

Which ADC Architecture Is Right for Your Application?

By Walt Kester [walt.kester@analog.com]

INTRODUCTION

Selecting the proper ADC for a particular application appears to be a formidable task, considering the thousands of converters currently on the market. A direct approach is to go right to the selection guides and parametric search engines, such as [those available](#)¹ on the Analog Devices website. Enter the sampling rate, resolution, power supply voltage, and other important properties, click the “find” button, and hope for the best. But it’s usually not enough. How does one deal with a multiplicity of apparent “best choices”? Is there a way to approach the task with greater understanding—and better results?

Most ADC applications today can be classified into four broad market segments: (a) *data acquisition*, (b) *precision industrial measurement*, (c) *voiceband and audio*, and (d) “*high speed*” (implying sampling rates greater than about 5 MSPS). A very large percentage of these applications can be filled by *successive-approximation* (SAR), *sigma-delta* (Σ - Δ), and *pipelined* ADCs. A basic understanding of these, the three most popular ADC architectures—and their relationship to the market segments—is a useful supplement to the selection guides and search engines.

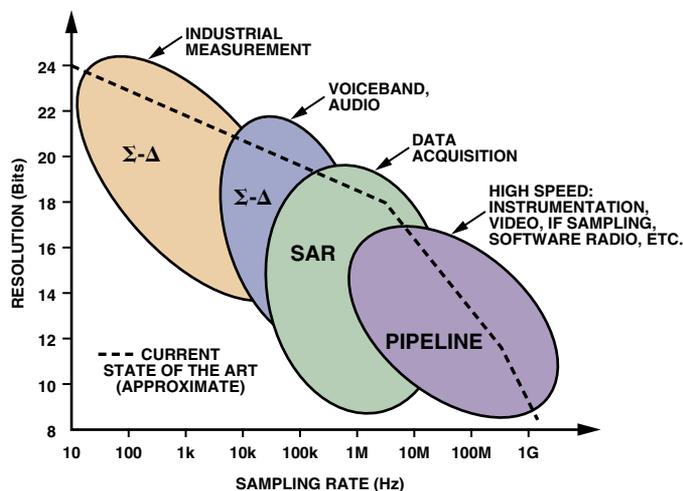


Figure 1. ADC architectures, applications, resolution, and sampling rates.

The classification in Figure 1 shows in a general way how these application segments and the associated typical architectures relate to ADC resolution (vertical axis) and sampling rate (horizontal axis). The dashed lines represent the approximate state of the art in mid-2005. Even though the various architectures have specifications with a good deal of overlap, the applications themselves are key to choosing the specific architecture required.

Successive-Approximation ADCs for Data Acquisition

The successive-approximation ADC is by far the most popular architecture for data-acquisition applications, especially when multiple channels require input multiplexing. From the modular and hybrid devices of the 1970s to today’s modern low-power ICs, the successive-approximation ADC has been the workhorse of data-acquisition systems. The architecture was first utilized in experimental *pulse-code-modulation* (PCM)

systems by Bell Labs in the 1940s. Bernard Gordon, at Epsco, introduced the first commercial vacuum-tube SAR ADC in 1954—an 11-bit, 50-kSPS ADC that dissipated 500 watts.

Modern IC SAR ADCs are available in resolutions from 8 bits to 18 bits, with sampling rates up to several MHz. At this writing, state-of-the-art performance of available devices is 16 bits at 3 MSPS (AD7621)² and 18 bits at 2 MSPS (AD7641)³. Output data is generally provided via a standard serial interface (I²C[®] or SPI[®], for example), but some devices are available with parallel outputs (at the obvious expense of increased pin count and package size).

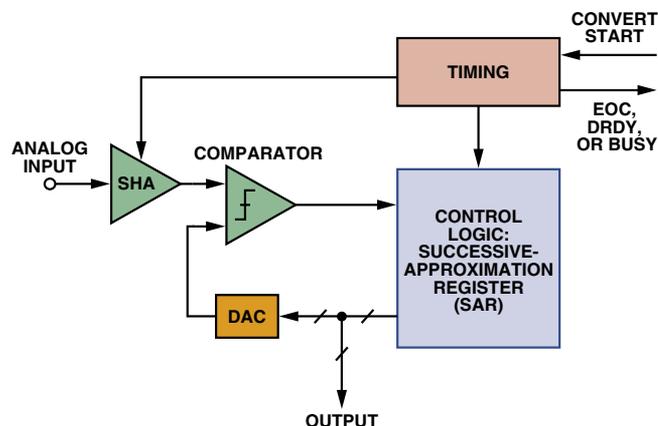


Figure 2. Basic successive-approximation (SAR) ADC.

The basic successive-approximation architecture is shown in Figure 2. In order to process rapidly changing signals, SAR ADCs have an input *sample-and-hold* (SHA) to keep the signal constant during the conversion cycle. The conversion starts with the internal *D/A converter* (DAC) set to midscale. The comparator determines whether the SHA output is greater or less than the DAC output, and the result (the *most-significant bit* (MSB) of the conversion) is stored in the *successive-approximation register* (SAR) as a 1 or a 0. The DAC is then set either to ¼ scale or ¾ scale (depending on the value of the MSB), and the comparator makes the decision for the second bit of the conversion. The result (1 or 0) is stored in the register, and the process continues until all of the bit values have been determined. At the end of the conversion process, a logic signal (EOC, DRDY, BUSY, etc.) is asserted. The acronym, *SAR*, which actually stands for *successive-approximation register*—the logic block that controls the conversion process—is universally understood as an abbreviated name for the entire architecture.

The timing diagram for a typical SAR ADC is shown in Figure 3. The functions shown are generally present in most SAR ADCs, but their exact labels can differ from device to device. Note that the data corresponding to that specific sample is available at the end of the conversion time, with no “pipeline” delay or “latency.” This makes the SAR ADC easy to use in single-shot, burst-mode, and multiplexed applications.

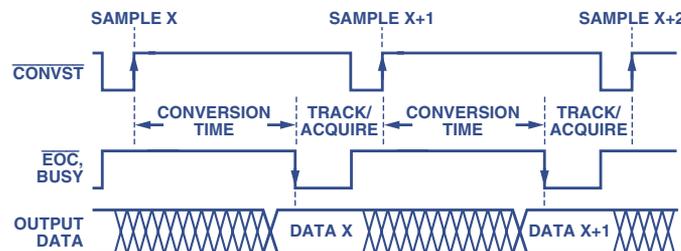


Figure 3. Simplified timing diagram of a SAR A/D converter.

