High Speed System Applications

SECTION 1
HIGH SPEED DATA CONVERSION OVERVIEW

Data Converter Selection: It's Not Just Bits and Speed

- **Making trade-offs involves many variables**
- **AC and DC performance**
- **Power consumption**
- **Degree of integration**
- **Ease of use**
  - **Output data formatting**
  - **Supply voltages**
  - **Package size**
  - **Integrated Functionality**
- **Cost**
- **Reputation of IC vendor**
  - **Design tools**
  - **Applications expertise**
  - **Clear, concise documentation**
  - **Help with product selection**

Selecting the proper ADC for a particular application can be a formidable task, especially considering the thousands of converters currently on the market. A traditional approach is to go right to the selection guides and parametric search engines, such as those available on the Analog Devices' website. Enter the sampling rate, resolution, power supply voltage, etc., click the "find" button, and hope for the best. Could there be a more productive way to approach the task?

Today's data converters are differentiated by much more than simply resolution and speed (sampling rate). This complicates the selection process even more. This section discusses the basic architectures, performance, and applications for high speed converters—understanding these fundamentals will greatly assist in the selection and application of the devices.

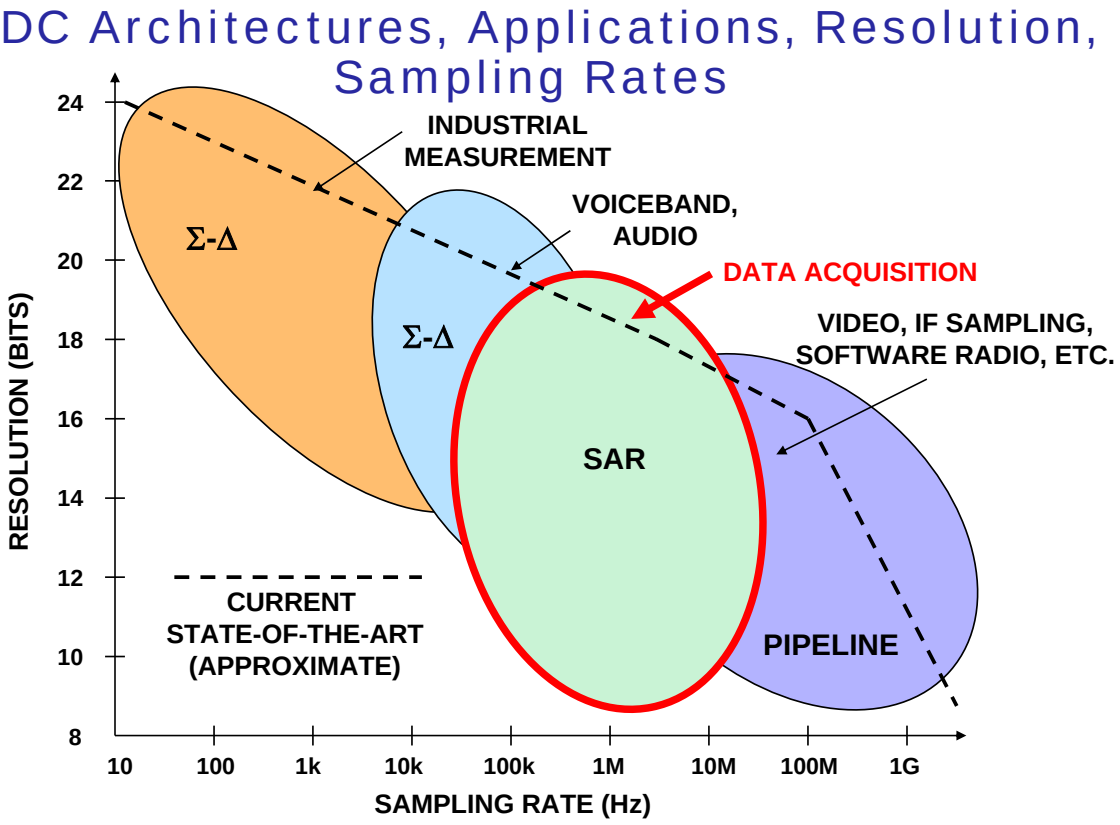

Most ADC applications today can be divided into four broad market segments: *Data Acquisition*, *Precision Industrial Measurement*, *Voiceband and Audio*, and *High Speed ("High Speed" implying sampling rates greater than approximately 10MSPS—although this line of demarcation is somewhat arbitrary. For instance, a 2MSPS sampling rate certainly qualifies as "high speed" for an 18-bit SAR ADC)*. A very large percentage of these applications can be filled with either the *Successive Approximation (SAR)*, *Sigma-Delta (Σ-Δ)*, or *Pipelined* ADC. A basic understanding of the three most popular ADC architectures is therefore valuable in selecting the proper ADC for a given application.

This figure shows approximately how these application segments and architectures relate to ADC resolution (vertical axis) versus sampling rate (horizontal axis). The dotted line represents the approximate state-of-the art today (2006). Even though there is considerable overlap between the coverage of the various architectures, the applications themselves differentiate between the specific architecture required.

The Sigma-Delta architecture dominates the precision industrial measurement, voiceband, and audio application space. This architecture is discussed in detail in the references below. The focus of this discussion will be on the SAR and Pipeline ADC architectures.

The SAR architecture will be discussed first, and it dominates the data acquisition application space, especially where multiple channels must be digitized.

Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN-0916550273 Chapter 3. Also available as *Data Conversion Handbook*, Elsevier-Newnes, 2005, ISBN: 0750678410, Chapter 3.

Successive Approximation ADCs

www.analog.com/pulsar

www.analog.com/icmos

Typical Muxed Data Acquisition System

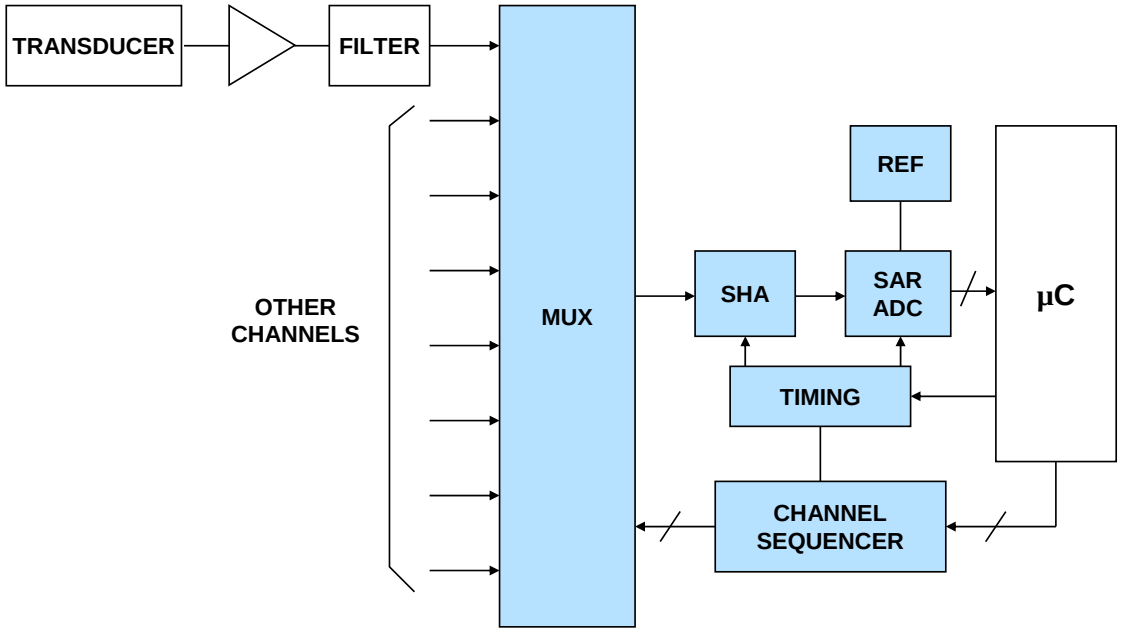

This figure shows a typical multiplexed data acquisition system. The successive approximation (SAR) ADC is the building block.

Rather than utilize a separate ADC per channel, the use of an analog multiplexer allows the conversion process to be accomplished with a single ADC. In the early days of integrated circuits, separate ICs were used for the multiplexer, sample-and-hold, reference, and SAR ADC. The user had to design the necessary timing and channel sequencing circuitry.

Modern IC technology allows all these functions (shaded) to be integrated into single packages, thereby providing complete data acquisition on a chip.

The following discussion will illustrate why the SAR ADC provides the optimum architecture for these systems.

Further description of the SAR architecture can be found in the following references.

---

1. Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN-0916550273 Chapter 1 and 3. Also available as *Data Conversion Handbook*, Elsevier-Newnes, 2005, ISBN: 0750678410, Chapter 1 and 3.

2. Tutorial MT-021, *Successive Approximation ADCs*, www.analog.com.

Basic Successive Approximation ADC (Feedback Subtraction ADC)

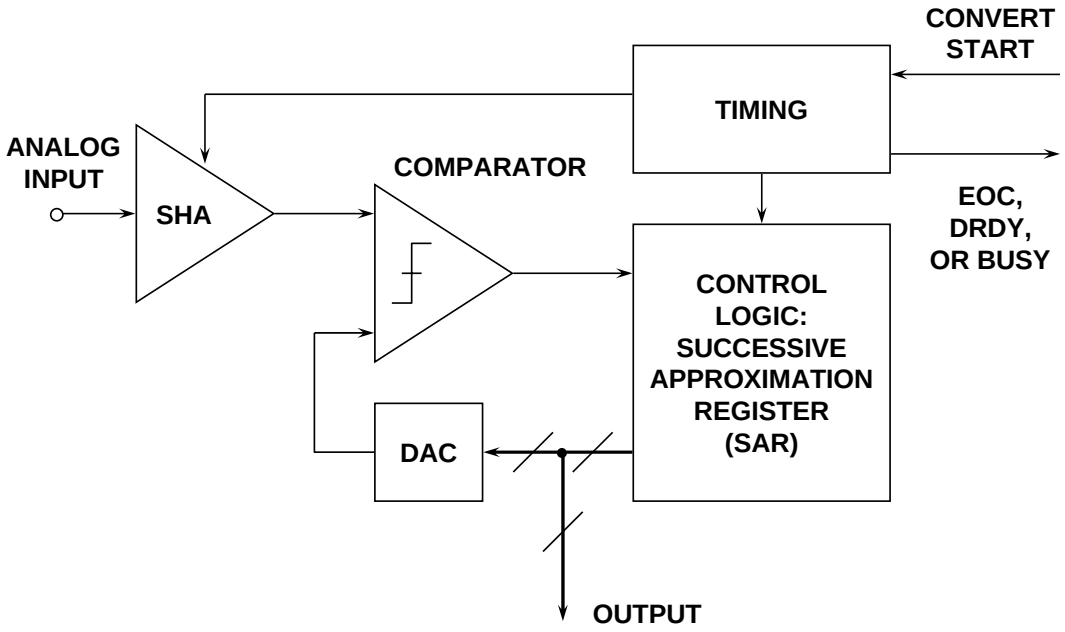

The SAR ADC performs conversions on command. On the assertion of the CONVERT START command (note that this function may actually be named something else or be combined with another control line), the sample-and-hold (SHA) is placed in the *hold* mode, and all the bits of the successive approximation register (SAR) are reset to "0" except the MSB which is set to "1". The SAR output drives the internal DAC. If the DAC output is greater than the analog input, this bit in the SAR is reset, otherwise it is left set. The next most significant bit is then set to "1". If the DAC output is greater than the analog input, this bit in the SAR is reset, otherwise it is left set. The process is repeated with each bit in turn. When all the bits have been set, tested, and reset or not as appropriate, the contents of the SAR correspond to the value of the analog input, and the conversion is complete. These bit "tests" can form the basis of a serial output version SAR-based ADC.

The basic accuracy of the SAR ADC is determined by the internal DAC.

---

1. Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN-0916550273 Chapter 1 and 3. Also available as *Data Conversion Handbook*, Elsevier-Newnes, 2005, ISBN: 0750678410, Chapter 1 and 3.

2. Tutorial MT-021, *Successive Approximation ADCs*, www.analog.com.

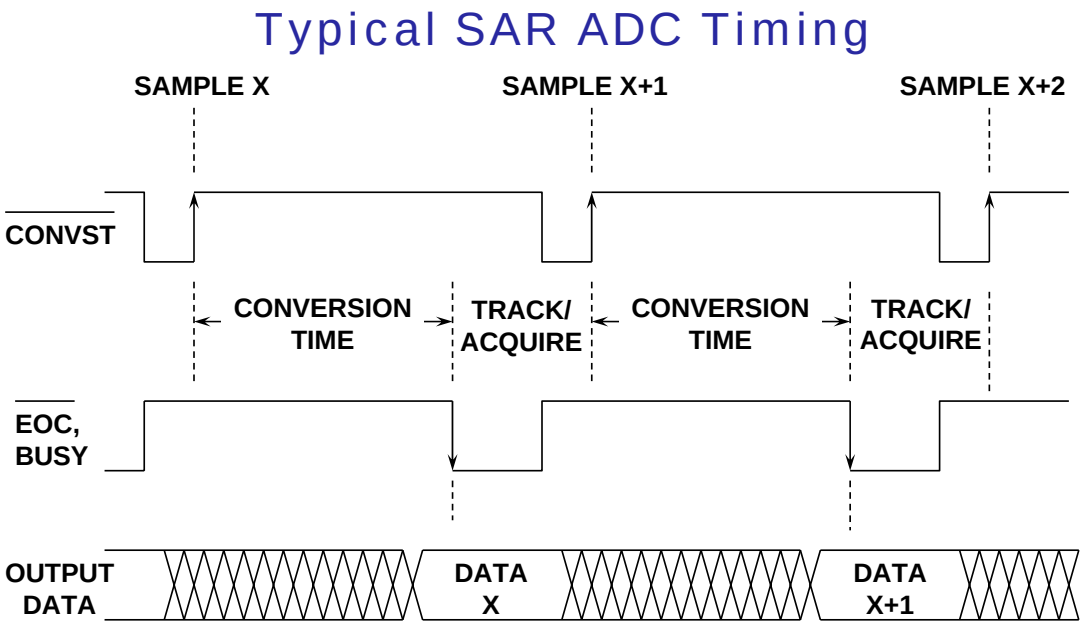

This figure shows the basic timing of a typical SAR ADC. The end of conversion is generally indicated by an end-of-convert (EOC), data-ready (DRDY), or a busy signal (actually, *not*-BUSY indicates end of conversion). The polarities and name of this signal may be different for different SAR ADCs, but the fundamental concept is the same. At the beginning of the conversion interval, the signal goes high (or low) and remains in that state until the conversion is completed, at which time it goes low (or high). The trailing edge is generally an indication of valid output data, but the data sheet should be carefully studied—in some ADCs additional delay is required before the output data is valid. An N-bit conversion takes N steps.

The exact labels assigned to these functions can vary from converter to converter, but are generally present in most SAR ADCs.

It should also be noted that some SAR ADCs require an external high frequency clock in addition to the CONVERT START command. In most cases, there is no need to synchronize the two. The frequency of the external clock, if required, generally falls in the range of 1MHz to 30MHz depending on the conversion time and resolution of the ADC. Other SAR ADCs have an internal oscillator which is used to perform the conversions and only require the CONVERT START command. Because of their architecture, SAR ADCs generally allow single-shot conversion at any repetition rate from dc to the converter's maximum conversion rate.

Note that at the end of the conversion time, the data corresponding to the sampling clock edge is available with no "pipeline" delay. Unlike most "pipelined" ADCs, the SAR ADC generally has no "minimum" specified sampling rate. SAR ADCs can be operated continuously, or in a "single-shot" mode. This feature is extremely useful in multiplexed applications.

The basic SAR is a serial output device. Although parallel output versions are available, the trend is toward the serial interface (SPI, I^2C, etc.) because of reduced pin count, package size, and cost.

3-Bit Switched Capacitor DAC

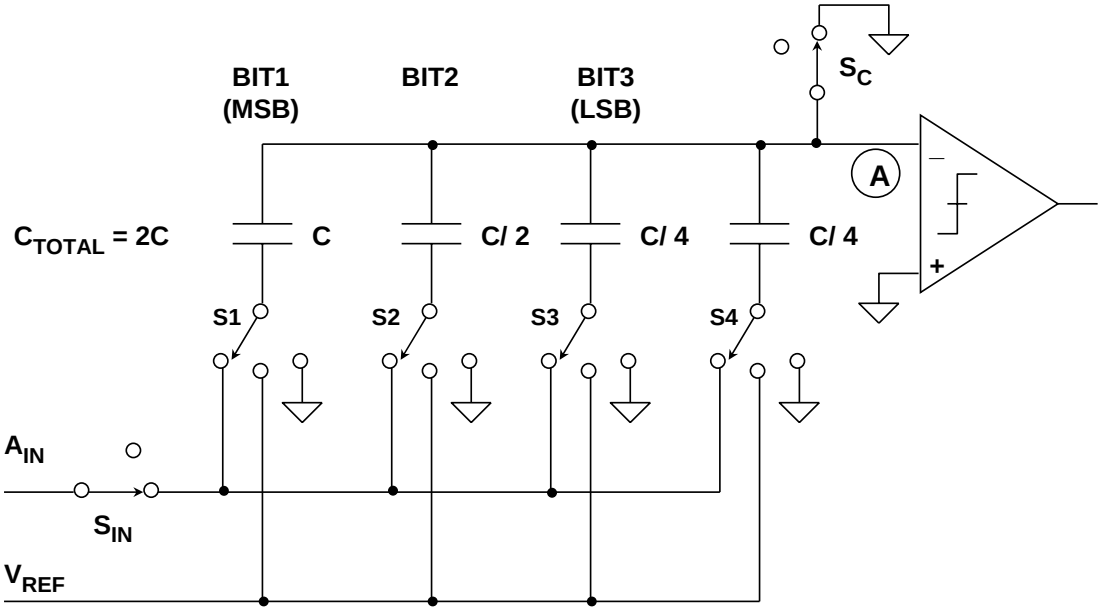

SWITCHES SHOWN IN TRACK (SAMPLE) MODE

The accuracy and linearity of the internal DAC determines the accuracy and linearity of the overall SAR ADC. Early SAR ADCs, such as the industry standard AD574, used thin film laser wafer trimmed internal DACs. Today, this approach has been replaced by switched capacitor (often called charge redistribution) CMOS DACs shown here. The capacitor matching is controlled by the precise lithography, and extra capacitors and switches can be added for trimming either at the factory or as part of autocalibration routines run at the system level after installation. The following describes the operation of the 3-bit DAC shown in the figure.

The switches are shown in the track, or sample mode where the analog input voltage, A_{IN}, is constantly charging and discharging the parallel combination of all the capacitors. The hold mode is initiated by opening S_{IN}, leaving the sampled analog input voltage on the capacitor array. Switch S_C is then opened allowing the voltage at node A to move as the bit switches are manipulated. If S1, S2, S3, and S4 are all connected to ground, a voltage equal to $-A_{IN}$ appears at node A. Connecting S1 to V_{REF} adds a voltage equal to $V_{REF}/2$ to $-A_{IN}$. The comparator then makes the MSB bit decision, and the SAR either leaves S1 connected to V_{REF} or connects it to ground depending on the comparator output (which is high or low depending on whether the voltage at node A is negative or positive, respectively). A similar process is followed for the remaining two bits. At the end of the conversion interval, S1, S2, S3, S4, and S_{IN} are connected to A_{IN}, S_C is connected to ground, and the converter is ready for another cycle.

Note that the extra LSB capacitor (C/4 in the case of the 3-bit DAC) is required to make the total value of the capacitor array equal to 2C so that binary division is accomplished when the individual bit capacitors are manipulated.

The operation of the capacitor DAC is similar to an R-2R resistive DAC. When a particular bit capacitor is switched to V_{REF}, the voltage divider created by the bit capacitor and the total array capacitance (2C) adds a voltage to node A equal to the weight of that bit. When the bit capacitor is switched to ground, the same voltage is subtracted from node A.

Modern 12-Bit 1.5MSPS SAR ADC with 8-Channel Input Multiplexer

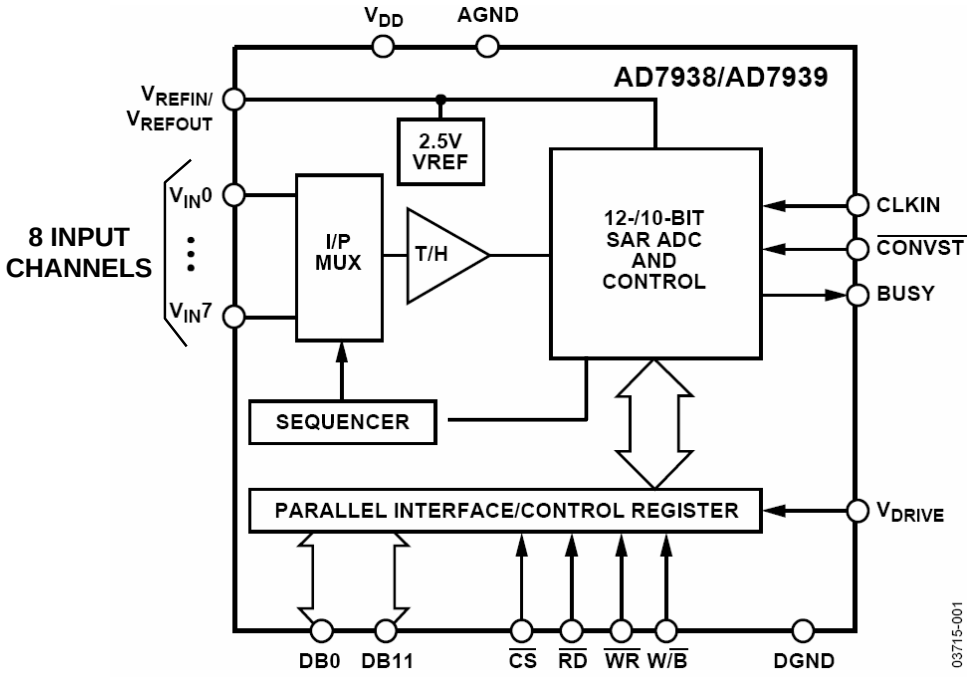

One drawback to the switched capacitor SAR architecture is the switching transient currents which can be injected on the analog input. This requires that the drive amplifier settle to these transient currents within approximately one-half of the conversion period. However, many SAR ADCs using the switched capacitor techniques can be driven directly from the signal source, provided the driving impedance is less than several kΩ.

Many modern general-purpose ADCs have on-chip multiplexers as shown here for the AD79xx series of 1MSPS ADCs. The AD7938/AD7939 are 12- and 10-bit, high speed, low power, successive approximation (SAR) ADCs which supply a parallel data output. A simplified block diagram is shown in the figure. The parts operate from a single 2.7V to 5.25V power supply and feature throughput rates up to 1.5MSPS. The parts contain a low noise, wide bandwidth, differential track/hold amplifier that can handle input frequencies up to 20MHz. The AD7938/AD7939 feature eight analog input channels with a channel sequencer to allow a pre-programmed selection of channels to be converted sequentially. These parts can operate with either single-ended, fully differential or pseudo-differential analog inputs. The analog input configuration is chosen by setting the relevant bits in the on-chip Control Register.

The AD7938/AD7939 has an accurate on-chip 2.5V reference that can be used as the reference source for the analog to digital conversion. Alternatively, this pin can be overridden to provide an external reference in the range 100mV to 3.5V. The pin can be optionally used for additional noise filtering if desired.

These parts use advanced design techniques to achieve very low power dissipation at high throughput rates. They also feature flexible power management options. An on-chip Control Register allows the user to set up different operating conditions including analog input range and configuration, output coding, power management, and channel sequencing. The parts are available in a 32-lead LFCSP package.

SAR ADCs are popular in multichannel applications because they have no "pipeline" delay as in the case of many of the other ADC architectures.

AD7641 18-Bit, 2 MSPS PulSAR® ADC

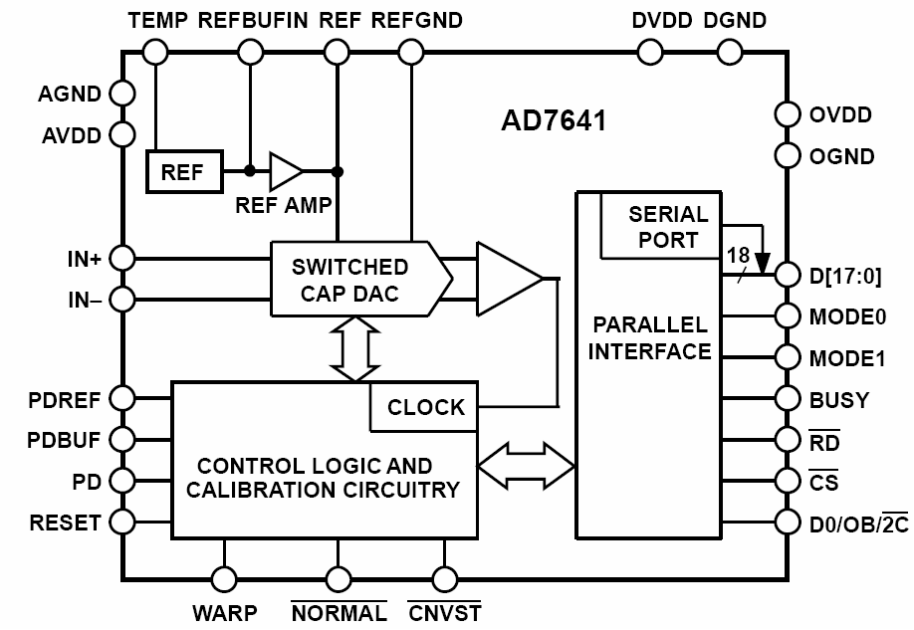

The technology in SAR ADCs has pushed the sampling rate to 2MSPS and resolution to 18-bits, as illustrated in the AD7641 PulSAR® shown in the figure.

