurements of that particular ADC. As in phase-plane compensation, this table is stored in RAM and is used to provide the correct ADC output for any given input signal amplitude. Tests using this technique also show as much as a 16-dB improvement in the SFDR over the **entire** 0 to $f_s/2$ frequency band for the particular sampling rate.

While compensation techniques require additional hardware and testing of individual ADCs, they can improve the SFDR of the ADC significantly without increasing its resolution (number of bits). In essence, they bring the characteristics of the ADC closer to the theoretical expectation

of its performance. An underlying assumption in these techniques, however, is that the ADC characteristics are static. Testing of ADCs has shown that for most ADCs this is a valid assumption [14].

4.4 Sampling Downconverters

The 6300 series sampling downconverters manufactured by Watkins-Johnson are an example of a technology potentially useful in the RF front-end of radio receivers. These sampling downconverters are microwave devices that use bandpass sampling techniques to downconvert an RF signal (typically in the 2- to 18-GHz range) to an IF signal (typically at 70 MHz up to 500 MHz). In its current configuration, the sampling downconverter functionally performs in the same manner as the conventional mixer. However, the device is actually a sample-and-hold circuit that uses a step recovery diode (SRD) to generate a sampling pulse train from a sinusoidal "clock" frequency. The sample-and-hold circuits in these downconverters are designed to work in the 2- to 18-GHz range, far higher in frequency than the sample-and-hold circuits integral to most ADCs. In the future, this technology may be coupled with high-resolution quantizers to produce bandpass sampling ADCs with a high analog input frequency capability and the high resolution of quantizers that would operate at much lower frequencies. Use of these sampling downconverters for radio receiver applications requires some careful considerations, however. The frequency content of the input must be bandlimited properly to prevent spectrum overlap in the desired signal output as with any bandpass sampling scheme. These downconverters do suffer a higher conversion loss than conventional mixers. The maximum specified conversion loss for the 6300 series sampling downconverters can be anywhere from 13-25 dB depending on the specific model. Typical conversion loss for conventional mixers ranges from roughly 5-9 dB. As with any downconverter, spurious suppression must be considered carefully. The specifications for the Watkins-Johnson 6300 series sampling downconverters state a minimum 15-dBc spurious suppression at a -10-dBm RF input level [36]. (This spurious suppression does not include the local oscillator (LO) signal leakage at the IF port or the 2^{nd} harmonic of the IF signal.) More typically, even a bad spur would be 25 dBc with the -10-dBm RF input level. Operation at lower RF input levels will provide even better spurious suppression [37]. These types of devices may be very useful for future radio receiver front-ends and ADCs.

It is tempting to compare the spurious suppression of the sampling downconverter to that of conventional mixers. While this comparison could be made for a specific mixer, it is not possible to make any general conclusions. While there are several types of conventional mixers, double-balanced mixers are the industry standard [38]. Even within the double-balanced mixer type, there are several different classes of mixers. These classes of mixers include class 1, class 2, and class 3 mixers. Class 2 mixers require more LO power and have better spurious suppression than class 1 mixers. Similarly, class 3 mixers require more LO power and have better spurious suppression than class 2 mixers. The higher the mixer class, the more LO power is required but the better the spurious suppression is. In addition to the different classes of mixers, spurious suppression is a function of LO and RF input power levels and frequencies. While there are specifications on mixers that help predict spurious suppression, spurious suppression for a specific mixer should be determined by measurement over a well-defined set of conditions [38].

4.5 Specialized Integrated Circuits

As digitization in radio receivers moves closer to the receive antenna, more of the traditionally analog receiver functions will be replaced with digital signal-processing algorithms. In addition, the area of digital communications requires extensive digital signal processing. While the speed of general digital signal processors continues to increase, specialized integrated circuits (ICs) can, in general, digitally implement radio receiver functions faster. Use of these specialized ICs should be quite prevalent in radio receiver applications of the future.

A current example of this type of specialized IC is the single chip digital downconverter. Digital downconversion offers a clear advantage over conventional analog downconversion by providing more precise frequency control of the LO and by providing a more ideal mixing operation. Current digital downconverters can control LO frequencies to within less than .01 Hz out of several MHz and can provide mixing with over 100 dB of SFDR. In addition, on-chip filters can be programmed to produce almost any desired bandwidth while preserving linear phase. In contrast to analog equivalents, these chips provide reproducible component accuracies that do not degrade over time or temperature.