The internal conversion process of most modern IC SAR ADCs is controlled by a high-speed clock (internal or external, depending on the ADC) that does not need to be synchronized to the CONVERT START input.

The basic algorithm used in the successive-approximation ADC conversion process can be traced back to the 1500s. It is related to the solution of a useful mathematical puzzle—the determination of an unknown weight by a minimal sequence of weighing operations (Reference 1). In this problem, as stated, the object is to determine the least number of weights which would serve to weigh an integral number of pounds from 1 lb to 40 lb using a balance scale. One solution put forth by the mathematician Tartaglia in 1556, was to use the binary series of weights 1 lb, 2 lb, 4 lb, 8 lb, 16 lb, and 32 lb (or 2^0 , 2^1 , 2^2 , 2^3 , 2^4 , and 2^5). The proposed weighing algorithm is the same one that is used in modern successive-approximation ADCs. (It should be noted that this solution will actually measure unknown weights up to 63 lb ($2^6 - 1$) rather than 40 lb as stated in the problem)*. The binary algorithm, using a balance scale, is shown in Figure 4 with an unknown weight of 45 lbs.

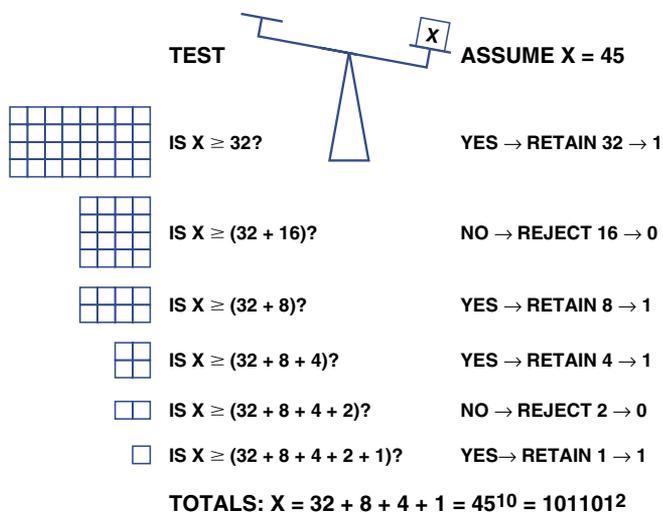


Figure 4. Successive-approximation ADC algorithm using balance scale and binary weights.

The overall accuracy and linearity of the SAR ADC are determined primarily by the internal DAC's characteristics. Early precision SAR ADCs, such as the industry-standard AD574,⁴ used DACs with laser-trimmed thin-film resistors to achieve the desired accuracy and linearity. However, the process of depositing and trimming thin-film resistors adds cost, and the thin-film resistor values may be affected after the device is subjected to the mechanical stresses of packaging.

For these reasons, switched-capacitor (or *charge-redistribution*) DACs have become popular in newer CMOS-based SAR ADCs. The principal advantage of the switched-capacitor DAC is that the accuracy and linearity are primarily determined by high-accuracy photolithography, which establishes the capacitor plate area, hence the capacitance and the degree of matching. In addition, small capacitors can be placed in parallel with the main capacitors—to be

*Note that if *ternary* (base-3: 1,0,-1) logic is permitted, the problem can be solved in four steps, with weights of 1, 3, 9, and 27 lbs applied on either side of the balance. Indeed, 40 lbs is then maximum with these weights.

switched in and out under control of autocalibration routines—to achieve high accuracy and linearity without the need for thin-film laser trimming. Because temperature tracking between the capacitors can be better than 1 ppm/°C, a high degree of temperature stability is achieved.

CMOS, the process of choice for modern SAR ADCs, is also the ideal process for *analog switches*. Therefore, input multiplexing can be added to the basic SAR ADC function relatively straightforwardly, thus allowing the integration of a complete data-acquisition system on a single chip. Additional digital functions are also easy to add to SAR-based ADCs, so features such as multiplexer sequencing, autocalibration circuitry, and more are becoming common.

Figure 5 illustrates the elements of the AD79x8 series of 1-MSPS SAR ADCs. The *sequencer* allows automatic conversion of the selected channels, or channels can be addressed individually if desired. Data is transferred via the serial port. SAR ADCs are popular in multichannel data-acquisition applications because they lack the “pipeline” delays typical in Σ - Δ and pipelined ADC architectures. The SAR ADC's conversion modes include “single-shot,” “burst,” and “continuous.”

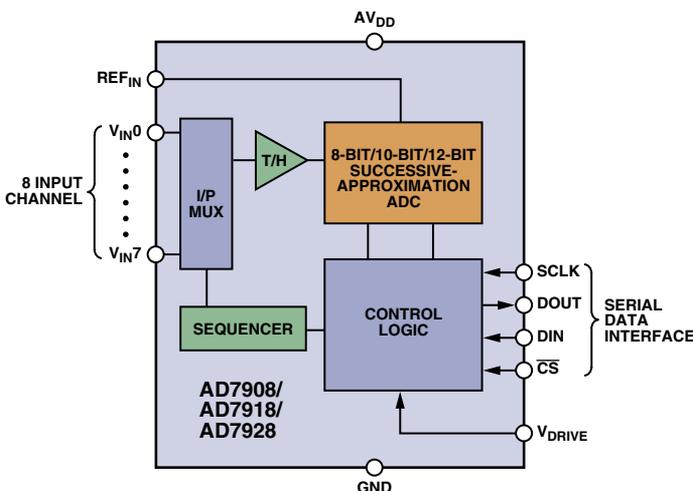


Figure 5. Functional block diagram of a modern 1-MSPS SAR ADC with 8-channel input multiplexer. Its family includes the AD7908⁵ (8 bits), AD7918⁶ (10 bits), and AD7928⁷ (12 bits).