The AD7641 is an 18-bit, 2MSPS, charge redistribution switched-capacitor SAR ADC, fully differential, analog-to-digital converter (ADC) that operates from a single 2.5V power supply. The part contains a high speed, 18-bit sampling ADC, an internal conversion clock, an internal reference (and buffer), error correction circuits, and both serial and parallel system interface ports. It features two very high sampling rate modes ("wideband warp" and "warp") and a fast mode (normal) for asynchronous rate applications. The AD7641 is hardware factory calibrated and tested to ensure AC parameters, such as signal-to-noise ratio (SNR), in addition to the more traditional dc parameters of gain, offset, and linearity, with operation specified from −40°C to +85°C.

SNR is 93.5dB typical, and THD is −112dB typical, for a 20kHz input (V_{REF} = 2.5V).

Other members of the PulSAR family of SAR ADCs can be found at www.analog.com/pulsar.

Interfacing Industrial-Level Bipolar Signals to Low-Voltage ADCs

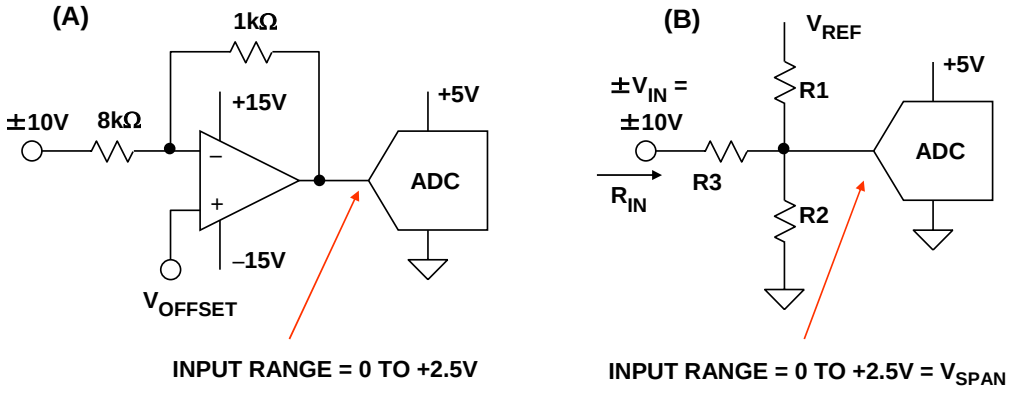

Many industrial applications still require ADCs that can handle ±10V signals

In recent years the trend in SAR ADCs has been toward lower power and lower supply voltages (typically less than 5V). Input signal ranges have decreased proportionally, and 2.5V fullscale is common using the lower supplies.

There is, however, a large industrial application base which requires digitization of signals up to ±10V. This figure shows two somewhat suboptimal approaches for interfacing the large signal to the single-supply ADCs.

In the circuit of (A) an op amp is used to level shift and attenuate the ±10V signal so that it "fits" the input span of the ADC, 0V to +2.5V. The obvious disadvantages here are the extra components and the load presented to the source by the feedforward resistor—in this case, 8kΩ. Another less obvious problem with the circuit is that since the op amp operates on ±15V supplies, it can overdrive the ADC unless some type of clamping circuit is added. Also, care must be taken that the +5V ADC supply is brought up before the op amp supplies; otherwise the ADC may latch up, depending on its input structure.

In the circuit of (B) a resistor network is used to accomplish the attenuation and level shifting. This approach requires three resistors as well as a voltage reference. It also presents a load to the source.

The resistors can be calculated using the following equations:

$$R_{IN} = R3 + R1||R2$$

$$(R1||R2) / (R3 + R1||R2) = V_{SPAN}/2V_{IN}$$

$$\{(R2||R3) / (R1 + R2||R3)\}\ V_{REF} = V_{SPAN}/2$$

AD7328 13-Bit, 1MSPS *i*CMOS™ ADC wih True Bipolar Inputs

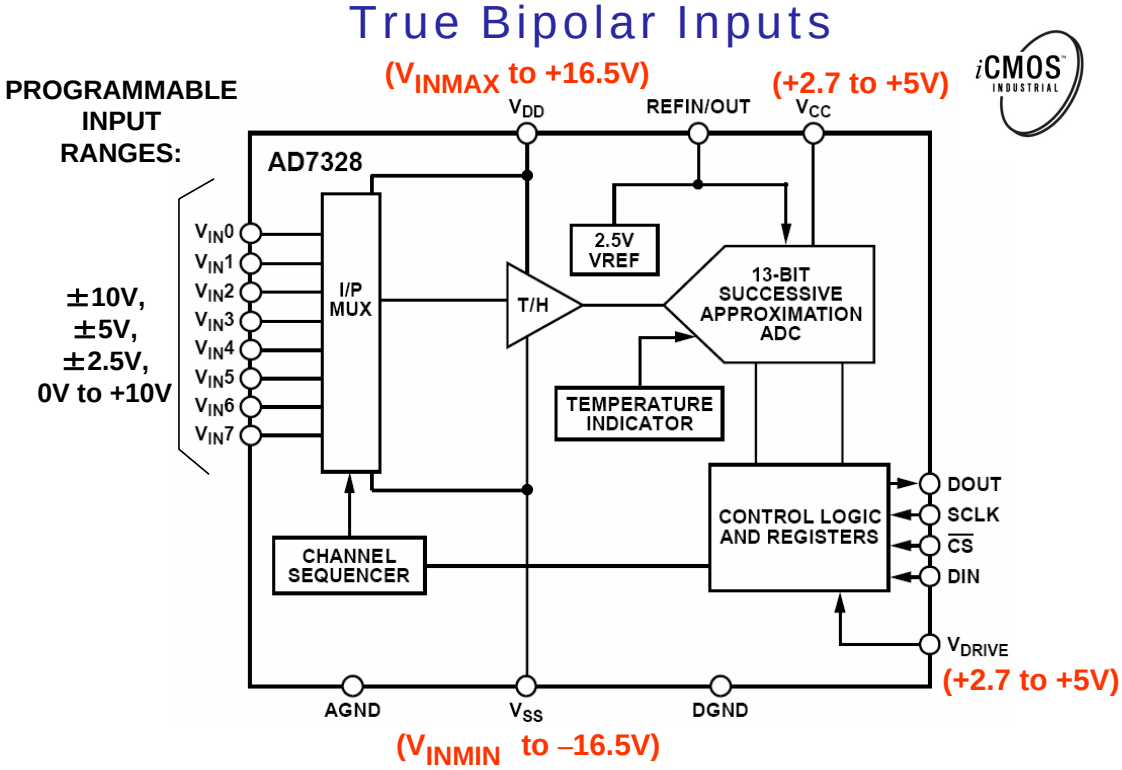

A much better solution available from Analog Devices uses a proprietary Industrial CMOS (*i*CMOS™) process which allows the input circuitry to operate on standard industrial ±15V supplies, while operating the ADC core on the low voltage supply (5V or less). This figure shows the AD7328 13-bit 8-channel input ADC, one of a number of *i*CMOS ADCs.

*i*CMOS is a process combining high voltage CMOS and low voltage CMOS. It enables the development of a wide range of high performance analog ICs capable of 33V operation in a footprint that no previous generation of high voltage parts could achieve. Unlike analog ICs using conventional CMOS processes, iCMOS components can accept bipolar input signals while providing increased performance, dramatically reducing power consumption, and having a reduced package size.

The AD7328 can accept true bipolar analog input signals. The AD7328 has four software-selectable input ranges, ±10V, ±5 V, ±2.5V, and 0V to 10V. Each analog input channel can be independently programmed to one of the four input ranges. The analog input channels on the AD7328 can be programmed to be single-ended, true differential, or pseudo differential. The ADC contains a 2.5V internal reference. The AD7328 also allows for external reference operation. If a 3V external reference is applied to the REFIN/OUT pin, the AD7328 can accept a true bipolar ±12V analog input. Minimum ±12V V_{DD} and V_{SS} supplies are required for the ±12V input range.

The low voltage core of the AD7328 operates on the V_{CC} supply which should be 5V nominal (4.75V to 5.5V) for specified performance. For V_{CC} between 2.7V and 4.75V, the AD7328 will meet its typical specifications. The AD7328 has a separate V_{DRIVE} pin which sets the I/O logic interface voltage (2.7V to 5.5V). The V_{DRIVE} voltage should not exceed V_{CC} by more than 0.3V.

The AD7328 has a high speed serial interface that can operate at throughput rates up to 1MSPS.

Other *i*CMOS products can be found at www.analog.com/icmos.

Pipelined ADCs

www.analog.com/adcs

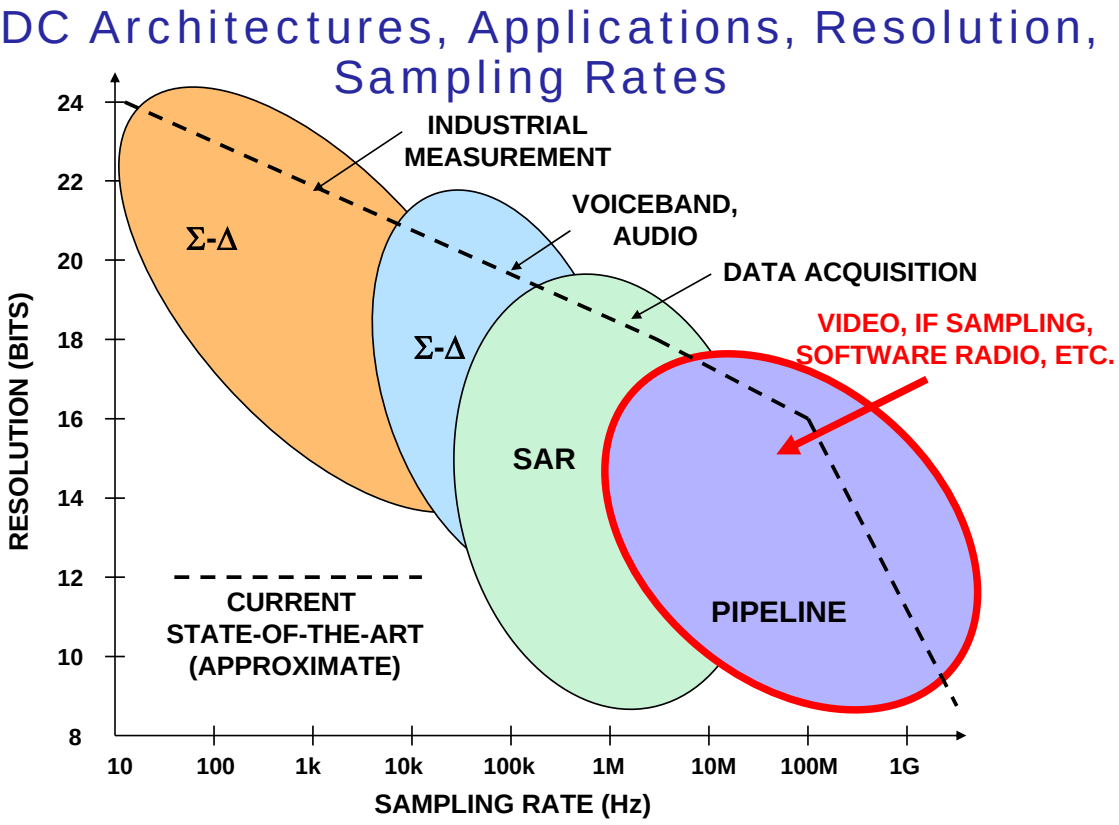

For the purposes of this discussion, "high speed" will be defined as sampling rates greater than about 10MSPS. Included in this application space are various types of instrumentation (digital oscilloscopes, digital spectrum analyzers), medical imaging, radar, IF sampling (including software radio), etc.

These applications are most often served by the "pipelined" ADC. Although there is some overlap between SAR ADCs and pipelined ADCs in the region of 1MSPS to 10MSPS, the applications themselves usually help determine which architecture is more appropriate.

3-Bit All-Parallel (Flash) Converter

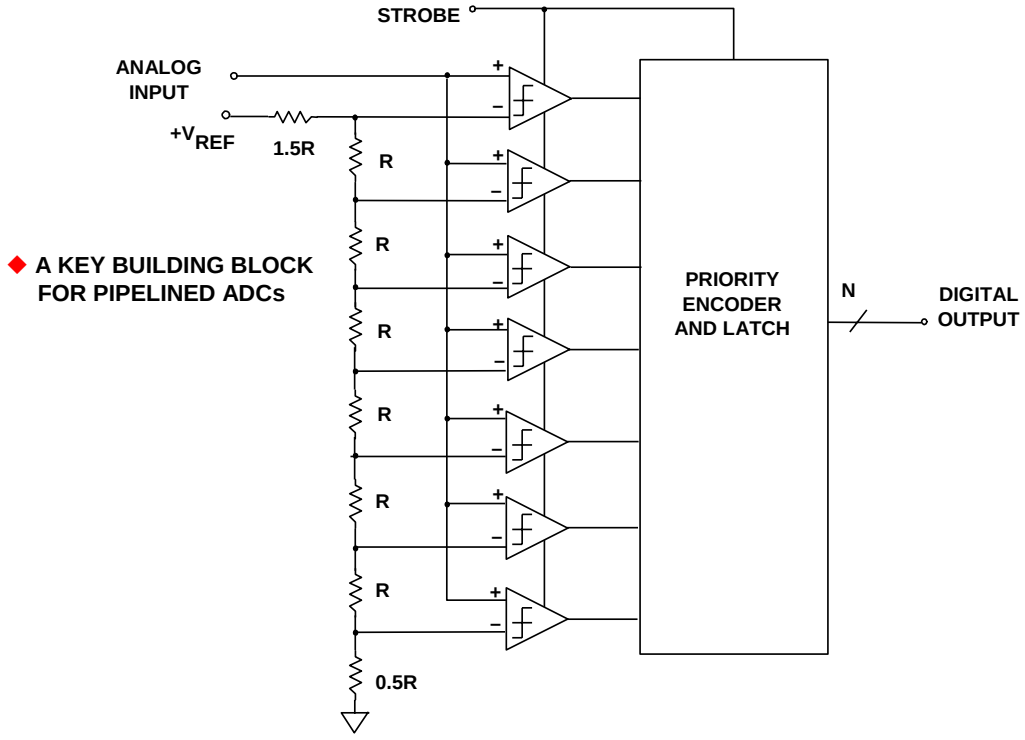

Because of its importance as a building block in subranging pipelined ADCs, one must first understand the basic flash converter.

The flash converter makes use of parallel comparators, each operating at a slightly different reference voltage determined by the resistor ladder network. An N-bit flash converter requires $2^N - 1$ latched comparators. Therefore, the technique is rarely used beyond 8-bits because of power dissipation and die size (cost).

The comparators are latched simultaneously; therefore, a separate SHA is not generally required. However, mismatches in timing between the comparators may require an external SHA for optimum performance at high input slew rates.

The output of the comparator bank is a thermometer code, which is decoded into the proper binary code by the decoding logic. Conceptually, the decoding logic is a priority encoder, but it may be more complicated to correct for comparator metastable state errors.

IC flash ADCs became extremely popular in the 1980s especially in the 8-bit, 20MSPS to 100MSPS sampling range. Today, however, standalone flash ADCs are mostly used at sampling rates of 1GSPS or higher for six to eight bits of resolution, and are high-power GaAs devices.

Other techniques, such as the subranging pipelined architecture, use lower power, lower cost CMOS processes to accomplish 8-bit to 14-bit resolution up to several hundreds of MSPS.

Low resolution flash ADCs are still used as building blocks in various subranging pipelined ADCs and in multi-bit sigma-delta ADCs.

Tutorial MT-024, *Pipelined Subranging ADCs*, Analog Devices, www.analog.com.

6-Bit Two-Stage Subranging ADC

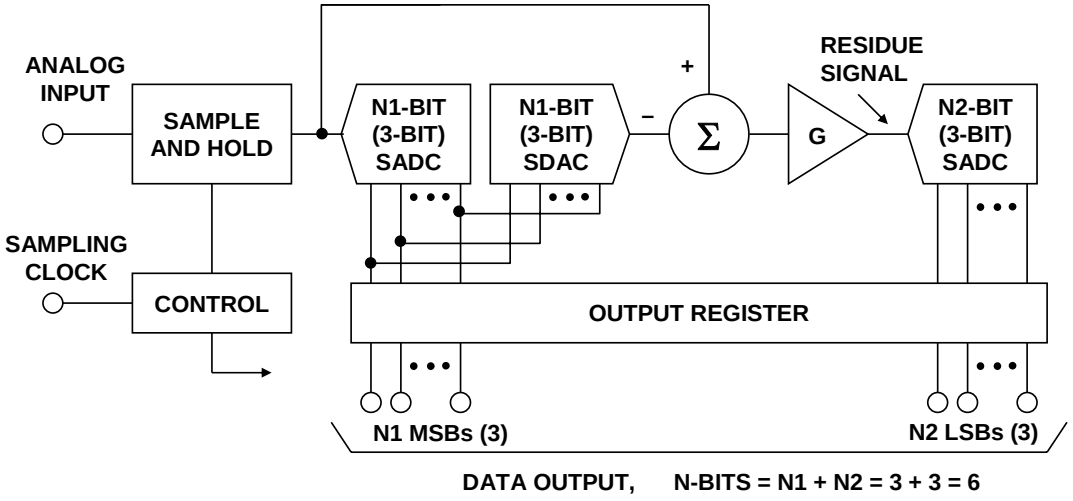

See: R. Staffin and R. Lohman, "Signal Amplitude Quantizer,"
U.S. Patent 2,869,079, Filed December 19, 1956, Issued January 13, 1959

Subranging ADCs were first documented in the mid 1950s as shown by this patent reference. This diagram shows a two-stage subranging ADC, but the concept can be continued to more than two stages. A single stage can also be used a number of times by "recirculating" the analog data using switches and a PGA.

There is a "coarse" conversion of N1 bits followed by a "fine" conversion of N2 bits. The individual sub-ADCs (labeled SADC) are generally flash converters, but do not have to be.

Subranging ADCs do not necessarily have to exhibit pipeline delay, but most do, in practice. On the other hand, a pipelined ADC is almost always subranging.

The N1-bit coarse conversion is converted back to analog by an N1 bit SDAC, subtracted from the held analog signal, amplified, and applied to the N2-bit SADC.

Note that the N1 bit SADC and SDAC must be accurate to better than N1 + N2 bits, even though their resolution is less. This type of ADC can be analyzed better by examining the "residue signal" into the second stage.

Residue Waveforms at Input of N2 SADC

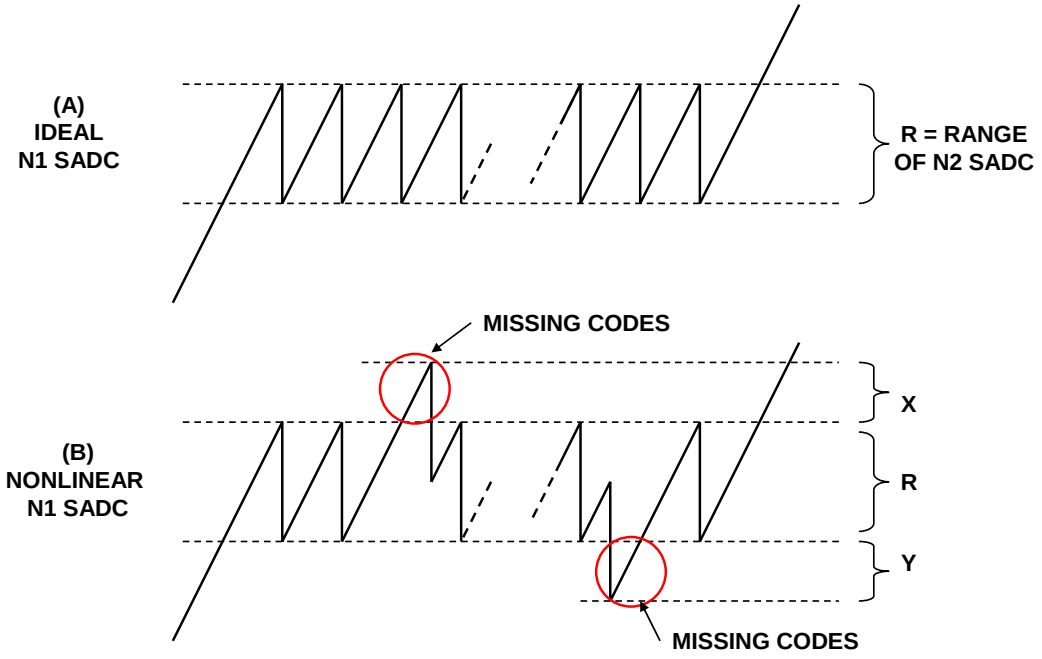

The residue waveform into the second N2 SADC must exactly fill the range of the N2SADC as shown in (A). Otherwise, as shown in (B), there will be nonlinearities in the overall transfer function, and possibly missing codes.

These nonlinearities can come from the N1SADC, the N1SDAC, or gain or offset errors in the summation amplifier, G. It is difficult to construct two-stage subranging ADCs with overall resolutions of greater than eight bits because of the effects of the first stage errors.

We will see shortly how expanding the resolution of the second stage ADC and the use of digital error correction techniques can minimize the effects of the first stage conversion errors on the overall ADC transfer function.

6-Bit Subranging Error Corrected ADC
N1 = 3, N2 = 4

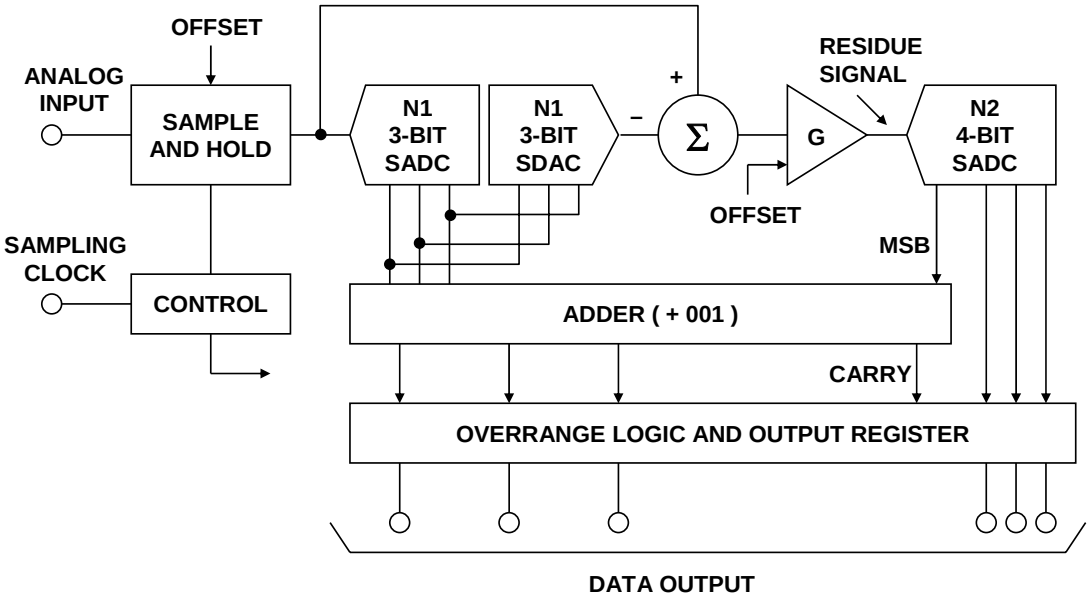

SEE: T. C. Verster, "A Method to Increase the Accuracy of Fast Serial-Parallel Analog-to-Digital Converters," *IEEE Transactions on Electronic Computers*,EC-13, 1964, pp. 471-473

The concept of digital error correction in a subranging ADC was implemented in the mid-1960s as shown in this reference.

In practice, rather than adding or subtracting 001 to the MSBs, an offset can be added to the residue signal so that the MSBs are either passed through to the output unmodified, or with 001 added to them. This simplifies the logic.

This figure shows a 6-bit subranging error corrected ADC with three bits in the first stage and four bits in the second stage. The extra bit in the second stage adds the additional range. The second stage MSB controls the digital adder.

There is no theoretical reason why more bits can't be added to the second stage, thereby allowing more errors in the first stage, but practical design considerations and tradeoffs come into play here.

Generalized Pipeline Stages in a Subranging ADC with Error Correction

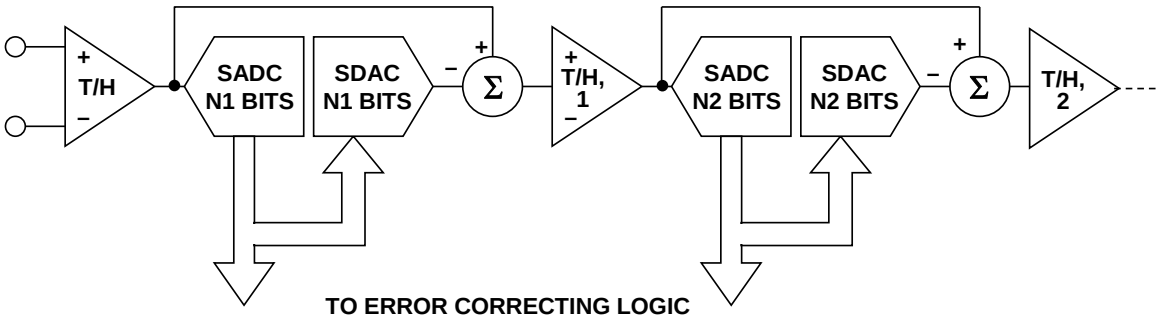

The pipelined architecture shown in this figure is a digitally corrected subranging architecture in which each stage operates on the data for one-half the sampling clock cycle and then passes its residue output to the next stage in the pipeline, prior to the next half cycle. The interstage track-and-hold (T/H) serves as an analog delay line—timing is set such that it enters the hold mode when the first stage conversion is complete. This gives 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 non-pipelined version.

The term "pipelined" architecture refers to the ability of one stage to process data from the previous stage during any given phase of the sampling clock cycle. At the end of each phase of a particular clock cycle, the output of a given stage is passed on to the next stage using the T/H functions, and new data is shifted into the stage. Of course this means that the digital outputs of all but the last stage in the "pipeline" must be stored in the appropriate number of shift registers so that the digital data arriving at the correction logic corresponds to the same sample.

Pipelined subranging ADCs usually have a number of identical stages in the pipeline. The first reference explains some of the more popular ones in much more detail, including the 1.5 bit/stage pipeline architecture.