The Harris HSP50016 is a programmable monolithic digital downconverter capable of extracting a narrow baseband signal from a wideband input bandpass signal (RF or IF) [39], [40]. The digital downconverter multiplies a digitized wideband input signal by samples representing both an in-phase LO and a quadrature-phase LO to produce in-phase and quadrature-phase baseband signals. Next, the digital downconverter filters the in-phase and quadrature-phase signals to extract the channel of interest. It then decimates the output to a sampling rate of twice the maximum frequency of the signal of interest. The digital downconverter can also be configured to provide a downconverted version of the input bandpass signal at a new IF. The HSP50016 can operate on 16-bit data at a rate of up to 52 Msamples/s [41].

5 EXAMPLES OF RADIOS USING RF OR IF DIGITIZATION

This section describes several radio receivers (currently existing or in development) employing direct digitization at the RF or IF. Digitization at the IF is becoming increasingly popular. Many radio receivers using this digitization currently are being developed. While direct digitization of the RF is a future design goal, few of these systems are in existence today except perhaps for frequencies below the HF band. In general, since many radio receivers using digitization at the RF or IF are currently in development, information about them is quite difficult to obtain. The receivers discussed below are a small sampling of these systems.

The dynamic signal analyzer is an excellent example of technology that uses direct digitization of the RF. The dynamic signal analyzer digitizes the input signal and provides various processing, display, and data storage options. A common signal-processing function is performing an FFT on an input signal and thereby providing a frequency domain display of the input signal. The HP 3587S is an example of this type of analyzer. This particular signal analyzer is implemented as a VXIbus system; the VXIbus is a standard modular instrument architecture that allows custom configuration of various measurement devices in a compact frame under computer control. The configuration of the HP 3587S includes a 10-Msample/s, 18-bit digitizer. This high-resolution ADC provides a maximum SFDR of 110 dB. An amplifier is used before the digitizer to set the system (random) noise power slightly above the quantization noise power. The system noise figure is then determined by this amplifier and is specified to be 14 dB. The signal analyzer contains a DSP module that can compute 250 MFLOPS. With this processor, a 1-MHz real-time bandwidth is achieved [42].

The HP1486A is a VXIbus signal-processing module currently in development that can process 9 billion math operations per second. It achieves this high processing speed by using application-specific integrated circuits (ASICs) along with performing computations as 28-bit fixed-point arithmetic. This module was designed for communication signal processing and can provide high-speed demodulation of both analog and digital communications signals. Processing speed is claimed to be fast enough to provide real-time demodulation of signals in the proposed HDTV digital transmission format. The HP1486A accepts data input up to 30 Msamples/s using a high-speed local bus. The receiver functions it can perform include: fixed and adaptive filtering; downconversion; extraction of signal amplitude, phase, and frequency information; resampling; adaptive clock and carrier synchronization; and symbol decoding. A general purpose DSP block is also available for more customized processing. These functions are individually programmable and can be used in a wide variety of combinations to implement various types of radio receivers [43].

A few years ago, the Novotel Communications Corporation (Calgary, Canada) was working on a global positioning system (GPS) receiver at 1.575 GHz using direct digitization of the RF. In addition to the challenges presented by direct digitization at this input frequency, the receiver design needed to be small, lightweight, inexpensive, and consume little power. A receiver using bandpass sampling techniques was considered; however, Novotel found that the receiver design was not feasible from a cost standpoint. Additionally, a sample-and-hold amplifier that would perform well at 1.575 GHz was not available commercially at that time so the technical feasibility was also in question.

The Honeywell Technology Center is developing an integrated avionics radio that digitizes and processes wideband IF signals between 100 and 1300 MHz [44]. The frequency bands that the system operates over include:

- 108-118 MHz VHF Omni Range and Instrument Landing System (ILS) Locator
- 118-137 MHz VHF Voice (Including Air Traffic Control Operations)
- 329-335 MHz Instrument Landing System Glide Path
- 960-1215 MHz Electronic Aids to Air Navigation (DME/Mode S)

Postdigitization filtering is used to isolate individual signals along with proprietary hardware to extend receiver dynamic range. Due to the developmental nature of this system, any further details (such as the IF bandwidth, dynamic range, sampling rate, ADC, processor hardware, and types of processing) are considered proprietary.