***Sigma-Delta* (Σ - Δ) ADCs for Precision Industrial Measurement and Instrumentation**

Modern Σ - Δ ADCs have virtually replaced the integrating-type ADCs (dual-slope, triple-slope, quad-slope, etc.) for applications requiring high resolution (16 bits to 24 bits) and effective sampling rates up to a few hundred hertz. High resolution, together with on-chip programmable-gain amplifiers (PGAs), allows the small output voltages of sensors—such as weigh scales and thermocouples—to be digitized directly. Proper selection of sampling rate and digital filter bandwidth also yields excellent rejection of 50-Hz and 60-Hz power-line frequencies. Σ - Δ ADCs offer an attractive alternative to traditional approaches using an instrumentation amplifier (in-amp) and a SAR ADC.

The basic concepts behind the Σ - Δ ADC architecture originated at Bell Labs in the 1950s—in work done on experimental digital transmission systems utilizing delta modulation and differential PCM. By the end of the 1960s, the Σ - Δ architecture was well understood. However, because digital filters (then a rarity) were

an integral part of the architecture, practical IC implementations did not appear until the late 1980s, when signal processing in digital CMOS became widely available. The basic concepts used in Σ - Δ —*oversampling*, *noise shaping*, *digital filtering*, and *decimation*—are illustrated in Figure 6.

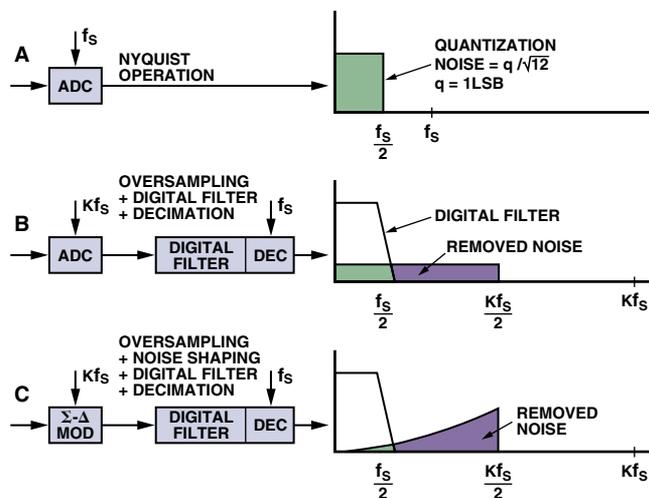


Figure 6. Noise-spectrum effects of the fundamental concepts used in Σ - Δ : oversampling, digital filtering, noise shaping, and decimation.

Figure 6A shows a noise spectrum for traditional “Nyquist” operation, where the ADC input signal falls between dc and $f_s/2$, and the quantization noise is uniformly spread over the same bandwidth. In Figure 6B, the sampling frequency has been increased by a factor, K , (the *oversampling* ratio), but the input signal bandwidth is unchanged. The quantization noise falling outside the signal bandwidth is then removed with a digital filter. The output data rate can now be reduced (*decimated*) back to the original sampling rate, f_s . This process of oversampling, followed by digital filtering and decimation, increases the SNR within the Nyquist bandwidth (dc to $f_s/2$). For each doubling of K , the SNR within the dc-to- $f_s/2$ bandwidth increases by 3 dB. Figure 6C shows the basic Σ - Δ architecture, where the traditional ADC is replaced by a Σ - Δ modulator. The effect of the modulator is to shape the quantization noise so that most of it occurs outside the bandwidth of interest, thereby greatly increasing the SNR in the dc-to- $f_s/2$ region.

The basic first-order Σ - Δ ADC is shown in Figure 7, with the Σ - Δ modulator shown in some detail.

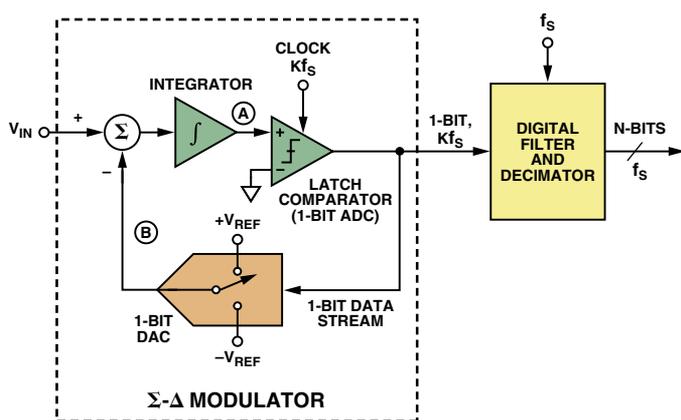


Figure 7. First-order sigma-delta ADC.

The heart of this basic modulator is a 1-bit ADC (comparator) and a 1-bit DAC (switch). Although there are a number of multibit Σ - Δ ADCs, those using the single-bit modulator have the obvious advantage of inherently excellent differential linearity.

The output of the modulator is a 1-bit stream of data. Because of negative feedback around the integrator, the average value of the signal at B must equal V_{IN} . If V_{IN} is zero (i.e., midscale), there are an equal number of 1s and 0s in the output data stream. As the input signal goes more positive, the number of 1s increases, and the number of 0s decreases. Likewise, as the input signal goes more negative, the number of 1s decreases, and the number of 0s increases. The ratio of the 1s in the output stream to the total number of samples in the same interval—the *ones density*—must therefore be proportional to the dc value of the input.

The modulator also accomplishes the *noise-shaping* function by acting as a low-pass filter for the signal and a high-pass filter for the quantization noise. Note that the digital filter is an integral part of the Σ - Δ ADC, and it can be optimized to give excellent 50-Hz/60-Hz power-frequency rejection. However, the digital filter does introduce inherent *pipeline delay*, which definitely must be considered in multiplexed and servo applications. If signals are multiplexed into a Σ - Δ ADC, the digital filter must be allowed to settle to the new value before the output data is valid. Several output clock cycles are generally required for this settling. Because of the pipeline delay of the digital filter, the Σ - Δ converter cannot be operated in a “single-shot” or “burst” mode.

Although the simple first-order single-bit Σ - Δ ADC is inherently linear and monotonic because of the 1-bit ADC and 1-bit DAC, it does not provide sufficient noise shaping for high-resolution applications. Increasing the number of integrators in the modulator (similar to adding poles to a filter) provides more noise shaping at the expense of a more complex design—as shown in Figure 8 for a second-order 1-bit modulator. Note the improvement in the noise shaping characteristic compared to a first-order modulator. Higher-order modulators (greater than third order) are difficult to stabilize and present significant design challenges.