---

1. Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN-0916550273 Chapter 1 and 3. Also available as *Data Conversion Handbook*, Elsevier-Newnes, 2005, ISBN: 0750678410, Chapter 1, 3.

2. Tutorial MT-024, *Pipelined Subranging ADCs*, www.analog.com.

Clock Issues in Pipelined ADCs

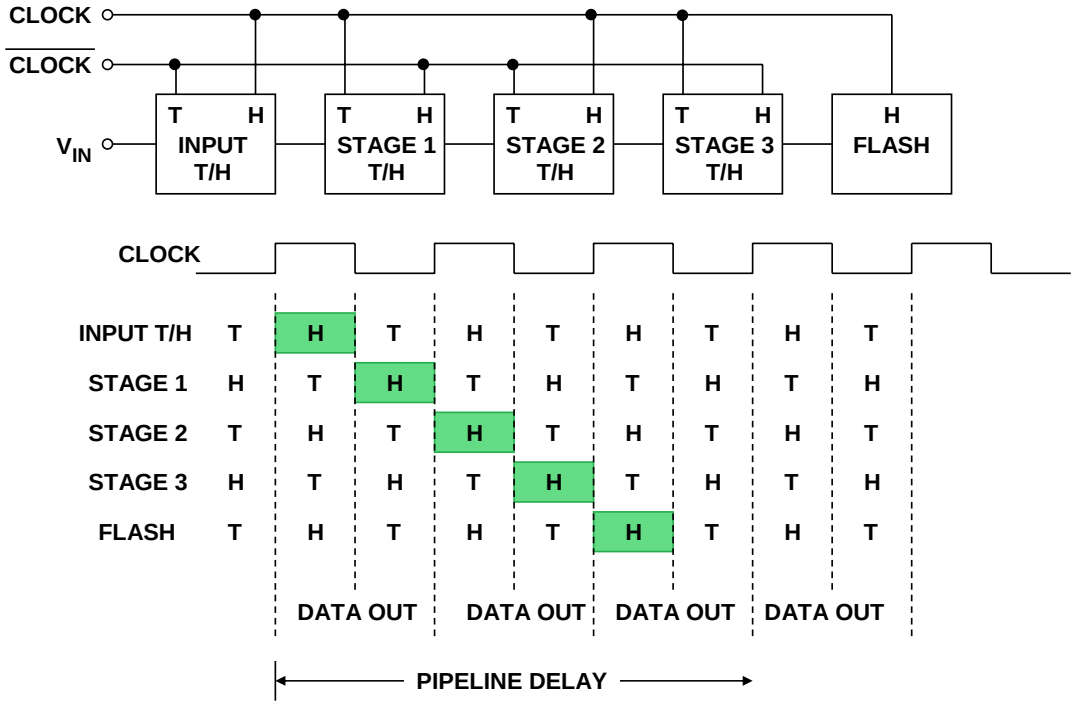

This figure shows a timing diagram of a typical pipelined subranging ADC. Notice that the phases of the clocks to the T/H amplifiers are alternated from stage to stage such that when a particular T/H in the ADC enters the hold mode it holds the sample from the preceding T/H, and the preceding T/H returns to the track mode. The held analog signal is passed along from stage to stage until it reaches the final stage in the pipelined ADC—in this case, a flash converter. When operating at high sampling rates, it is critical that the differential sampling clock be kept at a 50% duty cycle for optimum performance. Duty cycles other than 50% affect all the T/H amplifiers in the chain—some will have longer than optimum track times and shorter than optimum hold times; while others suffer exactly the reverse condition. Many newer pipelined ADCs have on-chip clock conditioning circuits to control the internal duty cycle and maintain rated performance even if there is some variation in the external clock duty cycle.

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.

Also, when the pipelined ADC is first powered up and the sampling clock applied, it takes a number of clock cycles to stabilize the clock circuitry and flush out the initial data in the pipeline.

The "latency" (pipeline delay) of pipelined ADCs make them somewhat difficult to use in traditional multiplexed data acquisition systems. SAR ADCs are much better here.

Typical Pipelined ADC Timing for AD9235 12-Bit, 65-MSPS ADC

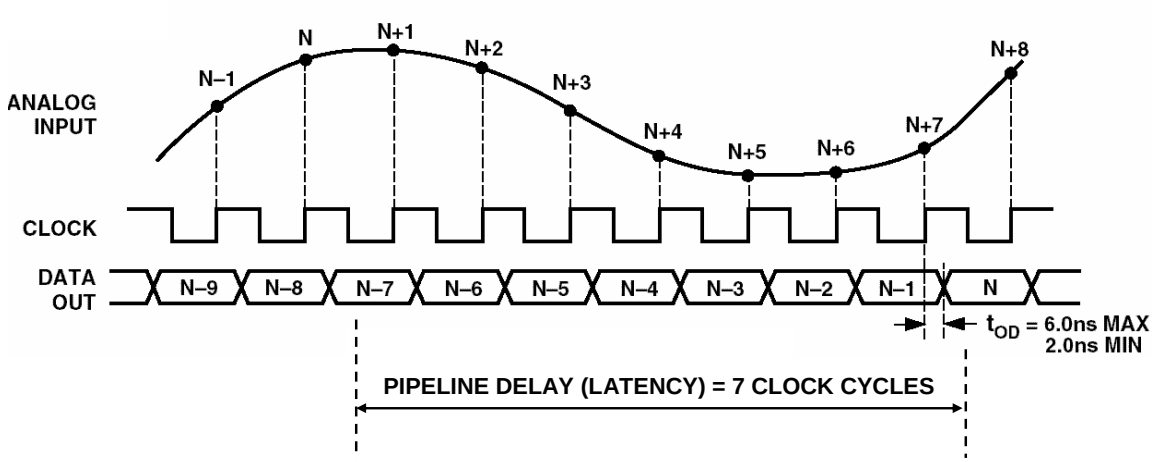

The effects of the "pipeline" delay (sometimes called "latency") in the output data are shown in in this figure for the AD9235 12-bit 65-MSPS ADC where there is a 7-clock cycle pipeline delay.

Note that the pipeline delay is a function of the number of stages and the particular architecture of the ADC under consideration—the data sheet should always be consulted for the exact details of the relationship between the sampling clock and the output data timing. In many applications the pipeline delay will not be a problem, but if the ADC is inside a feedback loop the pipeline delay may cause instability. The pipeline delay can also be troublesome in multiplexed applications or when operating the ADC in a "single-shot" mode. Other ADC architectures—such as successive approximation—are better suited to these types of applications.

It is often erroneously assumed that all subranging ADCs are pipelined, and that all pipelined ADCs are subranging. While it is true that most modern subranging ADCs are pipelined in order to achieve the maximum possible sampling rate, they don't necessarily have to be pipelined if designed for use at much lower speeds. For instance, the leading edge of the sampling clock could initiate the conversion process, and any additional clock pulses required to continue the conversion could be generated internal to the ADC using an on-chip timing circuit. At the end of the conversion process, an end-of-conversion or data-ready signal could be generated as an external indication that the data corresponding to that particular sampling edge is valid. This "no latency" approach is not often used for the obvious reason that the overall sampling rate is greatly reduced by eliminating the pipelined structure.

Summary: SAR vs. Pipelined ADCs

SAR ADCs

- Resolution to 18 bits
- Sample Rates to 3 MSPS
- Excellent DC Specifications
- Single-Shot Operation
- No Minimum Sample Rate
- No Latency (Pipeline Delay)
- Ideal for Muxed Applications
- Complete AC Specifications
- Easy to Use
- Key Applications:
  - Data Acquisition
  - Instrumentation
  - Industrial Process Control
  - Spectral Analysis
  - Medical Imaging
  - ATE

PIPELINED ADCs

- Resolution to 16 bits
- Sample Rates to 250 MSPS
- More Emphasis on AC Specifications
- Must Sample Continuously
- Minimum Sample Rate Specified
- Pipeline Delay
- Not Suitable for Muxed Systems
- Complete AC Specifications
- Easy to Use
- Key Applications
  - Wideband Multichannel Communications Receivers
  - Spectral Analysis
  - Medical Imaging
  - Display Electronics
  - Radar

This figure summarizes the differences between SAR and pipelined ADCs. As mentioned, there is some overlap in the sampling frequency range between 1MSPS and 10MSPS, but the application will generally dictate the appropriate architecture in this region.

In multiplexed data acquisition systems, the SAR ADC dominates because of its ease of use and lack of pipeline delay. For most other high speed ADC applications, the pipelined architecture dominates.

Measures of ADC Dynamic Performance

Important AC Performance Specifications for ADCs

- **Signal-to-Noise and Distortion Ratio (SINAD)**
- **Effective Number of Bits (ENOB)**
- **Signal-to-Noise Ratio (SNR)**
- **Single and Multitone Spurious Free Dynamic Range (SFDR)**
- **Total Harmonic Distortion (THD)**
- **Second Order Intermodulation Distortion (IMD2)**
- **Third Order Intermodulation Distortion (IMD3)**
- **Input Bandwidth**

- **Must Also Remember These**
  - **Minimum Sampling Frequency**
  - **Pipeline Delay (Latency)**

In order for an ADC to be useful in modern signal processing applications, it must meet system performance requirements, especially those associated with the frequency domain.

This list of AC specifications has evolved over the years, and today most customers and manufacturers agree on their basic definitions. This section will discuss them in more detail as well as how they influence overall system performance.

---

1. Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN-0916550273 Chapter 2 and 5. Also available as *Data Conversion Handbook*, Elsevier-Newnes, 2005, ISBN: 0750678410, Chapter 2 and 5.

2. Tutorial MT-003, *Understand SINAD, ENOB, SNR, THD, THD + N, and SFDR so You Don't Get Lost in the Noise Floor*, Analog Devices, www.analog.com.

Ideal N-bit ADC Quantization Noise

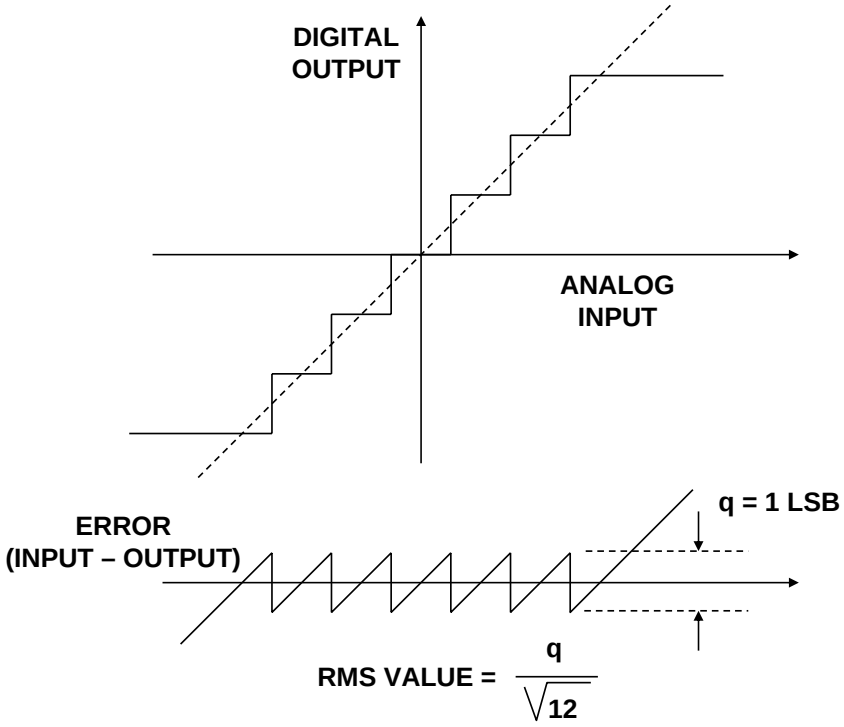

The maximum error an ideal converter makes when digitizing a signal is ±½ LSB as shown in the transfer function of an ideal N-bit ADC. The quantization error for any ac signal which spans more than a few LSBs can be approximated by an uncorrelated sawtooth waveform having a peak-to-peak amplitude of q, the weight of an LSB. Another way to view this approximation is that the actual quantization error is equally probable to occur at any point within the range ±½ q. Although this analysis is not precise, it is accurate enough for most applications. It can be shown (see References below) that the rms value of this sawtooth is $q/\sqrt{12}$.

The sawtooth error waveform produces harmonics which extend well past the Nyquist bandwidth of dc to $f_s/2$. However, all these higher order harmonics must fold (alias) back into the Nyquist bandwidth and sum together to produce an rms noise equal to $q/\sqrt{12}$.

The quantization noise is approximately Gaussian and spread more or less uniformly over the Nyquist bandwidth dc to $f_s/2$. The underlying assumption here is that the quantization noise is not correlated to the input signal. In other words, the error waveform is completely random with respect to the input signal. Under certain conditions, however, where the sampling clock and the signal are harmonically related, the quantization noise becomes correlated, and the energy is concentrated in the harmonics of the signal—however, the rms value remains approximately $q/\sqrt{12}$. The theoretical signal-to-noise ratio can now be calculated assuming a full-scale input sinewave. The result is:

$$SNR = 6.02N + 1.76dB$$

---

1. Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN-0916550273 Chapter 2. Also available as *Data Conversion Handbook*, Elsevier-Newnes, 2005, ISBN: 0750678410, Chapter 2.

2. Tutorial MT-001, *Taking the Mystery out of the Infamous Formula, "SNR = 6.02N + 1.76dB," and Why You Should Care*, Analog Devices, www.analog.com.

Quantization Noise Spectrum

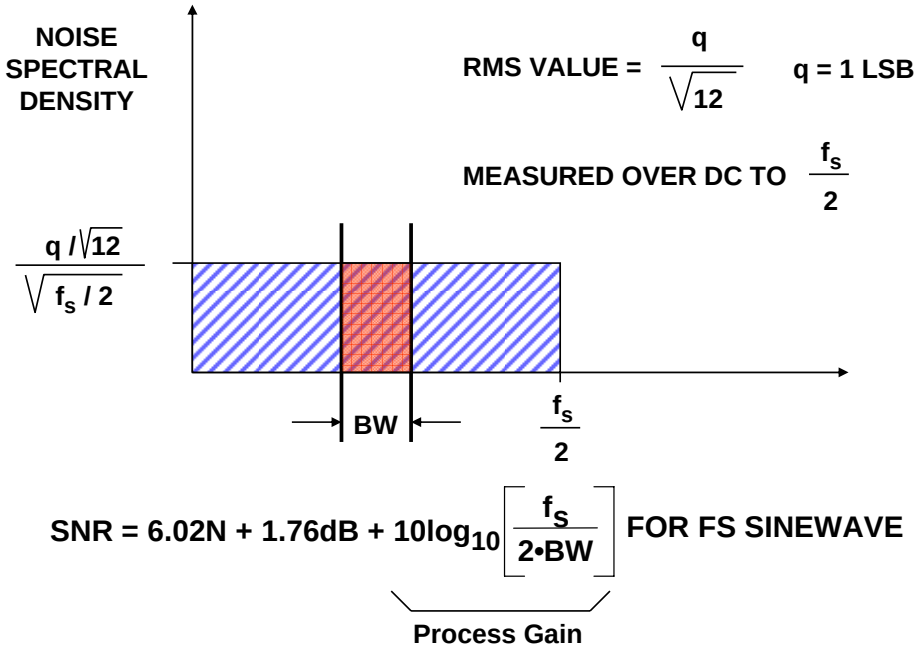

In many applications, the actual signal of interest occupies a smaller bandwidth, BW, which is less than the Nyquist bandwidth as shown in this figure. If digital filtering is used to filter out noise components outside the bandwidth BW, then a correction factor (called *process gain*) must be included in the equation to account for the resulting increase in SNR:

$$SNR = 6.02N + 1.76dB + 10\log(f_S/2BW),$$

where the term $10\log(f_S/2BW)$ is the process gain.

As an example, consider a multichannel GSM system which is sampled at a rate of 78MSPS. The individual bandwidth of each channel is 200kHz; therefore, the process gain is given by:

$$\text{Process Gain} = 10\log(f_S/2BW) = 10\log(78\times10^6/400\times10^3) = 22.9dB.$$

The process of sampling a signal at a rate which is greater than twice its bandwidth is referred to as *oversampling*. Oversampling in conjunction with quantization noise shaping and digital filtering are the key concepts in sigma-delta converters, although oversampling can be used with any ADC architecture.

SINAD, ENOB, SNR, and THD

- **SINAD (Signal-to-Noise-and-Distortion Ratio):**
  - **The ratio of the rms signal amplitude to the mean value of the root-sum-squares (RSS) of all other spectral components, including harmonics and noise, but excluding dc.**

- **ENOB (Effective Number of Bits):**

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

- **SNR (Signal-to-Noise Ratio, or Signal-to-Noise Ratio Without Harmonics:**
  - **The ratio of the rms signal amplitude to the mean value of the root-sum-squares (RSS) of all other spectral components, excluding the first 5 harmonics and dc**

- **THD (Total Harmonic Distortion):**
  - **The ratio of the rms signal amplitude to the mean value of the root-sum-square of its harmonics (generally only the first 5 harmonics are significant) and excluding dc**

SINAD, ENOB, and SNR, and THD are key dynamic ADC specifications and are defined in this figure.

SINAD is the ratio of the rms signal amplitude to the mean value of the rss values of all other spectral components INCLUDING harmonics and noise, but excluding dc.

If SINAD is substituted in the formula SNR = 6.02N + 1.76dB, and the equation solved for N, we get ENOB, or effective number of bits.

SNR is the ratio of the rms signal to the mean value of the RSS values of all other spectral components, EXCLUDING harmonics and dc. Usually, only the first five harmonics are significant.

SINAD and ENOB include all error sources and are good specifications when comparing various ADCs.

SNR is also important in many applications.

Carefully read the datasheets to be sure if the manufacturer is referring to SINAD or SNR, as some manufacturers don't differentiate clearly between them. SNR is generally greater than SINAD. Also, some manufacturers may use SNR when they really mean SINAD.

THD is the ratio of the rms signal amplitude to the mean value of the root-sum-square of its harmonics (generally only the first five harmonics are significant) and excluding dc.

AD9244 14-Bit, 65MSPS ADC
SINAD and ENOB for 1V and 2V Input Span

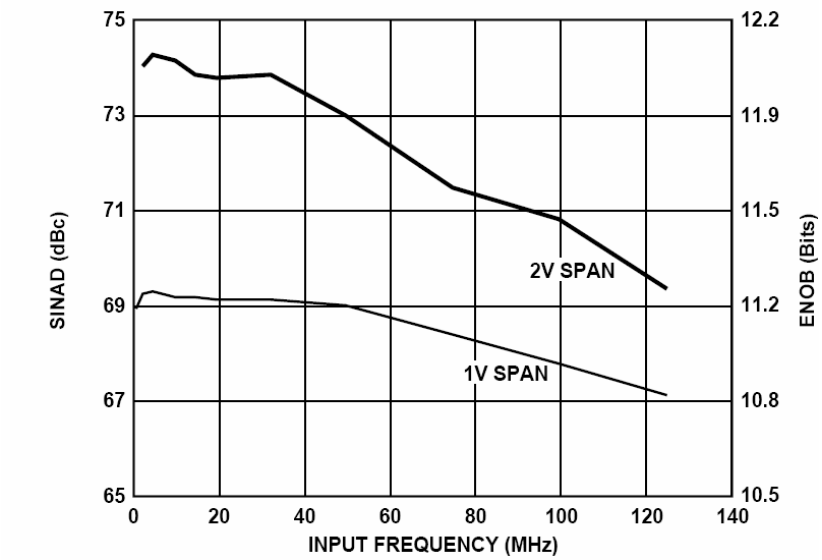

NOTE: AD9244 Full Power Input Bandwidth = 750MHz

SINAD, ENOB, and to a lesser degree SNR, all degrade as the analog input frequency to the ADC is increased. This is due to increased distortion and nonlinearities which occur at higher slew rates.

Some ADCs are designed to maintain good SINAD over the Nyquist bandwidth, others that are suitable for IF sampling maintain good SINAD in higher Nyquist zones.

This shows a plot for SINAD and ENOB for the AD9244 14-bit, 65MSPS ADC as a function of input frequency and input fullscale span. Values are given for differential input spans of 2V and 1V.

It is important to know SNR, SINAD, and ENOB at the analog input frequency of interest.

ADC bandwidth can be expressed in terms of full-power (FPBW) or small-signal bandwidth, as in the case of an amplifier. It is defined as the input frequency at which the FUNDAMENTAL signal in the FFT output drops 3dB. Therefore, the distortion may be very bad at the FPBW frequency.

ADC input bandwidth must always be examined in conjunction with ENOB, SNR, and SFDR to see if the bandwidth is really usable. Note that although the full power bandwidth of the AD9244 is 750MHz, the SINAD has begun to drop significantly at 120MHz.

Note that the above graph is for a fixed sampling rate with a varying input frequency. SINAD can also be a function of sampling rate for a fixed input frequency. In any case, the graphs included in the data sheet for an ADC should be studied with respect to system requirements—not just the numbers in the specification tables.

Relationship Between SINAD, SNR, and THD

- $\text{SNR} = 20 \log \left[\frac{S}{N} \right]$
- $\text{THD} = 20 \log \left[\frac{S}{D} \right]$
- $\text{SINAD} = 20 \log \left[\frac{S}{N + D} \right]$
- $\text{SINAD} = -10 \log \left[10^{-\text{SNR}/10} + 10^{-\text{THD}/10} \right]$

- **SNR, THD, and SINAD must be measured using the same signal amplitude.**
- **SINAD = THD + N if the measurement bandwidth is the same.**
- **Typically only the first 5 harmonics are included in the distortion term, D**

There is a known mathematical relationship between SINAD, SNR, and THD (assuming all are measured with the same input signal amplitude and frequency. In the above equations, SNR, THD, and SINAD are expressed in dB, and are derived from the actual numerical ratios S/N, S/D, and S/(N+D).

These equations are implemented in an easy to use SNR/THD/SINAD Calculator design tool on the Analog Devices' website, www.analog.com/designcenter.

It is important to emphasize again that these relationships hold true only if the input frequency and amplitude are equal for all three measurements.

Dynamic Performance Analysis of ADCs Using FFT Techniques

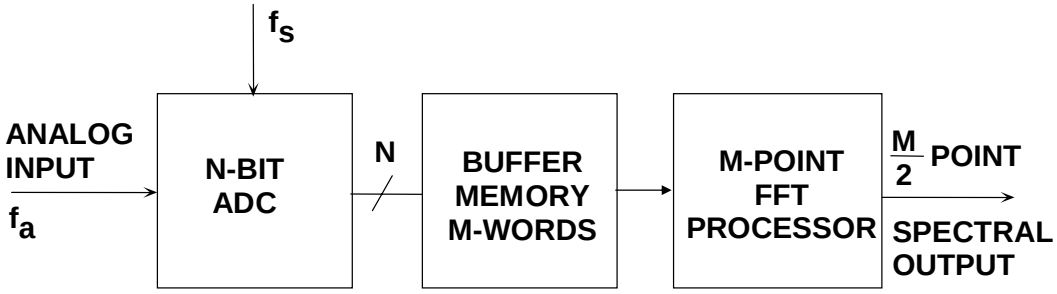

As previously discussed, there are a number of important frequency domain specifications for ADCs. All of them are based on an FFT analysis of the ADC output data using a generalized test setup such as shown in this figure.

Analog Devices has evaluation boards for all high speed ADCs, and the evaluation boards interface with a FIFO board which in turn interfaces to a PC. The details of the ADC evaluation hardware and software are described later in this section.

Ideal 12-Bit ADC, Input = 2.111MHz, f_S = 82MSPS, Average of 5 FFTs, M = 16,384

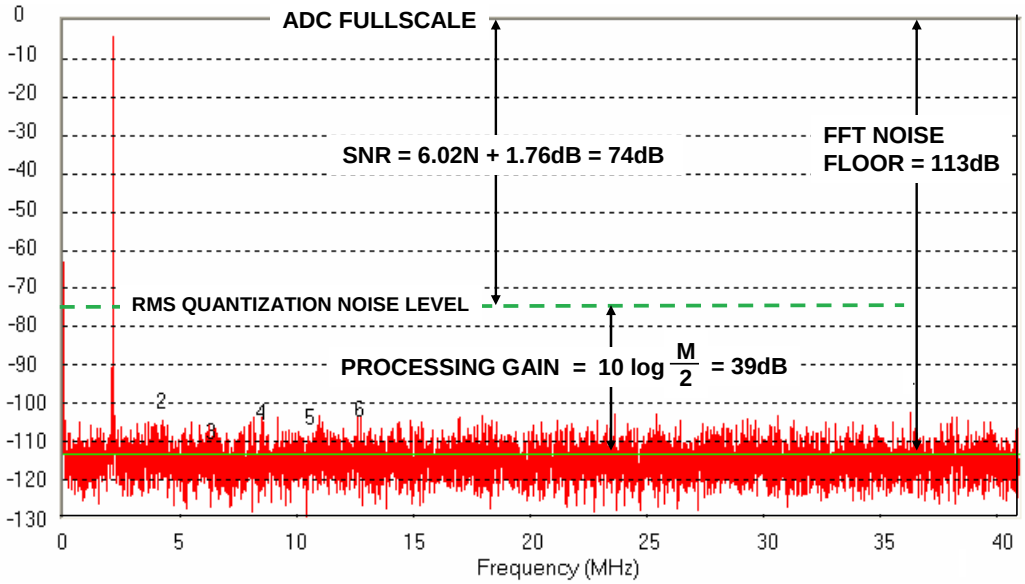

Data Generated Using ADIsimADC®

N = 12, M = 16,384 Bin Spacing = f_S / M

The spectral output of the FFT is a series of M/2 points in the frequency domain (M is the size of the FFT—the number of samples stored in the buffer memory). The spacing between the points is f_S/M, and the total frequency range covered is dc to f_S/2, where f_S is the sampling rate. The width of each frequency "bin" (sometimes called the *resolution* of the FFT) is f_S/M. This figure shows an FFT output for an ideal 12-bit ADC using the Analog Devices' ADIsimADC® program. Note that the theoretical noise floor of the FFT is equal to the theoretical SNR plus the FFT *process gain*, 10×log(M/2).