The SEA Corporation currently is producing a portable handheld transceiver operating in the 220- to 222-MHz band that uses digitization at the IF. SEA also is developing receivers for various applications in the 3- to 30-MHz band [45]. The 220- to 222-MHz band has 5-kHz channel bandwidths and therefore, requires sophisticated modulation techniques that generally depend on digital signal processing. Amplitude-companded single-sideband modulation (ACSSB) is being utilized with transmit tone in band (TTIB). However, discussions are underway to migrate to digital modulation. These receivers digitize IF signals at bandwidths slightly greater than the channel, and use bandpass sampling techniques for further downconversion. The system uses a 12-bit digitizer with a predigitization, 20-dB step AGC device.

Due to the very high dynamic range of input signal amplitudes present in the HF band, and limitations in the SFDR of high-speed ADCs currently available commercially, direct digitization at the RF in HF receivers is currently unrealistic [2]. There has been some recent activity, however, in digitization of the IF in HF radio receivers. An example design for an HF SSB radio receiver employing digitization at the IF is described in [10]. This SSB receiver uses a traditional superheterodyne front-end with two conversion stages to generate a 16-kHz wide IF signal centered at 456 kHz. An ADC, operating in a bandpass-sampling mode, samples the IF signal at 96 kHz thereby translating the desired replica of the spectrum to an IF centered at 24 kHz. Digital signal processing is then used to convert the IF signal to a baseband signal with an in-phase and quadrature-phase component. Decimation is then used to reduce the sample rate and allow the signal processing to take place in real time. The SSB signal is then digitally demodulated and converted to an analog signal for output.

Digitization at the wideband IF is being used in some receivers for cellular and PCS base stations. A potentially large market exists for these types of receivers. One example base station, the Watkins-Johnson $Base_2$, uses a receiver that digitizes at the wideband IF. The $Base_2$ base stations can be used for cellular or PCS applications by simply interchanging the RF front-end. The receiver in the $Base_2$ base station performs an initial downconversion and then digitizes up to a 15-MHz bandwidth at the IF. Digital signal processing is then used for channelization, filtering, tuning, and demodulation. The capability of digitizing up to a 15-MHz bandwidth permits digitization of either the entire mobile-to-base station cellular frequency block A or block B with a single receiver. Digitization of an entire, licensed, mobile-to-base station PCS

frequency block also can be achieved. The Steinbrecher and Airnet Corporations also produce cellular and PCS base stations that use wideband IF digitization.

There are several advantages of receivers in cellular or PCS base stations that digitize the wideband IF over traditional receivers used in analog cellular base stations. First, analog cellular base stations require one receiver per cellular channel. Only one wideband receiver is required for all of the cellular channels in a cell using a nonsectorized cell. Sectorized cells require one wideband receiver for all of the cellular channels used in a sector. Because of this, receivers that digitize at the wideband IF are cheaper, smaller, and consume less power than analog receivers in cellular base stations. Since filtering, demodulation, and signal processing are all performed digitally, the audio quality is claimed to be better than in analog receivers [46]. Perhaps one of the greatest advantages is that the receivers can receive multiple telephone calls with differing air-interface standards simultaneously. The receiver can be reconfigured with software and can operate with any air-interface standard. These standards include the advanced mobile phone service (AMPS), narrowband AMPS (N-AMPS), time-division multiple access (TDMA), cellular digital packet data (CDPD), and code-division multiple access (CDMA) [47].

Speakeasy is a new, joint-service military radio transceiver currently in development whose receiver uses digitization at the IF. The project is being managed by the Air Force Rome Laboratory. In Phase I of this project, the Speakeasy concept was demonstrated by implementing a VME-based radio operating over 2-2000 MHz using digitization of the IF at bandwidths up to approximately 10 MHz. General purpose, programmable DSPs were used for radio receiver operations whenever possible. When general-purpose DSPs were not fast enough to perform certain operations, application specific processors and dedicated digital hardware were used.

Two general modes of operation are available for the VME-based Speakeasy radio: a narrowband mode where an approximately 200-kHz IF bandwidth is digitized, and a wideband mode where an approximately 10-MHz bandwidth is digitized. Currently, two separate ADCs are used for these two modes. The narrowband mode uses a high-resolution ADC while the wideband mode uses a lower-resolution, faster ADC. The eventual goal is to use only one ADC for both modes when ADCs of sufficient resolution and sampling rate become available.