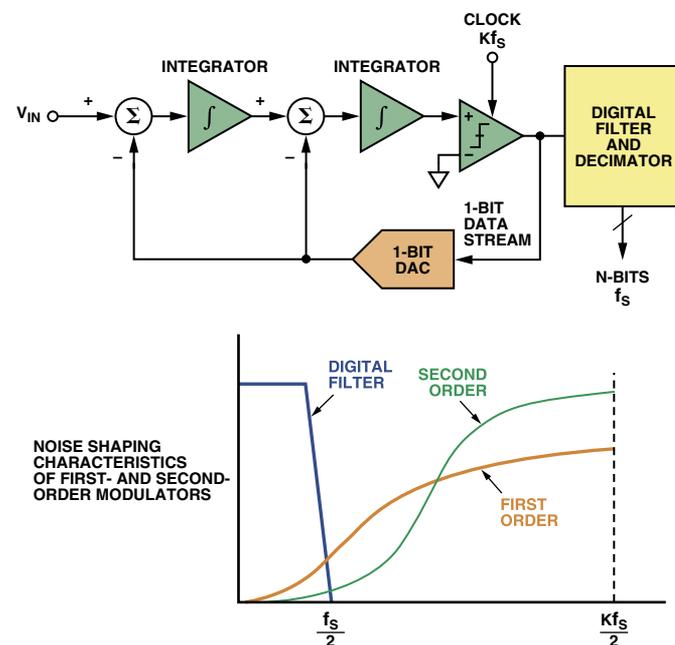


Figure 8. Second-order Σ - Δ modulator.

A popular alternative to higher-order modulators is to use a *multibit* architecture, where the 1-bit ADC (comparator) is replaced with an N -bit flash converter, and the single-bit DAC (switch) is replaced with a highly linear N -bit DAC. Expensive laser trimming in multibit Σ - Δ ADCs can be avoided by using techniques such as data scrambling to achieve the required linearity of the internal ADC and DAC.

While integrating architectures (dual-slope, triple-slope, etc.) are still used in applications such as digital voltmeters, the CMOS Σ - Δ ADC is the dominant converter for today's industrial measurement applications. These converters offer excellent power-line common-mode rejection and resolutions up to 24 bits as well as digital conveniences such as on-chip calibration. Many have programmable-gain amplifiers (PGAs), which allow small signals from bridge- and thermocouple transducers to be directly digitized without the need for additional external signal conditioning circuits and in-amps.

Figure 9 shows a simplified diagram of a precision load cell. This particular load cell produces 10-mV full-scale output voltage for a load of 2 kg with 5-V excitation.

The bridge's common-mode output voltage is 2.5 V. The diagram shows the bridge resistance values for a 2-kg load. The output voltage for any given load is directly proportional to the excitation voltage, i.e., it is ratiometric with the supply voltage.

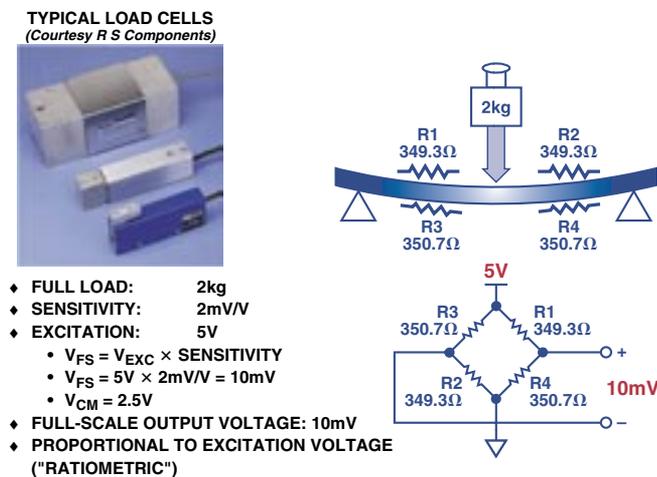


Figure 9. Load-cell signal-conditioning application.

A traditional approach to digitizing this low-level output would be to use an instrumentation amplifier to provide the necessary gain to drive a conventional SAR ADC of 14-bit to 18-bit resolution. Because of offset and drift considerations, an "auto-zero" in-amp such as the AD5555⁸ or AD8230⁹ is required. Appropriate filtering circuitry is needed due to the noise of the auto-zero in-amp. In addition, the output data from the SAR ADC is often averaged for further noise reduction.

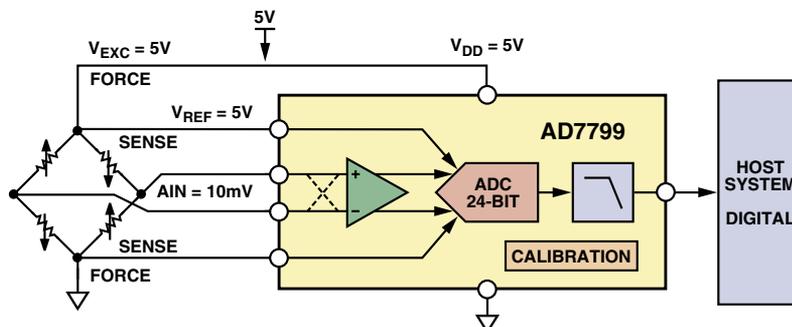


Figure 10. Load-cell signal conditioning using the AD7799 high-resolution Σ - Δ ADC.

An attractive alternative to the traditional in-amp/SAR ADC approach is shown in Figure 10, which uses a direct connection between the load cell and the AD7799¹⁰ high resolution Σ - Δ ADC. The full-scale bridge output of 10 mV is digitized to approximately 16 "noise-free" bits by the ADC at a throughput rate of 4.7 Hz. (For more discussion on input-referred noise and noise-free code resolution see for Further Reading 1). Ratiometric operation eliminates the need for a precision voltage reference.