It is important to remember that the value for noise used in the SNR calculation is the noise that extends over the entire Nyquist bandwidth (dc to f_S/2), but the FFT acts as a narrowband spectrum analyzer with a bandwidth of f_S/M that sweeps over the spectrum. This has the effect of pushing the noise down by an amount equal to the process gain—the same effect as narrowing the bandwidth of an analog spectrum analyzer.

The FFT data shown in this figure represents the average of five individual FFTs. Averaging a number of FFTs does not affect the average noise floor, it only acts to "smooth" the random variations in the amplitudes contained in each frequency bin.

Location of Distortion Products: Input Signal = 7MHz, Sampling Rate = 20MSPS

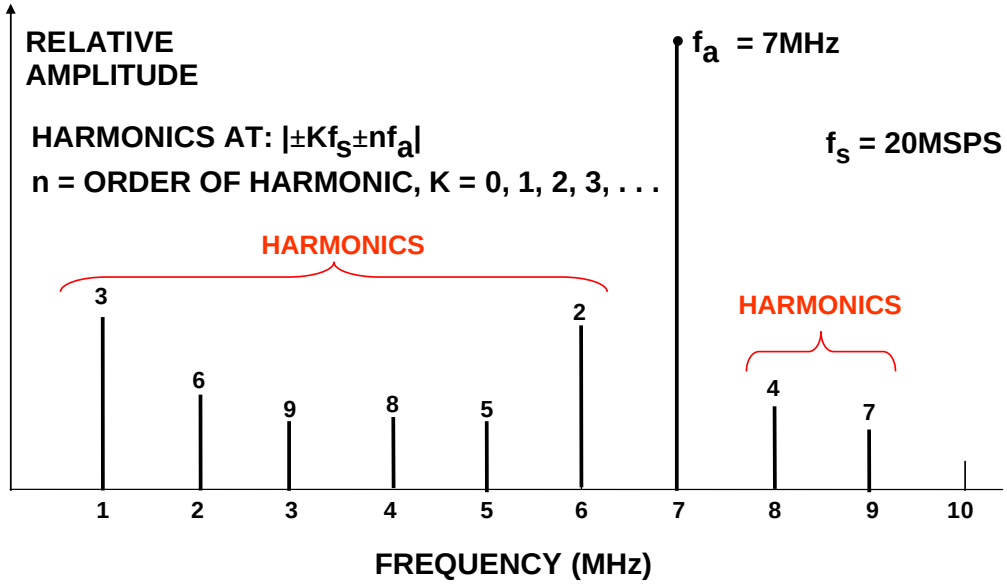

Harmonic distortion is normally specified in dBc (decibels below *carrier*), although in audio applications it may be specified as a percentage. It is the ratio of the rms signal to the rms value of the harmonic in question. Harmonic distortion is generally specified with an input signal near fullscale (generally 0.5 to 1dB below full-scale to prevent clipping), but it can be specified at any level. For signals much lower than full-scale, other distortion products due to the differential nonlinearity (DNL) of the converter—not direct harmonics—may limit performance.

Basic sampling theory says that harmonics of the fundamental signal which fall outside the Nyquist bandwidth do not go away, they all alias back into the dc to $f_s/2$ bandwidth of interest.

This figure shows a 7MHz input signal sampled at 20MSPS and the location of the first 9 harmonics. Aliased harmonics of f_a fall at frequencies equal to $|\pm Kf_S \pm nf_a|$, where n is the order of the harmonic, and K = 0, 1, 2, 3,.... The second and third harmonics are generally the only ones specified on a data sheet because they tend to be the largest, although some data sheets may specify the value of the *worst* harmonic.

In this figure, the fundamental signal is 7MHz, and the sampling frequency is 20MSPS. Note the location of the various harmonics which are aliased into the Nyquist bandwidth.

For instance, the second harmonic of 7MHz is 14MHz, but this is aliased back to 20 – 14 = 6MHz as shown.

There is an interactive Aliasing Suppression Assistant Tool on the ADI internet site which calculates the locations of the various harmonics for you. See www.analog.com/designcenter.

Spurious Free Dynamic Range (SFDR) in Communications Systems

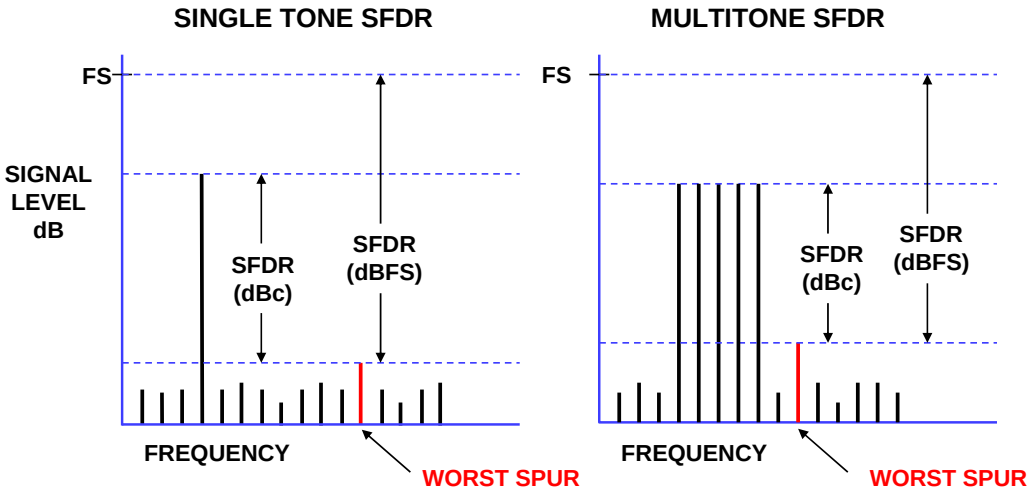

Spurious free dynamic range (SFDR) is the ratio of the rms value of the signal to the rms value of the worst spurious signal regardless of where it falls in the frequency spectrum or what the source. The worst spur may or may not be a harmonic of the original signal. SFDR is an important specification in communications systems because it represents the smallest value of signal that can be distinguished from a large interfering signal (blocker). SFDR can be specified with respect to full-scale (dBFS) or with respect to the actual signal amplitude (dBc).

SFDR can be specified for a single tone, two tone, or multitone signal. In testing using more than one tone, the amplitudes of the individual tones must be reduced so that the combined signal does not exceed the range of the ADC.

Second and Third-Order Intermodulation Products for f_1 = 5MHz, f_2 = 6MHz

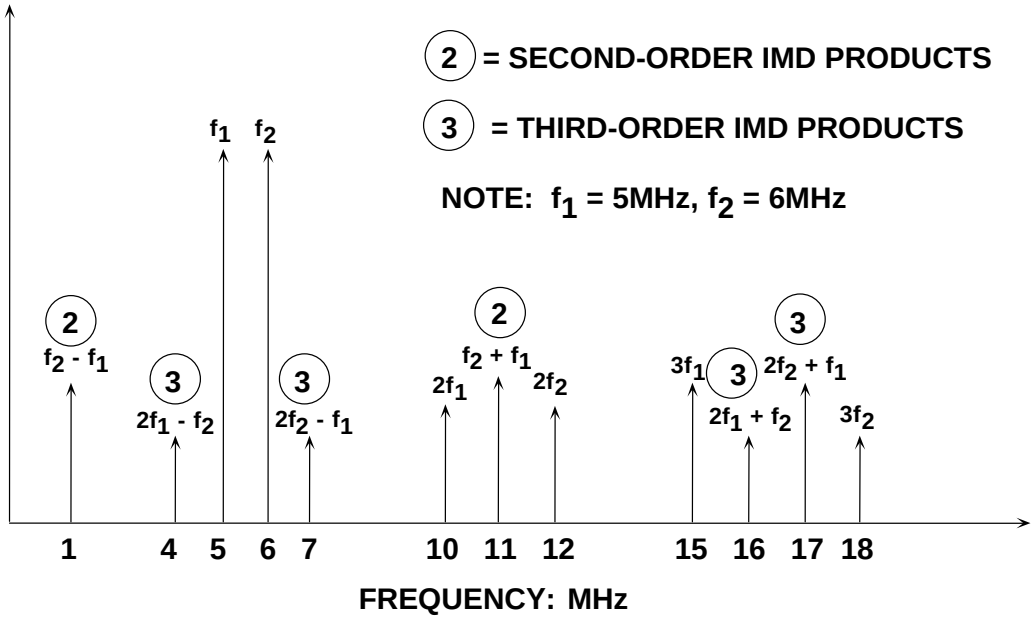

Two tone intermodulation distortion (IMD) is measured by applying two spectrally pure sinewaves to the ADC at frequencies f1 and f2, usually relatively close together. The amplitude of each tone is set slightly more than 6dB below fullscale so that the ADC does not clip when the two tones add in-phase.

The location of the second and third-order products are shown in the figure. Notice that the second-order products fall at frequencies which can be relatively easily removed by digital filters. However, the third-order products 2f2 – f1 and 2f1 – f2 are close to the original signals and are more difficult to filter. Unless otherwise specified, two-tone IMD refers to these "close-in" third-order products. It is customary to specify the value of the IMD product in dBc relative to the value of *either* of the two original tones, and not to their sum.

Note, however, that if the two tones are close to $f_S/4$, then the aliased third harmonics of the fundamentals can make the identification of the actual 2f2 – f1 and 2f1 – f2 products difficult. This is because the third harmonic of $f_S/4$ is $3f_S/4$, and the alias occurs at $f_S - 3f_S/4 = f_S/4$. Similarly, if the two tones are close to $f_S/3$, the aliased second harmonics may interfere with the measurement. The same reasoning applies here; the second harmonic of $f_S/3$ is $2f_S/3$, and its alias occurs at $f_S - 2f_S/3 = f_S/3$.

Two-Tone SFDR for AD9445 14-bit, 80/105MSPS ADC, Input Tones: 55.25MHz and 56.25MHz

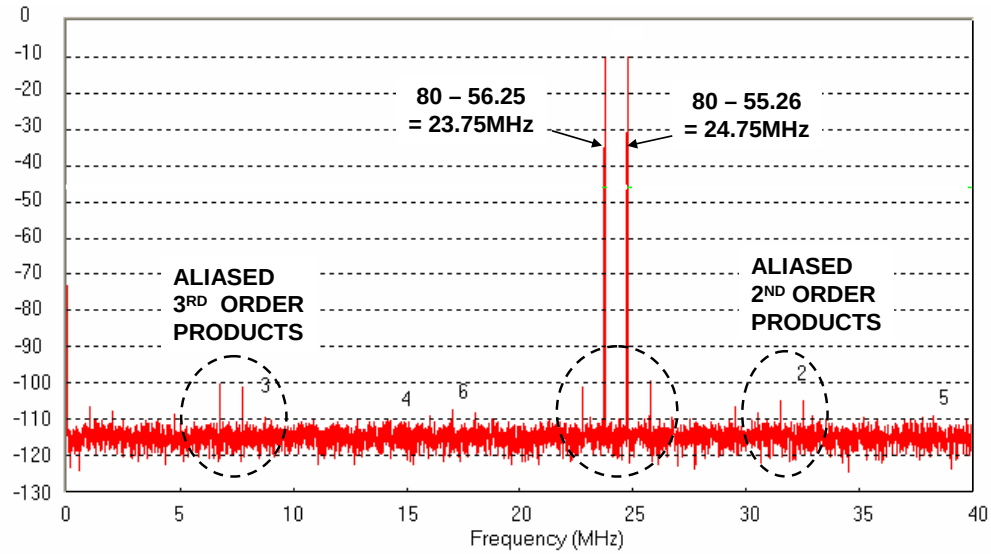

This figure shows the two-tone FFT output for the AD9445 sampling two tones of 55.25MHz and 56.25MHz at a sampling frequency of 80MSPS.

Note that the two tones must each be at least 6dB below FS in order to prevent saturating the ADC.

The two tones lie in the 2nd Nyquist zone, and are therefore aliased back to 80 – 56.25 = 23.75MHz and 80 – 55.25 = 24.75MHz.

The third-order products adjacent to the two tones are nearly 105dBFS. SFDR for two-tone inputs is the ratio of the amplitude of the tones to the worst spur, regardless of where the spur occurs.

Tradeoffs: SNR, SFDR, and Bandwidth Versus Power Level

	POWER	SNR @ 100 MHz	SFDR @ 100 MHz	BANDWIDTH
AD9446, 16 bits 100 MSPS	2.3 W	78.6 dB	82 dBc	540 MHz
AD9461, 16 bits 130 MSPS	2.2 W	76.0 dB	84 dBc	615 MHz
AD9445, 14 bits 125 MSPS	2.3 W	73.0 dB	95 dBc	615 MHz
AD9246, 14 bits 125 MSPS	0.395 W (1.8V)	71.5 dB	83 dBc	650 MHz

It is important to understand the ADC design tradeoffs that can be made between the two important parameters, SNR and SFDR.

The first three ADCs, the AD9446, AD9461, and the AD9445 are all designed on the same BiCMOS process and dissipate approximately the same amount of power (2.2W). Note that there is an inverse relationship between the SNR and the SFDR of the three converters.

In addition, the SNR is the highest for the device with the lowest bandwidth, as would be expected, since noise is proportional to the square root of the bandwidth.

The fourth ADC in the table, the AD9246, is designed on a 1.8V CMOS process for lower power dissipation (approximately one-fourth that of the other three). However, the SNR is 1.5dB worse and more importantly, the SFDR is 12dB worse than the AD9445. This illustrates the importance of knowing exact system requirements and not overspecifying the ADC.

The tradeoff between SNR and SFDR will be explained in more detail in the next figure.

SNR, SFDR, and Bandwidth Tradeoffs for Simplified Sample-and-Hold Model

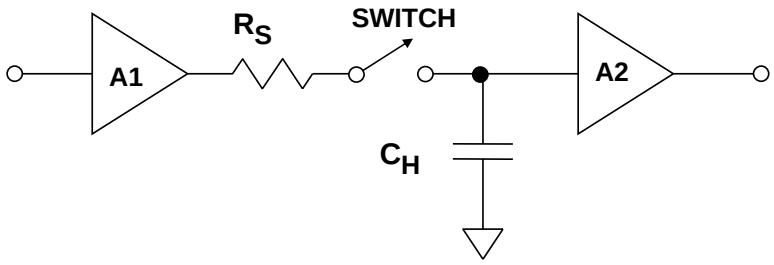

SMALLER C_H:

- **Higher Bandwidth**
- **Higher Noise (Lower SNR)**
- **Less Load on A1, Lower Distortion (Higher SFDR)**

LARGER C_H:

- **Lower Bandwidth**
- **Lower Noise (Higher SNR)**
- **More Load on A1, Higher Distortion (Lower SFDR)**

The fundamental tradeoff between SNR, SFDR, and bandwidth can be explained using this simple model for the sample-and-hold amplifier. Assume that A1 and A2 are wideband amplifiers and the only variable in the circuit is the hold capacitor, C_H. The resistor R_S is chosen to isolate A1 from C_H and maintain stability.

Small values of C_H result in higher bandwidth and higher noise (lower SNR). Small values of C_H also present a higher impedance load on the output of A1 (when the switch is closed in the track mode). The higher impedance load on the output of A1 results in lower distortion (higher SFDR).

Larger values of C_H result in lower bandwidth and lower noise (higher SNR). Larger values of C_H present a lower impedance load on the output of A1 which in turn results in higher distortion (lower SFDR).

Sampling Clock Jitter and Aperture Jitter Increase ADC Noise

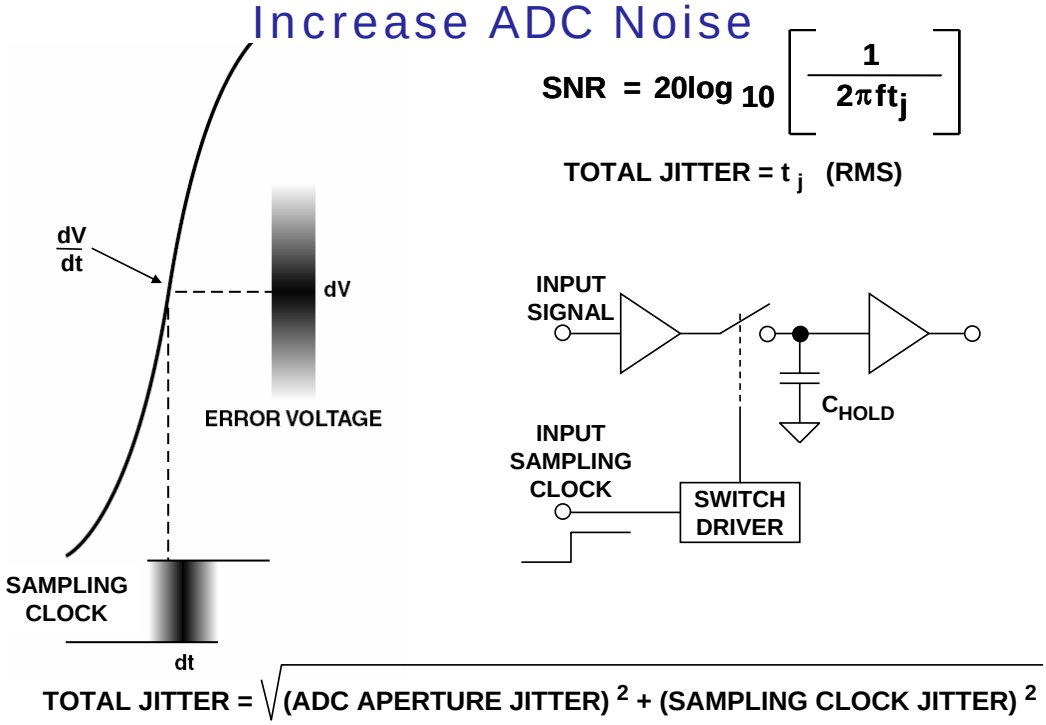

Aperture jitter is the sample-to-sample variation in the time at which the ADC sample is taken as shown in this model of a basic sample-and-hold. The total jitter in the sampling clock produces an error voltage, and the overall SNR of the ADC is limited by the formula as shown:

$$SNR = 20\log_{10}[1/2\pi ft_j],$$

where f is the fullscale analog input frequency and t_j is the total clock jitter.

Some modern IF sampling ADCs have aperture jitter specifications less than 100fs rms, but one must remember that the total jitter is the root-sum-square (rss) value of the ADC aperture jitter and the external sampling clock jitter. In most cases the sampling clock jitter will be the dominant source of error.

Note that the noise produced by jitter is directly proportional to the slew rate of the analog input signal. Jitter becomes especially critical in IF sampling applications where the input signal can be as high as several hundred MHz. The next figure graphically illustrates the effects of jitter on SNR.

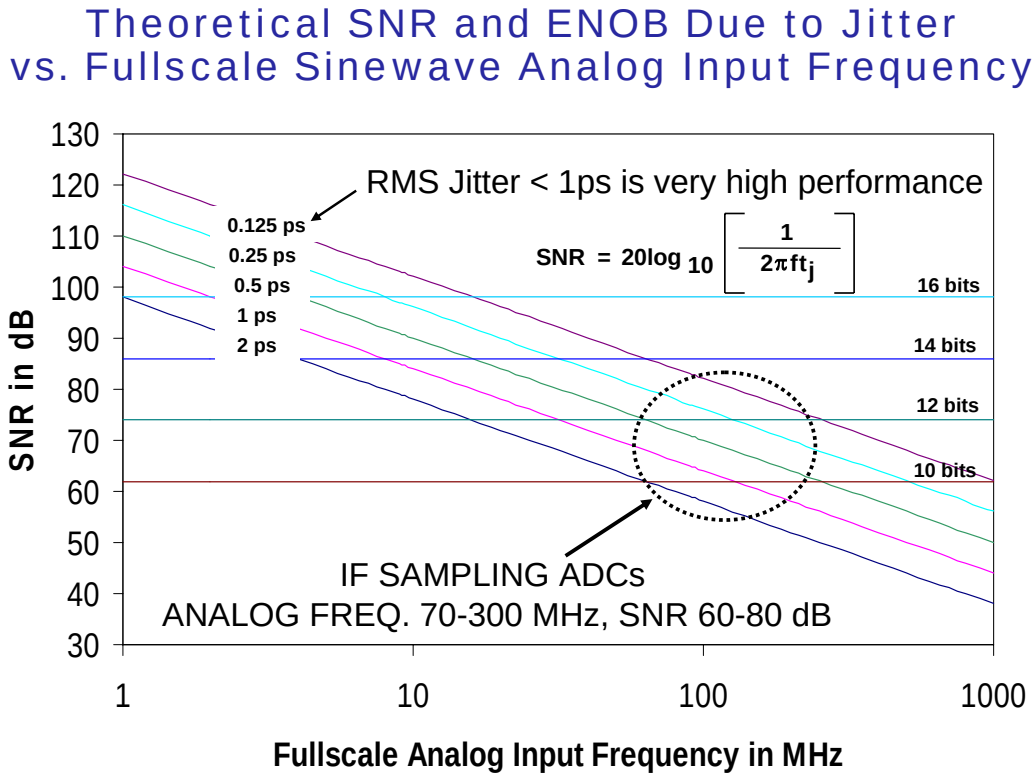

This figure shows the theoretical SNR (left side) due to total jitter versus fullscale analog input frequency. The ADC is assumed to have infinite resolution, and the only noise source is that produced by the jitter.

The effective number of bits (ENOB) is shown on the right and is related to SNR by the well known equation, SNR = 6.02N + 1.76dB, where N = ENOB.

IF sampling ADCs typically operate with analog frequencies between 70MHz and 300MHz, with system SNR requirements of 60dB to 80dB (the circled region). The range of allowable jitter is from about 0.1ps to 2ps, depending upon the IF frequency and the SNR requirement.

These stringent requirements mean that special care must be taken with the ADC sampling clock, which is often derived from other clocks in the system. The entire subject of high speed system clock generation and distribution is covered in more detail in Section 3 of this book.

High Speed ADC Applications in Software Radios

A software radio receiver uses an ADC to digitize the analog signal in the receiver as close to the antenna as practical, generally at an intermediate frequency (IF). Hence the term, *IF sampling* came into being. Once digitized, the signals are filtered, demodulated, and separated into individual channels often using specialized DSPs called *receive signal processors* (RSPs). Similarly, a software radio transmitter can perform coding, modulation, etc., in the digital domain—and near the final output IF stage, a DAC is used to convert the signal back to an analog format for transmission. The DSP which precedes the DAC is referred to as the transmit signal processor (TSP).

Ideally, the software radio eliminates quite a bit of expensive analog signal processing circuitry and performs these functions in low-cost DSPs. The software radio also allows the same hardware to handle various wireless air standards by making changes to the various DSP programs.

Wideband IF sampling places high demands on the ADCs and DACs in terms of SNR and SFDR. However, converter technology has progressed to the point that software radio is practical for most of the popular wireless air standards. For high volume applications, such as cellular telephone basestations and handsets, software radio has become a reality and a necessity.

U.S. Advanced Mobile Phone Service (AMPS) Superheterodyne Analog Receiver

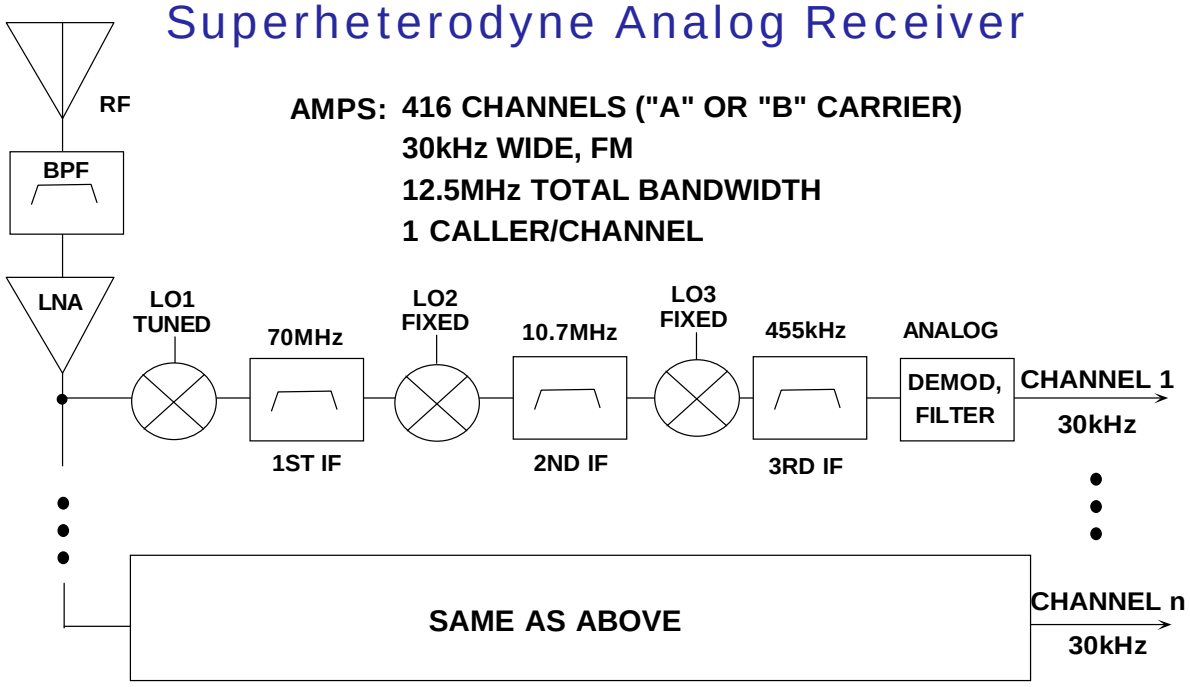

In order to understand the evolution of software radio, consider the analog superheterodyne receiver invented in 1917 by Major Edwin H. Armstrong. The frequencies shown in this figure correspond to the AMPS (Advanced Mobile Phone Service) analog cellular phone system currently being phased out in favor of new digital standards. The receiver is designed for AMPS signals at 900MHz RF. The signal bandwidth for the "A" or "B" carriers serving a particular geographical area is 12.5MHz (416 channels, each 30kHz wide). The receiver shown uses triple conversion, with a first IF frequency of 70MHz and a second IF of 10.7MHz, and a third IF of 455kHz. The image frequency at the receiver input is separated from the RF carrier frequency by an amount equal to twice the first IF frequency (illustrating the point that using relatively high first IF frequencies makes the design of the image rejection filter easier).