The VME-based Speakeasy radio allows selection (from a menu) of the types of signals that a user wants to transmit and receive. Currently, up to two types of signals can be processed (received or transmitted) simultaneously. The ultimate goal is to provide a capability to process four types of signals simultaneously. Speakeasy can process various types of amplitude, frequency, and phase-modulated signals as well as spread spectrum signals for analog and digital voice and data communication.

The goal of Phase II of the Speakeasy project is to provide a more compact and enhanced version of the VME-based radio developed in Phase I that will emulate more than 15 existing military radios [48].

Both direct broadcast satellite (DBS) and digital audio broadcast (DAB) systems transmit various forms of digitally modulated signals. Receivers for these systems can use either traditional demodulators (employing analog circuitry) after downconversion to IF, or digitization at the RF or IF before demodulation. Traditional demodulators require no analog-to-digital conversion

before demodulation. When using digitization at the RF or IF, demodulation is performed using a digital signal processor or dedicated digital-processing hardware.

The Advanced Wireless Technologies AWT2002 is a variable data rate BPSK/QPSK demodulator on a single IC. It is designed to operate on IF signals that have been digitized with a 6-bit ADC. Demodulation is accomplished digitally in this IC as opposed to demodulation requiring analog circuitry. The currently available version of this IC operates at a 2-Msymbol/s rate; however, a future version targeted for DBS systems will support symbol rates up to 31.5 Msymbols/s. Digitization at the RF or IF and then digital demodulation, like that provided by the AWT2002, may prove to be very useful in applications such as DBS and DAB.

Digitization at the RF or IF can be employed by receivers in radio propagation measurement systems. One example of a receiver that uses digitization at the wideband IF is in the recently patented digital sampling channel probe (DSCP). The DSCP was developed in a joint effort between the Institute for Telecommunication Sciences and Telesis Technologies Laboratory, Inc. The probe is ideal for making outdoor impulse response measurements to characterize wideband propagation in the radio channel [49].

The DSCP operates by transmitting an RF carrier modulated by a PN code over a radio propagation channel. The transmitted signal, modified by the propagation channel, then is received, downconverted to an IF, and digitized. After digitization at the IF, the complex impulse response of the radio propagation channel is generated via digital signal processing. Digitization at the IF enables the DSCP to measure an impulse response much faster than the traditional analog sliding correlator probe. This allows better characterization of rapidly changing propagation channels.

Many different system configurations are available with the DSCP. Current configurations use commercially available equipment such as RF signal generators, spectrum analyzers, digital oscilloscopes, and personal computers to achieve a high degree of flexibility and to minimize system setup time for specific field studies. Both the transmitter and receiver have a dual-channel capability allowing for transmission and reception with various combinations of two different PN codes, carrier frequencies, antenna polarizations, and antenna spacings. While the DSCP typically has been used to make impulse response measurements around 900 and 1850 MHz, it can be configured to accommodate measurements over an even broader range of frequencies.

In the typical configuration, the null-to-null bandwidth of the probe is 20 MHz, providing a delay resolution of 100 ns and a maximum measurable delay of 51 μs. The probe can be configured easily for wider bandwidths (finer time resolution) and different maximum delays. Recent improvements to the probe include the ability to measure absolute time and Doppler spread. Future plans include expanding the probe to multiple channels to help analyze the potential benefits of advanced antenna systems and antenna signal processing [49], [50].

6 SUMMARY AND RECOMMENDATIONS

The key factors in radio receivers where digitization occurs at the IF or RF are analog-to-digital conversion and digital signal processing. Although not discussed in this paper, digital-to-analog conversion is also a factor for applications requiring analog output (such as voice). Because of the rapid advances in hardware development of ADCs, DACs, digital signal processors, and specialized ICs, development of radio receivers using digitization at the IF (and in some cases at the RF) is becoming increasingly popular.

Hardware limitations of ADCs, digital signal processors, and DACs place constraints on digitization at the IF and RF in radio receivers. It was shown that digitization at the RF, in general, requires some type of bandlimiting (filtering) and amplification before the actual digitization takes place. The required amount of filtering and amplification is application-specific. ADC performance is improving rapidly. However, there is a tradeoff; one can get either high sampling rates or high resolution, but not both simultaneously. Therefore, the high sampling rate ADCs required for wide bandwidth applications may not have sufficient SFDR. Digitization at the RF is now being considered for satellite receivers since a large SFDR is not a necessity, very high sampling rate ADCs already exist, and even faster ADCs are being developed. For receiver applications requiring a large SFDR, such as HF communications, digitization at the IF is currently a more practical option.