The Σ - Δ ADC is an attractive alternative when very low-level signals must be digitized to high resolution—but the user should understand that the Σ - Δ ADC is more digitally intensive than the SAR ADC and may therefore require a somewhat longer development cycle. Evaluation boards and software can greatly assist in this process. Nevertheless, there are still many instrumentation and sensor signal-conditioning applications that can be efficiently solved with a traditional in-amp (for signal amplification and common-mode rejection) followed by a multiplexer and a SAR ADC.

Sigma-Delta ADCs for Voiceband and Audio

In addition to providing attractive solutions for a variety of industrial measurement applications—precision measurement, sensor monitoring, energy metering, and motor control—the Σ - Δ converter dominates modern voiceband and audio applications. A major benefit of the high oversampling rate inherent in Σ - Δ converters is that they simplify the input antialiasing filter for the ADC and the output anti-imaging filter for the DAC. In addition, the ease of adding digital functions to a CMOS-based converter makes features such as digital-filter programmability practical with only small increases in overall die area, power, and cost.

Digital techniques for voiceband audio began in the early days of PCM telecommunications applications in the 1940s. The early T-carrier systems used 8-bit companding ADCs and expanding DACs, and a sampling frequency of 8 kSPS became the early standard.

Modern digital cellular systems utilize higher-resolution oversampled linear Σ - Δ ADCs and DACs rather than the lower-resolution companding technique. Typical SNR requirements are 60 dB to 70 dB. If companding/expanding is required for compatibility with older systems, it is done in the DSP hardware or software. Voiceband "codecs"¹¹ (coder/decoders) having many applications other than PCM, such as speech processing, encryption, etc., are available in a variety of types.

Sigma-delta ADCs and DACs also dominate the more demanding audio markets, including, for example, FM stereo, computer audio, stereo compact disc (CD), digital audio tape (DAT), and DVD audio. *Total harmonic distortion plus noise* (THD + N) requirements range from 60 dB to greater than 100 dB, and sampling rates range from 48 kSPS to 192 kSPS. Modern CMOS Σ - Δ ADCs and DACs can meet these requirements and also provide the additional digital functions usually associated with such applications.

Pipelined ADCs for High-Speed Applications (Sampling Rates Greater than 5 MSPS)

In this article, we arbitrarily define any application requiring a sampling rate of greater than 5 MSPS as “high speed.” Figure 1 shows that there is an area of overlap between SAR and pipelined ADCs for sampling rates between approximately 1 MSPS and 5 MSPS. Except for this small region, the applications considered *high speed* are most often served by a pipelined ADC. Today, the low-power CMOS pipelined converter is the ADC of choice, not only for the video market but for many others as well. This contrasts strongly with the 1980s, when these markets were served by either the IC flash converter (which dominated the 8-bit video market with sampling rates between 15 MSPS and 100 MSPS) or the higher-resolution, more expensive modular/hybrid solutions. Although low-resolution flash converters remain an important building block for the pipelined ADC, they are rarely used by themselves, except at extremely high sampling rates—generally greater than 1 GHz or 2 GHz—requiring resolutions no greater than 6 bits to 8 bits.

Today, markets that require “high speed” ADCs include many types of *instrumentation* applications (digital oscilloscopes, spectrum analyzers, and medical imaging). Also requiring high-speed converters are video, radar, *communications* (IF sampling, software radio, base stations, set-top boxes, etc.), and *consumer electronics* (digital cameras, display electronics, DVD, enhanced-definition TV, and high-definition TV).

The pipelined ADC has its origins in the *subranging* architecture, first used in the 1950s. A block diagram of a simple 6-bit, two-stage subranging ADC is shown in Figure 11.

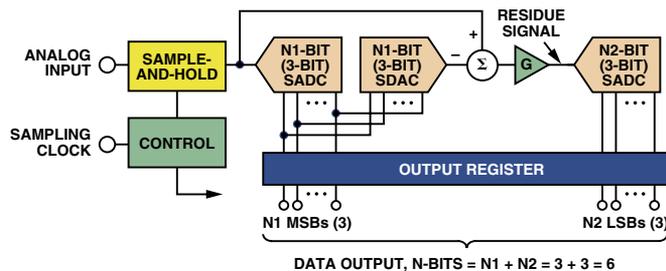


Figure 11. 6-bit, two-stage subranging ADC.

The output of the SHA is digitized by the first-stage 3-bit sub-ADC (SADC)—usually a flash converter. The coarse 3-bit MSB conversion is converted back to an analog signal using a 3-bit sub-DAC (SDAC). Then the SDAC output is subtracted from the SHA output, the difference is amplified, and this “residue signal” is digitized by a second-stage 3-bit SADC to generate the three LSBs of the total 6-bit output word.

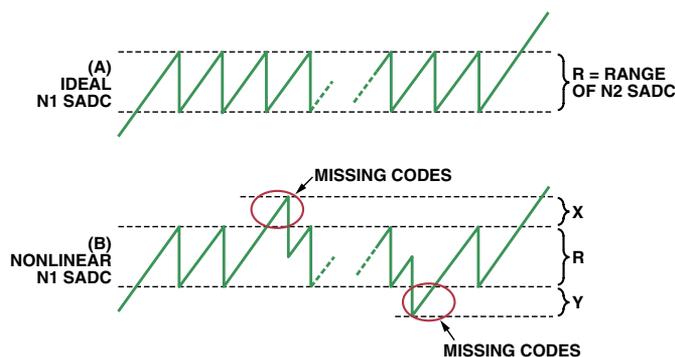


Figure 12. Residue waveform at input of second-stage SADC.

This subranging ADC can best be evaluated by examining the “residue” waveform at the input to the second-stage ADC, as shown in Figure 12. This waveform is typical for a low-frequency ramp signal applied to the analog input of the ADC. In order for there to be no missing codes, the residue waveform must not exceed the input range of the second-stage ADC, as shown in the ideal case of Figure 12A. This implies that both the N_1 -bit SADC and the N_1 -bit SDAC must be accurate to better than $N_1 + N_2$ bits. In the example shown, $N_1 = 3$, $N_2 = 3$, and $N_1 + N_2 = 6$. The situation shown in Figure 12B will result in missing codes when the residue waveform goes outside the range of the N_2 SADC, “R,” and falls within the “X” or “Y” regions—which might be caused by a nonlinear N_1 SADC or a mismatch of interstage gain and/or offset. The ADC output under such conditions might appear as in Figure 13.