The output of the third IF stage is demodulated using analog techniques (discriminators, envelope detectors, synchronous detectors, etc.). In the case of AMPS, the modulation is FM. An important point to notice about the above scheme is that there is *one receiver required per channel*, and only the antenna, prefilter, and LNA can be shared.

It should be noted that in order to make the receiver diagrams more manageable, the interstage amplifiers and other circuits are not shown. They are, however, an important part of the receiver, and the reader should be aware that they must be present.

Receiver design is a complicated art, and there are many tradeoffs that can be made between IF frequencies, single-conversion vs. double-conversion or triple conversion, filter cost and complexity at each stage in the receiver, demodulation schemes, etc. There are many excellent references on the subject, and the purpose of this section is only to acquaint the design engineer with some of the emerging architectures, especially in the application of ADCs and DACs in the design of advanced communications receivers.

Generic IF Sampling Wideband Software Radio Receiver and Transmitter

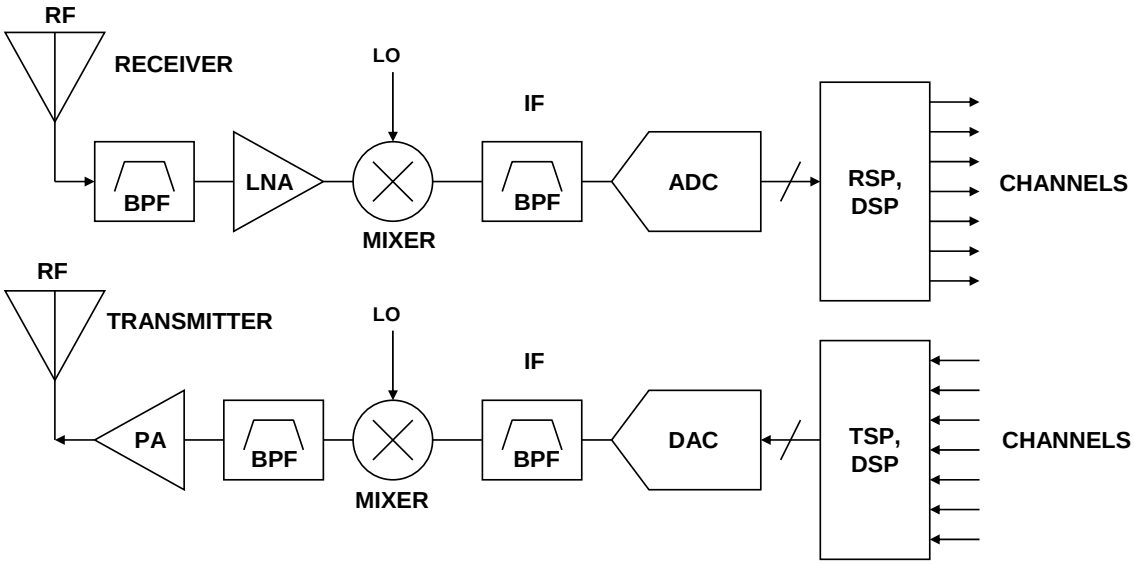

The standard superhetrodyne receiver has several stages of mixing down to baseband as previously described. There may be multiple parallel stages required for multichannel systems—one for each channel if all channels must be received simultaneously.

A generic wideband software radio receiver digitizes an entire bandwidth of multicarrier signals after only one or possibly two stages of downconversion. A basestation would be a typical application of software radio. Typical signal bandwidth can range from 5MHz to 30MHz. The receive signal processor (RSP) is a digital function which separates the individual channels, etc.

A similar approach is used in the transmitter. The transmit signal processor (TSP) combines the digital channel data and formats it suitable for driving a DAC.

Software radios are flexible and can handle multiple air standards by software changes. Digitizing the signal at high IF reduces analog component count (mixers, SAW filters, amplifiers, etc.) as well as the cost.

The example shown in this figure is for a basestation which handles multichannel wideband signals. The software radio technology can also be applied to signal channel applications, such as handsets, where multiple standards must be processed by the same hardware.

Undersampling and Frequency Translation Between Nyquist Zones

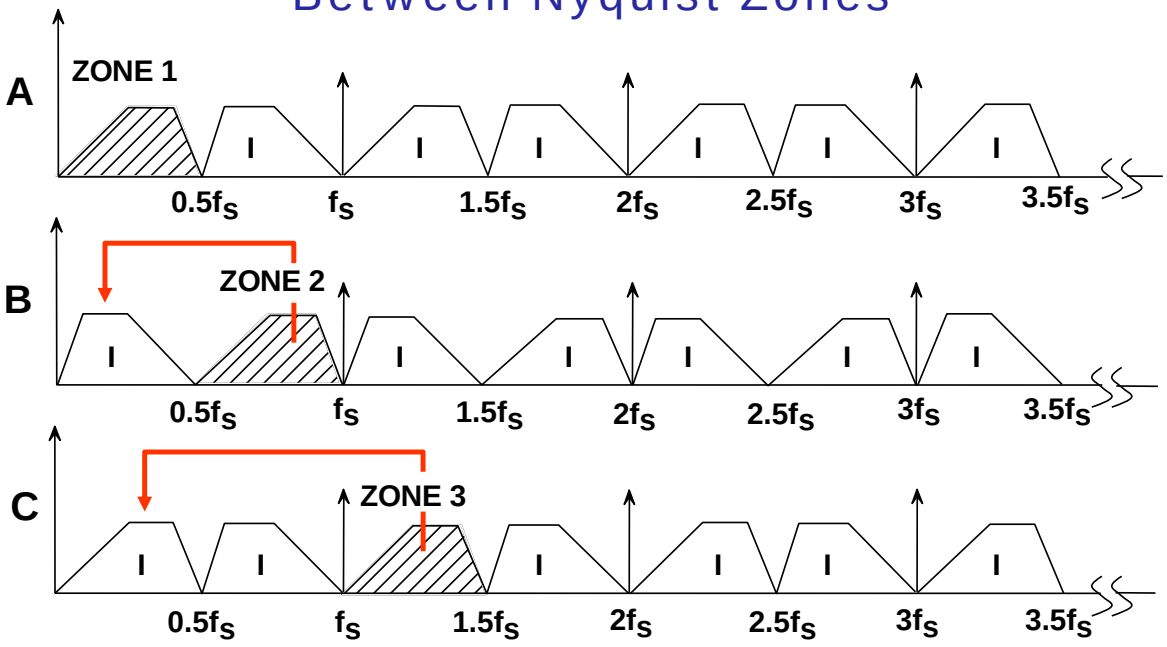

This figure illustrates the principles used in IF sampling, or direct IF-to-digital conversion. The signal of interest is bandlimited to a single Nyquist zone (not necessarily the first zone), and its image will always show up in the first Nyquist zone because of aliasing associated with the basic sampling process.

Note that the sample rate and the placement of the signal must be such that it is isolated to a single Nyquist zone (with appropriate filtering), and the sample rate must be at least twice the bandwidth of the signal.

Using higher IF frequencies can eliminate analog downconverter stages in the receiver as previously mentioned.

However, higher IF frequencies place more demands on the ADC in terms of bandwidth, SNR, SFDR.

Undersampling and Oversampling Combined Results in Process Gain

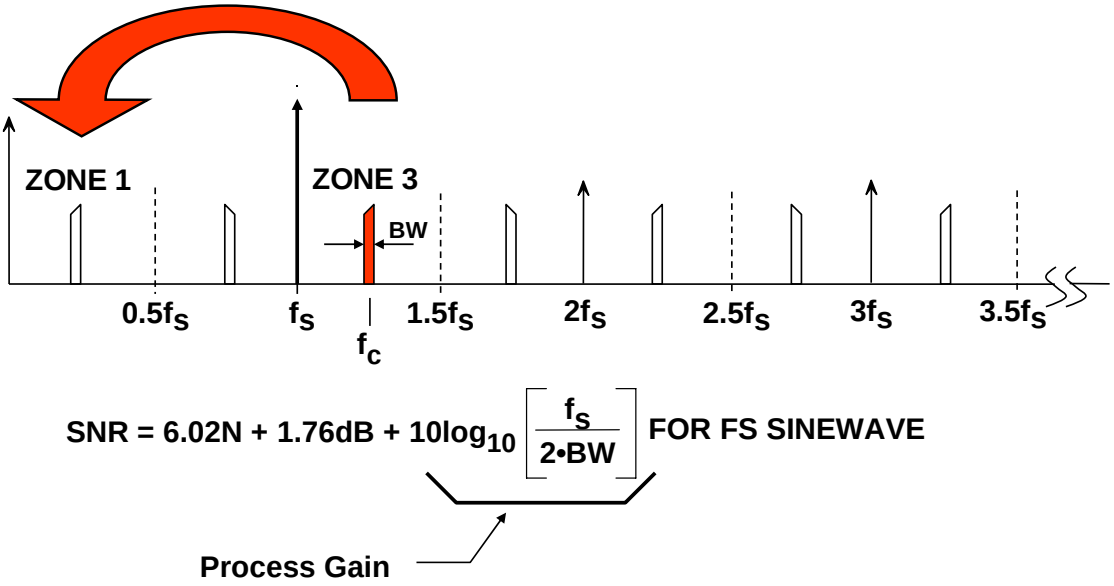

$$SNR = 6.02N + 1.76dB + 10\log_{10}\left[\frac{f_S}{2 \bullet BW}\right] \quad \text{FOR FS SINEWAVE}$$

Process Gain

This figure shows a signal of bandwidth BW which lies in the 3rd Nyquist zone sampled at frequency f_S.

The actual signal is "undersampled," since it lies outside the first Nyquist zone; but since the bandwidth of the signal is much less than $f_S/2$, it is also being "oversampled."

This often occurs in IF sampling applications, where the signal bandwidth is much less than $f_S/2$, but the signal of interest lies above the first Nyquist zone.

Note the increase in theoretical SNR due to the process gain as has been previously discussed.

AD9444 Sampling at 61.44MSPS with Four WCDMA Inputs Centered at 46.08MHz

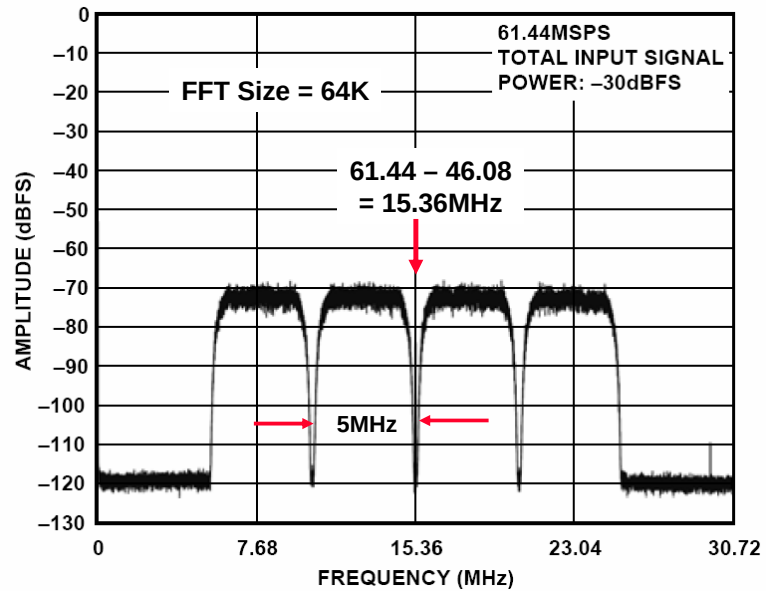

This figure shows how the direct IF to digital conversion process works for four wideband CDMA carriers which occupy a 20MHz bandwidth centered at 46.08MHz and sampled at 61.44MSPS.

The FFT output shows how the 46.08MHz center frequency which lies in the second Nyquist zone is aliased back to 15.36MHz.

The data was taken using the AD9444 14-bit, 80MSPS ADC.

Typical RF Spectrum of a Multicarrier CDMA2000 Receiver

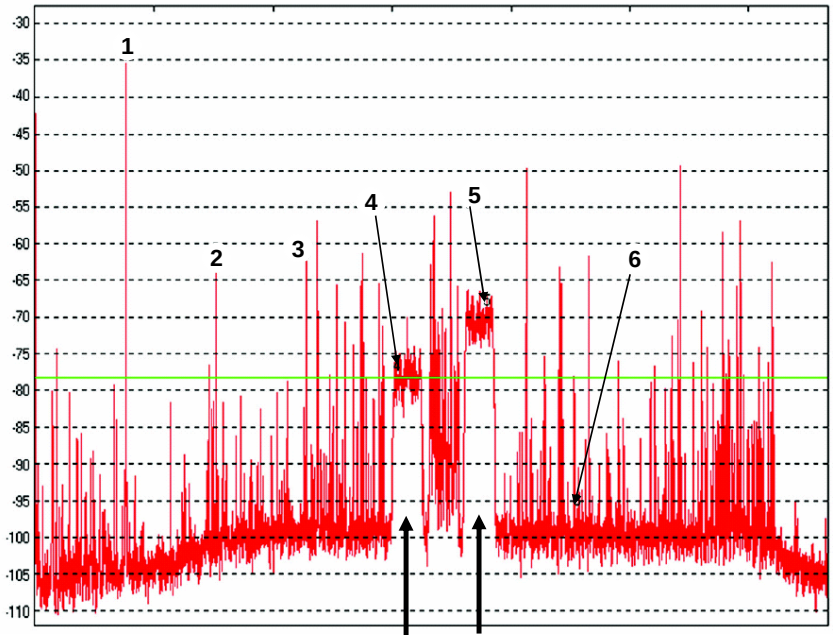

CDMA2000 CHANNELS, BW = 1.25MHz EACH, SAMPLING RATE = 61.44MSPS

The typical RF spectrum is very crowded, and contains large out-of-band signals whose harmonics can fall into the bandwidth of interest. These unwanted in-band signals are referred to as "interferers" or "blockers." The software radio (including the ADC) must be able to distinguish these interferers from the desired signals.

Shown here are two CDMA2000 channels, each having a bandwidth of 1.25MHz. Since the sampling frequency is 61MSPS, we are looking at a portion of the RF spectrum which is approximately 30.5MHz wide. This band has been down converted by at least one mixer stage.

Note that the 4th and 5th harmonics of the fundamental signal (labeled "1") act as "interferers" to the desired CDMA2000 signals.

I/Q Digital Modulation Schemes

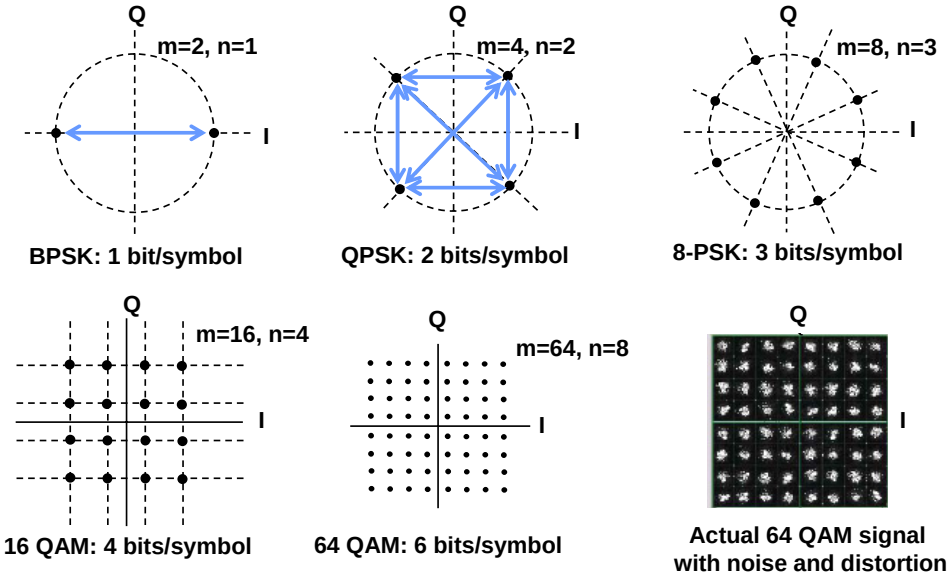

- **Higher Order Modulation Schemes → Higher Data Rate.**
- **But Symbols are closer together → Requires higher Signal-to-Noise Ratio for demodulation**
- **Increasing "Symbol Rate" increases data rate but widens spectrum**

Nearly all digital systems use some form of I/Q modulation in order to increase the amount of data that can be transmitted. This figure shows some of the more popular modulation types.

Regardless of the exact modulation scheme used, the receiver must determine which individual symbol in the I/Q constellation is being transmitted for each cycle of the symbol rate. Noise and distortion increase the probability of making an error in the receiver, especially for the higher order modulation schemes. The lower right-hand diagram shows an actual 64 QAM signal containing distortion and noise.

There is a relationship between the system bit-error-rate (BER) and the ratio of the carrier to the system noise (C/N) as well as the ratio of the carrier to the interferer (unwanted spur) level (C/I).

The higher order modulation schemes obviously place higher requirements on the ADC in terms of SNR, etc.

Bit Error Rate Versus Carrier-to-Noise Ratio (C/N)

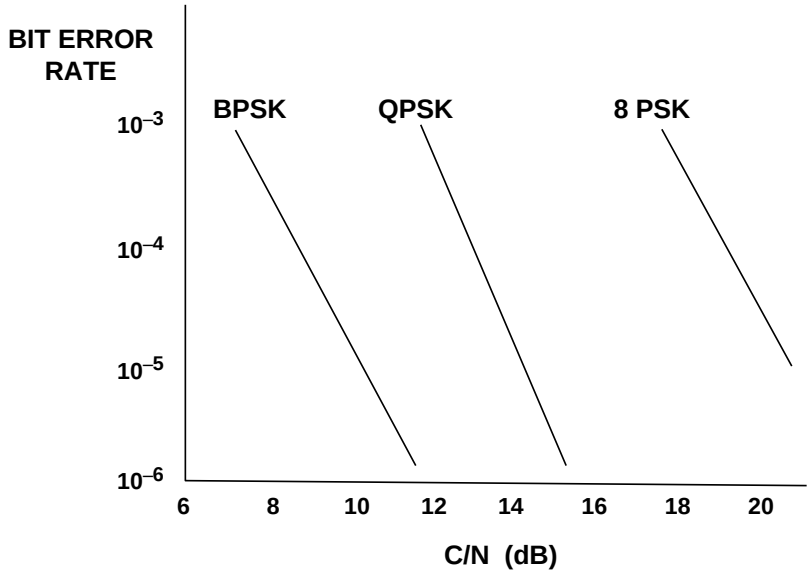

The figure shows the bit error rate as a function of the carrier-to-noise ratio (C/N) for three popular modulation schemes: BPSK, QPSK, and 8 PSK. Note that the more complex schemes (QPSK and 8 PSK) require a higher C/N ratio to achieve the same bit error rate as the simple BPSK.

The system bit error rate and the modulation method therefore determine the required C/N ratio. The C/N ratio can then be used to determine the SNR required of the ADC. In practice, the analysis is somewhat more complicated, but the concept is still valid.

GSM 1800/1900 MHz Wideband Receiver

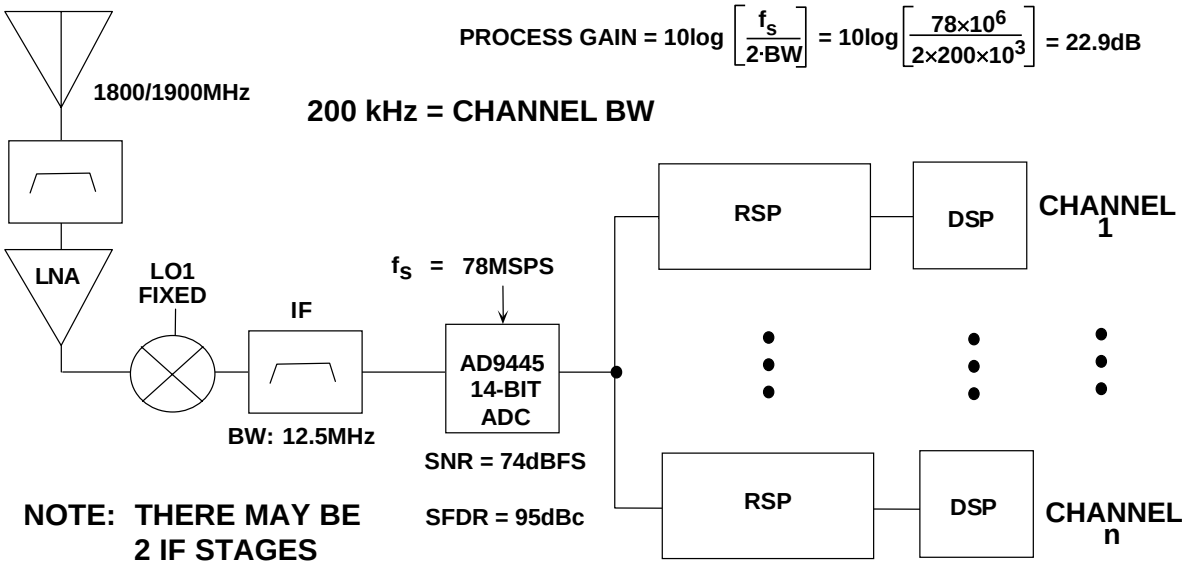

The noise analysis for a system depends upon the air standard and modulation method. For the GSM 1800/1900 MHz wideband receiver shown here, the AD9445 14-bit, 125MSPS ADC has 74dB SNR and 95dBc SFDR which is sufficient to meet the system bit error rate requirement (see next figure).

Details of the receiver noise analysis can be found in the reference.

Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN-0916550273 Chapter 8. Also available as *Data Conversion Handbook*, Elsevier-Newnes, 2005, ISBN: 0750678410, Chapter 8.

GSM 1800/1900 MHz Spurious Requirements

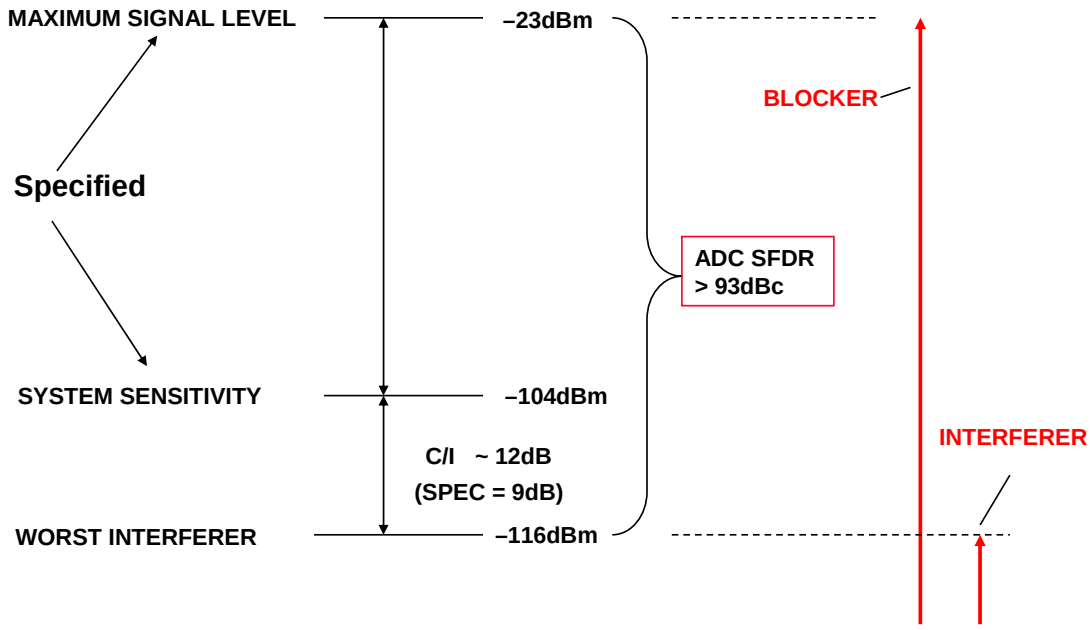

* NOTE: MAXIMUM SIGNAL LEVEL FOR GSM 1800MHz/1900MHz AND PCS SPEC = –23dBm

This figure shows how the system specifications can be used to determine the SFDR requirement for the ADC.

The amplitude of the "blocker" is equal to the maximum signal level of –23dBm. The maximum interferer (which would be an ADC spur) can be no greater than –116dBm in order to achieve a C/I ratio of 12dB.

This is equivalent to an SFDR of 93dBc for a fullscale single tone input. This specification can be met with the AD9445 14-bit ADC which has an SFDR of 95dBc.

Two-Tone Intermodulation Distortion in Multichannel System (GSM-1800/1900 MHz)

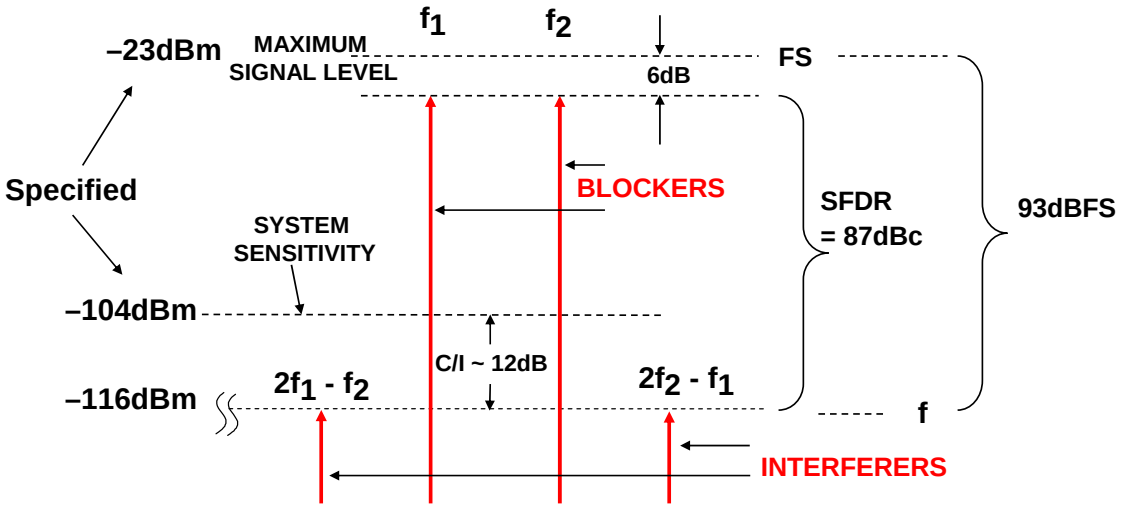

NOTE: MAXIMUM SIGNAL LEVEL FOR GSM 1800/1900 SPEC = –23dBm

A similar analysis can be used to determine the two-tone SFDR requirement for the ADC. In this case, the third order IMD products must be no greater than –116dBm for tone amplitudes of 6dB below fullscale. This requires a two-tone SFDR of at least 87dBc which is also met by the AD9445 ADC.