Digital signal processors may present an even greater limitation than ADCs in radio receivers using digitization at the RF or IF. The speed, size, and cost of these processors are important for a particular radio receiver application. The requirement for real-time operation for many radio receivers places a heavy burden on these processors. It is difficult to discuss limitations of digital signal processing in general terms since many algorithms can be implemented in radio receivers depending on the specific application. The amount of time that signal processing requires is a function of the bandwidth of the signal, the speed of the processor, and the number and complexity of the algorithms required to perform the needed radio receiver functions.

Several potential methods and devices that are expected to be useful in radio receivers employing digitization of the RF or IF are discussed in this report. Various types of quantization techniques include uniform, μ-law, adaptive, and differential quantization. These techniques can be used to improve the dynamic range of ADCs as well as make the SNR insensitive to input signal amplitude. Nonlinear compression devices such as log amplifiers and automatic gain control amplifiers also can be used to improve the dynamic range of ADCs. Postdigitization algorithms for improving the SFDR provide extended dynamic range for presently available ADCs. Sampling downconverters, based on the theory of bandpass sampling, may be coupled with ADC technology to provide improved ADC performance for bandpass sampling applications. Specialized ICs for digital signal-processing tasks required in radio receiver applications show great potential in increasing processing speed; this allows more receiver functions to be executed in real time. Bandpass sampling is expected to be important for radio receiver applications since it allows the use of more readily available ADCs with lower sampling rates and higher SFDR.

This report provides information on receivers implementing digitization at the RF or IF. The topic of EMC analysis for these radio receivers needs to be investigated further. Due to the tremendous variety of possible receiver implementations employing digitization at RF or IF,

focusing on some specific RF front-end configurations probably would be the best approach for determining EMC analysis procedures. In general, however, RF front-end specifications such as SFDR and intermodulation distortion should be important, just as in EMC analyses of traditional receivers not employing digitization of the RF or IF.

The potential for spectrum overlap due to the sampling process is probably the major difference between receivers that employ digitization at the RF or IF and those that do not. Effects of this spectrum overlap are best determined by considering the specific radio system details such as the type of source information (i.e., voice, data, or video); desired signal bandwidth; modulation and coding techniques; undesired signal characteristics (bandwidth, power, and type of signal); and the performance criterion used to evaluate the quality of the reception of the desired signal. The sampling rate in relationship to the maximum frequency content of the signal and the signal bandwidth is of utmost importance in evaluating spectrum overlap. Computer simulation is probably the best tool available to evaluate the effects of spectrum overlap on receiver performance.

7 REFERENCES

[1] J. Mitola III, "Software radios survey, critical evaluation and future directions," *IEEE AES Systems Magazine*, Vol. 8, No. 4, pp. 25-36, Apr. 1993.

[2] R.M. Lober, "A DSP-based approach to HF receiver design: higher performance at a lower cost," *RF design*, Vol. 16, No. 8, pp. 92-100, Directory Issue, Aug. 1993.

[3] F.G. Stremler, *Introduction to Communication Systems*, Reading, MA: Addison-Wesley, Inc., 1977, pp. 112-120.

[4] E.O. Brigham, *The Fast Fourier Transform and its Applications*, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1988, pp. 83-86 and 320-337.

[5] D. Asta, "Recent dynamic range characterization of analog-to-digital converters for spectral analysis applications," Massachusetts Institute of Technology, Lincoln Laboratory, Project Report AST-14, Jul. 9, 1991.

[6] M. Bartz, "Large scale dithering enhances ADC dynamic range," *Microwaves & RF*, Vol. 32, No.5, pp. 192-198, May 1993.

[7] B. Brannon, "Overcoming converter nonlinearities with dither," Analog Devices Application Note AN-410, 1995.

[8] R. Groshong and S. Ruscak, "Exploit digital advantages in an SSB receiver," *Electronic Design*, Vol. 39, No. 11, pp. 89-96, Jun. 13, 1991.

[9] M. Amarandos and S. Andrezyk, "Considerations in the development of a low cost, high performance receiver based on DSP techniques," *DSP Applications*, Vol. 2, No. 12, pp. 1-14, Dec. 1993.

[10] R. Groshong and S. Ruscak, "Undersampling techniques simplify digital radio," *Electronic Design*, Vol. 39, No. 10, pp. 67-78, May 23, 1991.

[11] Harris Semiconductor Corporation, *Harris High-Speed Digital Signal Processing Seminar*, Melbourne, FL: Harris Semiconductor Corporation, 1994, pp. 8-7.

[12] Analog Devices, "AD9042 data specification sheet," Analog Devices, Inc., 1995.

[13] N.W. Spencer, "Comparison of state-of-the-art analog-to-digital converters," Massachusetts Institute of Technology, Lincoln Laboratory, Project Report AST-4, Mar. 4, 1988.

[14] F.H. Irons and T.A. Rebold, "Characterization of high-frequency analog-to-digital converters for spectral analysis applications," Massachusetts Institute of Technology, Lincoln Laboratory, Project Report AST-2, Nov. 1986.

[15] N.T. Thao and M. Vetterli, "Optimal MSE signal reconstruction in oversampled A/D conversion using convexity," in *Proc. ICASSP '92 Vol. 4*, 1992, pp. 165-168.

[16] S.R. Broadstone, "Evaluation of the Hughes Aircraft Company 14-bit analog-to-digital converter," Massachusetts Institute of Technology, Lincoln Laboratory, Project Report AST-28, Jul. 27, 1993.

[17] Harris Semiconductor Corporation, "Principles of data acquisition and conversion," Harris Semiconductor Application Note A002, Oct. 1986.

[18] Analog Devices, *System Applications Guide*, Norwood, MA: Analog Devices, Inc., 1993, pp. 14.1-14.16.

[19] Analog Devices, *Practical Analog Design Techniques*, Norwood, MA: Analog Devices, Inc., 1995, pp. 3.2-3.11.

[20] Analog Devices, *Linear Design Seminar*, Norwood, MA: Analog Devices, Inc., 1995, pp. 8.44-8.68.

[21] Analog Devices, *Analog-to-Digital Conversion Handbook*, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1986, pp. 559-560.

[22] B. Sklar, "A structured overview of digital communications - a tutorial review - part 1," *IEEE Communications Magazine*, Vol. 21, No. 8, pp. 4-17, Aug. 1983.

[23] B. Sklar, "A structured overview of digital communications - a tutorial review - part 2," *IEEE Communications Magazine*, Vol. 21, No. 10, pp. 6-21, Oct. 1983.

[24] A.V. Oppenheim and R.W. Schafer, *Discrete-Time Signal Processing*, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1989, pp. 119-123.

[25] J.J. Dongarra, "Performance of various computers using standard linear equations software," Computer Science Department, University of Tennessee, Knoxville, Tennessee, Sep. 26, 1994.

[26] Hewlett Packard Corporation, "The fundamentals of signal analysis," Hewlett Packard Application Note 243, Feb. 1985.

[27] L.R. Rabiner and R.W. Schafer, *Digital Processing of Speech Signals*, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1978.[6]

[28] B. Smith, "Instantaneous companding of quantized signals," *Bell System Tech. J.*, Vol. 36, No.3, pp. 653-709, May 1957.

[6] Reissue note: Between the first publication of this report and this reissue, *Digital Processing of Speech Signals* was superseded as a fundamental reference text by L.R. Rabiner and R.W. Schafer, *Theory and Applications of Digital Speech Processing*, Englewood Cliffs, NJ: Prentice-Hall, Inc., 2010. Although the authors consulted the 1978 textbook in preparing this report, we recommend that readers consult the 2010 textbook on any topics that refer to [27].

[29] P. Noll, "A comparative study of various schemes for speech encoding," *Bell System Tech. J.*, Vol. 54, No. 9, pp. 1597-1614, Nov. 1975.

[30] R.B. Hirsch, "Plain talk on log amps," *Microwaves*, Vol. 17, No. 3, pp. 40-44, Mar. 1978.

[31] R.S. Hughes, "Design automatic gain control loops the easy way," *EDN*, Vol. 23, No. 2, pp. 123-128, Oct. 5, 1978.

[32] R.S. Hughes, *Logarithmic Amplification*, Dedham, MA: Artech House, Inc., 1986.

[33] M. Parkyn, "Monolithic design shrinks logamp architecture," *Microwaves & RF*, Vol. 29, No. 9, pp. 179-183, Sep. 1990.

[34] J.G. Graeme and G.E. Tobey, *Operational Amplifiers*, New York: McGraw-Hill Book Co., 1971, pp. 258-268.