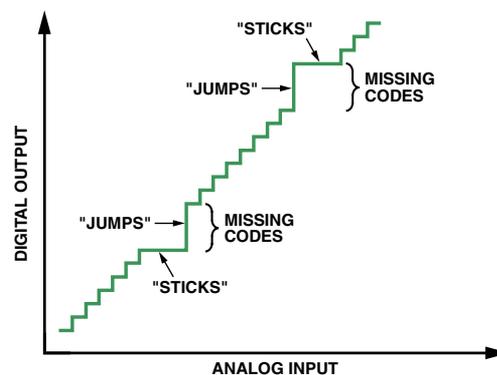


Figure 13. Missing codes due to MSB ADC nonlinearity or interstage misalignment.

This architecture, as shown, is useful for resolutions up to about 8 bits ($N_1 = N_2 = 4$); however maintaining better than 8-bit alignment between the two stages (over temperature variations, in particular) can be difficult. At this point it is worth noting that there is no particular requirement—other than certain design issues beyond the scope of this discussion—for an equal number of bits per stage in the subranging architecture. In addition, there can be more than two stages. Nevertheless, the architecture as shown in Figure 11 is limited to approximately 8-bit resolution unless some form of error correction is added.

The *error-corrected subranging* ADC architecture appeared in the mid-1960s as an efficient means to achieve higher resolutions, while still utilizing the basic subranging architecture. In the two-stage 6-bit subranging ADC, for example, an extra bit is added to the second-stage ADC which allows the digitization of the regions shown as “X” and “Y” in Figure 12. The extra range in the second-stage ADC allows the residue waveform to deviate from its ideal value—provided it does not exceed the range of the second-stage ADC. However, the internal SDAC must still be accurate to more than the overall resolution, $N_1 + N_2$.

A basic 6-bit subranging ADC with error correction is shown in Figure 14, with the second-stage resolution increased to 4 bits, rather than the original 3 bits. Additional logic, required to modify the results of the N_1 SADC when the residue waveform falls in the “X” or “Y” overrange regions, is implemented with a simple adder in conjunction with a dc offset voltage added to the residue waveform. In this arrangement, the MSB of the second-stage SADC controls whether the MSBs are incremented by 001 or passed through unmodified.

It’s worth noting that more than one correction bit can be used in the second-stage ADC, a trade-off—part of the converter design process—beyond the scope of this discussion.

The *error-corrected subranging* ADC shown in Figure 14 does not have a pipeline delay. The input SHA remains in the *hold* mode during the time required for the following events to occur: the first-stage SADC makes its decision, its output is reconstructed by the first-stage SDAC, the SDAC output is subtracted from the SHA output, amplified, and digitized by the second-stage SADC. After the digital data passes through the error correction logic and output registers, it is ready for use; and the converter is ready for another sampling-clock input.

In order to increase the speed of the basic subranging ADC, the “pipelined” architecture shown in Figure 15 has become very popular. This pipelined ADC has a *digitally corrected subranging* architecture—in which each of the two stages operates on the data for one-half of the conversion cycle, and then passes its residue output to the next stage in the “pipeline” prior to the next phase of the sampling clock. The interstage track-and-hold (T/H) serves as an analog delay line—it is timed to enter the hold mode when the first-stage conversion is complete. This allows more settling time for the internal SADCs, SDACs, and amplifiers, and allows the pipelined converter to operate at a much higher overall sampling rate than a nonpipelined version.

There are many design trade-offs that can be made in the design of a pipelined ADC, such as the number of stages, the number of bits per stage, number of correction bits, and the timing. In order to ensure that the digital data from the individual stages corresponding to a particular sample arrives at the error correction logic simultaneously, the appropriate number of shift registers must be added to each of the outputs of the pipelined stages. For example, if the first stage requires seven shift-register delays, the next stage will require six, the next five, etc. This adds the digital pipeline delay to the final output data, as shown in Figure 16, the timing for a typical pipelined ADC, the AD9235.¹²

For the 12-bit, 65-MSPS AD9235, there are seven clock cycles of pipeline delay (sometimes referred to as *latency*). This latency may or may not be a problem, depending upon the application. If the ADC is within a feedback control loop, latency may be a problem—in the overlap area, the successive-approximation architecture would be a better choice. Latency also makes pipelined ADCs difficult to use in multiplexed applications. However, in the bulk of applications for which frequency response is more important than settling time, the latency issue is not a real problem.

A subtle issue relating to most CMOS pipelined ADCs is their performance at *low* sampling rates. Because the internal timing generally is controlled by the external sampling clock, very low sampling rates extend the *hold* times for the internal track-and-holds to the point where excessive droop causes conversion errors. Therefore, most pipelined ADCs have a specification for *minimum* as well as *maximum* sampling rate. Obviously, this precludes operation in single-shot or burst-mode applications—where the SAR ADC architecture is more appropriate.

Finally, it is important to clarify the distinction between *subranging* and *pipelined* ADCs. From the discussions above, it can be seen that, although pipelined ADCs are generally subranging (with error correction, of course), subranging ADCs are not necessarily pipelined. As a matter of fact, the pipelined subranging architecture is predominant because of the demands for high sampling rates, where internal settling time is of utmost importance.

Pipelined ADCs are available today with resolutions of up to 14 bits and sampling rates over 100 MHz. They are ideal for many applications that require not only high sampling rates but high *signal-to-noise ratio* (SNR) and *spurious-free dynamic range* (SFDR). A popular application for these converters today is in software-defined radios (SDR) that are used in modern cellular telephone base stations.

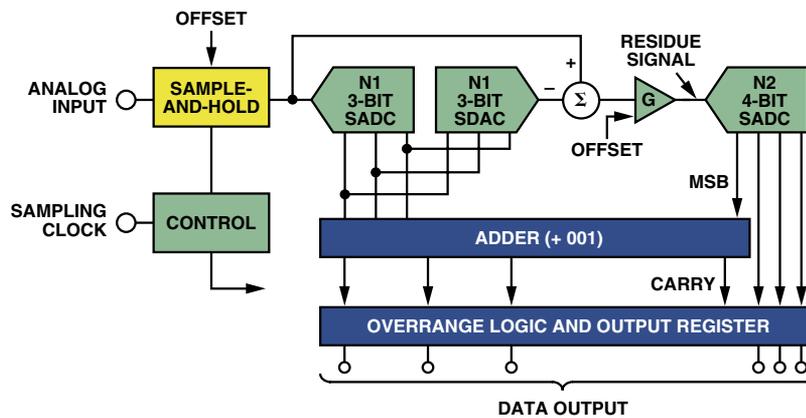


Figure 14. 6-bit subranging error-corrected ADC, N1 = 3, N2 = 4.

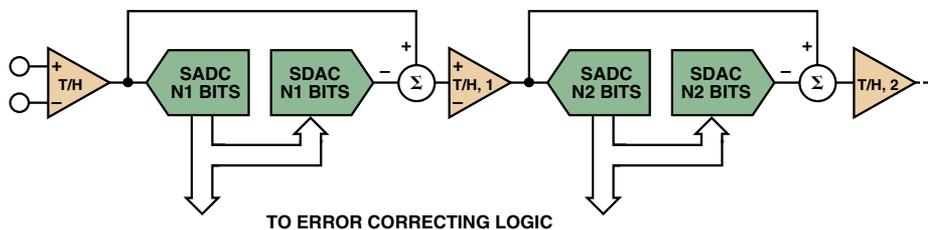


Figure 15. Generalized pipeline stages in a subranging ADC with error correction.

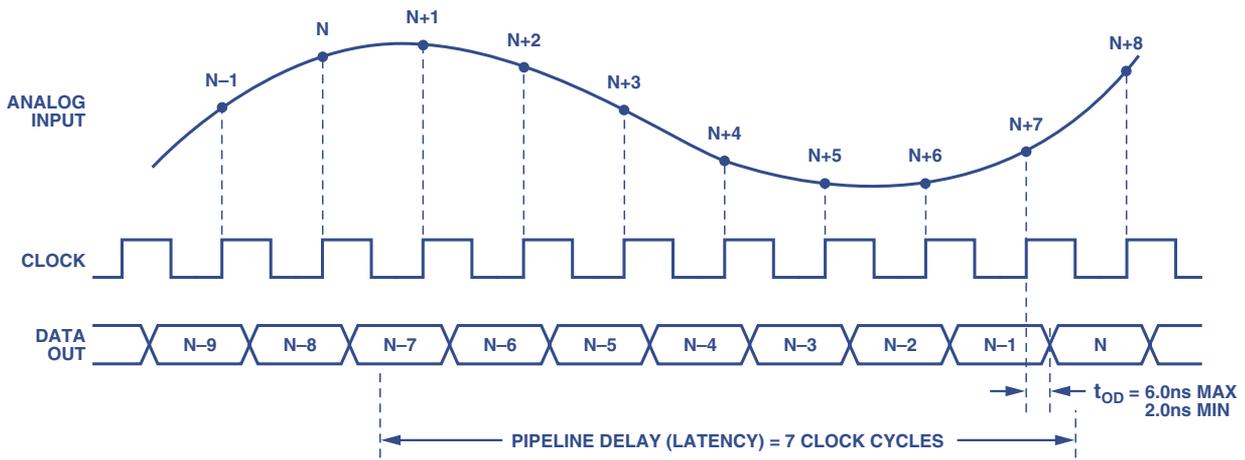


Figure 16. Timing of a typical pipelined ADC, the 12-bit, 65-MSPS AD9235.

Figure 17 shows a simplified diagram of a generic software radio receiver and transmitter. An essential feature is this: rather than digitize each channel separately in the receiver, the entire bandwidth containing many channels is digitized directly by the ADC. The total bandwidth can be as high as 20 MHz, depending on the air standard. The channel-filtering, tuning, and separation are performed digitally in the receive-signal processor (RSP) by a high-performance *digital signal-processor* (DSP).

Digitizing the frequency band at a relatively high intermediate frequency (IF) eliminates several stages of down-conversion. This leads to a lower-cost, more flexible solution in which most of the signal processing is performed digitally—rather than in the more complex analog circuitry associated with standard analog superheterodyne radio receivers. In addition, various air standards (GSM, CDMA, EDGE, etc.) can be processed by the same hardware simply by making appropriate changes in the software. Note that the transmitter in the software radio uses a transmit signal processor (TSP) and DSP to format the individual channels for transmission via the upstream DAC.

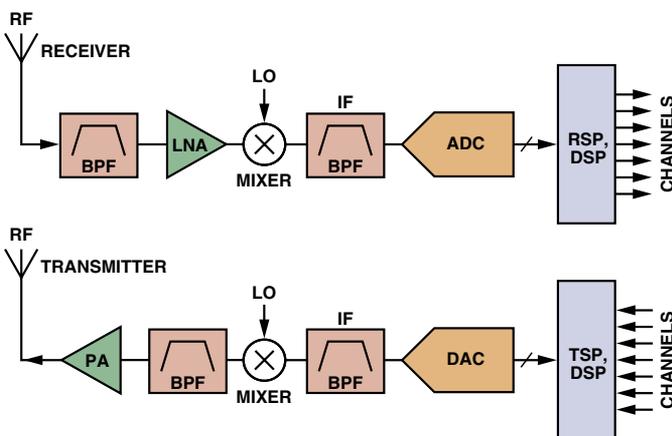


Figure 17. Generic IF-sampling wideband software radio receiver and transmitter.

The ADC requirements for the receiver are determined by the particular air standards the receiver must process. The frequencies in the bandwidth presented to the ADC consist of

the desired signals as well as large-amplitude “interferers” or “blockers.” The ADC must not generate *intermodulation products* due to the blockers, because these unwanted products can mask smaller desired signals. The ratio of the largest expected blocker to the smallest expected signal basically determines the required spurious-free dynamic range (SFDR). In addition to high SFDR, the ADC must have a signal-to-noise ratio (SNR) compatible with the required receiver sensitivity.

Another requirement is that the ADC meet the SFDR and SNR specifications at the desired IF frequency. The basic concept of IF sampling is shown in Figure 18, where a 20-MHz band of signals is digitized at a rate of 60 MSPS. Note how the IF sampling process shifts the signal from the third Nyquist zone to baseband without the need for analog down-conversion. The signal bandwidth of interest is centered in the third Nyquist zone at an IF frequency of 75 MHz. The numbers chosen in this example are somewhat arbitrary, but they serve to illustrate the concept of undersampling. These applications place severe requirements on the ADC performance, especially with respect to SNR and SFDR. Modern pipelined ADCs, such as the 14-bit, 80-MSPS AD9444,¹³ can meet these demanding requirements. For instance, the AD9444 has an SFDR of 97 dBc and an SNR of 73 dB with a 70-MHz IF input. The input bandwidth of the AD9444 is 650 MHz. Other 14-bit ADCs optimized for SFDR and/or SNR are the AD9445¹⁴ and AD9446.¹⁵

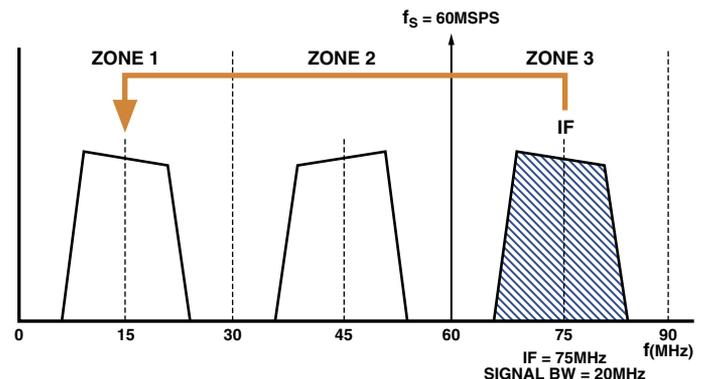


Figure 18. Sampling a 20-MHz BW signal with an IF frequency of 75 MHz at a sampling rate of 60 MSPS.

CONCLUSION

We have discussed here the successive approximation, Σ - Δ , and pipelined architectures—those most widely used in modern integrated circuit ADCs.

Successive-approximation is the architecture of choice for nearly all multiplexed data acquisition systems, as well as many instrumentation applications. The SAR ADC is relatively easy to use, has no pipeline delay, and is available with resolutions to 18 bits and sampling rates up to 3 MSPS.

For a wide variety of industrial measurement applications, the *sigma-delta* ADC is ideal; it is available in resolutions from 12 bits to 24 bits. Sigma-delta ADCs are suitable for a wide variety of sensor-conditioning, energy-monitoring, and motor-control applications. In many cases, the high resolution and the addition of on-chip PGAs allow a direct connection between the sensor and the ADC without the need for an instrumentation amplifier or other conditioning circuitry.

The Σ - Δ ADC and DAC, easily integrated into ICs containing a high degree of digital functionality, also dominate the voiceband and audio markets. The inherent oversampling in these converters greatly relaxes the requirements on the ADC antialiasing filter and the DAC reconstruction filter.

For sampling rates greater than approximately 5 MSPS, the *pipelined architecture* dominates. These applications typically require resolutions up to 14 bits with high SFDR and SNR at sampling frequencies ranging from 5 MSPS to greater than 100 MSPS. This class of ADCs is used in many types of instrumentation, including digital oscilloscopes, spectrum analyzers, and medical imaging. Other applications are video, radar, and communications applications—including IF sampling, software radio, base stations, and set-top boxes—and consumer electronics equipment, such as digital cameras, display electronics, DVDs, enhanced definition TVs, and high-definition TVs.

The use of manufacturers' selection guides and parametric search engines, coupled with a fundamental knowledge of the three basic architectures, should help the designer select the proper ADC for the application. The use of manufacturers' [evaluation boards](#)¹⁶ makes the process much easier. The Analog Devices' [ADIsimADC® program](#)¹⁷ allows the customer to evaluate the dynamic performance of the ADC without the need for any hardware. The required software and the ADC models (and many other analog and digital design aids) are

free downloads at <http://www.analog.com>.¹⁸ This tool can be extremely valuable in the selection process.

Not to be overlooked is the proper design of the ADC input-, output-, and sampling-clock circuitry. Data sheets and application notes should be consulted regarding these important issues. Finally, and equally critical to achieving a successful mixed-signal design, are layout, grounding, and decoupling. For a detailed treatment of these and other design issues, the reader is encouraged to consult the two comprehensive texts listed in the Further Reading as well as the Analog Devices website, <http://www.analog.com>. 

FURTHER READING

1. Walt Kester, Editor, *Data Conversion Handbook*, Published by Newnes, an imprint of Elsevier, 2005, ISBN: 0-7506-7841-0. See in particular Chapter 3, "Data Converter Architectures." In addition to detailed discussions of the various ADC and DAC architectures themselves, the chapter also includes historical aspects.
2. Walt Jung, Editor, *Op Amp Applications Handbook*, Published by Newnes, an imprint of Elsevier, 2005, ISBN: 0-7506-7844-5.
3. For more information on products or applications, please visit the Analog Devices, Inc., website: <http://www.analog.com>.

REFERENCES—VALID AS OF MAY 2005

- 1 <http://www.analog.com/en/cat/0,2878,760,00.html>
- 2 ADI website: www.analog.com (Search) AD7621 (GO)
- 3 ADI website: www.analog.com (Search) AD7641 (GO)
- 4 ADI website: www.analog.com (Search) AD574A (GO)
- 5 ADI website: www.analog.com (Search) AD7908 (GO)
- 6 ADI website: www.analog.com (Search) AD7918 (GO)
- 7 ADI website: www.analog.com (Search) AD7928 (GO)
- 8 ADI website: www.analog.com (Search) AD5555 (GO)
- 9 ADI website: www.analog.com (Search) AD8230 (GO)
- 10 ADI website: www.analog.com (Search) AD7799 (GO)
- 11 http://www.analog.com/audio_codec_ic.html
- 12 ADI website: www.analog.com (Search) AD9235 (GO)
- 13 ADI website: www.analog.com (Search) AD9444 (GO)
- 14 ADI website: www.analog.com (Search) AD9445 (GO)
- 15 ADI website: www.analog.com (Search) AD9446 (GO)
- 16 http://www.analog.com/en/cList/0,2880,760__15,00.html
- 17 http://www.analog.com/Analog_Root/static/designCenter/analogDigitalConverters/design.html#ADIsimADC
- 18 <http://www.analog.com>