Approximate Wideband ADC Requirements for Popular Wireless Air Interface Standards

	MULTIPLE ACCESS METHOD	CHANNEL SPACING (BW)	TYPICAL TOTAL BW	ADC SAMPLING RATE (TYP.)	ADC SFDR	ADC SNR
AMPS	FDMA	30kHz	12.5MHz	61.44MSPS	96dBc	72dBFS
IS-136	TDMA/FDM	30kHz	5-15MHz	61-92MSPS	88dBc	68dBFS
GSM 900 MHz	TDMA/FDM	200kHz	5-15MHz	61-92MSPS	106dBc	85dBFS
GSM 1800/1900MHz, PCS	TDMA/FDM	200kHz	5-15MHz	61-92MSPS	93dBc	75dBFS
IS-95	CDMA	1.25MHz	5-15MHz	61-92MSPS	83dBc	74dBFS
CDMA2000	CDMA	1.25MHz	5-15MHz	61-92MSPS	79dBc	74dBFS
WCDMA (UMTS)	CDMA	5MHz	5-20MHz	61-92MSPS	79dBc	69dBFS

This figure summarizes the basic characteristics and wideband ADC requirements for some of the various air standards. Note that ADCs are available which will meet these SFDR and SNR requirements, (with the possible exception of GSM 900).

CCD/CIS Imaging for Digital Cameras and Camcorders

www.analog.com/afe

Linear and Area CCD Arrays

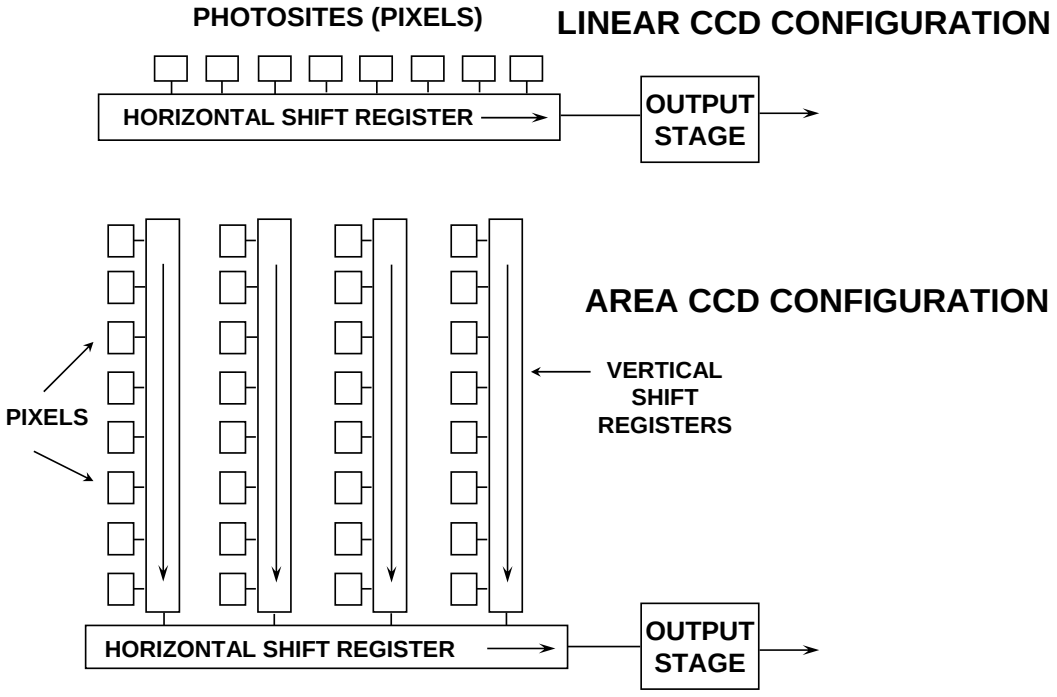

The *charge-coupled-device* (CCD) and *contact-image-sensor* (CIS) are widely used in consumer imaging systems such as scanners and digital cameras. The imaging sensor (CCD, CMOS, or CIS) is exposed to the image or picture much like film is exposed in a camera. After exposure, the output of the sensor undergoes some analog signal processing and then is digitized by an ADC. The bulk of the actual image processing is performed using fast digital signal processors. At this point, the image can be manipulated in the digital domain to perform such functions as contrast or color enhancement/correction, etc.

The building blocks of a CCD are the individual light sensing elements called pixels. A single pixel consists of a photo sensitive element, such as a photodiode or photocapacitor, which outputs a charge (electrons) proportional to the light (photons) that it is exposed to. The charge is accumulated during the exposure or integration time, and then the charge is transferred to the CCD shift register to be sent to the output of the device. The amount of accumulated charge will depend on the light level, the integration time, and the quantum efficiency of the photo sensitive element. A small amount of charge will accumulate even without light present; this is called dark signal or dark current and must be compensated for during the signal processing.

The pixels can be arranged in a linear or area configuration as shown in this figure. Clock signals transfer the charge from the pixels into the analog shift registers, and then more clocks are applied to shift the individual pixel charges to the output stage of the CCD.

Scanners generally use the linear configuration, while digital cameras use the area configuration.

The analog shift register typically operates at pixel frequencies between 1 and 10MHz for linear sensors, and 5 to 25MHz for area sensors.

Note that the size of the array for digital cameras can be quite large for megapixel resolutions.

Output Stage and Waveforms

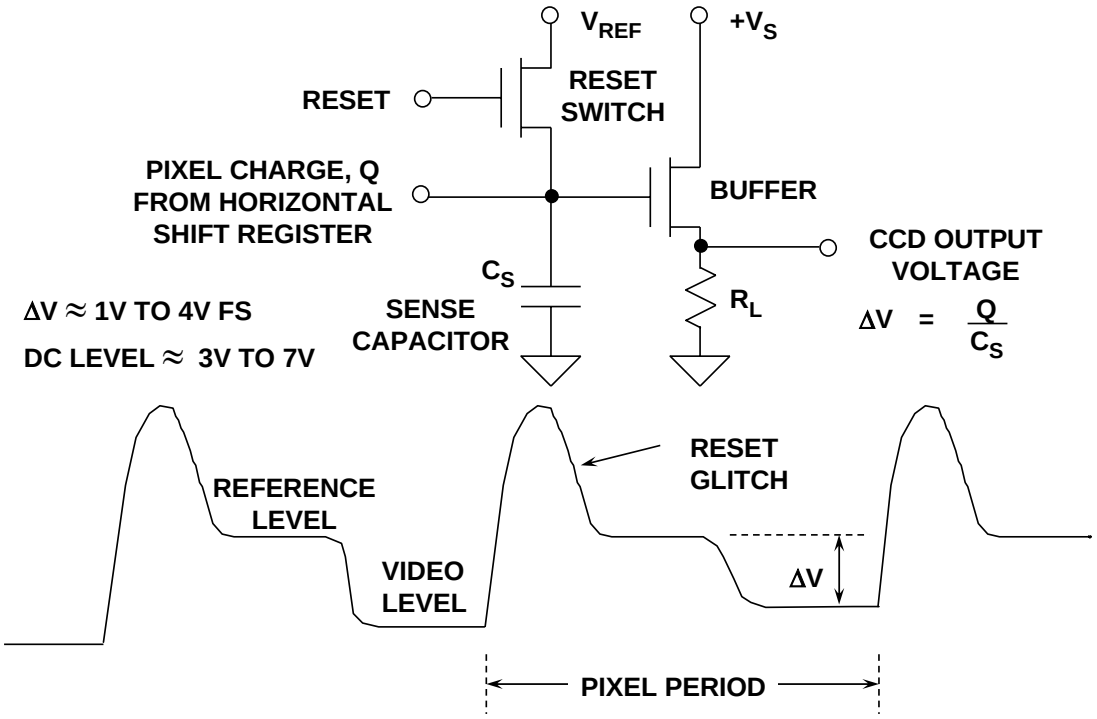

A typical CCD output stage is shown in this figure along with the associated voltage waveforms. The output stage of the CCD converts the charge of each pixel to a voltage via the sense capacitor, C_S. At the start of each pixel period, the voltage on C_S is reset to the reference level, V_{REF}, causing a reset glitch to occur. The amount of light sensed by each pixel is measured by the difference between the reference and the video level, ΔV. CCD charges may be as low as 10 electrons, and a typical CCD output has a sensitivity of 0.6μV/electron. Most CCDs have a saturation output voltage of about 500mV to 1V for area sensors and 2V to 4V for linear sensors. The dc level of the waveform is between 3V and 7V.

CCD processes generally have limited capability to perform on-chip signal conditioning. Therefore the CCD output is generally processed by external conditioning circuits. The nature of the CCD output requires that it be clamped before being digitized by the ADC. In addition, offset and gain functions are generally part of the analog signal processing.

CCD output voltages are small and quite often buried in noise. The largest source of noise is the thermal noise in the resistance of the FET reset switch (often called "KT/C" noise). During the reset interval, the storage capacitor C_S is connected to V_{REF} via a CMOS switch. Note that when the reset switch opens, the noise is stored on C_S and remains constant until the next reset interval. It therefore occurs as a *sample-to-sample* variation in the CCD output level and is common to both the reset level and the video level for a given pixel period.

Correlated Double Sampling (CDS)

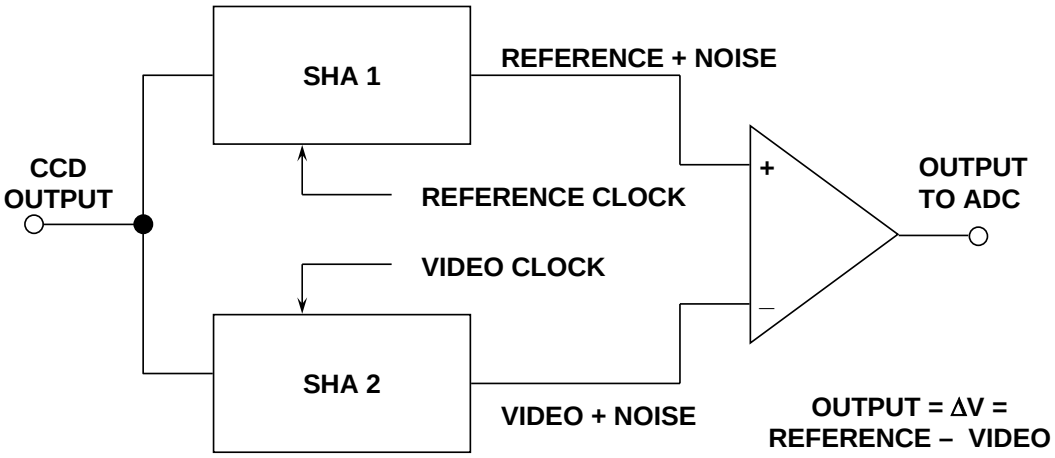

- **Typical ADC Resolution: 10 to 14 bits**
- **Typical ADC Sampling Rates: 12 to 65MSPS**

A technique called *correlated double sampling* (CDS) is often used to reduce the effect of this noise. This figure shows one circuit implementation of the CDS scheme, though many other implementations exist. The CCD output drives both SHAs. At the end of the reset interval, SHA1 holds the reset voltage level plus the kT/C noise. At the end of the video interval, SHA2 holds the video level plus the kT/C noise. The SHA outputs are applied to a difference amplifier which subtracts one from the other. In this scheme, there is only a short interval during which both SHA outputs are stable, and their difference represents ΔV, so the difference amplifier must settle quickly. Note that the final output is simply the difference between the reference level and the video level, ΔV, and that the kT/C noise is removed.

Contact Image Sensors (CIS) are linear sensors often used in facsimile machines and low-end document scanners instead of CCDs. Although a CIS does not offer the same potential image quality as a CCD, it does offer lower cost and a more simplified optical path. The output of a CIS is similar to the CCD output except that it is referenced to or near ground, eliminating the need for a clamping function. Furthermore, the CIS output does not contain correlated reset noise within each pixel period, eliminating the need for a CDS function. Typical CIS output voltages range from a few hundred mV to about 1V fullscale. Note that although a clamp and CDS is not required, the CIS waveform must be sampled by a sample-and-hold before digitization.

AD9898 CCD Signal Processor with Precision Timing™ Generator

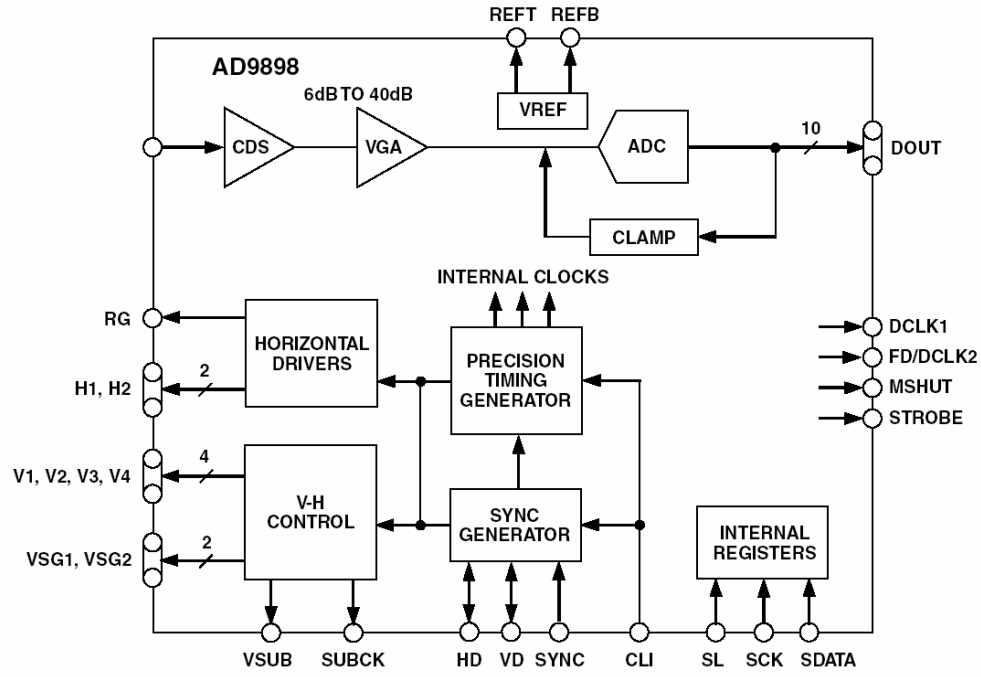

The AD9898 is a highly integrated CCD signal processor for digital still camera and digital video camera applications. A simplified block diagram is shown in this figure. It includes a complete analog front end with 10-bit A/D conversion combined with a full function programmable timing generator. A precision timing core allows adjustment of high speed clocks with 1ns resolution at 20MSPS operation. The AD9898 is specified at pixel rates as high as 20MHz. The analog front end includes black level clamping, CDS, VGA, and a 10-bit A/D converter. The timing generator provides all the necessary CCD clocks: RG, H-clocks, V-clocks, sensor gate pulses, substrate clock, and substrate bias pulse. Operation is programmed using a 3-wire serial interface.

Packaged in a space saving 48-lead LFCSP, the AD9898 is specified over an operating temperature range of –20°C to +85°C.

Other CCD image processing products can be found at www.analog.com/afe.

ADC Applications in Video

www.analog.com/video

Digital video had its beginnings in the early 1970s when 8-bit ADCs with sampling frequencies of 15MSPS to 20MSPS became available. Subjective tests showed that 8-bit resolution was sufficient for digitizing the composite video signal at sampling frequencies of 3 or 4 times the NTSC color subcarrier frequency (3.58 MHz). Digital techniques were first applied to "video black boxes" which replaced functions previously implemented using analog techniques. These early digital black boxes had an analog video input and an analog video output, and replaced analog-based equipment such as time-base correctors, frame stores, standards converters, etc. The availability of low-cost IC ADCs in the 1980s played a large role in the growth of digital video.

Digital videotape recorders (VTRs) emerged in the 1980s, based on CCIR recommendations. More digital black boxes proliferated, such as digital effects generators, graphic systems, and still stores—these devices operating in a variety of noncorrelated and incompatible standards. Digital connections between the black boxes were difficult or impossible, and the majority were connected with other equipment using analog input and output ports.

In the 1980s, the Society of Motion Picture and Television Engineers (SMPTE) developed a digital standard (SMPTE 244M) which defined the characteristics of $4f_{SC}$ sampled NTSC composite digital signals as well as the characteristics of a bit-parallel digital interface which allowed up to 10-bit samples. The digital interface consisted of 10 differential ECL-compatible data signals, 1 differential ECL-compatible clock signal, 2 system grounds, and 1 chassis ground, for a total of 25 pins. Later, digital systems using $4f_{SC}$ NTSC composite digital signals adopted a high-speed bit-serial interface, with a data rate of 143 Mb/s (defined in SMPTE 259M).

Model for Generating the Composite Video Signal from RGB Components

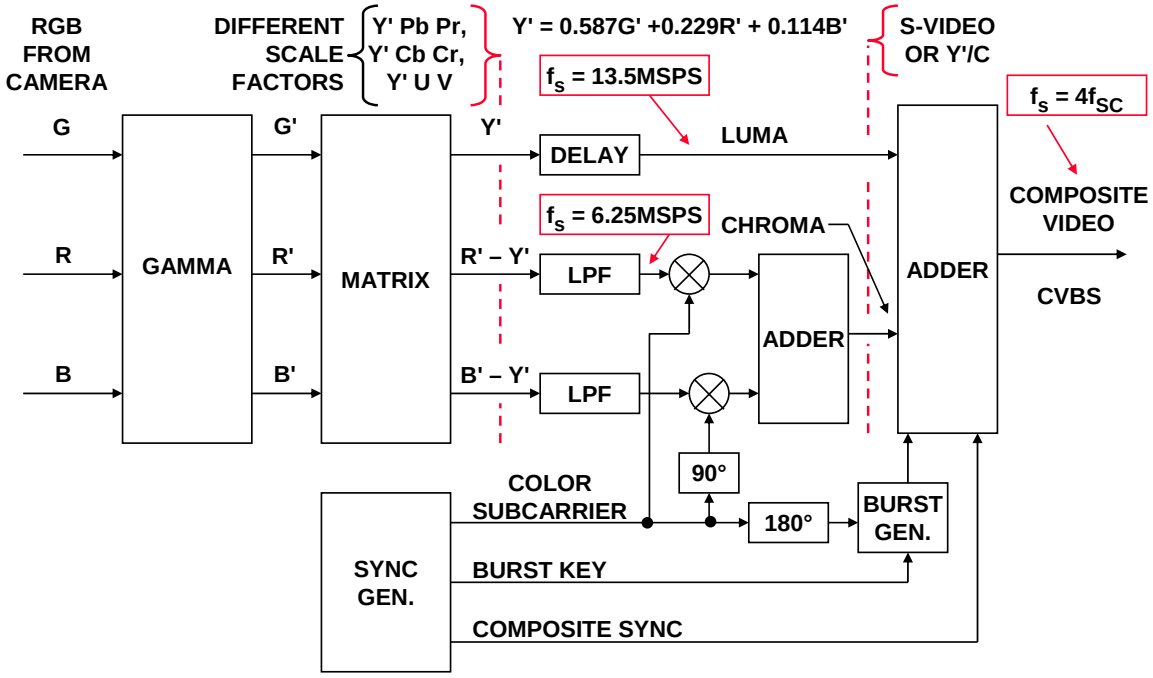

Even before the finalization of the $4f_{SC}$ composite digital standard, work was progressing on digital *component* systems, which offer numerous advantages over the composite digital systems. To understand the differences and advantages, this figure shows a generalized block diagram of how the composite broadcast video signal is constructed.

The native RGB signals from the color camera are first passed through a nonlinear *gamma* unit which compensates for the inherent nonlinearity in the receiving CRT. The R'G'B' outputs of the gamma unit then pass through a resistive *matrix* which generates a high-bandwidth *luma* signal (often incorrectly called *luminance*) and two reduced-bandwidth color difference signals. The luma signal, Y', is formed using the relationship Y' = 0.587G' + 0.229R' + 0.114B'. In addition, two *color difference* signals, designated R' – Y' and B' – Y' are formed. The color subcarrier is then used to modulate the color difference signals in quadrature, and they are summed to form the *chroma* signal (often incorrectly called *chrominance*). The color burst and composite sync signals are then combined with the luma and chroma signals to form the *composite* video signal, designated CVBS (composite video with burst and sync)—the composite signal is ultimately broadcast.

A reverse process occurs in the television receiver, where the composite signal is decomposed into the various components and finally into an RGB signal which ultimately drives the three color inputs to the CRT.

It should be noted that the NTSC system was developed to be backwards compatible with the existing black-and-white standards.

What the Analog Connectors Look Like on a High-End Receiver

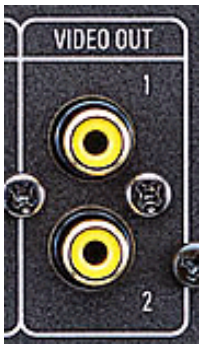

CVBS
Composite Video
with Burst and Sync

2 Outputs

S-Video
Y'/C

2 Outputs

Component Video
Y'PrPb or
Y'CrCb

1 Output

Note that each step in the construction of the composite video signal after the output of the resistive matrix has the potential of introducing artifacts in the signal. For this reason, engineers working in digital video soon realized that it would be advantageous to keep the digital video signal as close to the native R'G'B' format as possible. The first so-called *component* analog video standard developed was designated as *Y'PbPr* (note that the prime notation has been dropped from most modern nomenclature). The corresponding digital standard is designated *Y'CbCr*.

Another analog component standard is designated as Y'UV and is similar to Y'PbPr with different scaling factors for the color difference signals.

The final popular analog component standard to be discussed is the so-called *S-Video*, or simply *Y'/C*. This is a two-component analog system and is often used in high-end VCRs, DVDs, and TV receivers and monitors.

Details of these various standards can be found in the following reference:

Keith Jack, *Video Demystified, Fourth Edition*, Elsevier-Newnes, 2005, ISBN: 0-7506-7822-4.

Video Decoders and Encoders for Enhanced Definition and High Definition TV

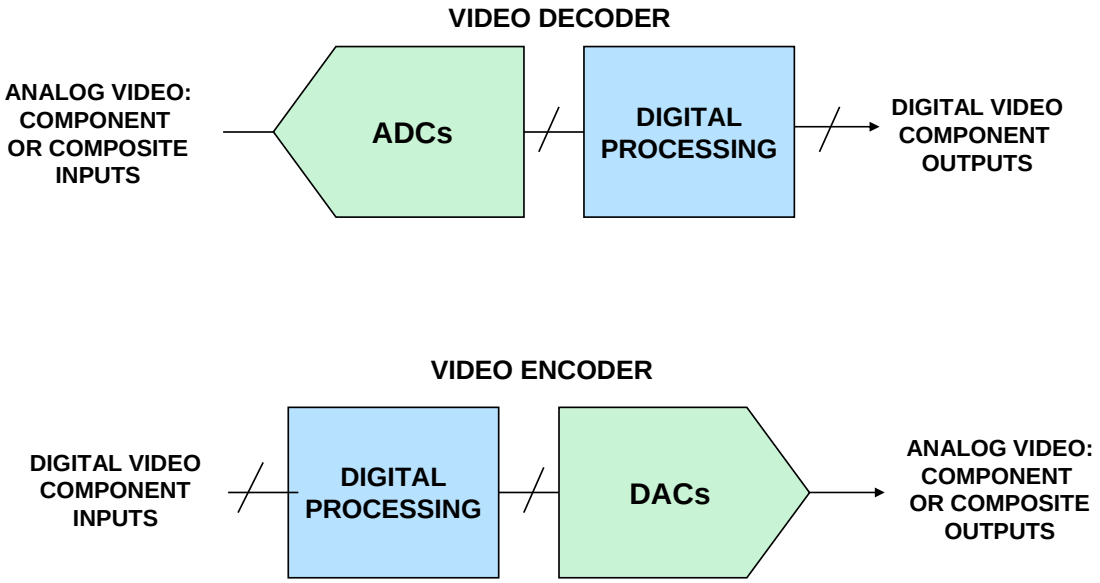

A video *decoder* converts *analog* composite video (CVBS), S-Video (Y'/C), Y'UV, or Y'PbPr signals into a digital video steam in the form of a Y'CbCr digital stream per ITU-R BT.656 4:2:2 component video compatible with NTSC, PAL B/D/G/H/I/ PAL M, PAL N, or others. An ADC function is implicit in the definition of the video decoder.

A digital video *encoder* converts digital component video (ITU-R BT.601 4:2:2, for example) into a standard composite analog baseband signal compatible with NTSC, PAL B/D/G/H/I, PAL M, or PAL N. In addition to the composite output signal, there is often the facility to output S-Video (Y'/C), RGB, Y'PbPr, or Y'UV component analog video.

In contrast to the digital video terminology, in ADC and DAC terminology, the terms *encoder* and *decoder* are used to refer to the ADC and the DAC function, respectively, and the combination is called a *codec* (coder-decoder).

The reason for this is that video engineers consider a composite color signal to have the chroma *encoded* on top of the luma signal. The video *decoder* (with the ADCs) *decodes* (separates) the chroma and luma signal and is referred to as the *decoder*. On the other hand, the video *encoder encodes* the chroma and the luma back into the composite signal.

Important Video ADC Specifications

- **Resolution, Sampling Rate, Linearity, Bandwidth**
- **Differential Gain (CVBS)**
- **Differential Phase (CVBS)**
- **SNR**
- **Chroma-Specific (Component Video)**
  - **Hue Accuracy**
  - **Color Saturation Accuracy**
  - **Color Gain Control Range**
  - **Analog Color Gain Range**
  - **Digital Color Gain Range**
  - **Chroma Amplitude Error**
  - **Chroma Phase Error**
  - **Chroma/Luma Intermodulation**
- **Luma-Specific (Component Video)**
  - **Luma Brightness Accuracy**
  - **Luma Contrast Accuracy**

This figure lists the important video specifications for the ADC used in the video decoder. These specifications are now easily met with 8 or 10-bit CMOS pipelined ADCs which are integrated into the overall decoder function.