[35] R.S. Hughes, "Automatic gain control: a practical approach to its analysis and design," Naval Weapons Center Publication 5948, Aug. 1977.

[36] Watkins-Johnson Company, "WJ-6300 series sampling downconverters," Technical Data Sheet, Apr. 1993.

[37] J. Bellantoni, "Watkins-Johnson Company 6300 series sampling downconverters," Watkins-Johnson Company Application Note AN-5, 1993.

[38] B.C. Henderson, "Mixers: part 2 (theory and technology)," Watkins-Johnson Company Tech-Notes, Vol. 8, No. 3, 1981.

[39] D.B. Chester and G. Phillips, "Single chip digital down converter simplifies RF DSP applications," *RF Design*, Vol. 15, No. 12, pp. 39-46, Nov. 1992.

[40] M. Petrowski, D.B. Chester, and W.R. Young, "Single chip digital down converter architecture," in *Proc. ICASSP 93 - 1993 IEEE International Conf. Acoustics, Speech, and Signal Processing*, Minneapolis, MN, 1993, pp. I-349 to I-352.

[41] Harris Semiconductor Corporation, *Digital Signal Processing Data Book*, Melbourne, FL: Harris Semiconductor Corporation, 1994, pp. 6-3.

[42] Hewlett Packard Corporation, *Test and Measurement Catalog 1994*, Santa Clara, CA: Hewlett Packard Corporation Test and Measurement Sector, 1993, pp. 230.

[43] H. Hilton, "DSP module for real-time demodulation," Hewlett Packard 1994 Signal Surveillance Forum, Washington, D.C., 1994.

[44] J. Kriz, Private Communication, Integrated Avionics Radio Program Manager, Honeywell Technology Center, Minneapolis, MN, Feb. 1996.

[45] J.F. Cleveland, Private Communication, Senior Project Engineer, SEA Corporation, Mountlake Terrace, WA, Sep. 1994 and Feb. 1996.

[46] T.T. Yang, "Base$_2$™ wireless network system," Watkins-Johnson Company, Aug. 7, 1995.

[47] Steinbrecher Corporation, "Corporate profile," Jun. 24, 1994.

[48] R.J. Lackey and D.W. Upmal, "Speakeasy: the military software radio," *IEEE Communications Magazine*, Vol. 33, No. 5, pp. 56-61, May 1995.

[49] J.A. Wepman, "Digital sampling channel probe," *Federal Lab Test & Measurement Tech Briefs*, p. 6a, May 1995.

[50] J.A. Wepman, J.R. Hoffman, and L.H. Loew, "Analysis of impulse response measurements for PCS channel modelling applications," *IEEE Trans. Vehic. Tech.*, Vol. 44, No. 3, pp. 613-620, Aug. 1995.

APPENDIX: ACRONYMS AND ABBREVIATIONS

ΣΔ	sigma-delta
ACSSB	amplitude companded single-sideband
ADC	analog-to-digital converter
AGC	automatic gain control
AMPS	advanced mobile phone service
ARPA	Advanced Research Projects Agency
ASIC	application specific integrated circuit
BW	bandwidth
CDMA	code-division multiple access
CDPD	cellular digital packet data
CW	continuous-wave
DAB	digital audio broadcast
DAC	digital-to-analog converter
DBS	direct broadcast satellite
DME	distance measuring equipment
EMC	electromagnetic compatibility
ENOB	effective number of bits
FFT	Fast Fourier Transform
FIR	finite impulse response
FSR	full scale range
GPS	global positioning system
HF	high frequency
IC	integrated circuit
IF	intermediate frequency
IIR	infinite impulse response
ILS	Instrument Landing System
IMD	intermodulation distortion
LO	local oscillator
LSB	least significant bit
MA	magnitude amplifier
MFLOPS	millions of floating-point operations per second
MS	mean squared
MSB	most significant bit
N-AMPS	narrowband AMPS
NF	noise figure
NPR	noise power ratio

NTIA	National Telecommunications and Information Administration
PCS	personal communications services
RF	radio frequency
SFDR	spurious free dynamic range
SHA	sample-and-hold amplifier
SINAD	signal-to-noise plus distortion ratio
SNR	signal-to-noise ratio
SNR_{aj}	signal-to-noise ratio due to aperture jitter
SRD	step recovery diode
SSB	single-sideband
TDMA	time-division multiple access
THA	track-and-hold amplifier
TTIB	transmit tone in band
VHF	very high frequency

FORM **NTIA-29**
(4-80)

U.S. DEPARTMENT OF COMMERCE
NAT'L. TELECOMMUNICATIONS AND INFORMATION ADMINISTRATION

BIBLIOGRAPHIC DATA SHEET

1. PUBLICATION NO.	2. Government Accession No.	3. Recipient's Accession No.
96-328		

4. TITLE AND SUBTITLE	5. Publication Date
	Mar. 1996, Reissued. Feb. 2014
	6. Performing Organization Code NTIA/ITS
7. AUTHOR(S)	9. Project/Task/Work Unit No. 6 90 6105
8. PERFORMING ORGANIZATION NAME AND ADDRESS Institute for Telecommunication Sciences National Telecommunications & Information Administration U.S. Department of Commerce 325 Broadway Boulder, CO 80305	10. Contract/Grant Number.
11. Sponsoring Organization Name and Address National Telecommunications & Information Administration Herbert C. Hoover Building 14th & Constitution Ave., NW Washington, DC 20230	12. Type of Report and Period Covered

14. SUPPLEMENTARY NOTES

15. ABSTRACT (A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here.)

Hardware development of analog-to-digital converters (ADCs) and digital signal processors, including specialized integrated circuits, has advanced rapidly within the last few years. These advances have paved the way for development of radio receivers using digitization at the IF and in some cases at the RF. Applications for these receivers are expected to increase rapidly in areas such as cellular mobile, satellite, and personal communications services (PCS) systems. The constraints placed on these receivers due to hardware limitations of these devices are investigated in this paper. Some examples of state-of-the-art ADCs, signal processors, and specialized integrated circuits are listed. Various quantization techniques, nonlinear compression devices, postdigitization algorithms for improving dynamic range, sampling downconverters, and specialized integrated circuits are discussed as they are expected to be useful in the development of these types of receivers. Several examples of radio receivers employing digitization at the IF and RF are also presented.

16. Key Words (Alphabetical order, separated by semicolons)
analog-to-digital converters; automatic gain control devices; digital signal processors; digitization; logarithmic amplifiers; quantization; radio receivers; sampling downconverters; signal-to-noise ratio; spurious free dynamic range

17. AVAILABILITY STATEMENT	18. Security Class. (This report)	20. Number of pages
☒ UNLIMITED.	Unclassified	60
☐ FOR OFFICIAL DISTRIBUTION.	19. Security Class. (This page) Unclassified	21. Price:

NTIA FORMAL PUBLICATION SERIES

NTIA MONOGRAPH (MG)

A scholarly, professionally oriented publication dealing with state-of-the-art research or an authoritative treatment of a broad area. Expected to have long-lasting value.

NTIA SPECIAL PUBLICATION (SP)

Conference proceedings, bibliographies, selected speeches, course and instructional materials, directories, and major studies mandated by Congress.

NTIA REPORT (TR)

Important contributions to existing knowledge of less breadth than a monograph, such as results of completed projects and major activities. Subsets of this series include:

NTIA RESTRICTED REPORT (RR)

Contributions that are limited in distribution because of national security classification or Departmental constraints.

NTIA CONTRACTOR REPORT (CR)

Information generated under an NTIA contract or grant, written by the contractor, and considered an important contribution to existing knowledge.

JOINT NTIA/OTHER-AGENCY REPORT (JR)

This report receives both local NTIA and other agency review. Both agencies' logos and report series numbering appear on the cover.

NTIA SOFTWARE & DATA PRODUCTS (SD)

Software such as programs, test data, and sound/video files. This series can be used to transfer technology to U.S. industry.

NTIA HANDBOOK (HB)

Information pertaining to technical procedures, reference and data guides, and formal user's manuals that are expected to be pertinent for a long time.

NTIA TECHNICAL MEMORANDUM (TM)

Technical information typically of less breadth than an NTIA Report. The series includes data, preliminary project results, and information for a specific, limited audience.

For information about NTIA publications, contact the NTIA/ITS Technical Publications Office at 325 Broadway, Boulder, CO, 80305 Tel. (303) 497-3572 or e-mail info@its.bldrdoc.gov.

The original version of this report is for sale as Accession Number PB97-185581 by the National Technical Information Service, 5285 Port Royal Road, Springfield, VA 22161,Tel. (800) 553-6847.