Definitions of the above specifications and a good explanation of digital video can be found in the following reference:

Keith Jack, *Video Demystified, Fourth Edition*, Elsevier-Newnes, 2005, ISBN: 0-7506-7822-4.

ADV7180 10-Bit SDTV Video Decoder

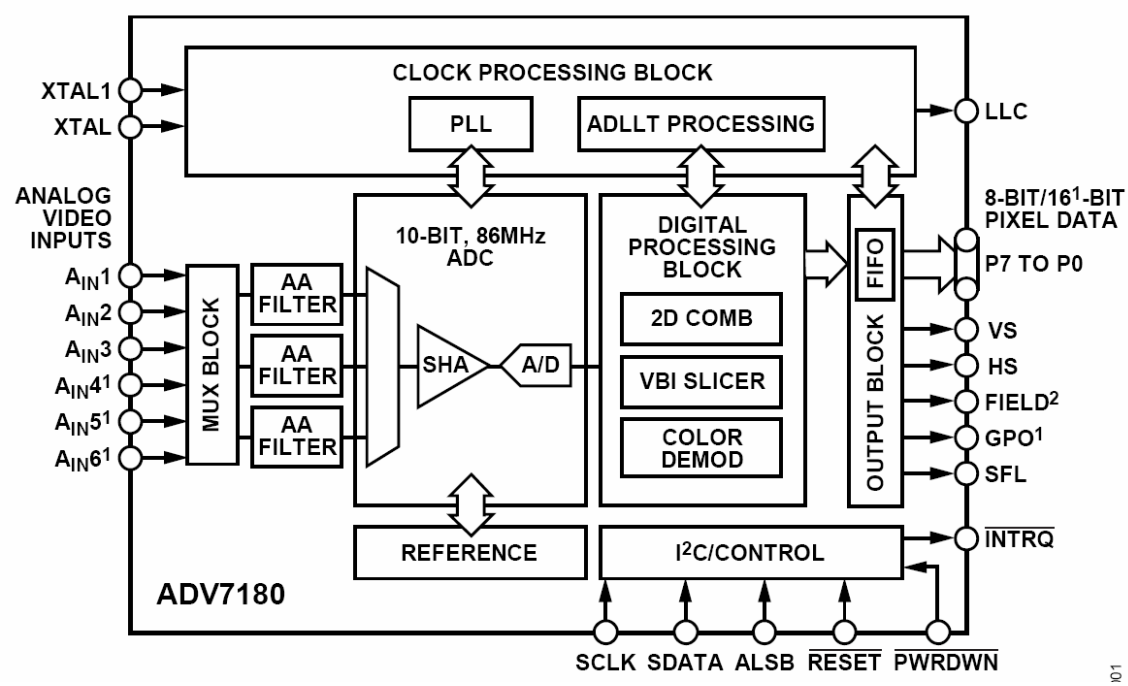

The ADV7180 automatically detects and converts SDTV (standard definition TV) analog baseband television signals compatible with worldwide NTSC, PAL, and SECAM standards into 4:2:2 component video data compatible with the 8-bit ITU-R BT.656 interface standard. The simple digital output interface connects gluelessly to a wide range of MPEG encoders, codecs, mobile video processors, and ADI digital video encoders, such as the ADV7179.

External HS, VS, and FIELD signals provide timing references for LCD controllers and other video ASICs, if required. The accurate 10-bit analog-to-digital conversion provides professional quality video performance for consumer applications with true 8-bit data resolution.

Three analog video input channels accept standard composite, S-video, or component video signals, supporting a wide range of consumer video sources. AGC and clamp-restore circuitry allow an input video signal peak-to-peak range up to 1.0V. Alternatively, these can be bypassed for manual settings. The line-locked clock output allows the output data rate, timing signals, and output clock signals to be synchronous, asynchronous, or line locked even with ±5% line length variation. Output control signals allow glueless interface connections in many applications.

The ADV7180 is programmed via a 2-wire, serial, bidirectional port (I^2C compatible). The ADV7180 is fabricated in a 1.8V CMOS process. Its monolithic CMOS construction ensures greater functionality with lower power dissipation. A chip-scale, 40-lead, LFCSP package option makes the decoder ideal for space constrained portable applications. A 64-lead LQFP package is also available (pin compatible with ADV7181B).

Further details on other video encoders and decoder products can be found at www.analog.com/video.

Flat Panel Display Interface Electronics

www.analog.com/display

Flat Panel Analog and Digital Interfaces

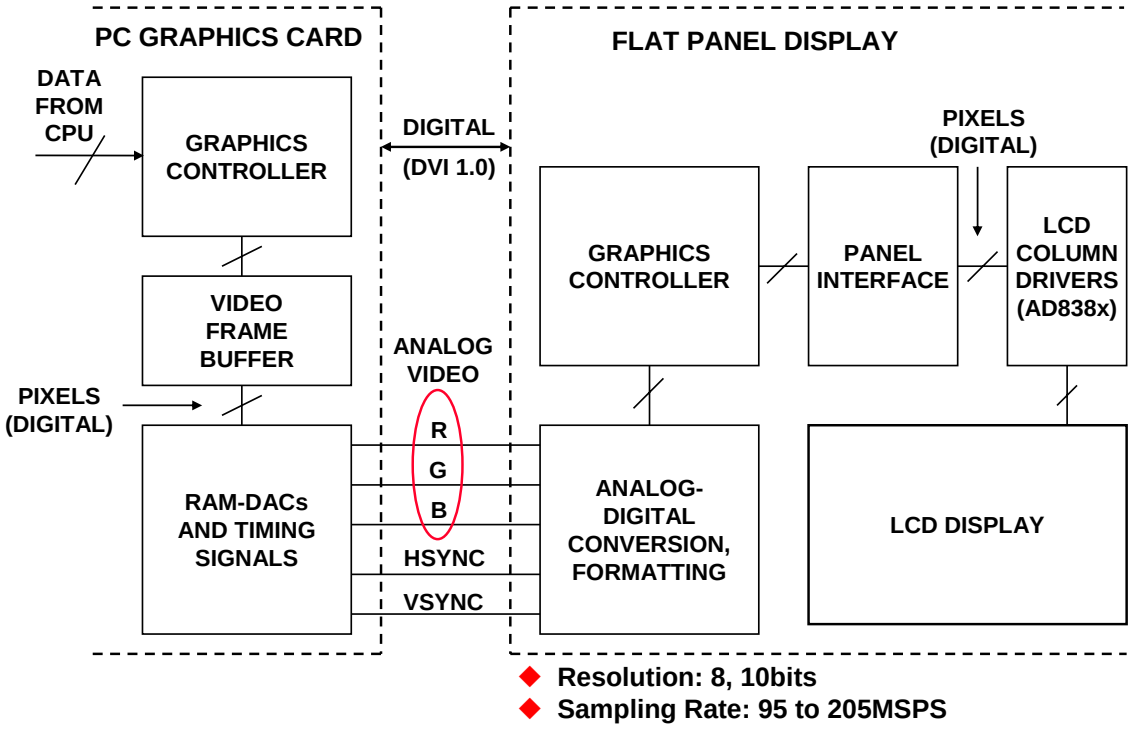

The popularity of flat panel LCD-based displays has steadily increased over the last few years, and they are rapidly replacing CRT-based monitors in desktop computer systems. In addition, LCD projectors have virtually replaced 35mm slide and overhead projectors as a means of delivering presentation material.

The graphics card in a typical desktop computer system converts the digital pixel data to an analog RGB signal for driving an external monitor. In a laptop computer the built-in LCD display is generally driven directly with the digital data, and it is also converted to analog RGB video using video DACs, where it is available on an output connector for driving an external monitor or projector.

The analog RGB interface to the CRT is the primary workhorse in the display of computer-generated graphics data. A large legacy of PC-graphics adapters currently exist that use RAM-DACs (video DACs with on-chip color lookup table memory) to convert digital graphics data to analog RGB signals. The new flat panel displays must therefore be able to interface with this conventional technology to achieve market penetration and fast acceptance.

In an effort to establish an industry-wide standard for the next-generation flat panel displays, the Digital Display Working Group (DDWG) developed the Digital Video Interface (DVI 1.0) specification. This specification describes how designers should implement the analog and digital interfaces. Analog timing is described in the Video Electronics Standards Association (VESA) standard for monitors, and the digital interface uses Transition Minimized Differential Signaling (TMDS) format.

AD9888 100/140/170/205MSPS Analog Flat Panel Interface

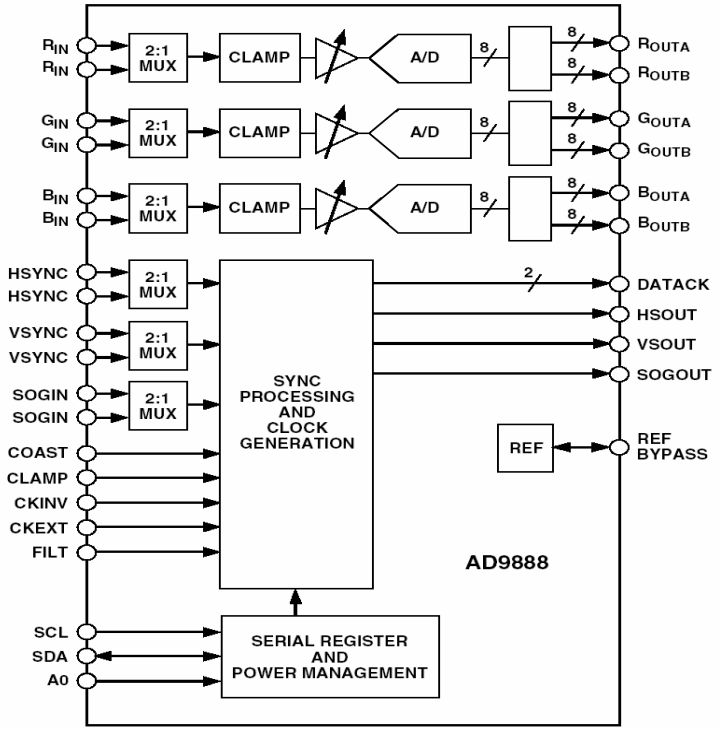

The AD9888 is an example of a highly integrated flat panel interface product. It is a complete 8-bit, 205MSPS monolithic analog interface optimized for capturing RGB graphics signals from personal computers and workstations. Its 205MSPS sampling rate capability and full-power analog bandwidth of 500MHz supports resolutions up to UXGA (1600 × 1200 at 75Hz).

For ease of design and to minimize cost, the AD9888 is a fully integrated interface solution for flat panel displays. The AD9888 includes an analog interface with a 205MHz triple ADC with internal 1.25V reference, PLL to generate a pixel clock from HSYNC and COAST, midscale clamping, and programmable gain, offset, and clamp control. The user provides only a 3.3V power supply, analog input, and HSYNC and COAST signals.

Three-state CMOS outputs may be powered from 2.5V to 3.3V. The AD9888's on-chip PLL generates a pixel clock from HSYNC and COAST inputs. Pixel clock output frequencies range from 10MHz to 205MHz. PLL clock jitter is typically less than 450ps p-p at 205MSPS. When the COAST signal is presented, the PLL maintains its output frequency in the absence of HSYNC. A sampling phase adjustment is provided. Data, HSYNC, and clock output phase relationships are maintained. The PLL can be disabled and an external clock input can be provided as the pixel clock.

The AD9888 also offers full sync processing for composite sync and Sync-on-Green applications. A clamp signal is generated internally or may be provided by the user through the CLAMP input pin. This interface is fully programmable via a 2-wire serial interface. Fabricated in an advanced CMOS process, the AD9888 is provided in a space-saving 128-lead MQFP surface-mount plastic package and is specified over the 0°C to 70°C temperature range.

Details of other display interface products can be found at www.analog.com/display.

High Speed ADC Applications in Ultrasound

www.analog.com/adcs

www.analog.com/amps

www.analog.com/vga

Ultrasound Application: Analog Versus Digital Beamforming

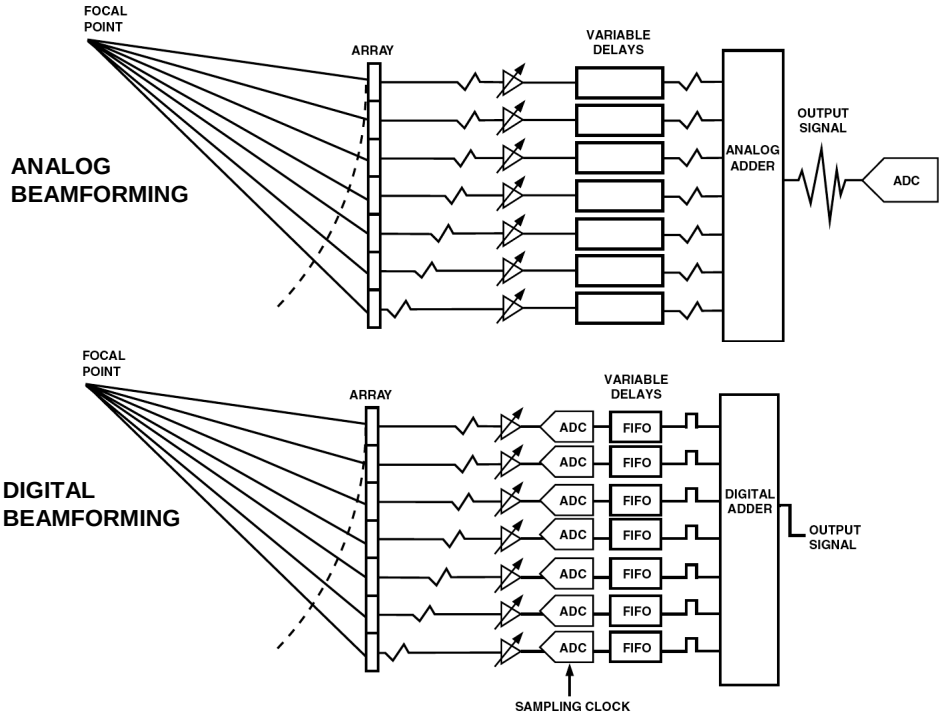

Medical ultrasound machines are among the most sophisticated signal processing machines in widespread use today. As in any complex machine, there are many trade-offs in implementation due to performance requirements, physics, and cost. Some system-level understanding is necessary to fully appreciate the desired front-end IC functions and performance levels, especially for: the low-noise amplifier (LNA); time gain compensation amplifier (TGC); and analog-to-digital converters (ADCs).

In analog beamforming (ABF) and digital beamforming (DBF) ultrasound systems, the received pulses reflected from a particular focal point along a beam are stored for each channel, then aligned in time, and coherently summed—this provides spatial processing gain because the noise of the channels is uncorrelated. Images may be formed as either a sequence of analog levels that are delayed with analog delay lines, summed, and converted to digital after summation (ABF)—or digitally by sampling the analog levels as close as possible to the transducer elements, storing them in a memory (FIFO), and then summing them digitally (DBF).

This figure show basic respective block diagrams of ABF and DBF systems. Both types of systems require perfect channel-to-channel matching. Note that the variable-gain amplifiers (VGAs) are needed in both implementations—and will continue to be in the digital case until ADCs with a large enough dynamic range become available at reasonable cost and low enough power.

Note that an ABF imaging system needs only one very high resolution and high speed ADC, but a DBF system requires many high speed, high resolution ADCs. Sometimes a logarithmic amplifier is used in the ABF systems to compress the dynamic range before the ADC.

Time Gain Amplifier Compensates for Signal Attenuation in Body

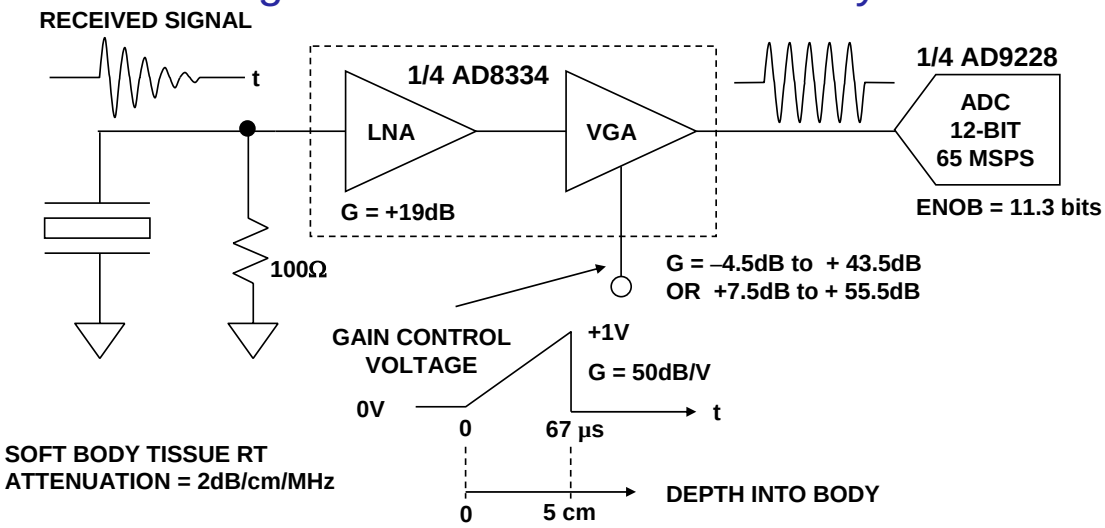

AT 5 MHz:

- **RT PROP. DELAY = 13.4 µs/cm FOR SOFT BODY TISSUE**
- **RT ATTENUATION = 10 dB/cm**
- **TGA RAMP SLOPE = (10 dB/cm)/(13.4 µs/cm) = 0.746 dB/µs = 15 mV/µs**

Operating frequencies for medical ultrasound are in the 1MHz to 40MHz range, with external imaging machines typically using frequencies of 1MHz to 15MHz, while intravenous cardiovascular machines use frequencies as high as 40MHz. Higher frequencies are in principle more desirable, since they provide higher resolution—but tissue attenuation limits how high the frequency can be for a given penetration distance. However, one cannot arbitrarily increase the ultrasound frequency to get finer resolution, since the signal experiences an attenuation of about 1dB/cm/MHz; i.e., for a 5MHz ultrasound signal and a penetration depth of 5 cm, the round-trip signal has been attenuated by 5 × 2 × 5 = 50dB. To handle an instantaneous dynamic range of about 60dB at any location, the required dynamic range would be 110dB.

A single channel of a DBF ultrasound system is shown in this figure. The AD8334 time gain amplifier provides a variable gain range of approximately 48dB that is proportional to the control voltage (G = 50dB/V). A ramp control voltage of the proper slope allows the AD8334 to compensate for the attenuation through body tissue and to present a signal to the ADC which has a more or less constant amplitude for the duration of the burst.

For the example shown in the figure, the round trip propagation delay is 13.4µs/cm. The round trip attenuation is 10dB/cm. The gain compensation should therefore be 10dB/cm ÷ 13.4µs/cm = 0.746dB/µs. The AD8334 gain-to-control voltage ratio is 50dB/V. Therefore, a ramp slope of 0.746dB/µs ÷ 50dB/V = 15mV/µs compensates for the body tissue attenuation at 5MHz.

The ADC selected for this application is the quad 12-bit, 65MSPS AD9228.

Multichannel Ultrasound Application for AD8334 VGA and AD9228 Uses 264mW/Channel

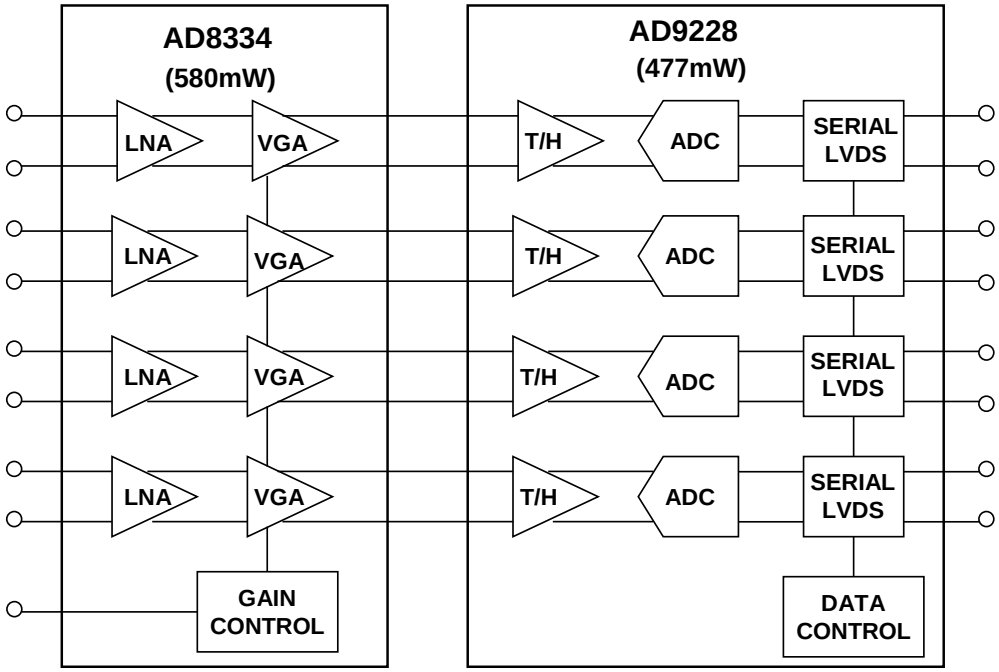

This figure shows a block diagram of how the quad AD8334 amplifier and the quad AD9228 ADC form an ultrasound system dissipating only 264mW/channel. The AD8331/AD8332/AD8334 are single-, dual-, and quad-channel ultralow noise, linear-in-dB, variable gain amplifiers (VGAs). Optimized for ultrasound systems, they are usable as a low noise variable gain element at frequencies up to 120MHz. Included in each channel are an ultralow noise preamplifier (LNA), an X-AMP® VGA with 48dB of gain range, and a selectable gain postamplifier with adjustable output limiting. The LNA gain is 19dB with a single-ended input and differential outputs. Using a single resistor, the LNA input impedance can be adjusted to match a signal source without compromising noise performance. The 48dB gain range of the VGA makes these devices suitable for a variety of applications. Excellent bandwidth uniformity is maintained across the entire range. The gain control interface provides precise linear-in-dB scaling of 50dB/V for control voltages between 40mV and 1V. Factory trim ensures excellent part-to-part and channel-to-channel gain matching.

Differential signal paths result in superb second- and third-order distortion performance and low crosstalk. The VGA's low output-referred noise is advantageous in driving high speed differential ADCs. The gain of the postamplifier can be pin selected to 3.5dB or 15.5dB to optimize gain range and output noise for 12-bit or 10-bit converter applications. The output can be limited to a user-selected clamping level, preventing input overload to a subsequent ADC. An external resistor adjusts the clamping level. The operating temperature range is −40°C to +85°C. The AD8334 is available in a 64-lead LFCSP package.

The AD9228 is a quad 12-bit 65MSPS ADC which dissipates only 477mW and has 11.3 effective bits at ultrasound frequencies.

Eberhard Brunner, "How Ultrasound System Considerations Influence Front-End Component Choice," *Analog Dialogue* 36-03, 2002.

Key Specifications for Ultrasound ADCs

- **Sampling Rate: 65MSPS**
- **Resolution: 12 bits**
- **Bandwidth: 100MHz**
- **Effective Number of Bits (ENOB): Greater than 10 at the maximum ultrasound burst frequency**
- **Power per channel**
- **Cost per channel**
- **Package size per channel**
- **Data outputs are generally serial LVDS to conserve pins and reduce noise)**

This figure shows the key specifications for ultrasound ADCs. Serial LVDS data outputs are used to conserve pin count, reduce package size, and minimize digital switching noise.

ADC Evaluation Hardware and ADIsimADC® Modeling Tool

www.analog.com/fifo

www.analog.com/adisimadc

www.analog.com/designtools

Analog Devices' High Speed ADC FIFO Evaluation Kit

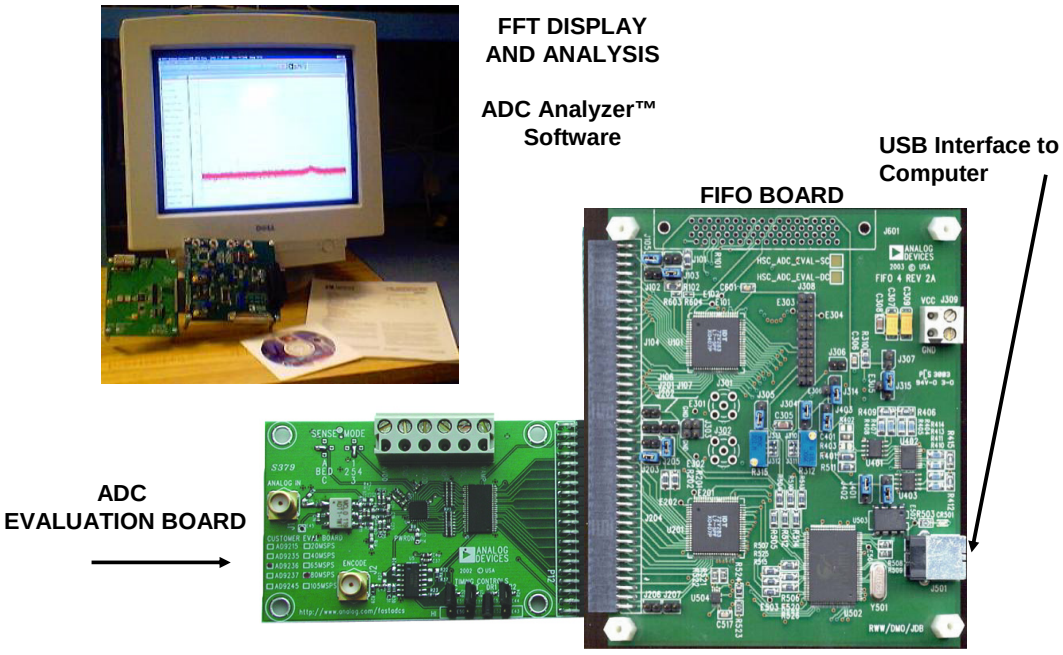

Analog Devices offers complete hardware and software tools for evaluating data converters. In order to make FFT measurements, a suitable memory must capture the ADC output data. The FIFO Board shown in the figure is designed to mate with the product-specific evaluation board and store the ADC data. The interface to the PC is via a standard USB port.

The ADC Analyzer™ software provides a convenient way to generate the FFT of the converter output, and the software also computes SNR, SFDR, harmonic distortion, etc.

The evaluation board schematic, layout, and parts list is contained in the data sheet for the individual ADCs. The evaluation represents an optimum layout and can be used as a guide when laying out the actual system board that contains the ADC. Gerber files for the PCB layout are available upon request.

FPGA Board Allows Easy Conversion from Serial LVDS

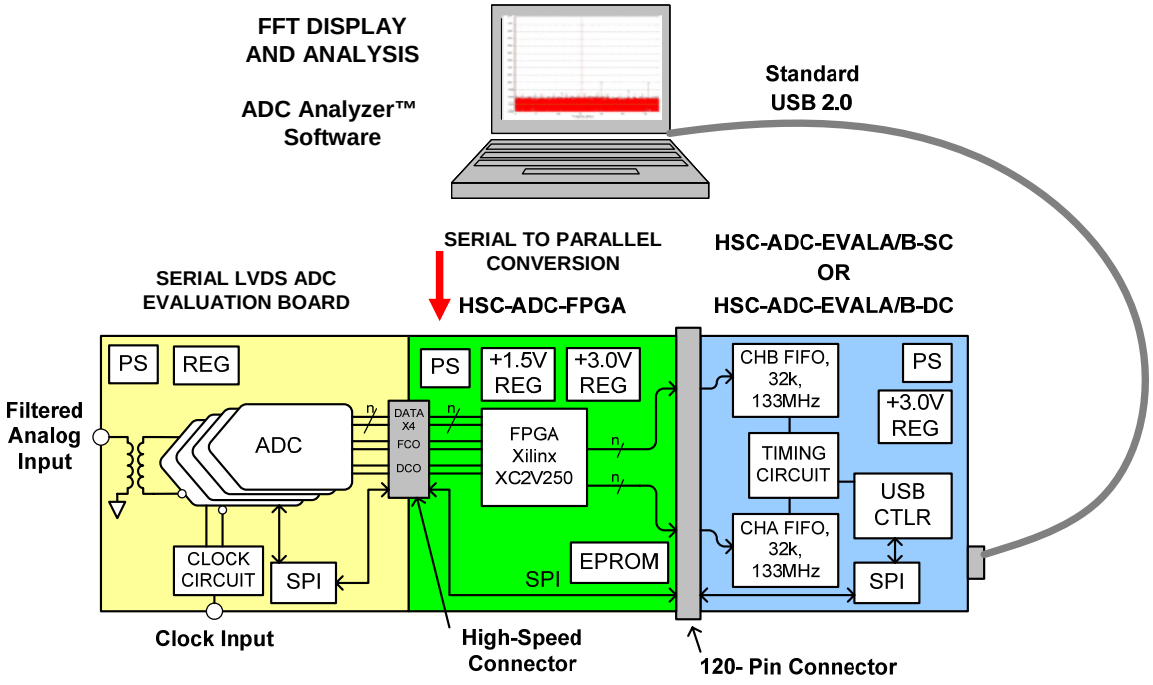

In order to accommodate ADCs with LVDS outputs, the HSC-ADC-FPGA interface board is added to perform the serial-to-parallel conversion required to interface LVDS ADCs to the FIFO Board.

Further details on the operation of the FIFO Board as well as the ADC evaluation boards can be found at www.analog.com/fifo.

ADIsimADC® Modeling Tool

- **Virtual Evaluation Boards**
  - **Rapidly Preview Many Converters Without Bench Setup**
  - **With ADC Analyzer You Can Directly Compare the ADIsimADC Models to the Hardware Performance**
- **Live Data Sheets**
  - **Don't Like the Conditions in the Data Sheet? User Can Pick the Input Conditions**
- **Tool Support**
  - **Matlab Documentation Available. Others Released on an Ongoing Basis**

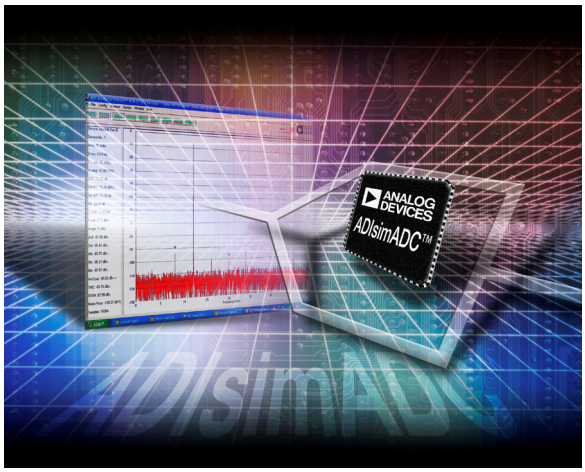

ADIsimADC® simulates behavior using many critical specifications of the data converters, including offset, gain, sample rate, bandwidth, jitter, latency, both ac and dc nonlinearity.

These models provide a way to validate performance of the converter at the system level to determine the applicability of a selected device.

Users can run real-time FFT and time-domain analysis; analyze SNR, SINAD, SFDR, and harmonics; and export data for additional analysis.

The ADIsimADC program and models can be downloaded from www.analog.com/adisimadc at no charge. The model library is continually updated as new products are released.

An interactive web-based version of ADIsimADC is available at www.analog.com/designcenter.

ADC Modeling: ADIsimADC®

- ADC behavioral modeling allows the behavior of the ADC to be included in system simulations.
- Behavioral Model
  - Not bit exact model
  - Not algorithmic (i.e., doesn't attempt to model the signal flow within the ADC)
  - Transfer function based on actual INL data
  - Modeling allows for many design tradeoffs to be made before the system is committed to hardware.
- Many ADCs can quickly be evaluated, thereby reducing the uncertainty of selecting the correct ADC.
- Modeling accounts for:
  - Gain, offset, pipelines, jitter, DNL & INL, frequency dependent transfer characteristics, thermal noise, and many other characteristics

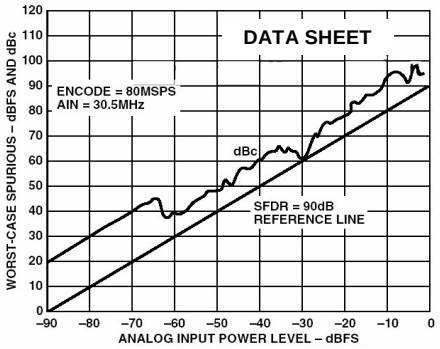

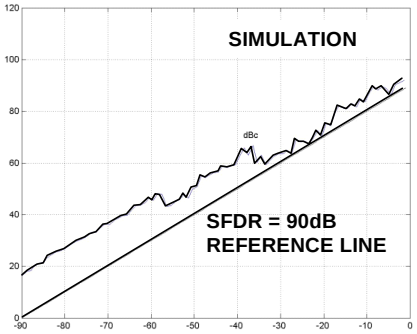

ADIsimADC is a behavioral model which utilizes a transfer function based on characterization data for the ADC under consideration. The models are posted on the ADI website and are updated when new products are released.

This figure shows how close the simulation matches the actual data sheet plot of SFDR versus input power level for the AD6645 ADC. The "dips" in SFDR occur at predictable places across the dynamic range of the converter and are typical of the nonlinearities associated with the transfer function of the pipelined architecture.

Setting Up ADIsimADC®

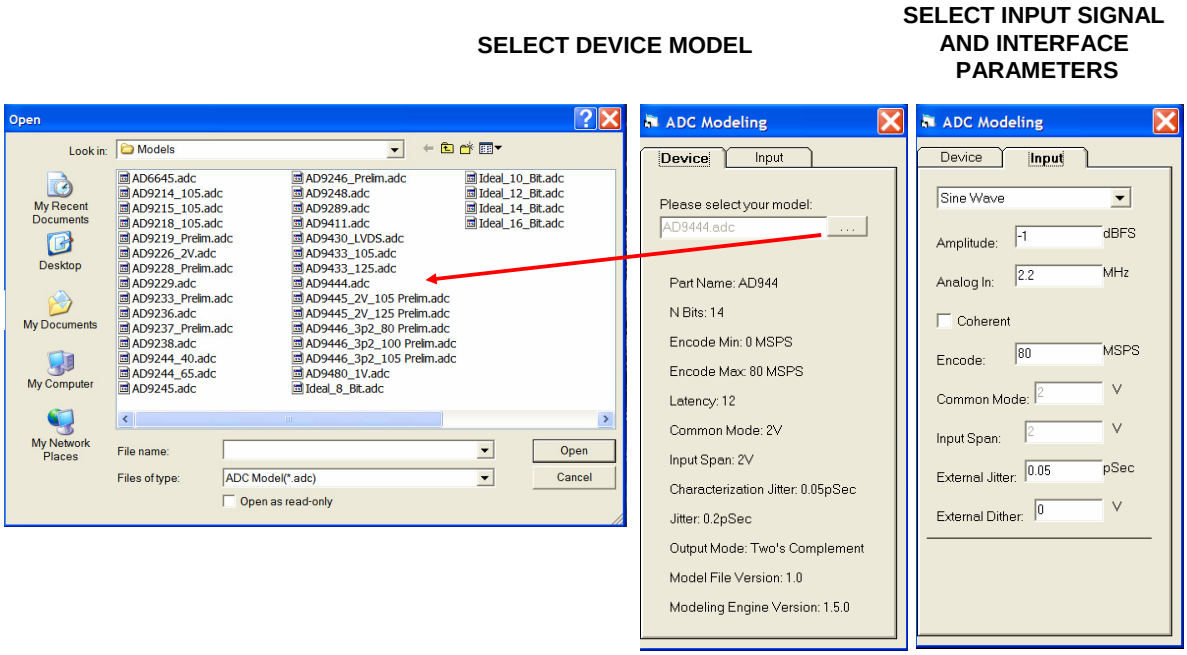

The ADIsimADC software is easy to run. This figure shows how the "Device" tab is first selected. The latest model set can be downloaded from www.analog.com/adisimadc, and each ADC model shows up in the list.

The "Input" tab is then used to select the input signal frequency, amplitude, and sample rate. Note that external jitter as well as dither can be added if desired.

In addition to individual product models, models are provided for *ideal* 8, 10, 12, 14, and 16-bit ADCs. These ideal models are useful as educational tools, as well as comparing ideal and actual performance.

AD9246 14-Bit, 105/125MSPS ADC
f_{in} = 2.211MHz, f_s = 125MSPS, FFT Output

Device: AD9246_Prelim
Device No.: 1
Avcc: 5. Volts
Dvcc: 3.3 Volts
Encode: 125. MSPS
Analog: 2.22 MHz
SNR: 71.03 dB
SNRFS: 72.03 dBFS
UDSNR: 92.23 dB
NF: 30.71 dB
SINAD: 70.85 dB
Fund: -0.999 dBfs
Image: 0. dBc
2nd: -90.48 dBc
3rd: -86.86 dBc
4th: -101.11 dBc
5th: -101.36 dBc
6th: -93.89 dBc
WoSpur: -89.78 dBc +
THD: -84.54 dBc
SFDR: 86.87 dBc
Noise Floor: -111.16 dBFS
Samples: 16384
Windowing: Bh4

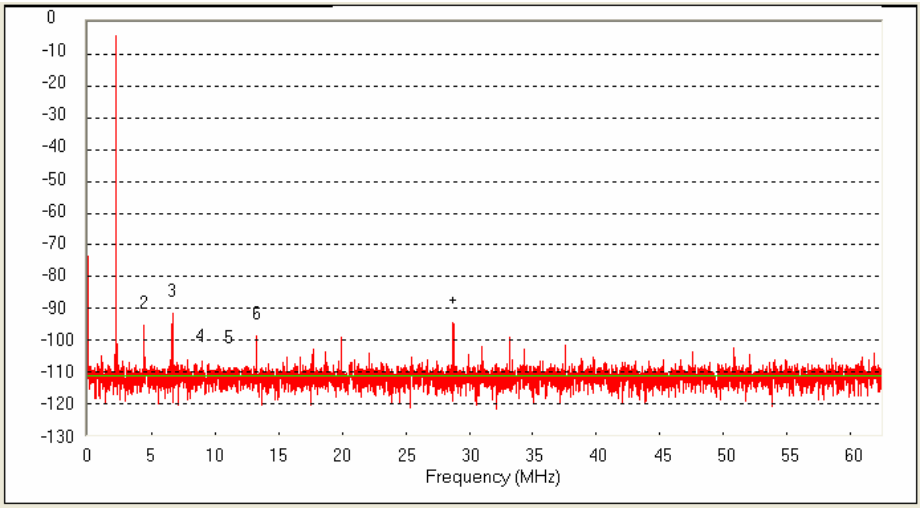

DATA TABLE **FFT OUTPUT**

This shows an ADIsimADC FFT output for the AD9246 14-bit, 105/125MSPS ADC for an input frequency of 2.211MHz and a sample rate of 125MSPS. Note that the program labels the position of the first five harmonic products

The data table on the left shows the input conditions as well as the calculated values for SNR, SINAD, harmonic distortion products, THD, and SFDR. The noise floor of the output is also calculated.

Note that the largest non-harmonic spur is labeled with a "+" symbol.

AD9246 14-Bit, 105/125MSPS ADC
f_{in} = 170.111MHz, f_s = 125MSPS, FFT Output

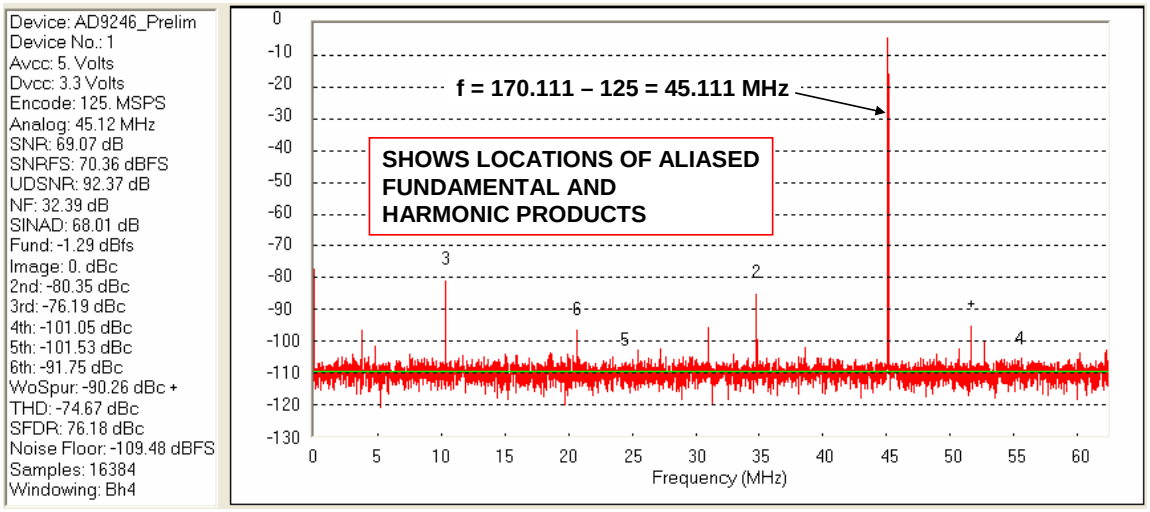

The ADIsimADC program works for IF sampled signals which are outside the first Nyquist zone. In this figure, the sample rate is 125MSPS and the input signal is located at the IF frequency of 170.111MHz. The FFT output shows the aliased fundamental of this signal at 45.111MHz.

Note that the position of the aliased harmonic products are also labeled in the FFT output.

AD9246 14-Bit, 105/125MSPS ADC SFDR Amplitude Sweep f_{in} = 170.111MHz, f_s = 125MSPS

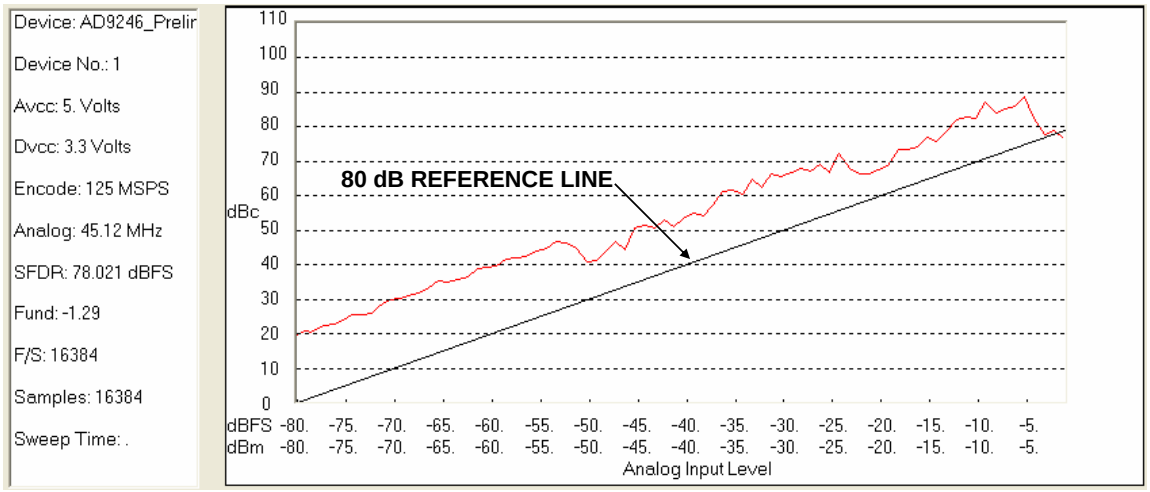

The FFT data generated by ADIsimADC can be displayed in a standard FFT as previously shown, or as an amplitude "sweep" as shown in this figure for SFDR.

The input is 170.111MHz, the sampling rate is 125MSPS, and the SFDR is displayed as the input signal level is swept from –80dBFS to –1dBFS. The user defines the range as well as the step size.

AD9246 14-Bit, 105/125MSPS ADC SFDR Frequency Sweep f_{in} = 2 to 170MHz, f_s = 125MSPS

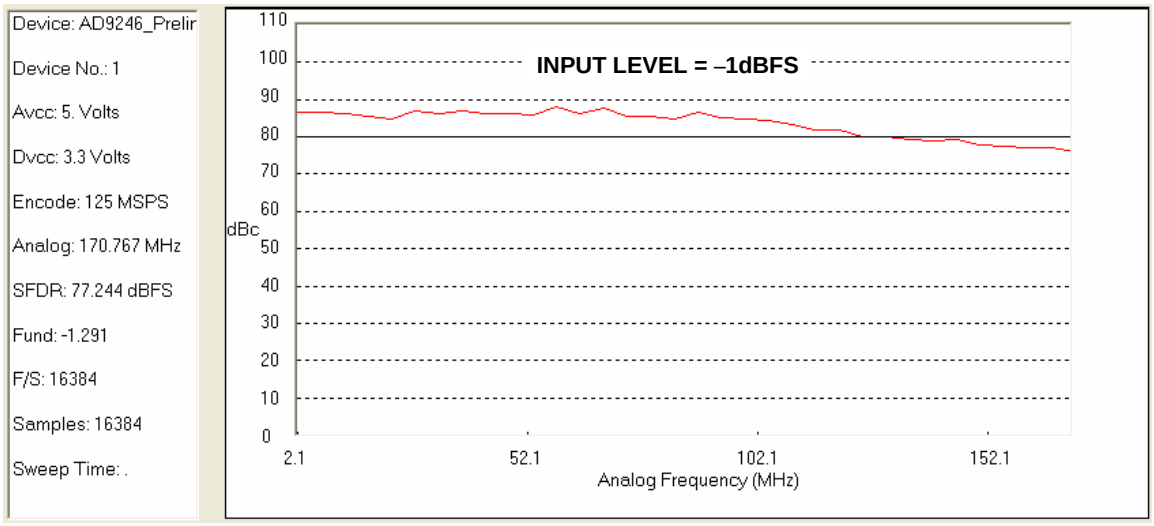

This figure shows a frequency "sweep" output, where the ADC is operated at a fixed sampling frequency and input amplitude, and the input frequency is "swept" over the desired range.

In this example, the sampling rate is set for 125MSPS, and the SFDR is plotted as the input frequency is varied from 2.1MHz to 170.1MHz.

SNR and SINAD can also be plotted as an amplitude or frequency sweep.

FFT Averaging and Windowing Options

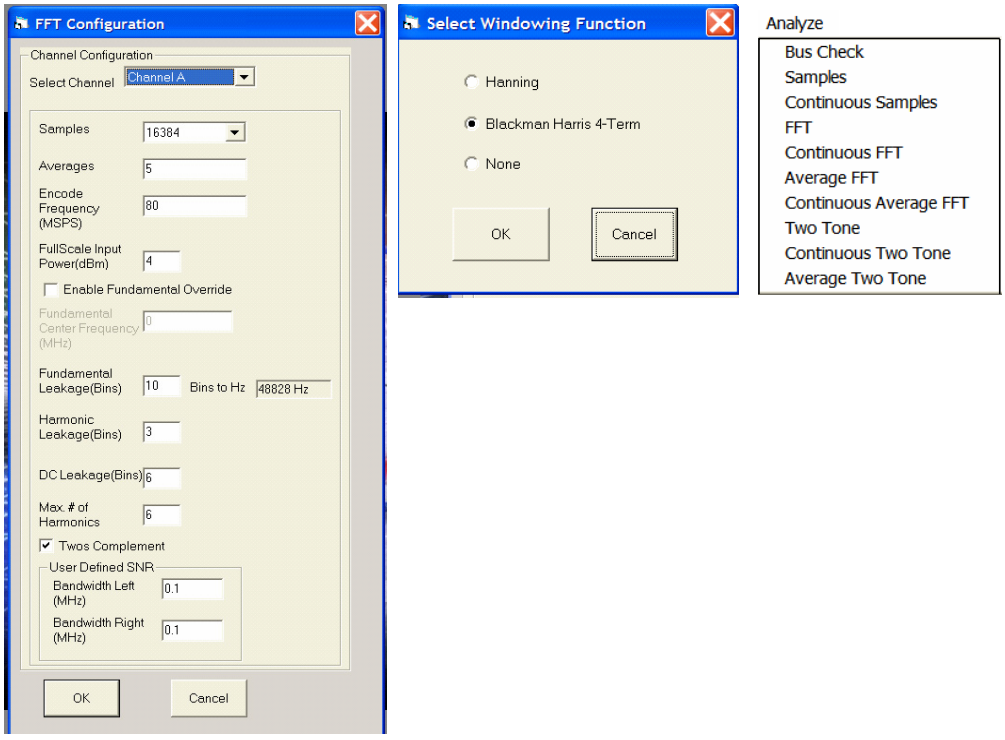

This figure shows some of the FFT and display options available in ADIsimADC.

The following can be set in the first menu: FFT Sample size, number of FFTs to be averaged, number of "bins" to include in calculating the energy of the fundamental, number of "bins" to include in calculating the energy in the harmonics, and the number of bins to include in the dc component.

The second window allows the selection of the FFT windowing function: Hanning, 4-Term Blackman-Harris, or no window.

The third menu determined what is displayed in frequency output: bus check, actual sample amplitudes, single FFT, continuous FFT, average of several FFTs, continuous average FFT, two tone input, continuous two tone input, and average two tone input.

ADIsimADC Web-Based Version

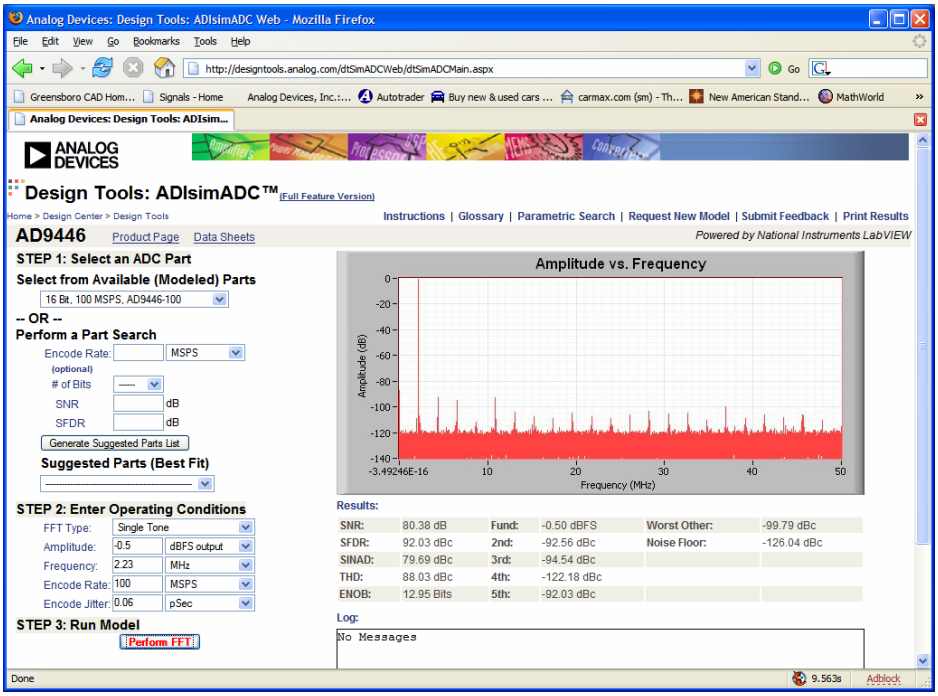

A web-based version of ADIsimADC is also available at www.analog.com/designtools. This on-line program is capable of running the basic FFT function as well as frequency and amplitude sweeps on the various ADCs under user-defined operating conditions.

The program also will perform a parts search based on sampling rate, resolution, SNR, and SFDR requirements. Suggested parts that fit the operating parameters are then displayed, and the simulation capability can then be used to assist in further selection.

The web-based version is based on the full-featured downloadable ADIsimADC, but with a more limited number of computational and display options.

Notes: