High Speed System Applications

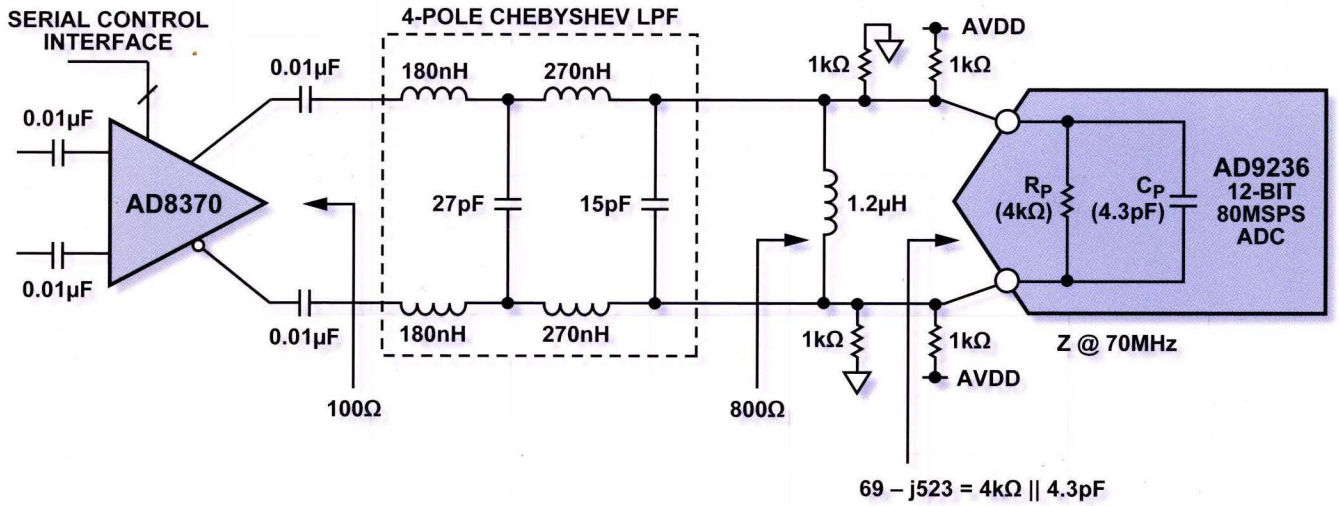

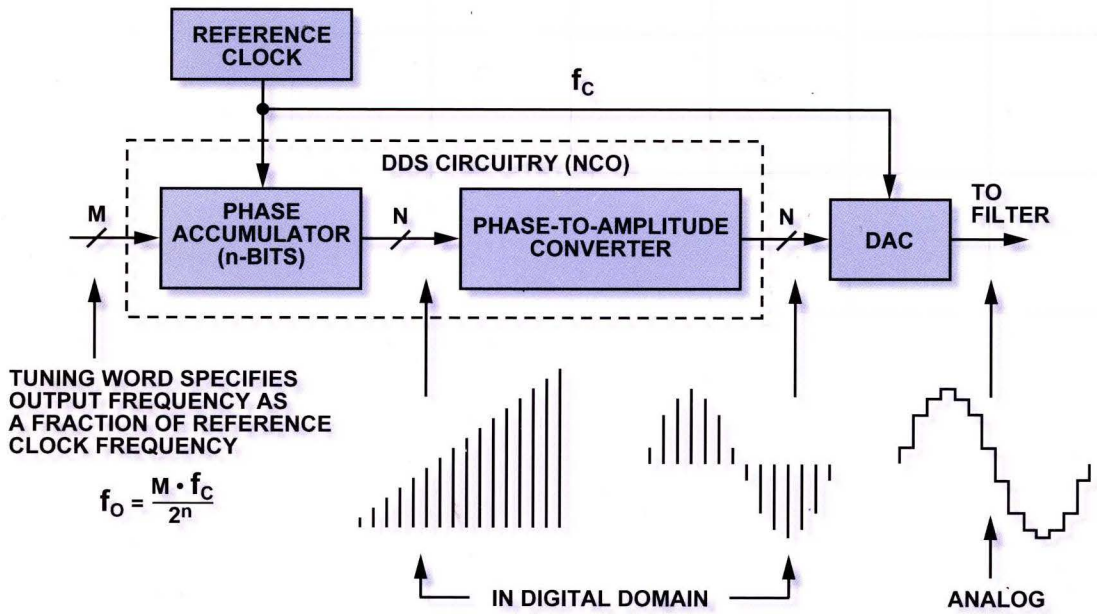

High Speed System Applications
ANALOG DEVICES

HIGH SPEED SYSTEM APPLICATIONS

Acknowledgments:

Thanks are due the many technical staff members of Analog Devices in Engineering, Marketing, and Applications who provided invaluable inputs during this project.

Special thanks go to Hank Zumbahlen who provided major contributions and also reviewed all the material.

John Galgay proofread the manuscript, and Bill Fennell prepared the artwork for the cover.

Walt Kester, August 2006
Analog Devices Central Applications Department

Direct questions to Linear.Apps@analog.com, with a subject line of "High Speed System Applications"

ISBN-10: 1-56619-909-3
ISBN-13: 978-1-56619-909-4

High Speed System Applications

Table of Contents

High Speed System Applications

Table of Contents

High Speed System Applications

Table of Contents

High Speed System Applications

Table of Contents

High Speed System Applications

Table of Contents

High Speed System Applications

Table of Contents

High Speed System Applications

Table of Contents

High Speed System Applications

Table of Contents

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.

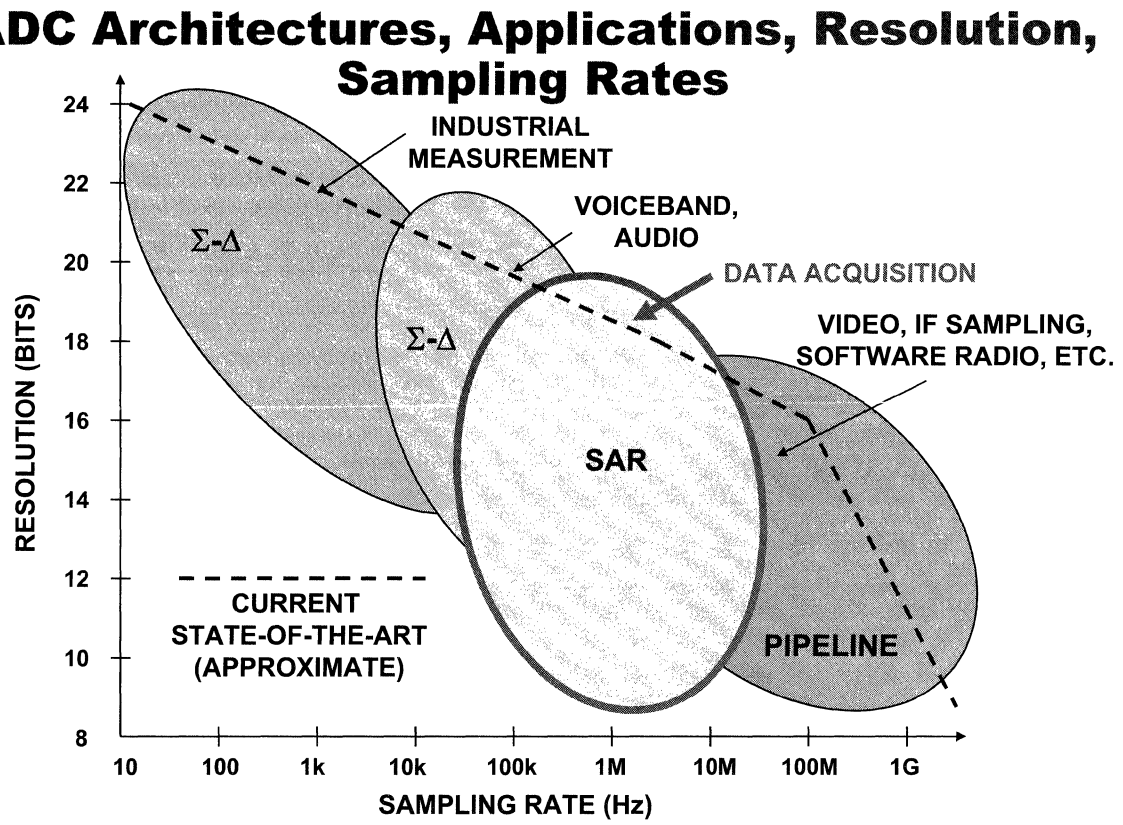

Most ADC applications today can be divided into four broad market segments: *Data Acquisition*, *Precision Industrial Measurement*, *Voiceband and Audio*, and *High Speed ("High Speed" implying sampling rates greater than approximately 10MSPS—although this line of demarcation is somewhat arbitrary. For 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

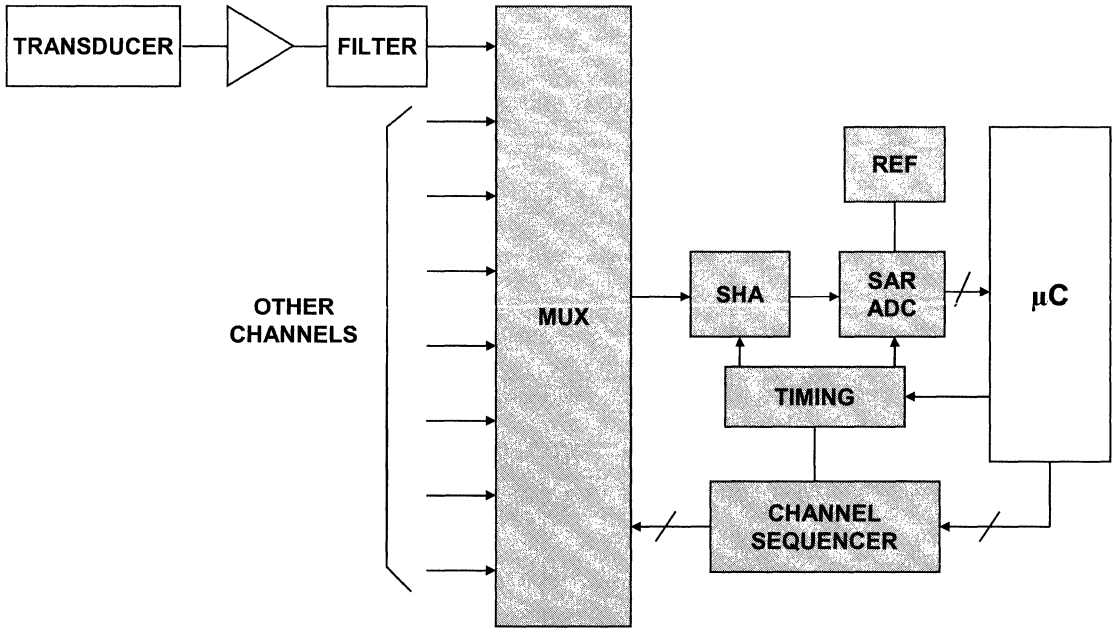

This figure shows a typical multiplexed data acquisition system. The successive approximation (SAR) ADC is the building block.

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)

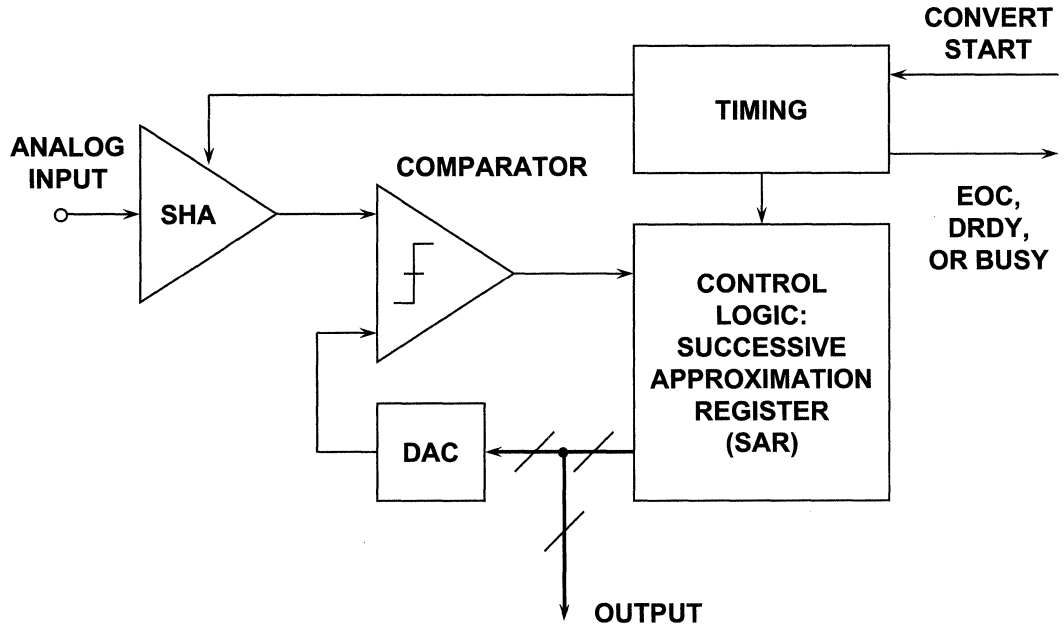

The SAR ADC performs conversions on command. On the assertion of the CONVERT START command (note that this function may actually be named something else or be combined with another control line), the sample-and-hold (SHA) is placed in the *hold* mode, and all the bits of the successive approximation 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.

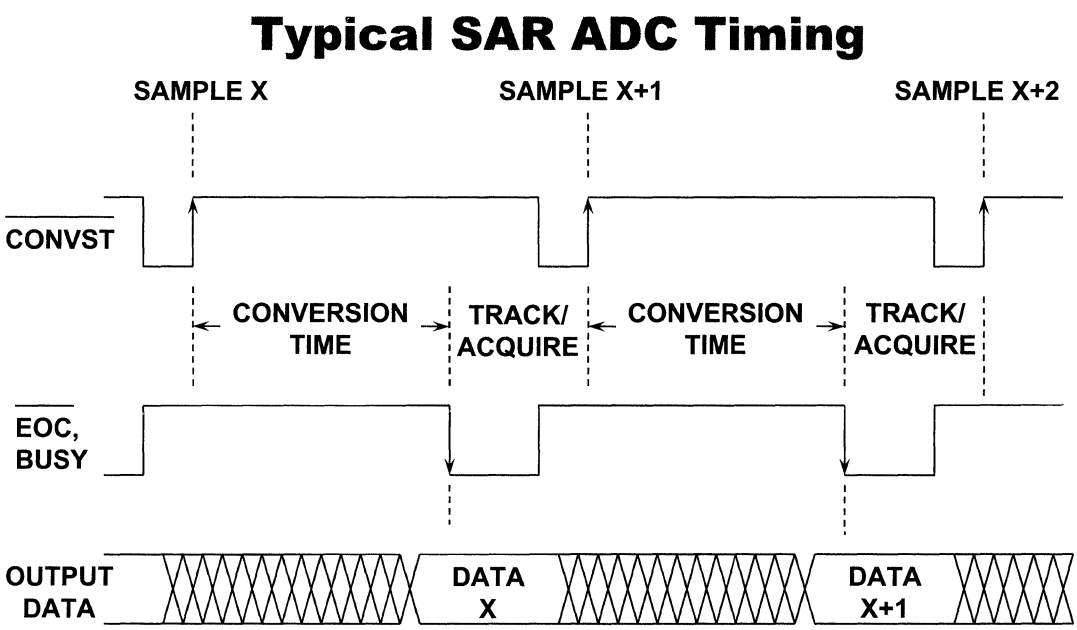

This figure shows the basic timing of a typical SAR ADC. The end of conversion is generally indicated by an end-of-convert (EOC), data-ready (DRDY), or a busy signal (actually, *not*-BUSY indicates end of conversion). The polarities and name of this signal may be different for different SAR ADCs, but 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

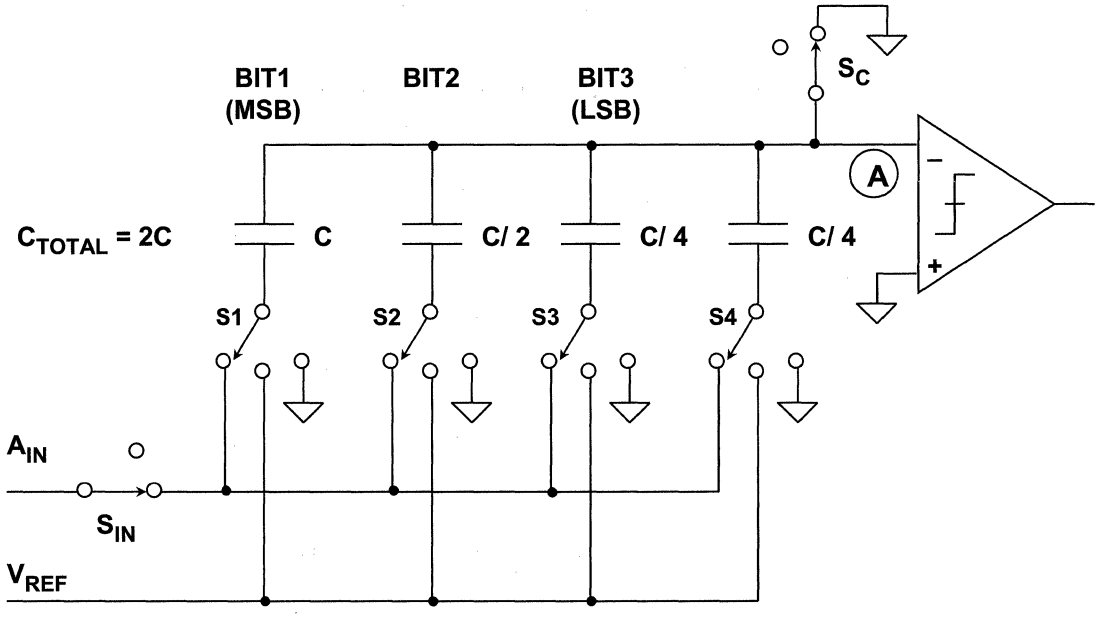

SWITCHES SHOWN IN TRACK (SAMPLE) MODE

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

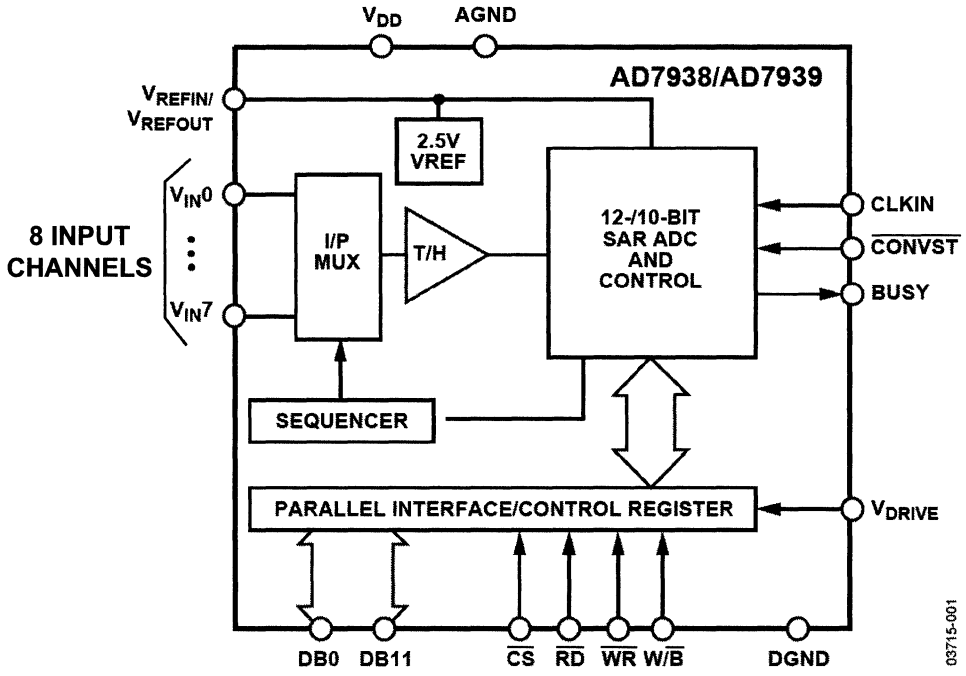

One drawback to the switched capacitor SAR architecture is the switching transient currents which can be injected on the analog input. This requires that the drive amplifier settle to these transient currents within approximately one-half of the conversion period. However, many SAR ADCs using the switched 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

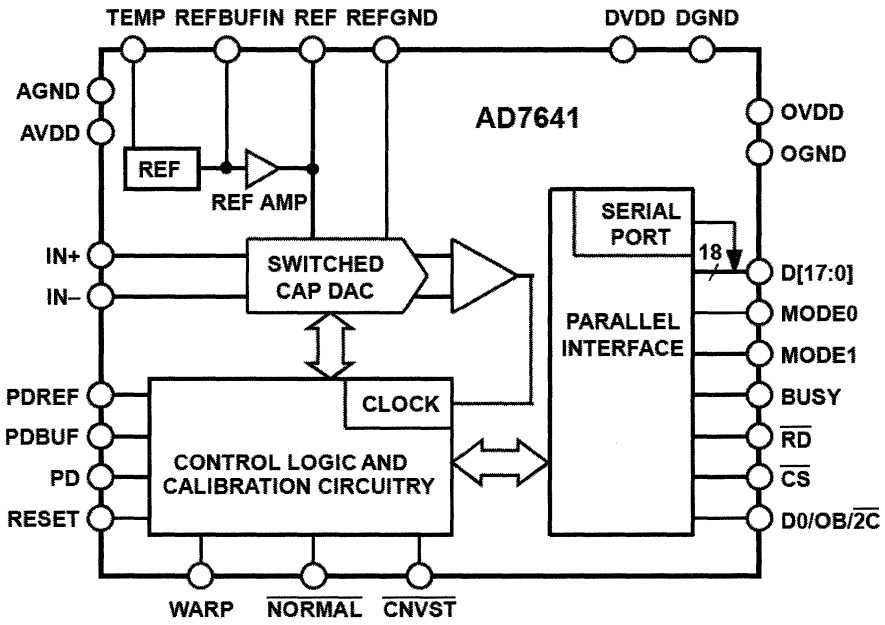

The technology in SAR ADCs has pushed the sampling rate to 2MSPS and resolution to 18-bits, as illustrated in the AD7641 PulSAR® shown in the figure.

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

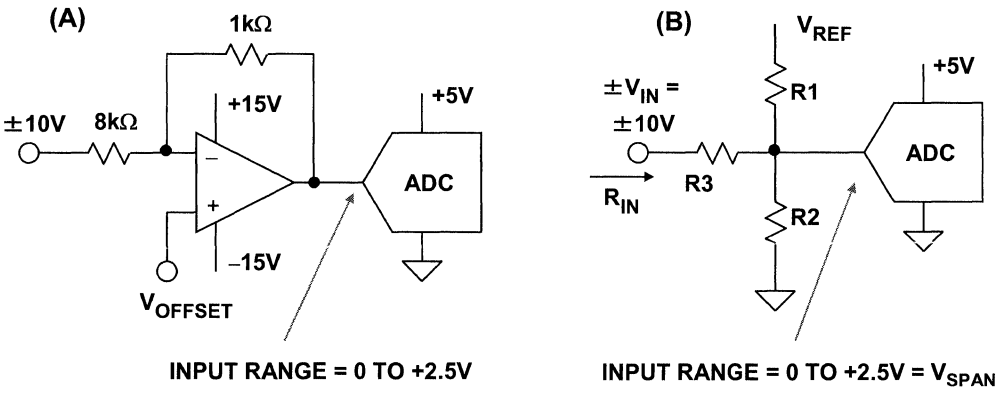

Many industrial applications still require ADCs that can handle ±10V signals

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

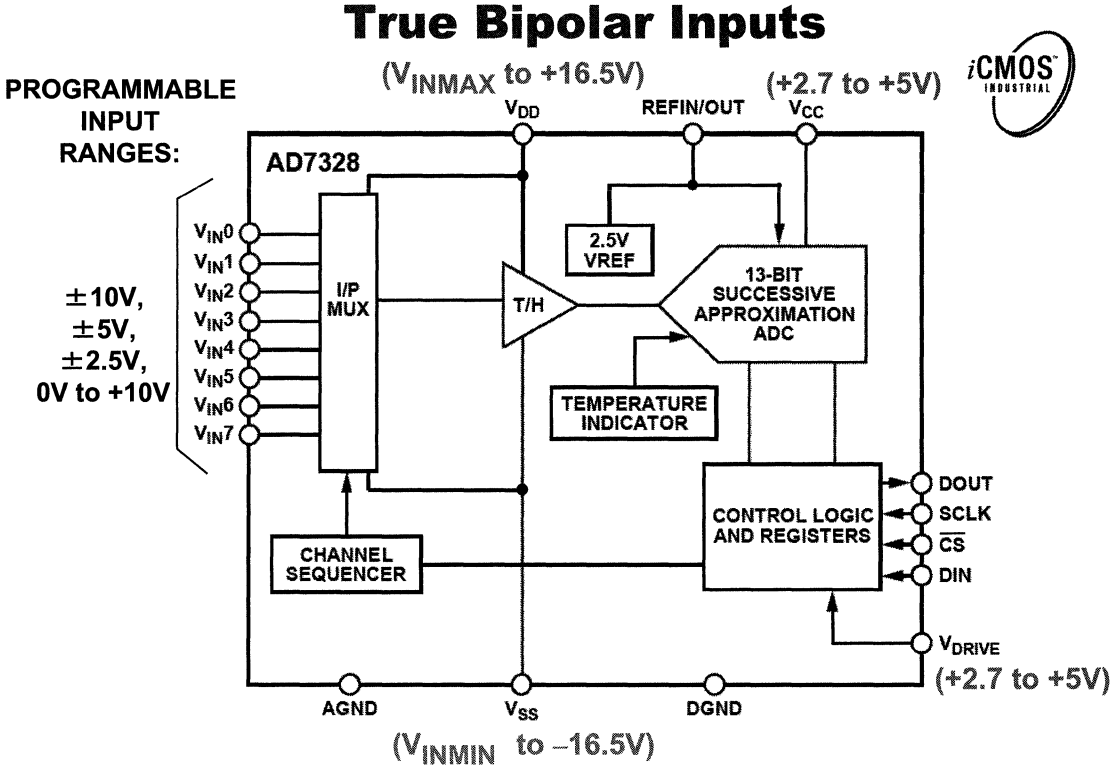

A much better solution available from Analog Devices uses a proprietary Industrial CMOS (*i*CMOS™) process which allows the input circuitry to operate on standard industrial ±15V supplies, while operating the ADC core on the low voltage supply (5V or less). This figure shows the AD7328 13-bit 8-channel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

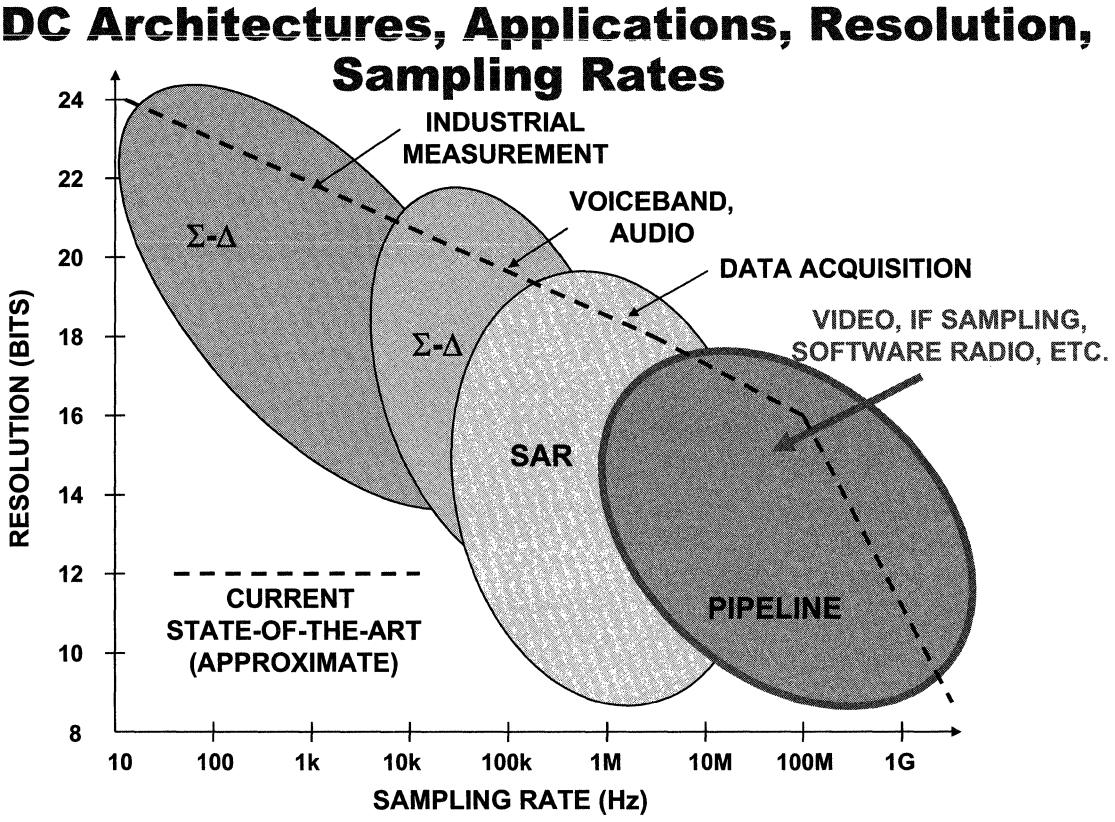

For the purposes of this discussion, "high speed" will be defined as sampling rates greater than about 10MSPS. Included in this application space are various types of instrumentation (digital oscilloscopes, digital spectrum analyzers), medical imaging, radar, IF sampling (including software radio), etc.

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

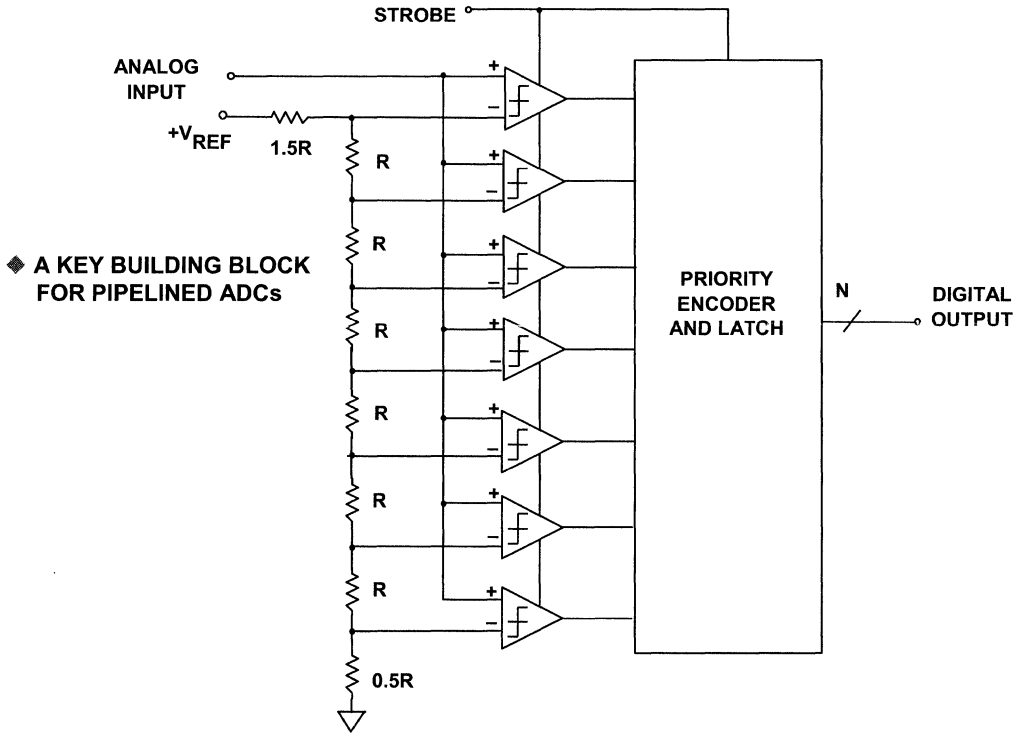

Because of its importance as a building block in subranging pipelined ADCs, one must first understand the basic flash converter.

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

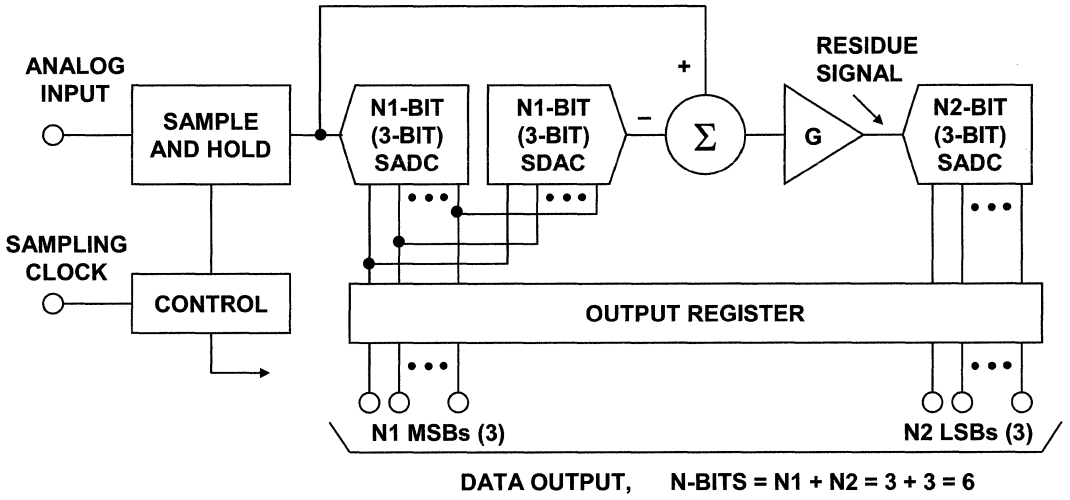

See: R. Staffin and R. Lohman, "Signal Amplitude Quantizer," U.S. Patent 2,869,079, Filed December 19, 1956, Issued January 13, 1959

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

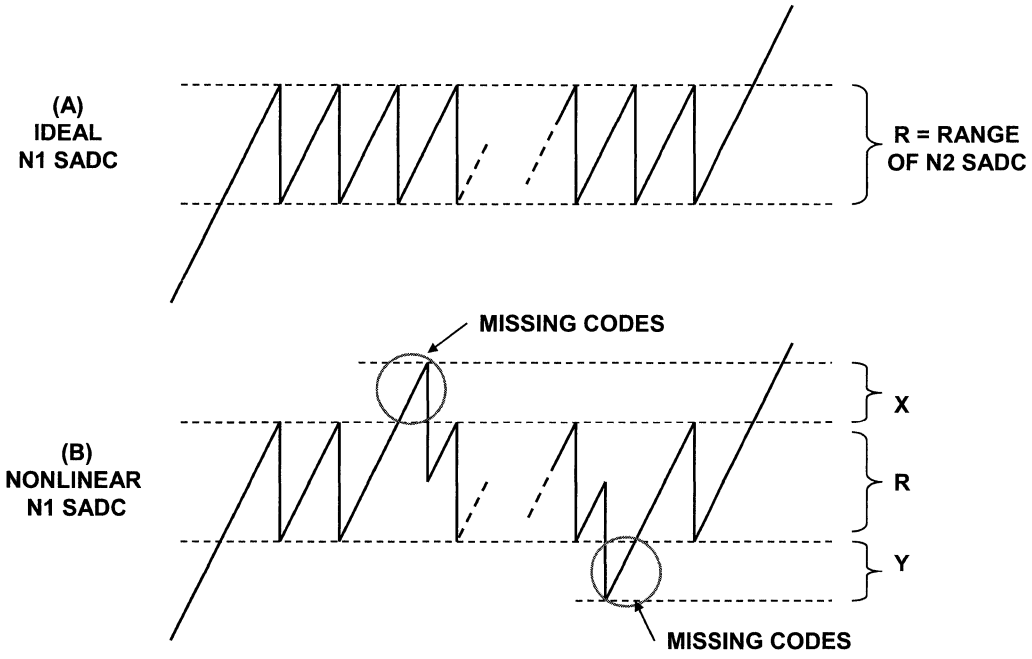

The residue waveform into the second N2 SADC must exactly fill the range of the N2SADC as shown in (A). Otherwise, as shown in (B), there will be nonlinearities in the overall transfer function, and possibly missing codes.

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

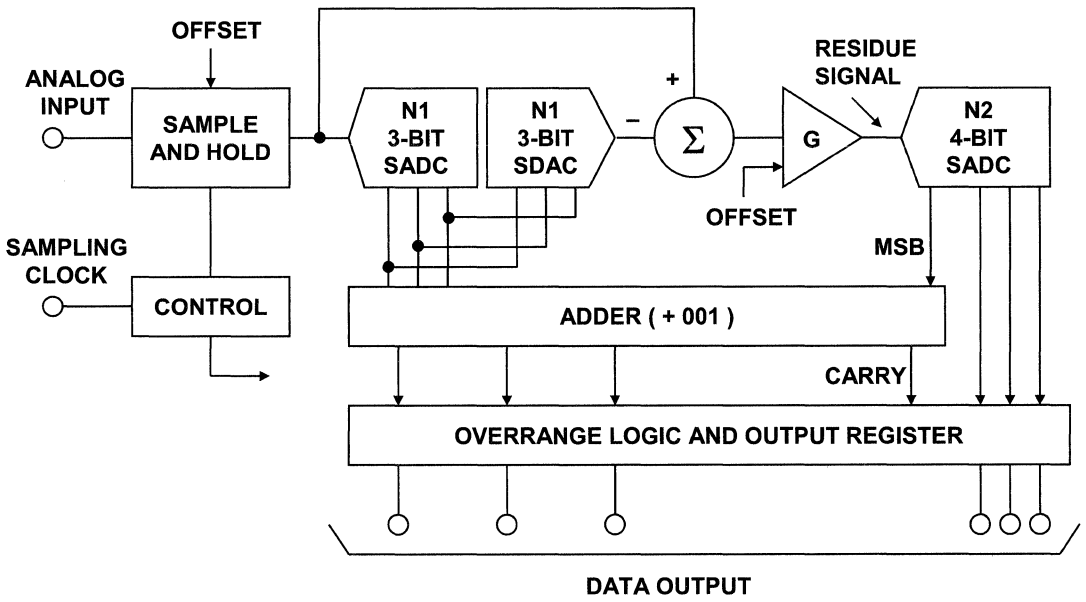

SEE: T. C. Verster, "A Method to Increase the Accuracy of Fast Serial-Parallel Analog-to-Digital Converters," *IEEE Transactions on Electronic Computers*,EC-13, 1964, pp. 471-473

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

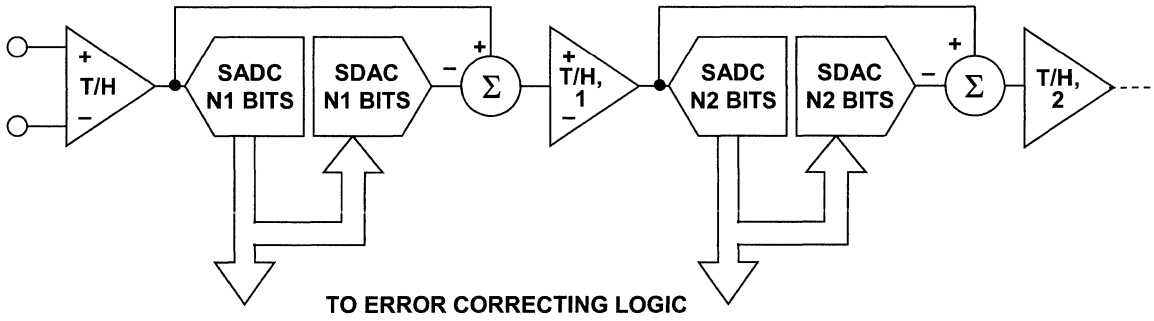

The pipelined architecture shown in this figure is a digitally corrected subranging architecture in which each stage operates on the data for one-half the sampling clock cycle and then passes its residue output to the next stage in the pipeline, prior to the next half cycle. The interstage track-and-hold (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

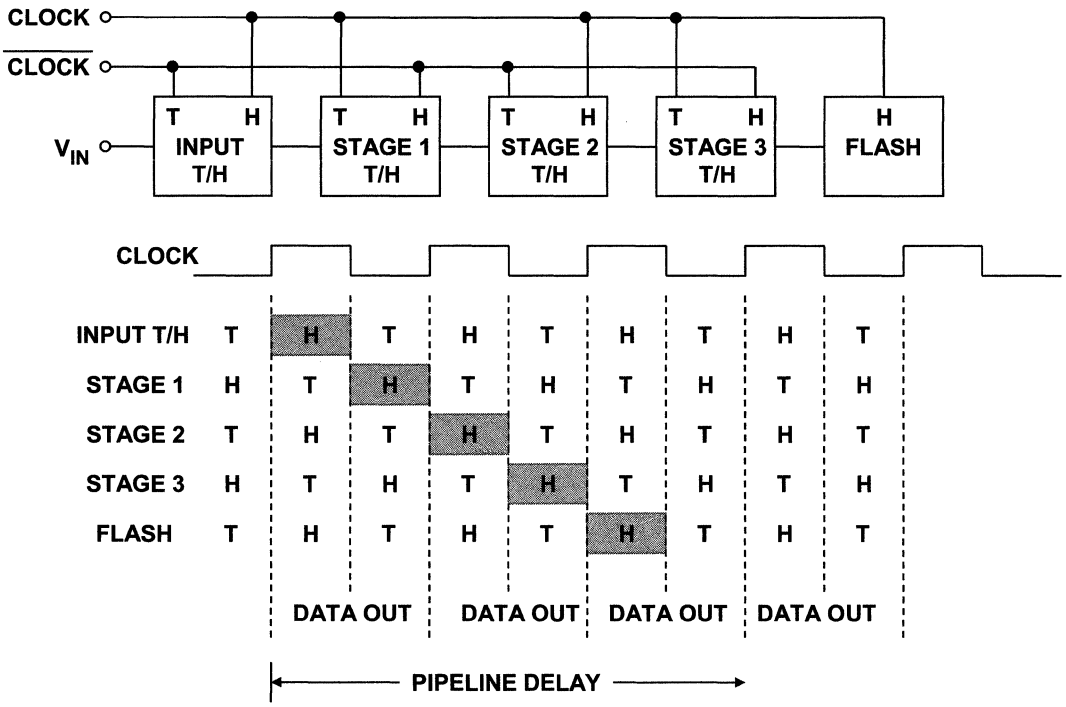

This figure shows a timing diagram of a typical pipelined subranging ADC. Notice that the phases of the clocks to the T/H amplifiers are alternated from stage to stage such that when a particular T/H in the ADC enters the hold mode it holds the sample from the preceding T/H, and the preceding T/H returns 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

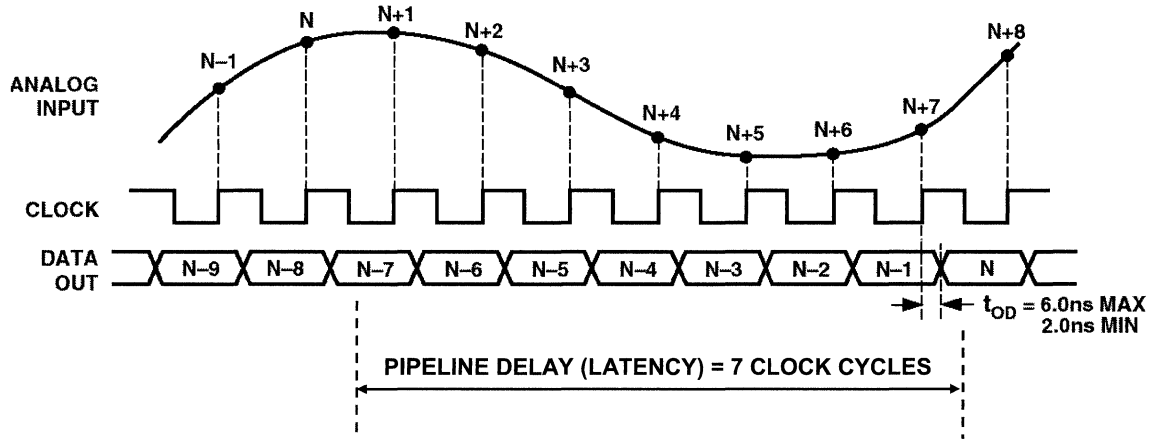

The effects of the "pipeline" delay (sometimes called "latency") in the output data are shown in in this figure for the AD9235 12-bit 65-MSPS ADC where there is a 7-clock cycle pipeline delay.

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

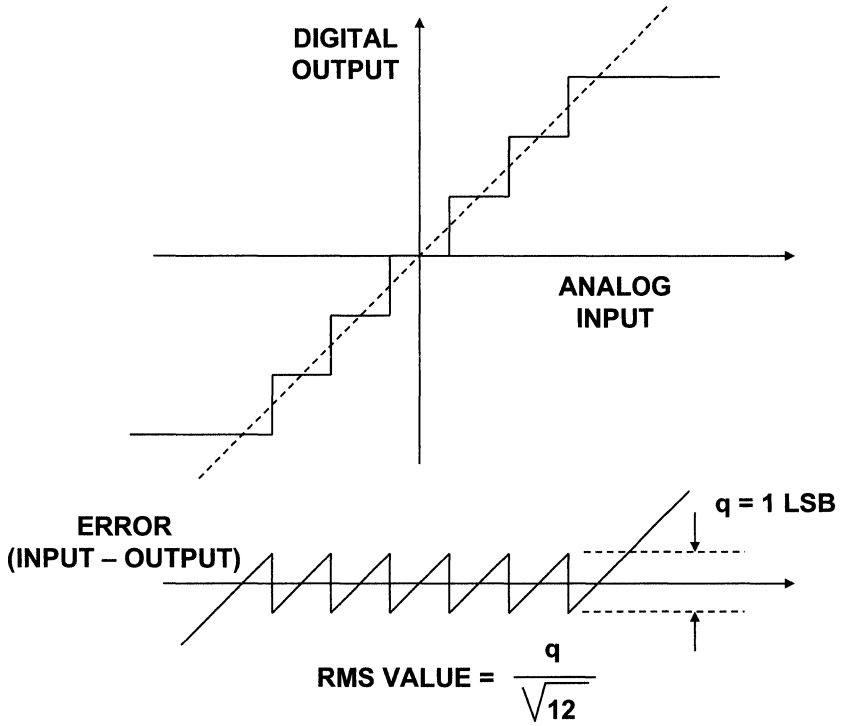

The maximum error an ideal converter makes when digitizing a signal is ±½ LSB as shown in the transfer function of an ideal N-bit ADC. The quantization error for any ac signal which spans more than a few LSBs can be approximated by an uncorrelated sawtooth waveform having a peak-to-peak amplitude of 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

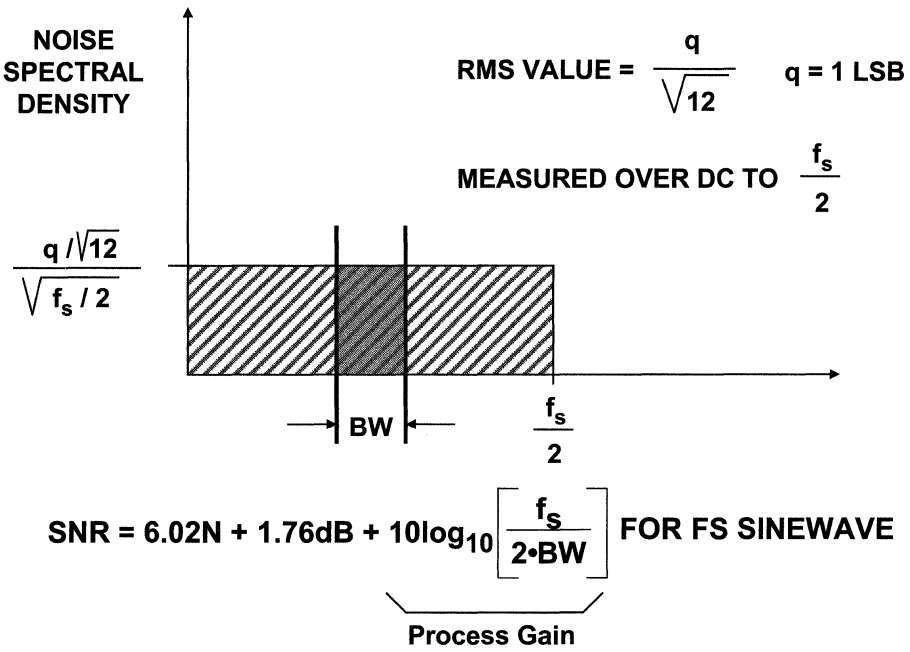

In many applications, the actual signal of interest occupies a smaller bandwidth, BW, which is less than the Nyquist bandwidth as shown in this figure. If digital filtering is used to filter out noise components outside the bandwidth BW, then a correction factor (called *process gain*) must be included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

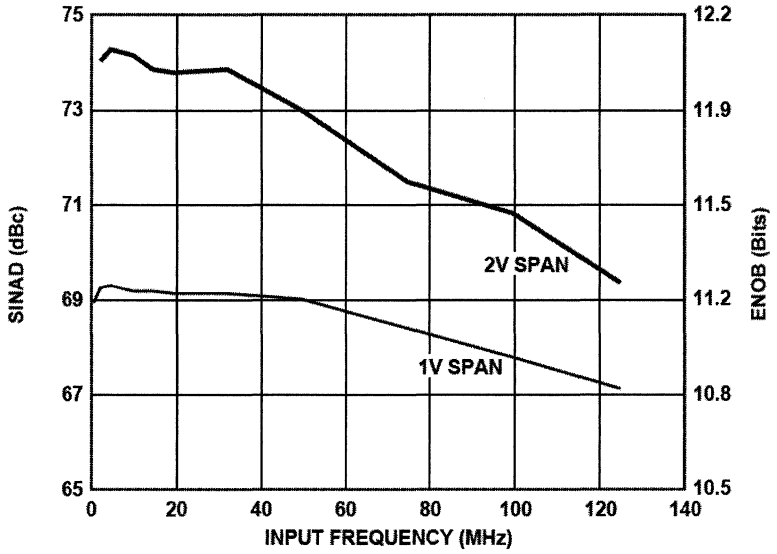

NOTE: AD9244 Full Power Input Bandwidth = 750MHz

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

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

- **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

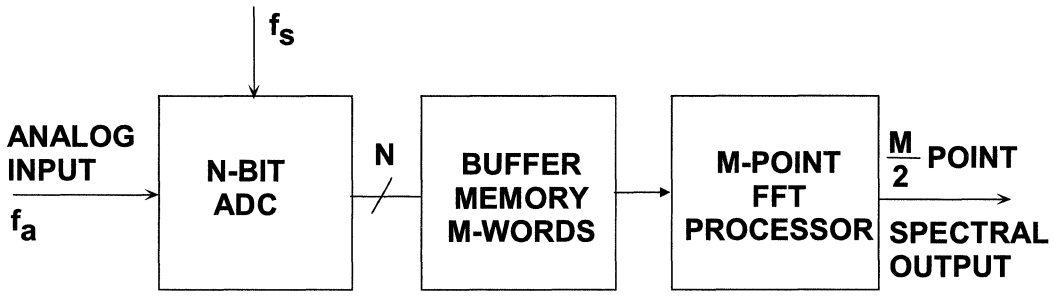

As previously discussed, there are a number of important frequency domain specifications for ADCs. All of them are based on an FFT analysis of the ADC output data using a generalized test setup such as shown in this figure.

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

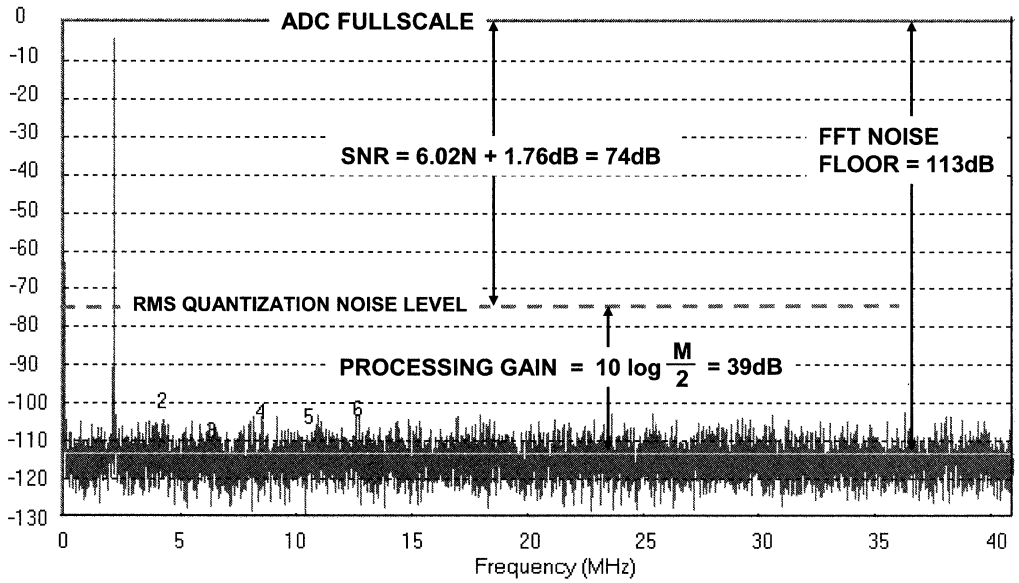

Data Generated Using ADIsimADC®

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

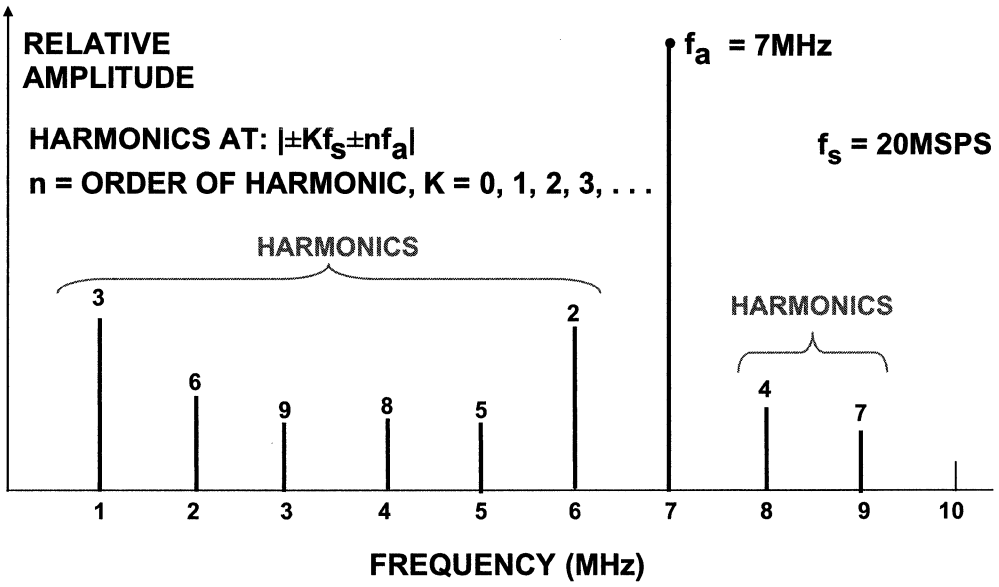

Harmonic distortion is normally specified in dBc (decibels below *carrier*), although in audio applications it may be specified as a percentage. It is the ratio of the rms signal to the rms value of the harmonic in question. Harmonic distortion is generally specified with an input signal near fullscale (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

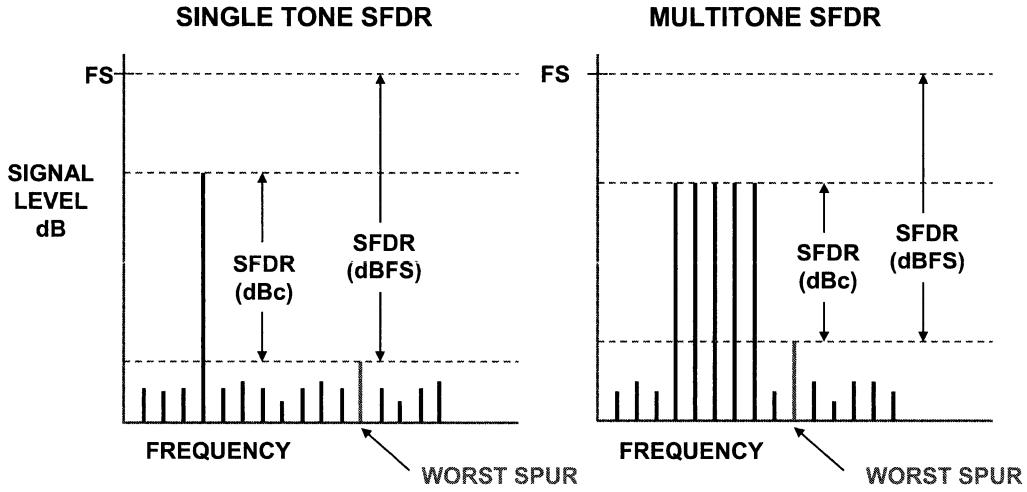

Spurious free dynamic range (SFDR) is the ratio of the rms value of the signal to the rms value of the worst spurious signal regardless of where it falls in the frequency spectrum or what the source. The worst spur may or may not be a harmonic of the original signal. SFDR is an important specification 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

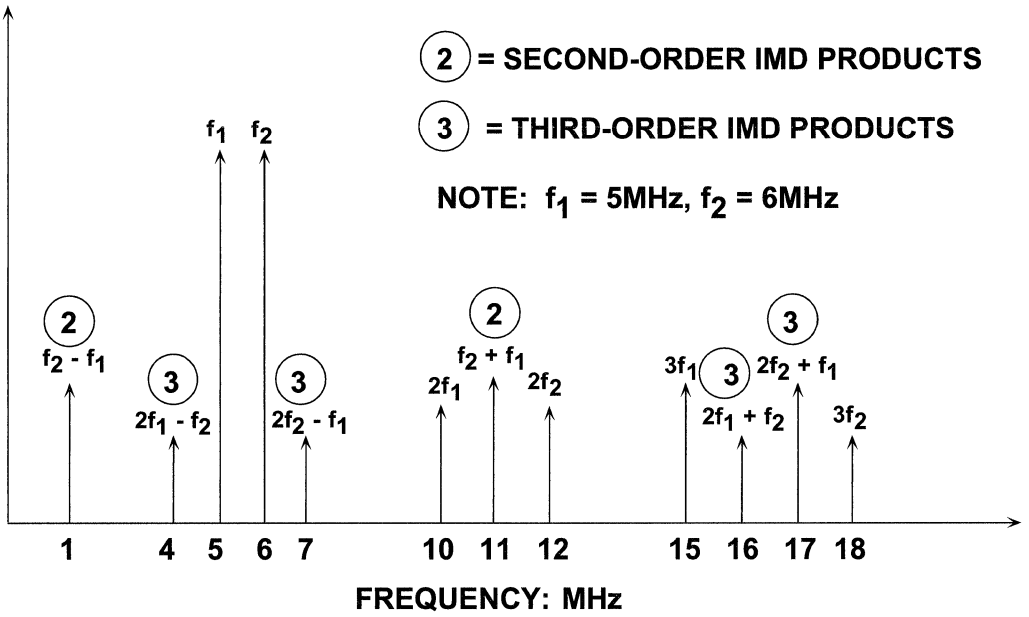

Two tone intermodulation distortion (IMD) is measured by applying two spectrally pure sinewaves to the ADC at frequencies f1 and f2, usually relatively close together. The amplitude of each tone is set slightly more than 6dB below fullscale so that the ADC does not clip when the two tones add in-phase.

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

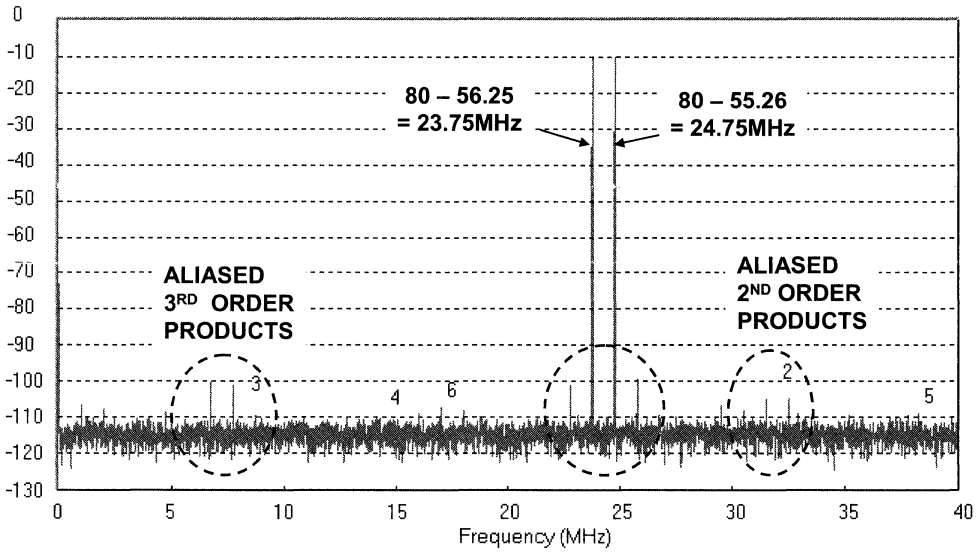

This figure shows the two-tone FFT output for the AD9445 sampling two tones of 55.25MHz and 56.25MHz at a sampling frequency of 80MSPS.

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

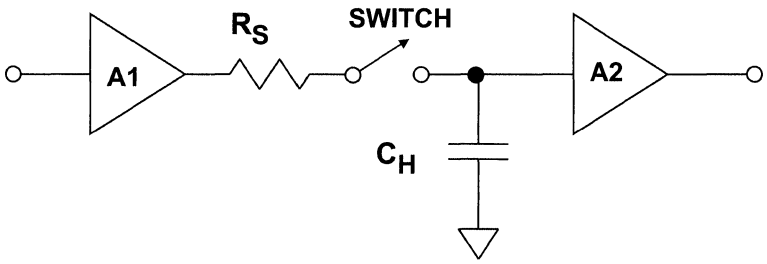

SMALLER C_H:

- **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).

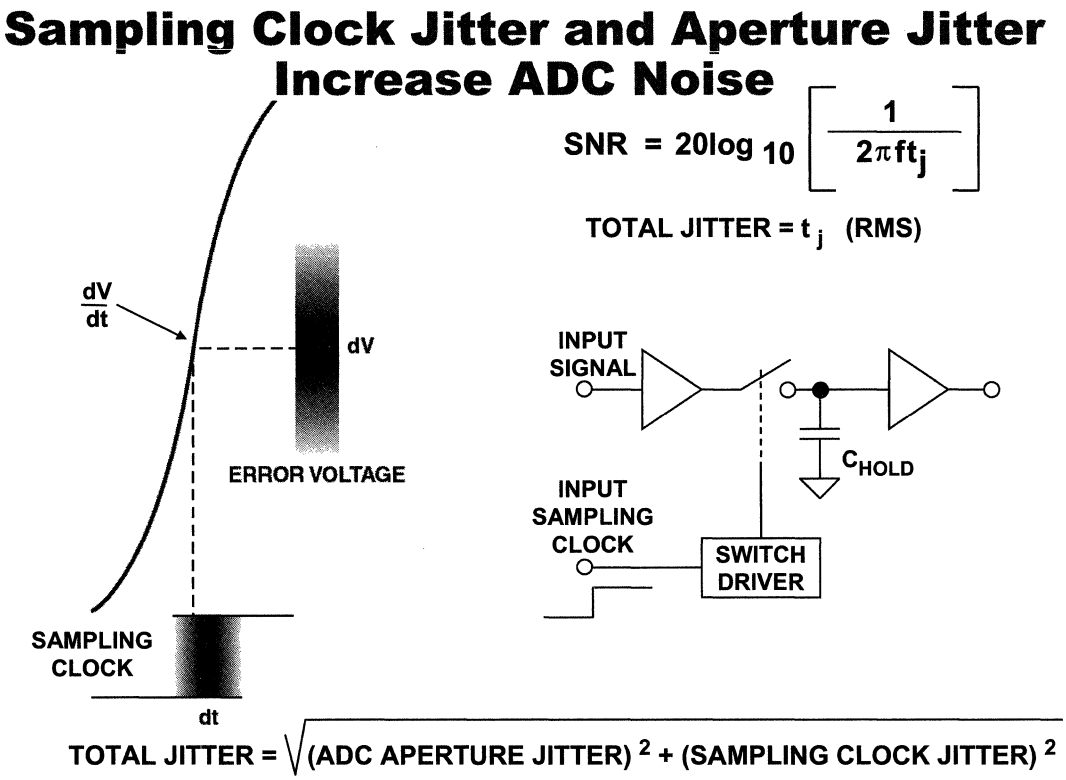

Aperture jitter is the sample-to-sample variation in the time at which the ADC sample is taken as shown in this model of a basic sample-and-hold. The total jitter in the sampling clock produces an error voltage, and the overall SNR of the ADC is limited by the formula as shown:

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

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.

Theoretical SNR and ENOB Due to Jitter vs. Fullscale Sinewave Analog Input Frequency

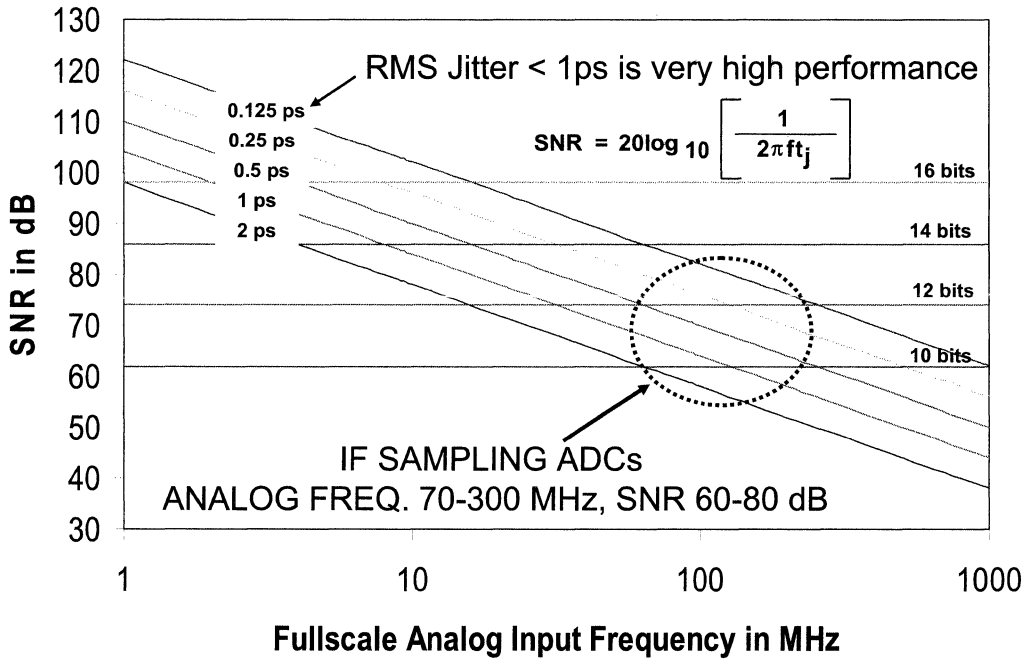

This figure shows the theoretical SNR (left side) due to total jitter versus fullscale analog input frequency. The ADC is assumed to have infinite resolution, and the only noise source is that produced by the jitter.

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

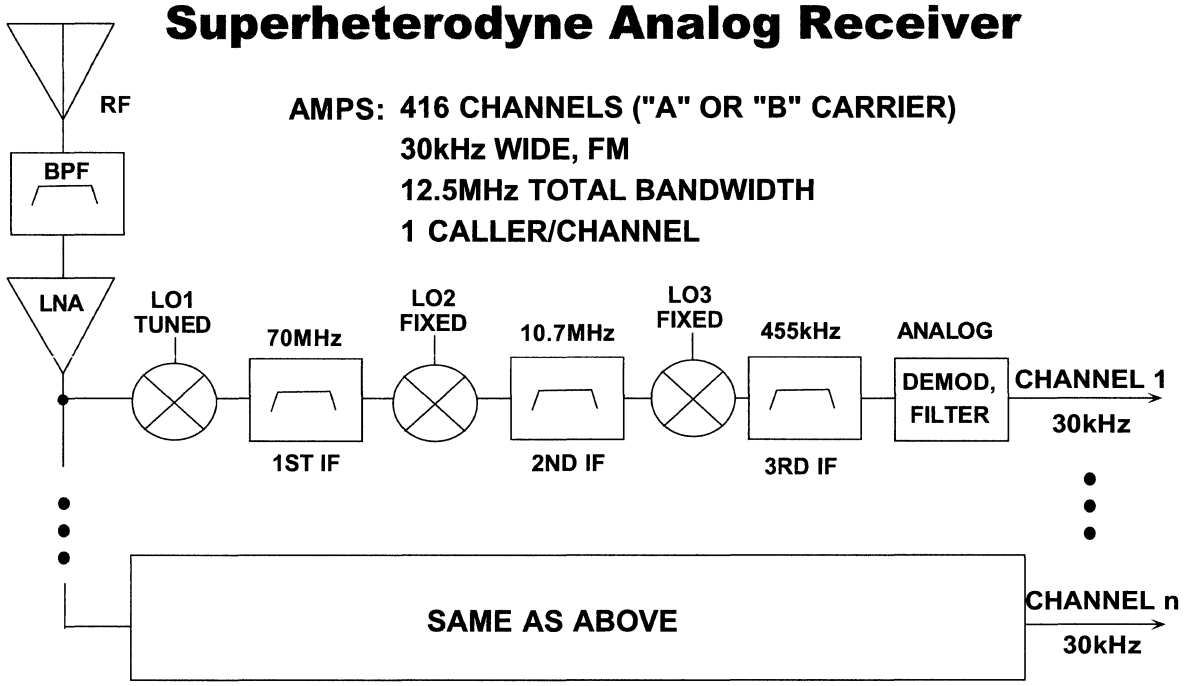

In order to understand the evolution of software radio, consider the analog superheterodyne receiver invented in 1917 by Major Edwin H. Armstrong. The frequencies shown in this figure correspond to the AMPS (Advanced Mobile Phone Service) analog cellular phone system currently being phased out in favor 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

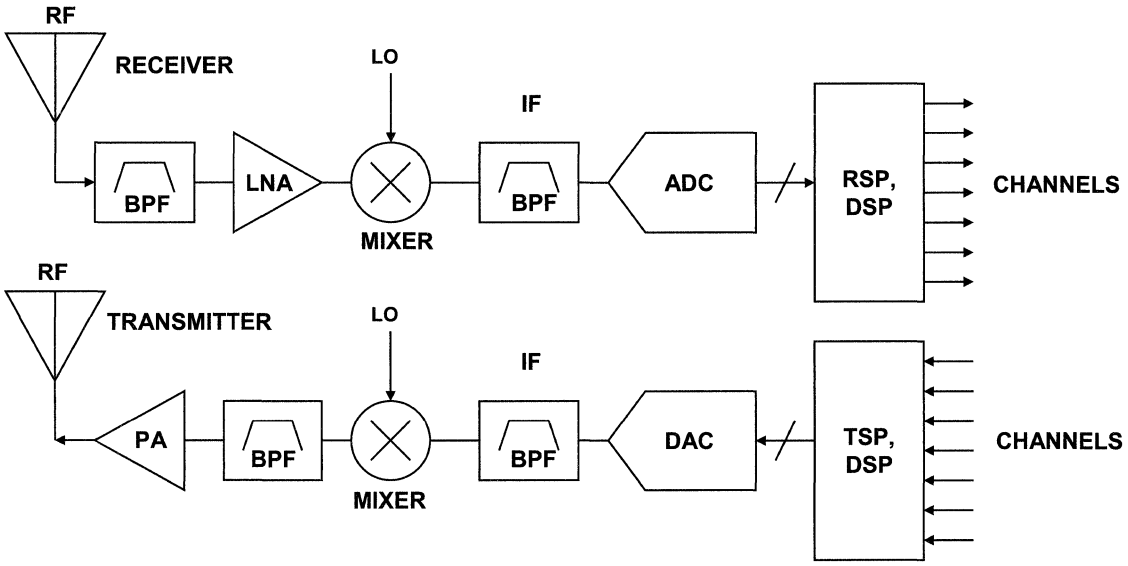

The standard superhetrodyne receiver has several stages of mixing down to baseband as previously described. There may be multiple parallel stages required for multichannel systems—one for each channel if all channels must be received simultaneously.

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.

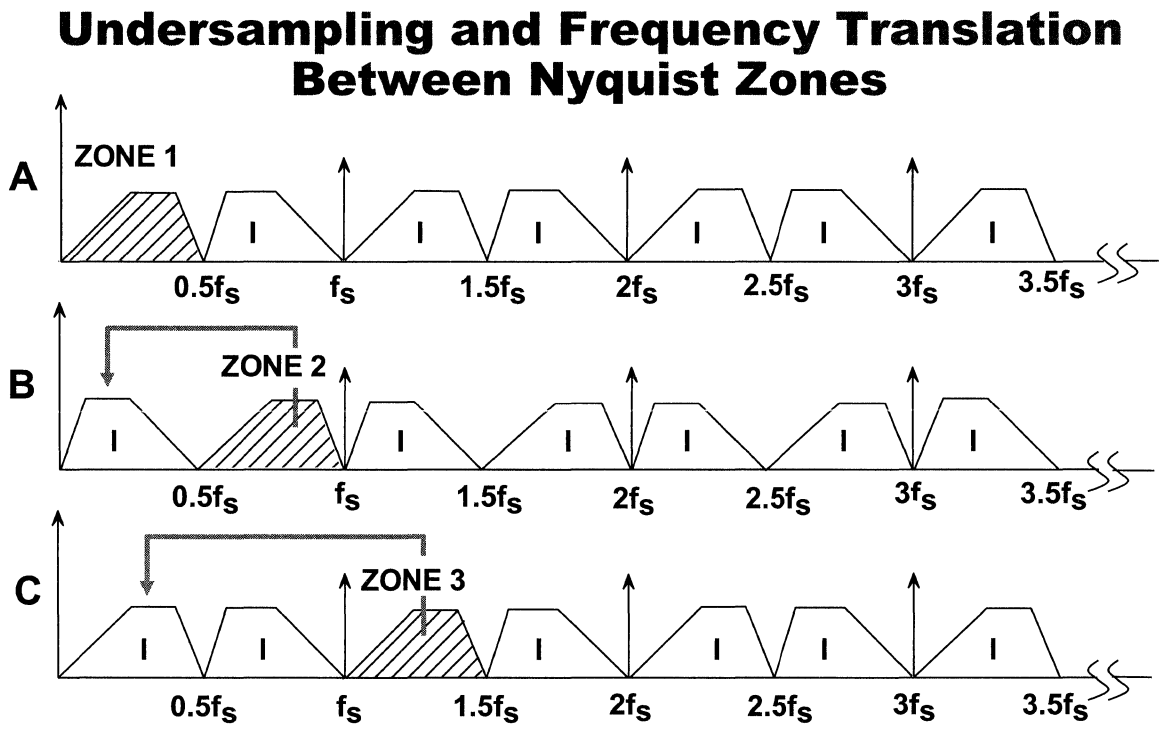

This figure illustrates the principles used in IF sampling, or direct IF-to-digital conversion. The signal of interest is bandlimited to a single Nyquist zone (not necessarily the first zone), and its image will always show up in the first Nyquist zone because of aliasing associated with the basic sampling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

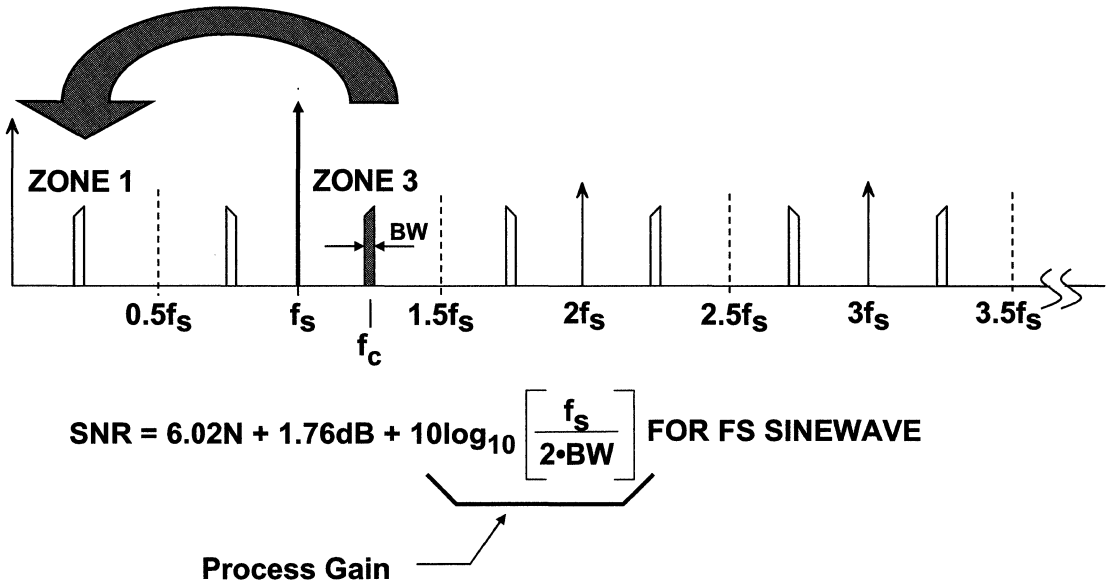

This figure shows a signal of bandwidth BW which lies in the 3rd Nyquist zone sampled at frequency f_S.

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

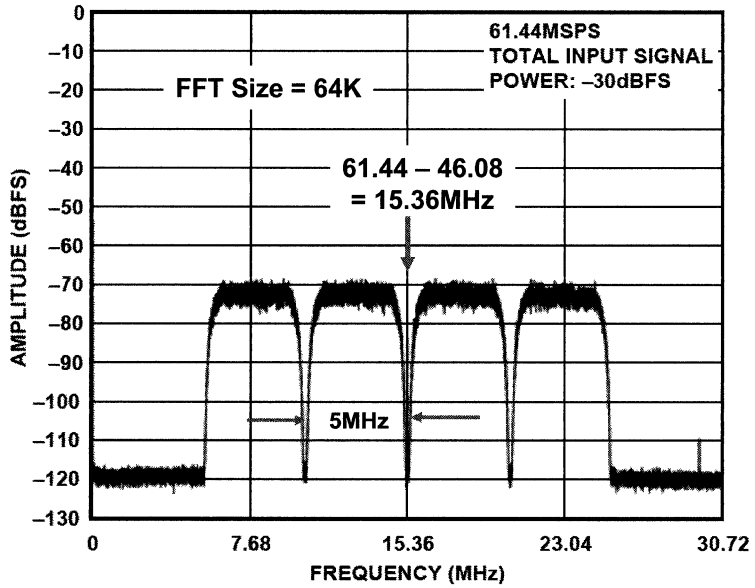

This figure shows how the direct IF to digital conversion process works for four wideband CDMA carriers which occupy a 20MHz bandwidth centered at 46.08MHz and sampled at 61.44MSPS.

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

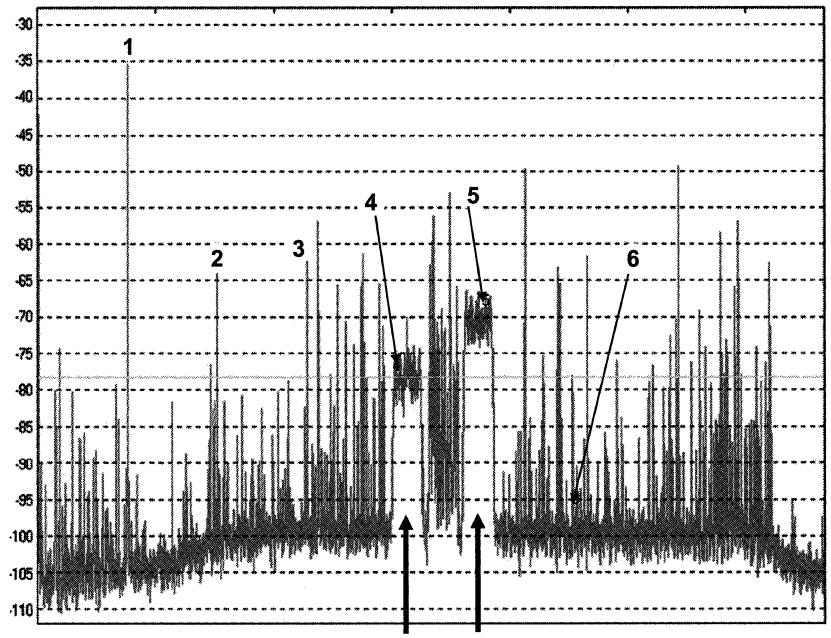

CDMA2000 CHANNELS, BW = 1.25MHz EACH, SAMPLING RATE = 61.44MSPS

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.

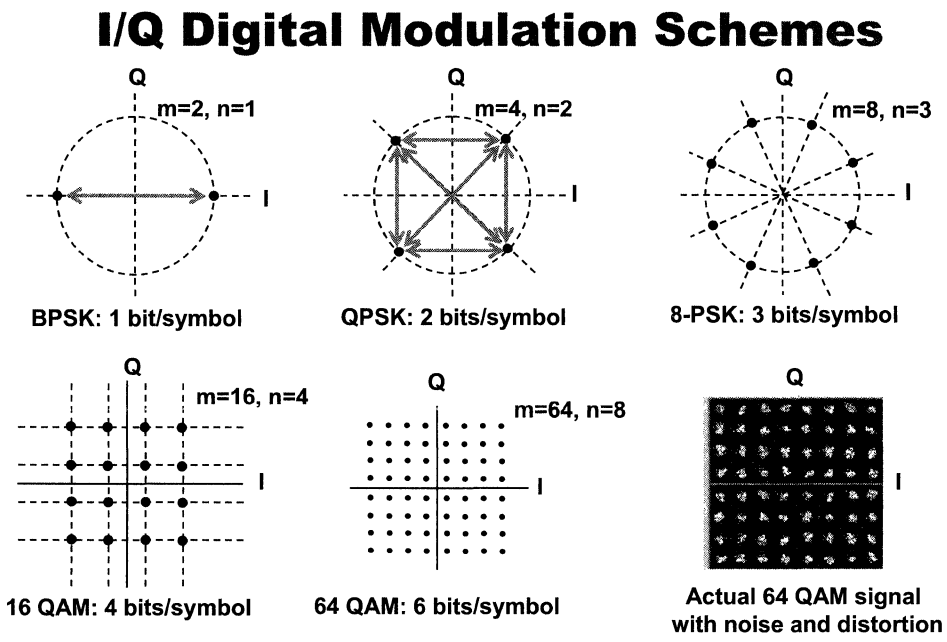

- Higher Order Modulation Schemes → Higher Data Rate.
- But Symbols are closer together → Requires higher Signal-to-Noise Ratio for demodulation
- Increasing "Symbol Rate" increases data rate but widens spectrum

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)

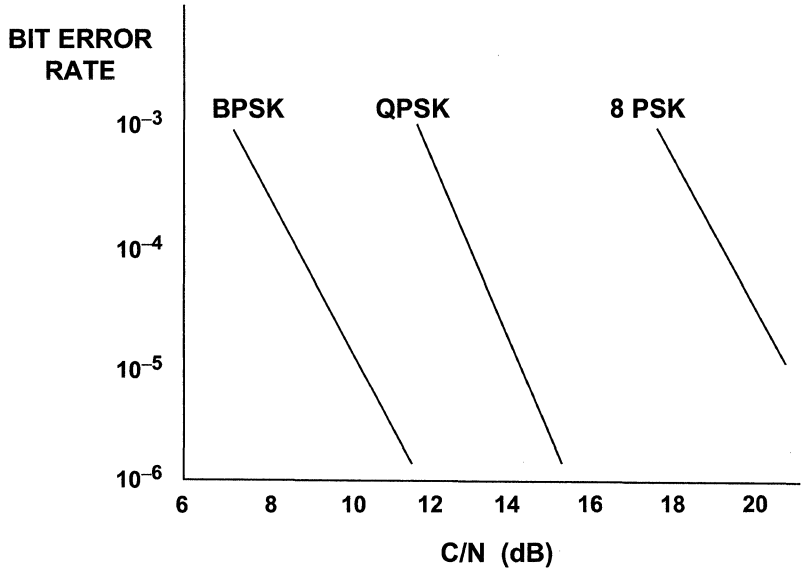

The figure shows the bit error rate as a function of the carrier-to-noise ratio (C/N) for three popular modulation schemes: BPSK, QPSK, and 8 PSK. Note that the more complex schemes (QPSK and 8 PSK) require a higher C/N ratio to achieve the same bit error rate as the simple BPSK.

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

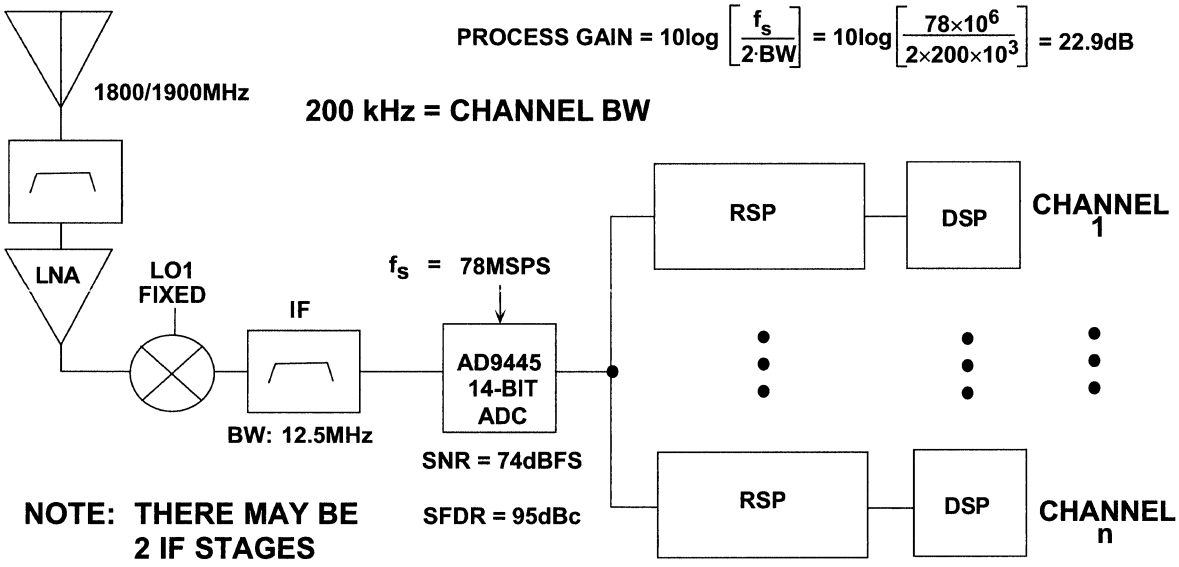

The noise analysis for a system depends upon the air standard and modulation method. For the GSM 1800/1900 MHz wideband receiver shown here, the AD9445 14-bit, 125MSPS ADC has 74dB SNR and 95dBc SFDR which is sufficient to meet the system bit error rate requirement (see next figure).

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

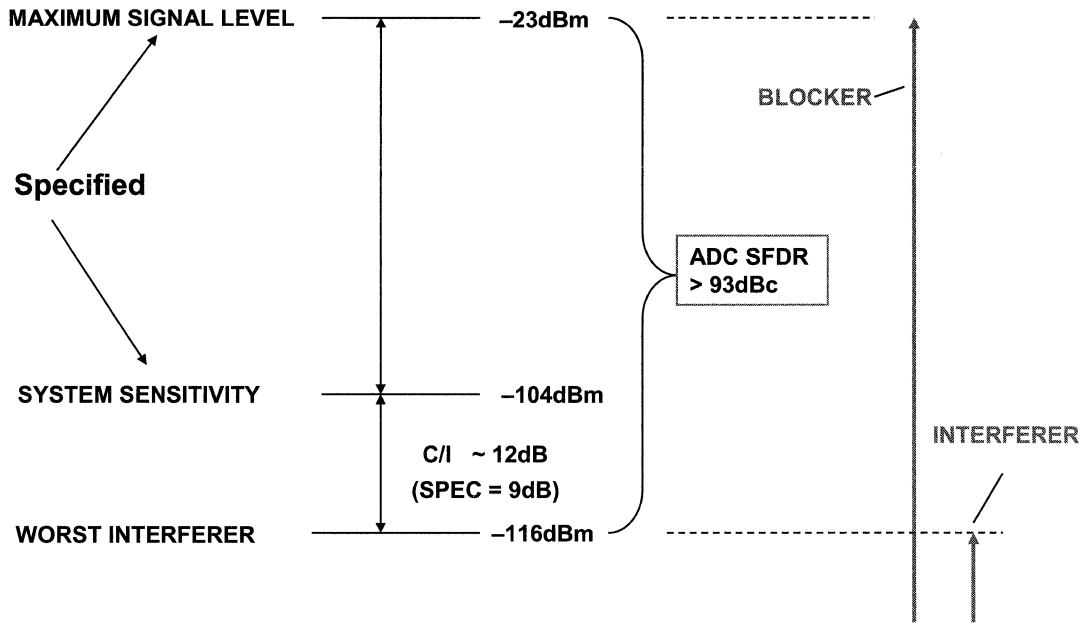

* NOTE: MAXIMUM SIGNAL LEVEL FOR GSM 1800MHz/1900MHz AND PCS SPEC = –23dBm

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)

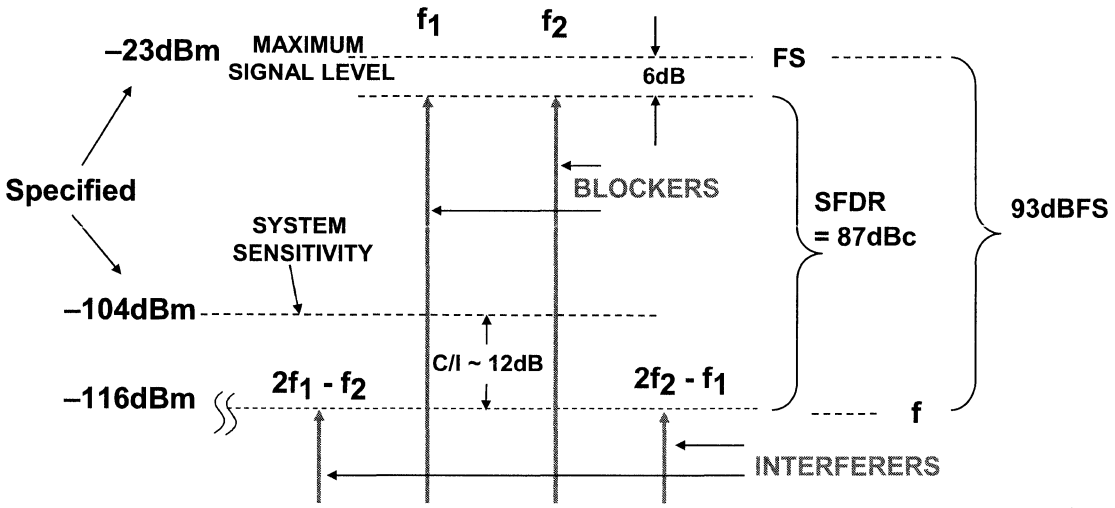

NOTE: MAXIMUM SIGNAL LEVEL FOR GSM 1800/1900 SPEC = –23dBm

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

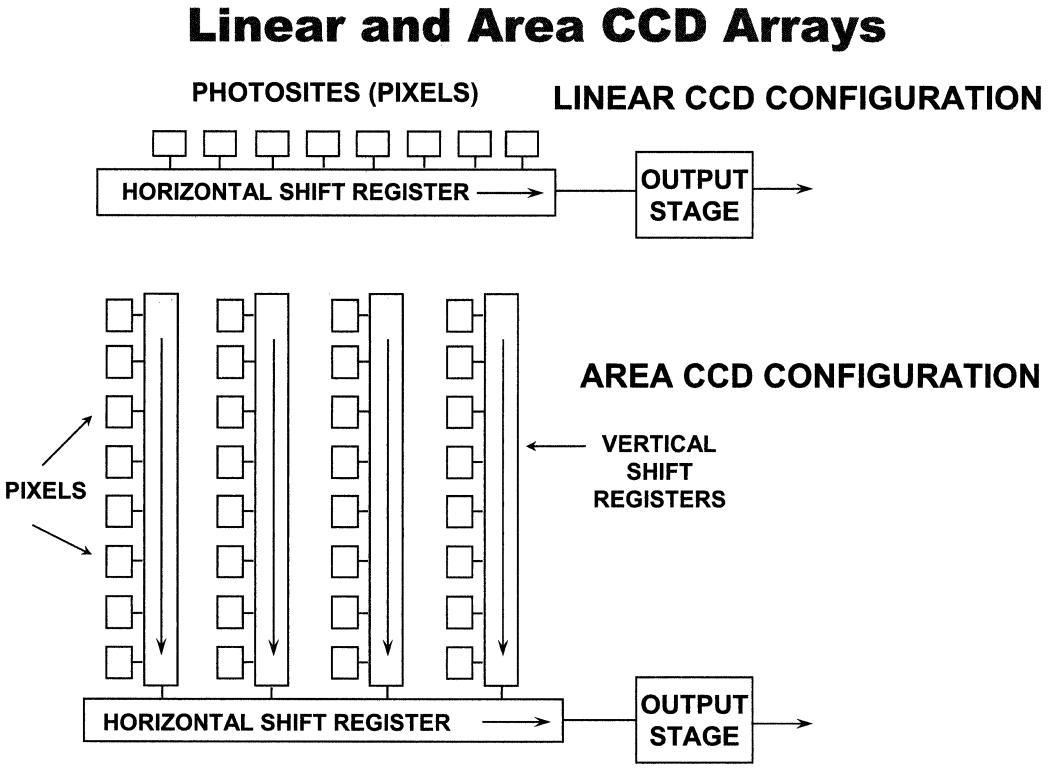

The *charge-coupled-device* (CCD) and *contact-image-sensor* (CIS) are widely used in consumer imaging systems such as scanners and digital cameras. The imaging sensor (CCD, CMOS, or CIS) is exposed to the image or picture much like film is exposed in a camera. After exposure, the output of the sensor 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

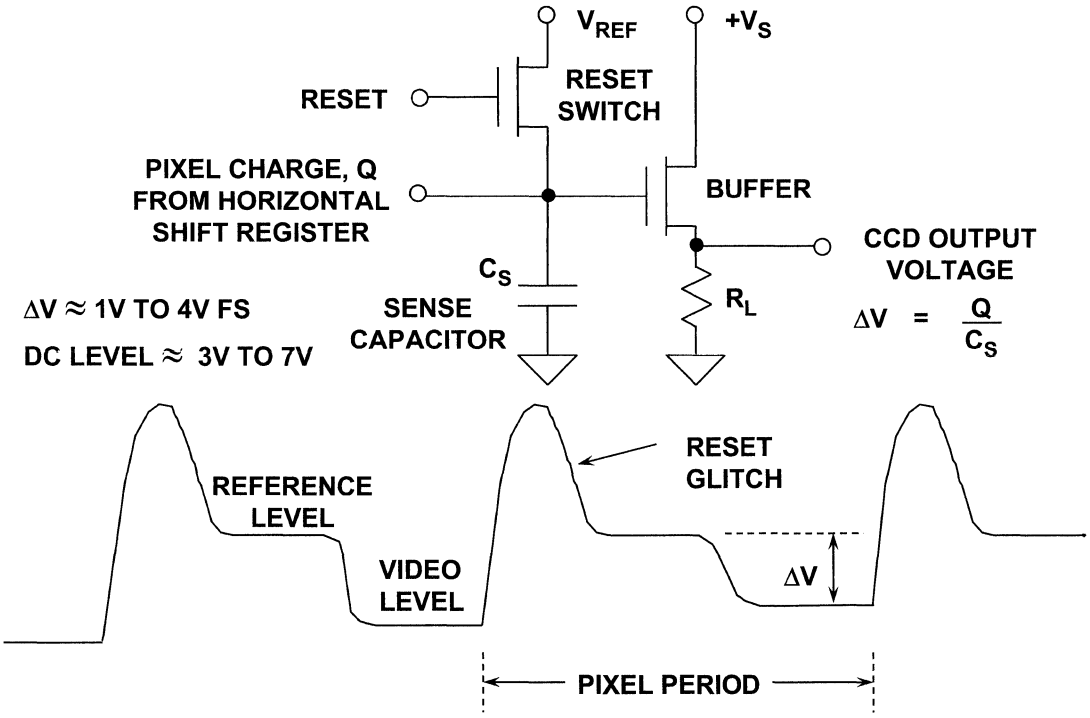

A typical CCD output stage is shown in this figure along with the associated voltage waveforms. The output stage of the CCD converts the charge of each pixel to a voltage via the sense capacitor, C_S. At the start of each pixel period, the voltage on C_S is reset to the reference level, V_{REF}, causing 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)

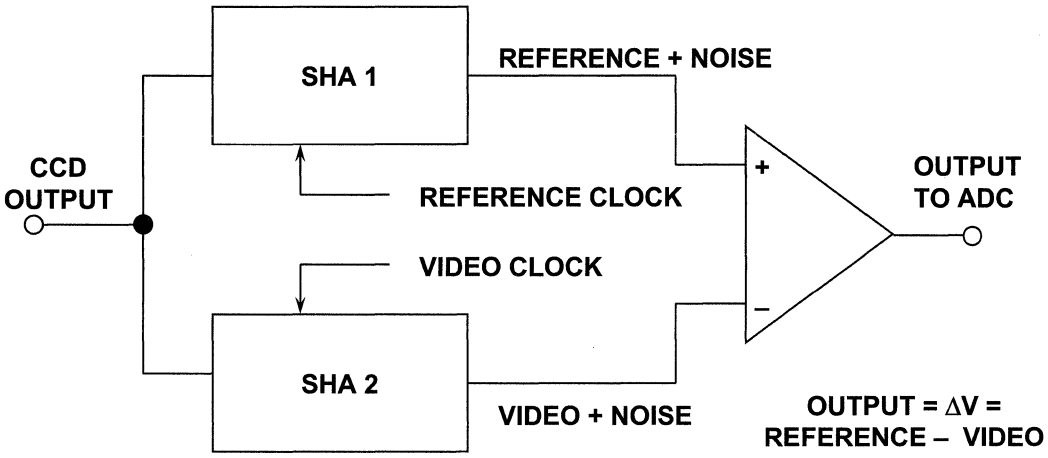

- Typical ADC Resolution: 10 to 14 bits
- Typical ADC Sampling Rates: 12 to 65MSPS

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

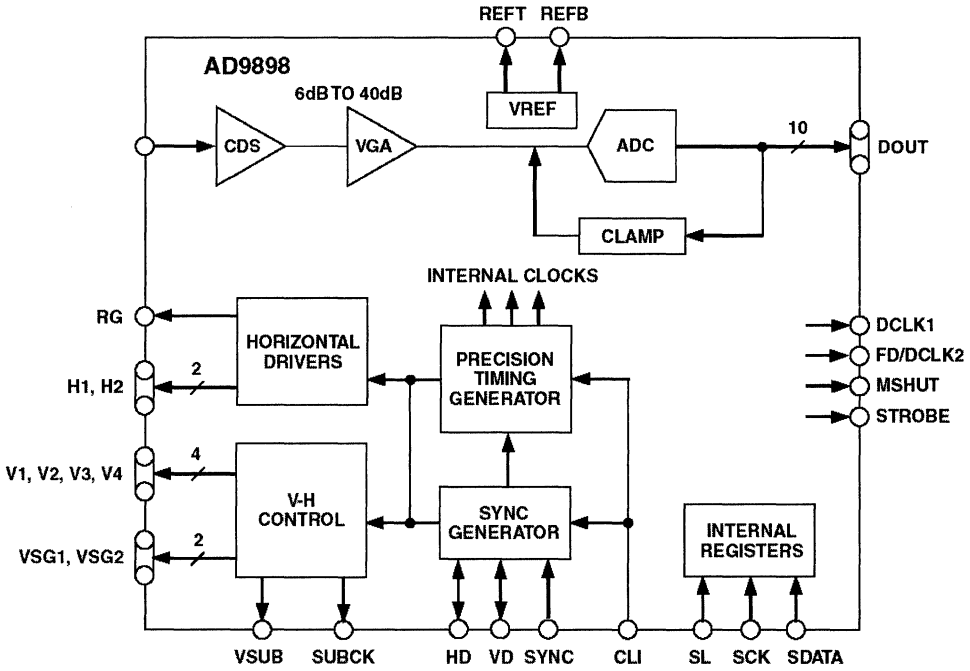

The AD9898 is a highly integrated CCD signal processor for digital still camera and digital video camera applications. A simplified block diagram is shown in this figure. It includes a complete analog front end with 10-bit A/D conversion combined with a full function programmable timing generator. A 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

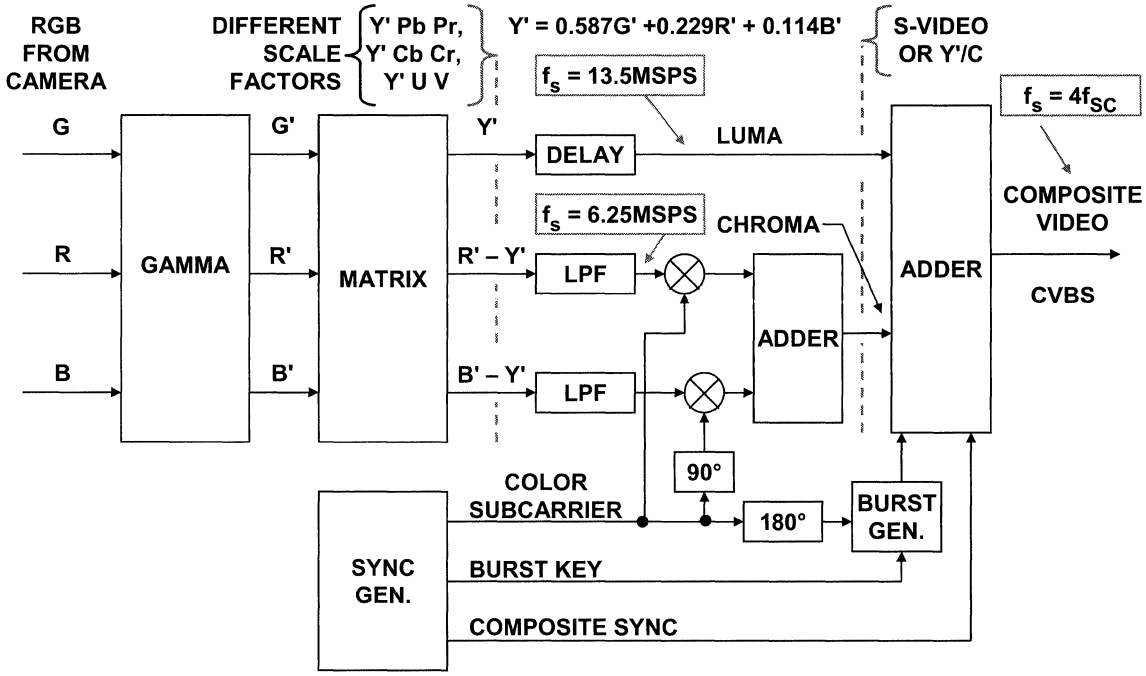

Even before the finalization of the $4f_{SC}$ composite digital standard, work was progressing on digital *component* systems, which offer numerous advantages over the composite digital systems. To understand the differences and advantages, this figure shows a generalized block diagram of how the composite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

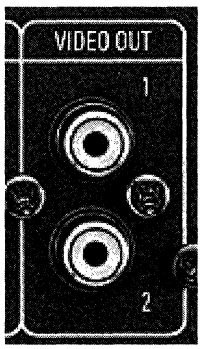

CVBS
Composite Video
with Burst and Sync

2 Outputs

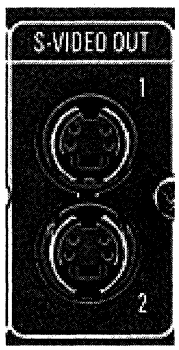

S-Video
Y'/C

2 Outputs

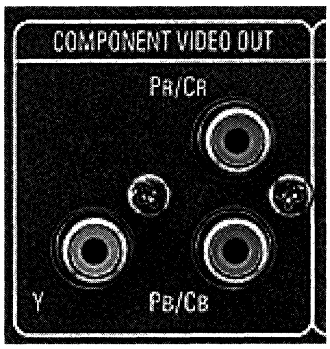

Component Video
Y'PrPb or
Y'CrCb

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

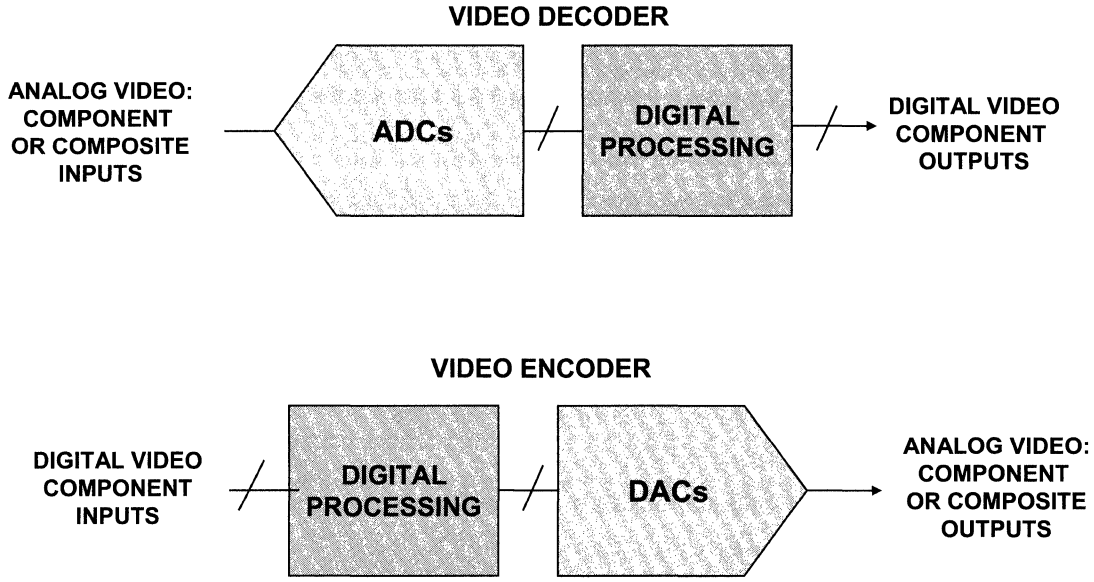

A video *decoder* converts *analog* composite video (CVBS), S-Video (Y'/C), Y'UV, or Y'PbPr signals into a digital video steam in the form of a Y'CbCr digital stream per ITU-R BT.656 4:2:2 component video compatible with NTSC, PAL B/D/G/H/I/ PAL M, PAL N, or others. An ADC function is implicit in the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

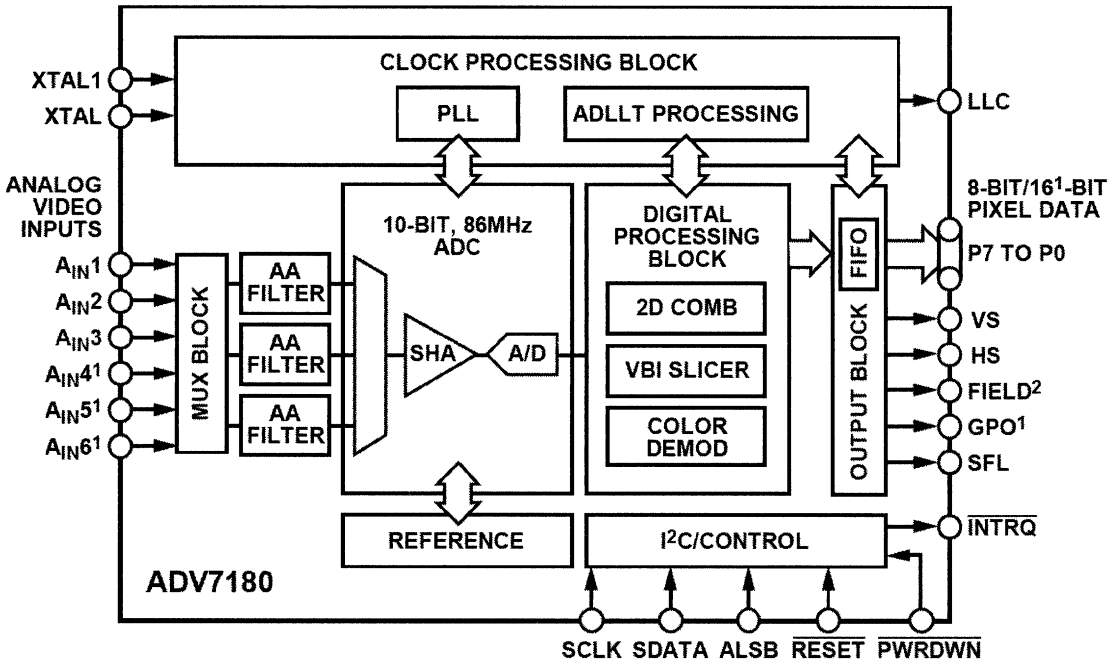

The ADV7180 automatically detects and converts SDTV (standard definition TV) analog baseband television signals compatible with worldwide NTSC, PAL, and SECAM standards into 4:2:2 component video data compatible with the 8-bit ITU-R BT.656 interface standard. The simple digital output interface connects 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

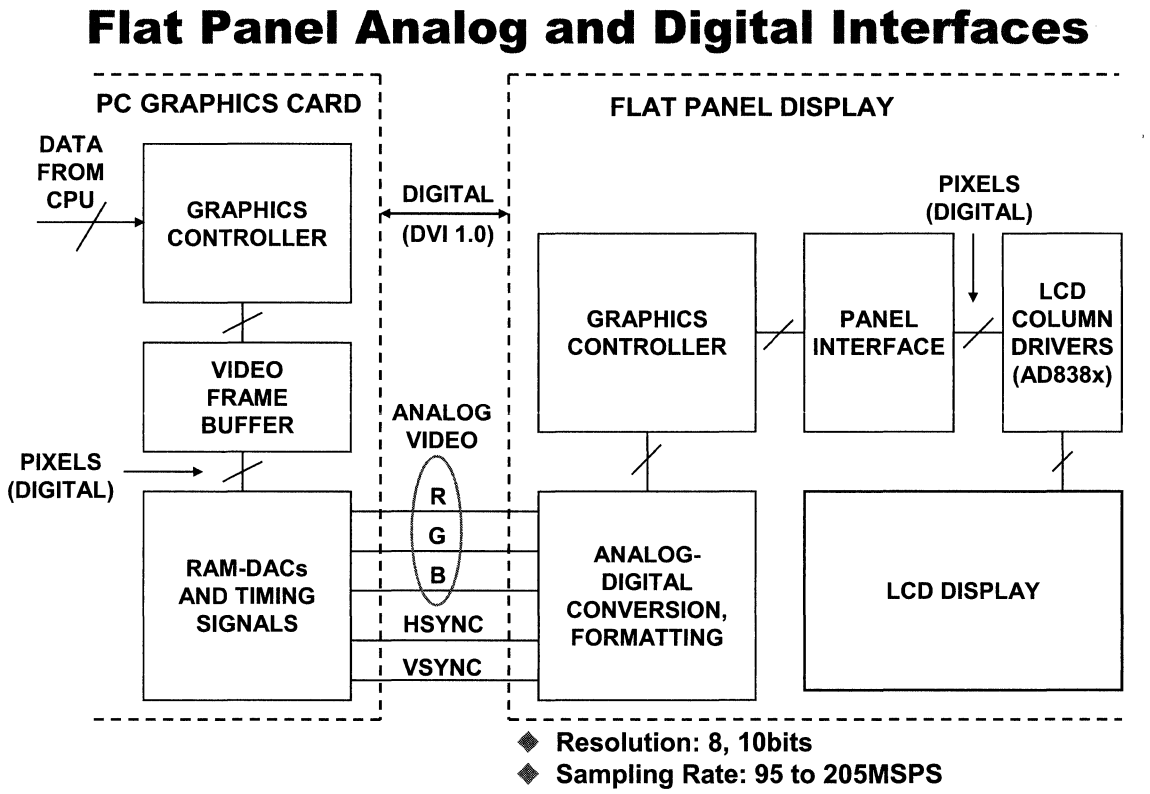

The popularity of flat panel LCD-based displays has steadily increased over the last few years, and they are rapidly replacing CRT-based monitors in desktop computer systems. In addition, LCD projectors have virtually replaced 35mm slide and overhead projectors as a means of delivering presentation material.

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

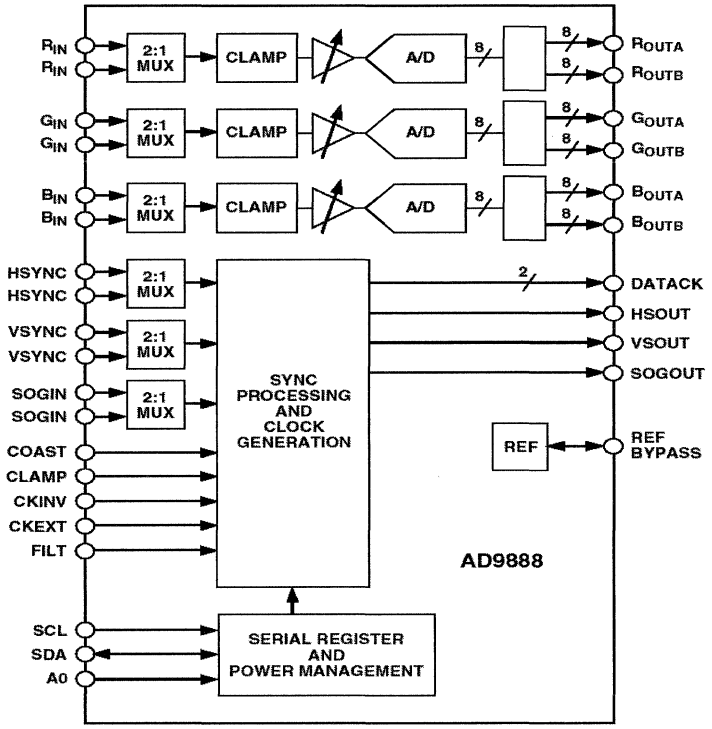

The AD9888 is an example of a highly integrated flat panel interface product. It is a complete 8-bit, 205MSPS monolithic analog interface optimized for capturing RGB graphics signals from personal computers and workstations. Its 205MSPS sampling rate capability and full-power analog bandwidth of 500MHz 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

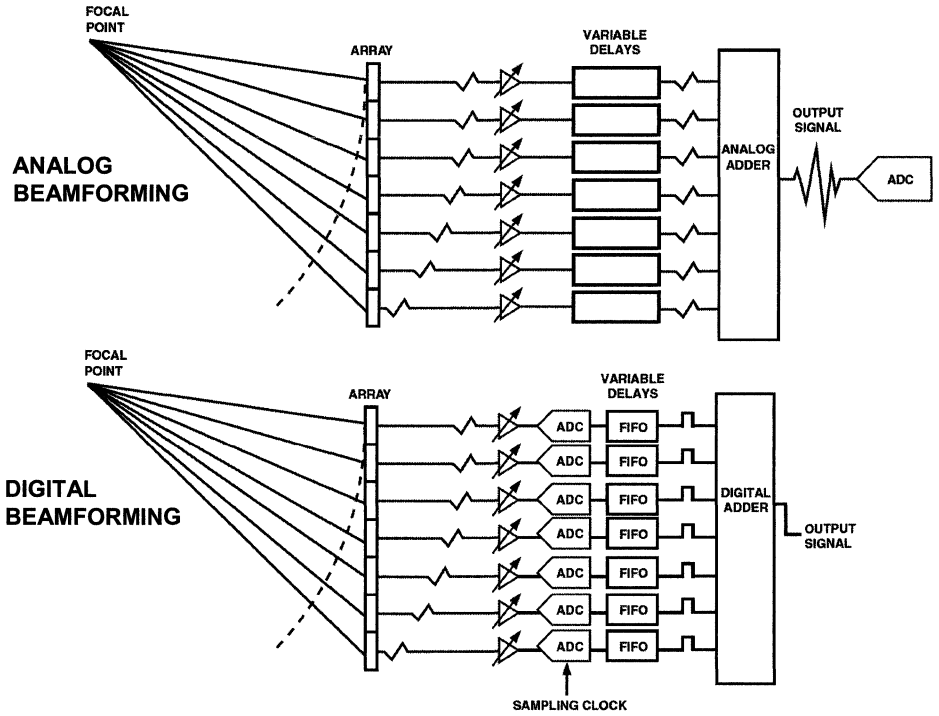

Medical ultrasound machines are among the most sophisticated signal processing machines in widespread use today. As in any complex machine, there are many trade-offs in implementation due to performance requirements, physics, and cost. Some system-level understanding is necessary to fully appreciate 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

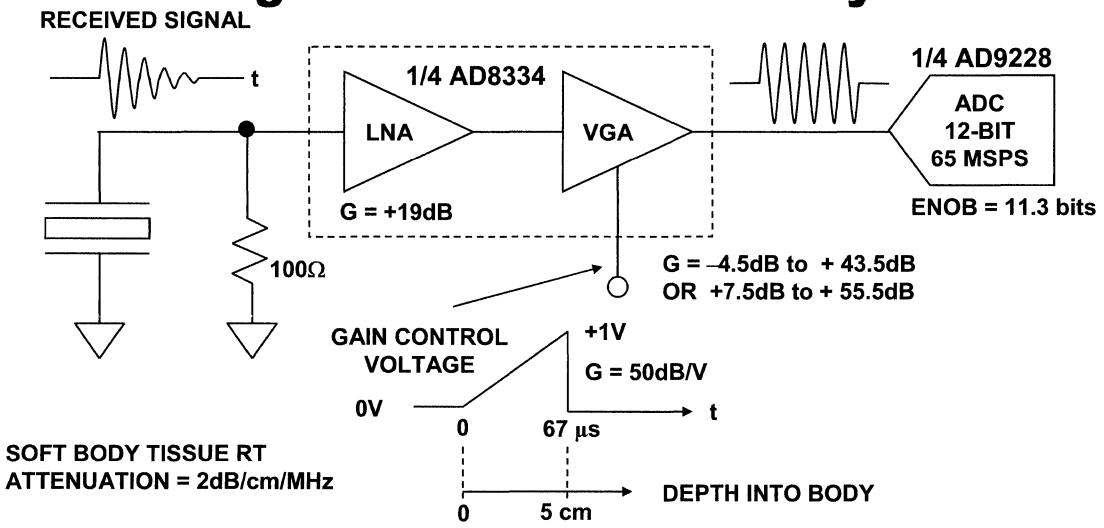

AT 5 MHz:

- 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

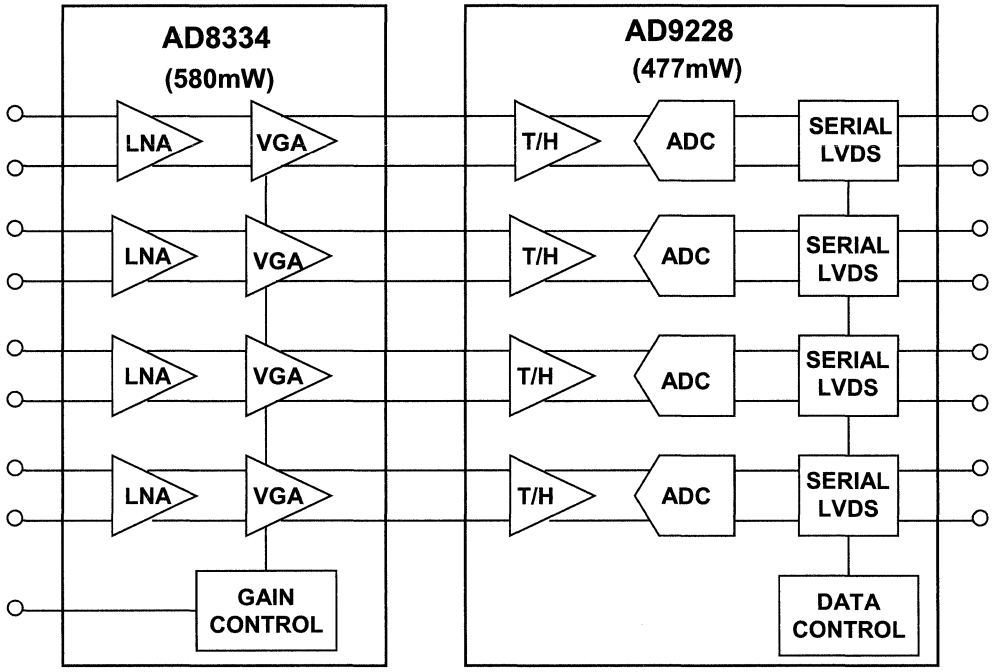

This figure shows a block diagram of how the quad AD8334 amplifier and the quad AD9228 ADC form an ultrasound system dissipating only 264mW/channel. The AD8331/AD8332/AD8334 are single-, dual-, and quad-channel ultralow noise, linear-in-dB, variable gain amplifiers (VGAs). Optimized for ultrasound systems, 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

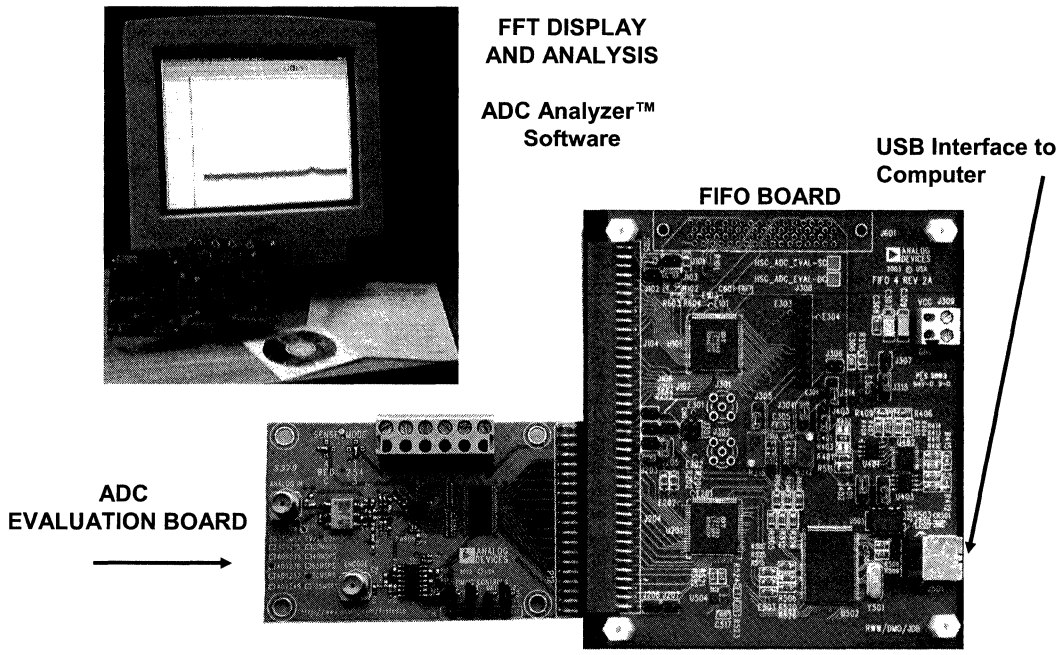

Analog Devices offers complete hardware and software tools for evaluating data converters. In order to make FFT measurements, a suitable memory must capture the ADC output data. The FIFO Board shown in the figure is designed to mate with the product-specific evaluation board and store the ADC data. The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

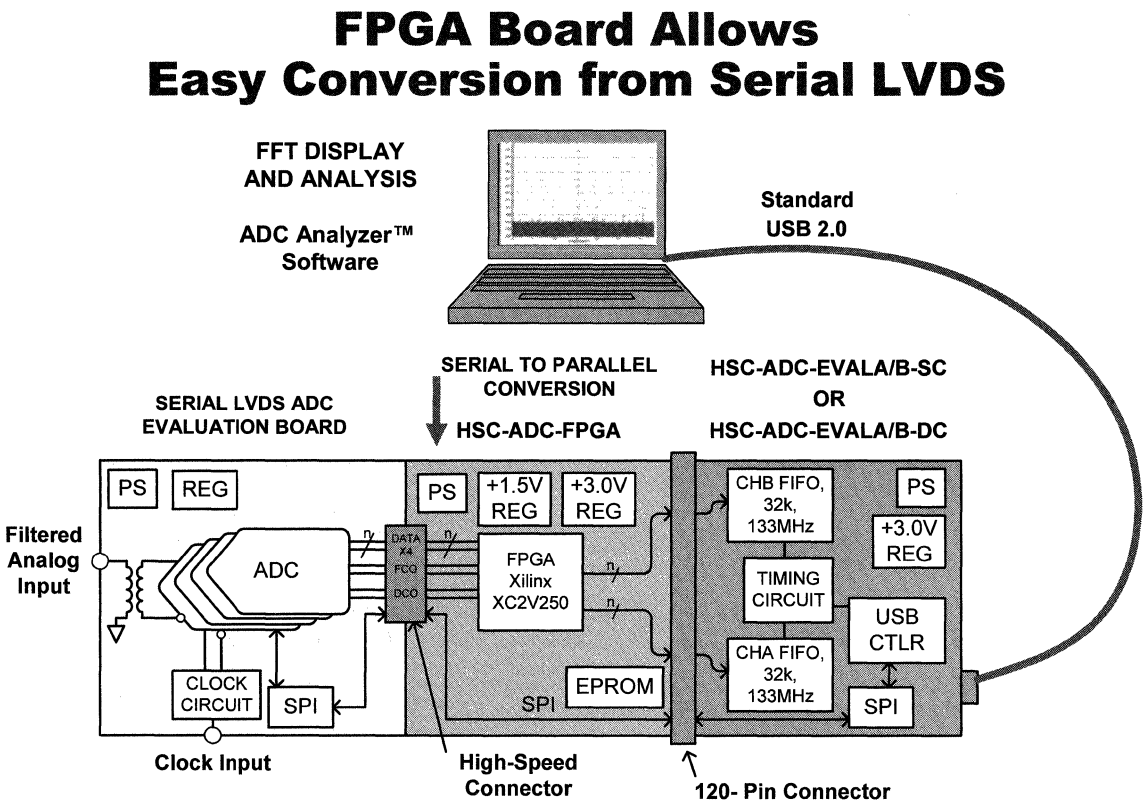

In order to accommodate ADCs with LVDS outputs, the HSC-ADC-FPGA interface board is added to perform the serial-to-parallel conversion required to interface LVDS ADCs to the FIFO Board.

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**

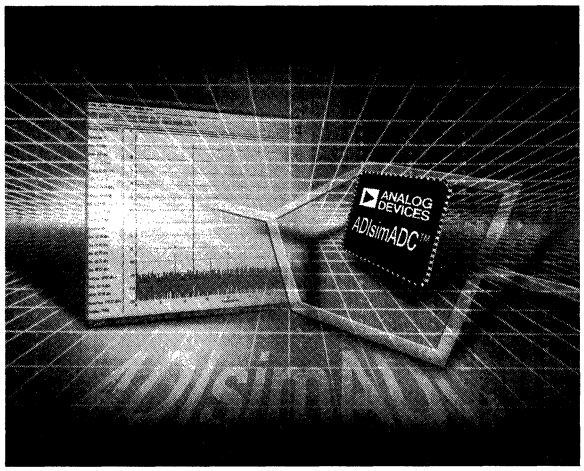

ADIsimADC® simulates behavior using many critical specifications of the data converters, including offset, gain, sample rate, bandwidth, jitter, latency, both ac and dc nonlinearity.

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

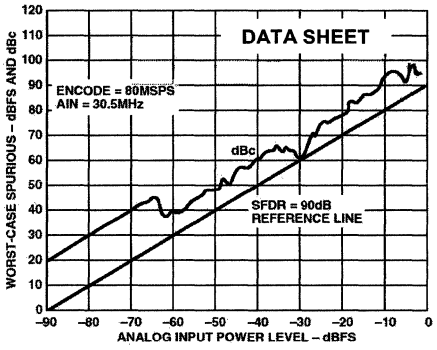

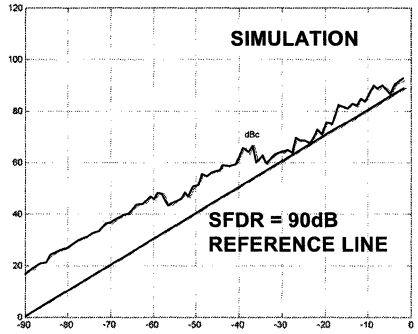

ADIsimADC is a behavioral model which utilizes a transfer function based on characterization data for the ADC under consideration. The models are posted on the ADI website and are updated when new products are released.

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®

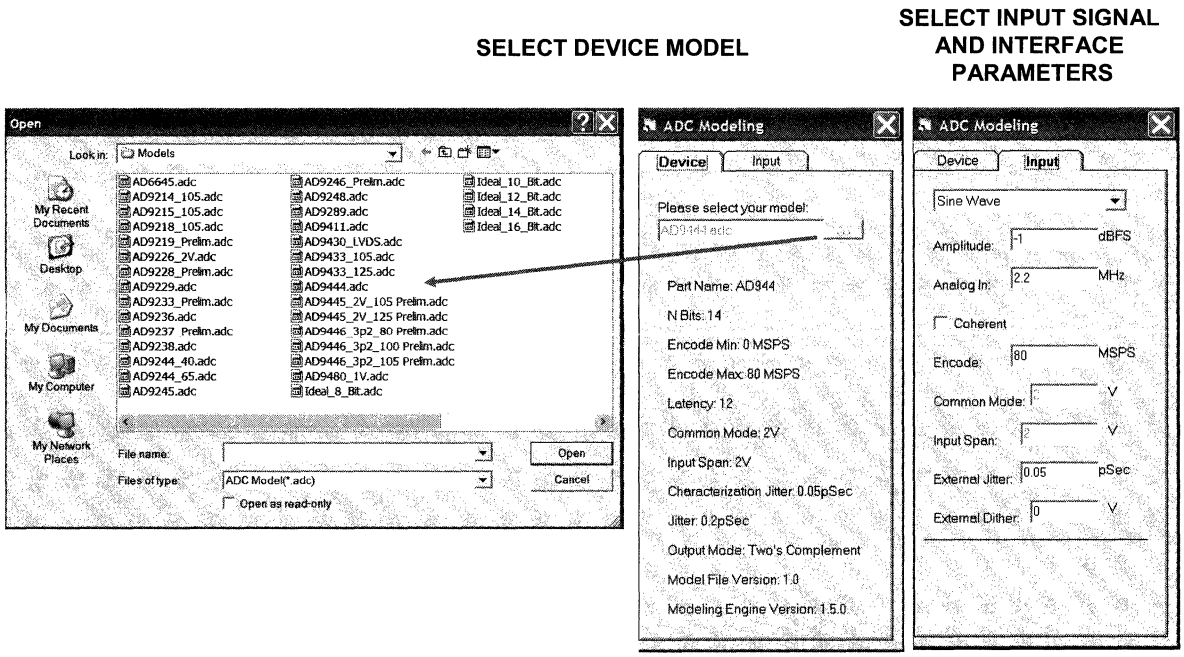

The ADIsimADC software is easy to run. This figure shows how the "Device" tab is first selected. The latest model set can be downloaded from www.analog.com/adisimadc, and each ADC model shows up in the list.

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

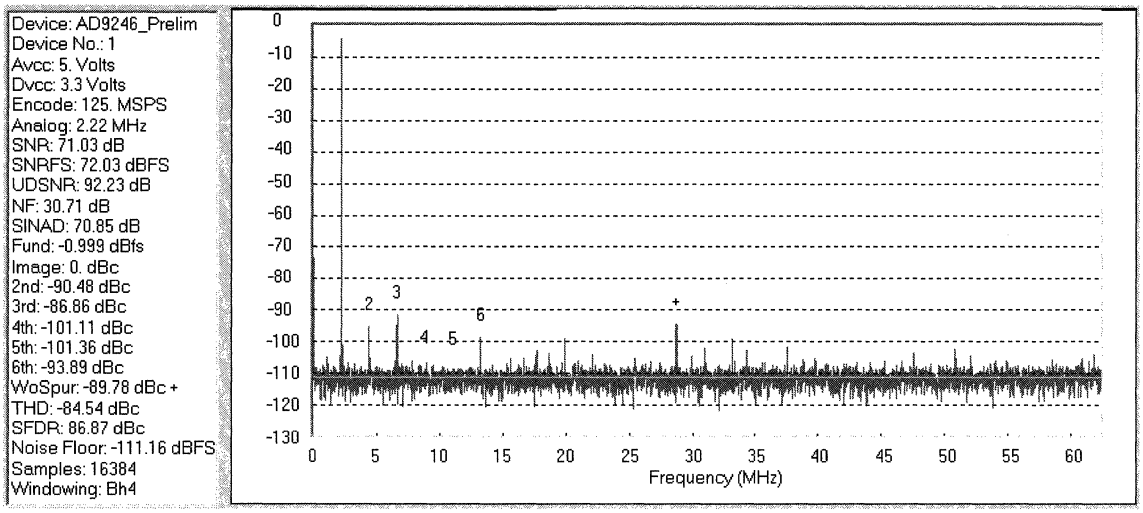

DATA TABLE **FFT OUTPUT**

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

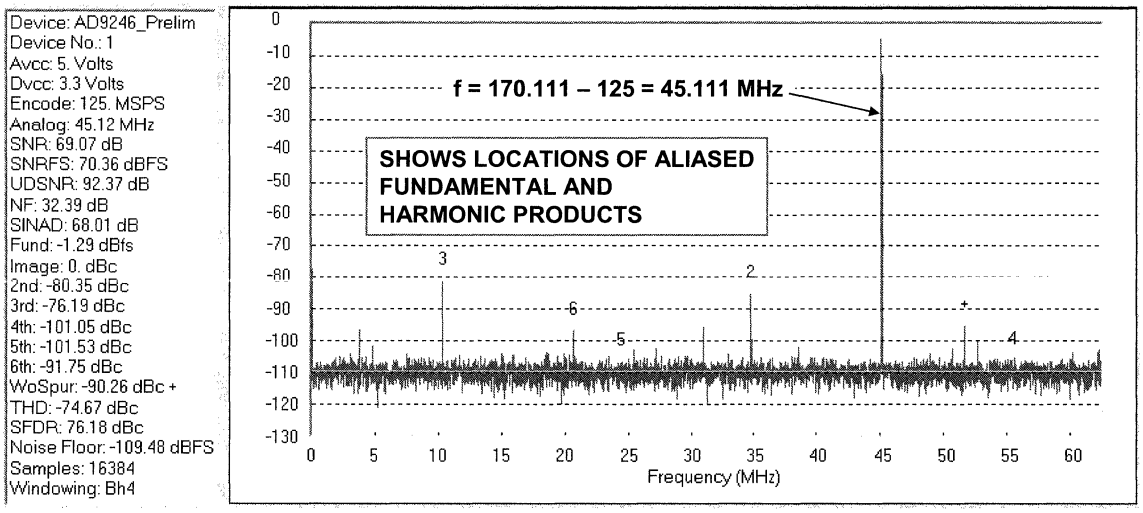

The ADIsimADC program works for IF sampled signals which are outside the first Nyquist zone. In this figure, the sample rate is 125MSPS and the input signal is located at the IF frequency of 170.111MHz. The FFT output shows the aliased fundamental of this signal at 45.111MHz.

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

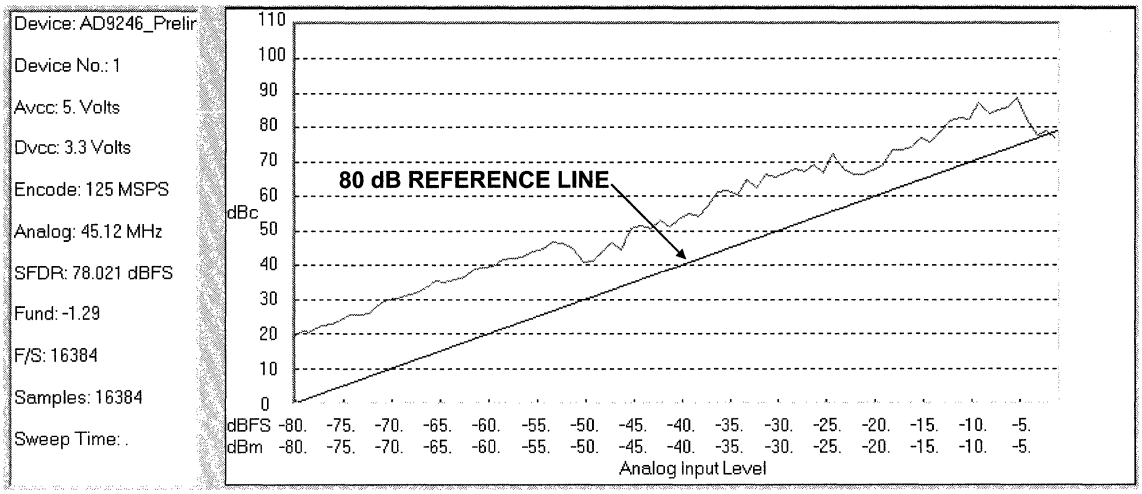

The FFT data generated by ADIsimADC can be displayed in a standard FFT as previously shown, or as an amplitude "sweep" as shown in this figure for SFDR.

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

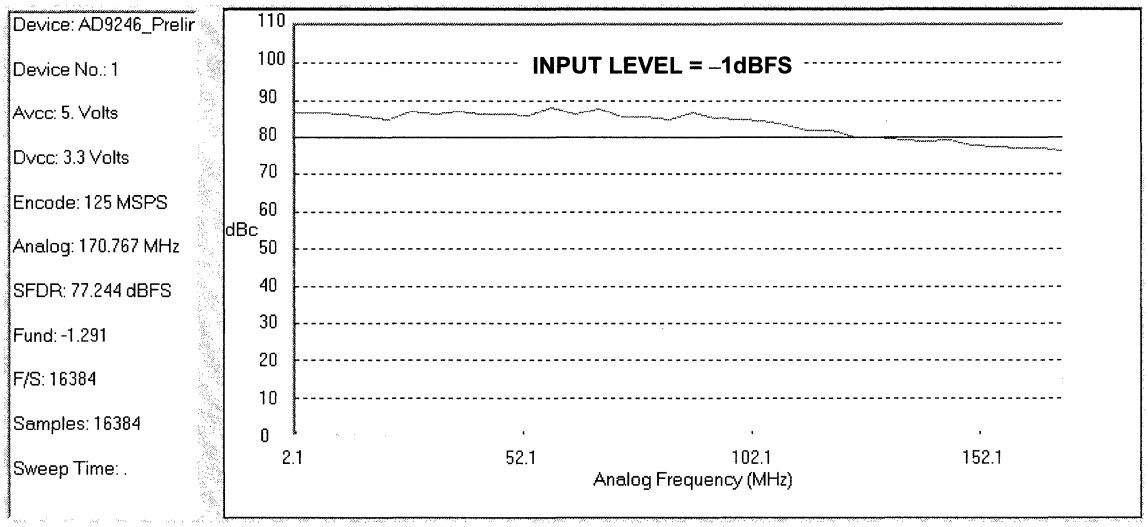

This figure shows a frequency "sweep" output, where the ADC is operated at a fixed sampling frequency and input amplitude, and the input frequency is "swept" over the desired range.

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

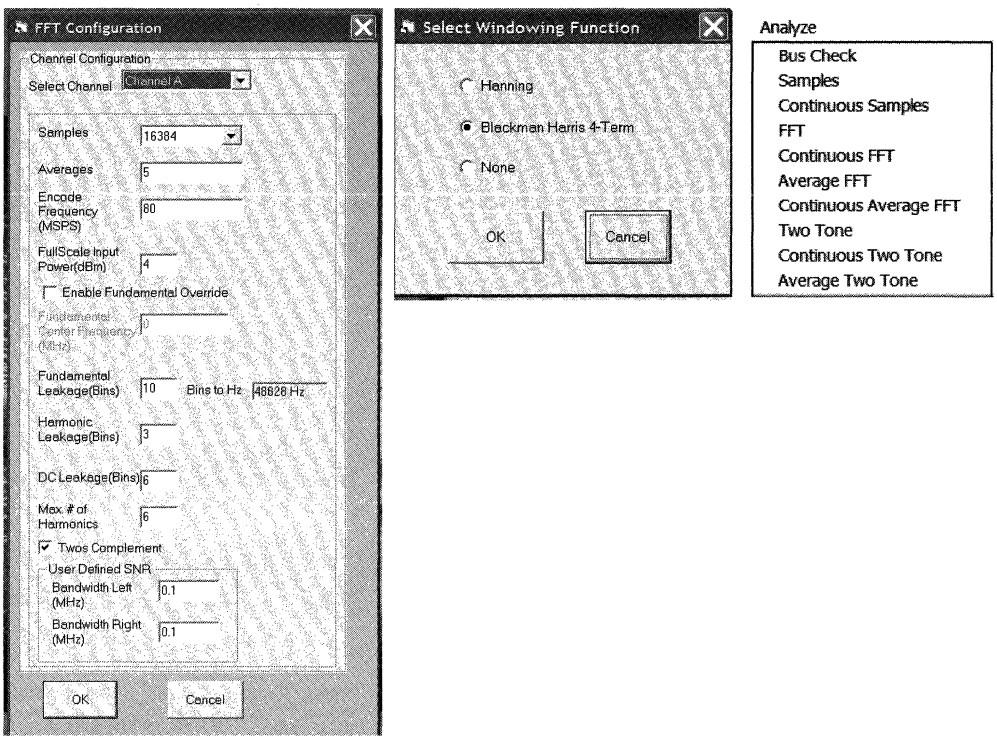

This figure shows some of the FFT and display options available in ADIsimADC.

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

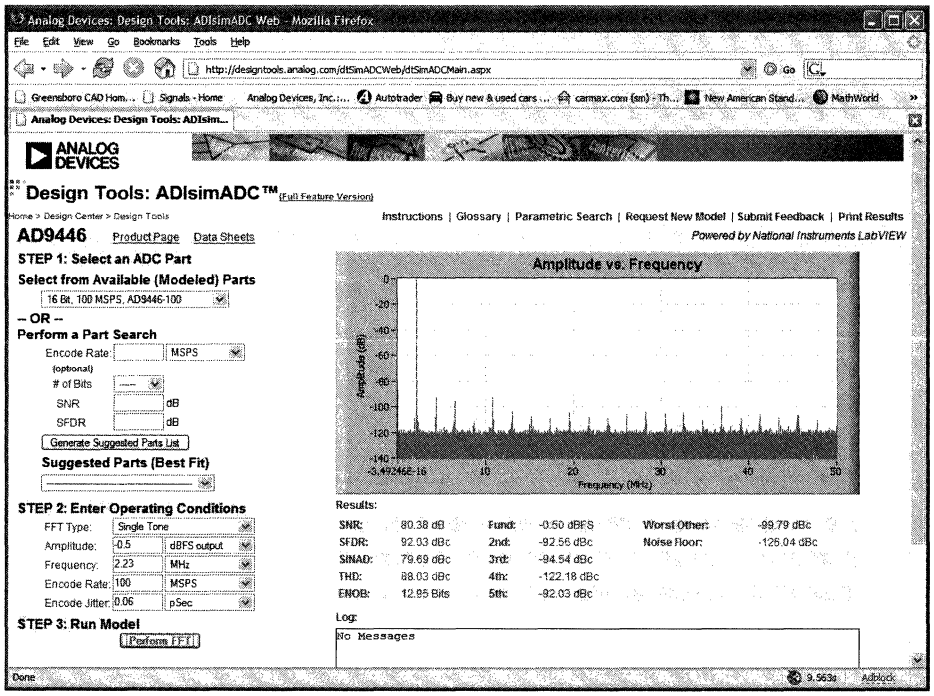

A web-based version of ADIsimADC is also available at www.analog.com/designtools. This on-line program is capable of running the basic FFT function as well as frequency and amplitude sweeps on the various ADCs under user-defined operating conditions.

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:

High Speed System Applications

SECTION 2
OPTIMIZING DATA CONVERTER INTERFACES

Data Converter Interface Overview

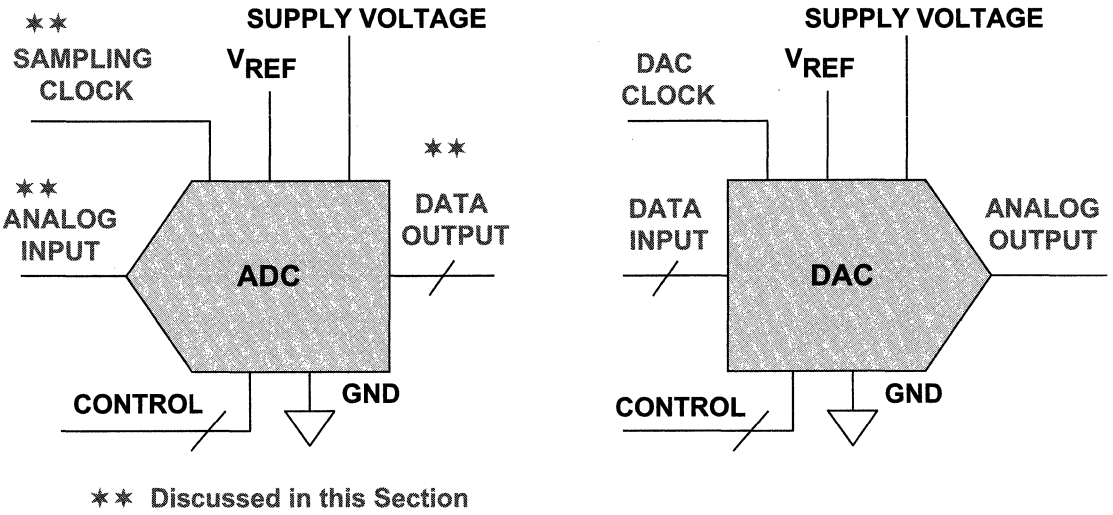

This figure shows the critical interfaces to an ADC or DAC:

Analog Input/Output
Sampling Clock/DAC Clock
Data Output/Input

The reference voltage, supply voltages, and ground are also important. Section 4 of this book will address the issues associated with grounding and decoupling, and this section discusses the ones listed above.

This section concentrates on ADCs; however, the same interface concepts apply equally to DACs, some of which are discussed in Section 3.

General Trends in Data Converters Affecting Interface Design

- Higher sampling rates, higher resolution
- Excellent ac performance
- Single-supply operation (e.g., +5V, +3V, +2.5V, +1.8V)
- Smaller input/output signal swings
- Differential inputs/outputs
- More sensitivity to noise
- Lower power, shutdown or sleep modes
- Maximize usage of low cost foundry CMOS processes
- Small surface mount packages

Several issues have complicated the design of data converter interfaces in recent years. The primary one is the trend to lower voltage, lower power, single-supply ICs which reduce signal swings proportionally. Smaller signal swings make modern data converters more sensitive to noise, grounding, and decoupling.

Many ADCs are now designed with differential inputs to reduce sensitivity to noise and also to get more signal swing for a given supply voltage. Selecting the proper drive amplifier is more complex because not only must it have differential outputs, but many times it must convert a single-ended signal to a differential one as well as perform a level shifting function to match the common-mode input voltage of the ADC.

These factors, added to the increased demand for higher sampling rates, resolutions, and excellent ac performance, make proper analog interface design critical to achieving the desired system performance.

Driving the ADC Analog Input

ADC Analog Inputs Are Not Ideal and Require Suitable Drivers

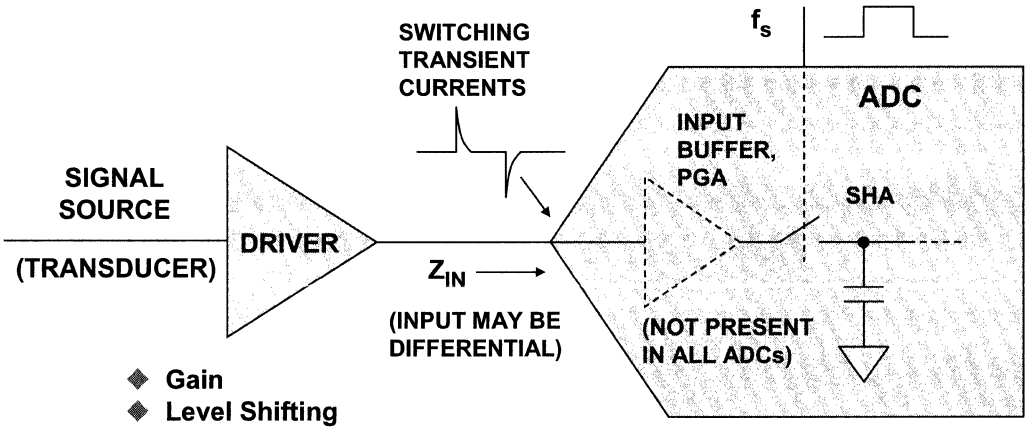

- Gain
- Level Shifting
- Isolation from Signal Source
- Impedance Transformation
- Single-Ended to Differential Conversion

- Driver Should Not Degrade DC or AC Performance of ADC
- A Few ADCs Are Designed to Interface Directly with Transducers (Some Σ-Δ and SAR ADCs)

An ideal ADC analog input circuit would have a constant resistive input impedance (in most cases, the resistance is several kΩ, but a few ADCs are designed with lower impedances) and an input range compatible with the signal source. Practical ADCs, however, present a finite complex (real and reactive components) input impedance to the driver, and may have transient currents on the input which are due to the switching action of the sample-and-hold function in the converter.

The input of the ADC may be buffered internally to minimize these transients, but CMOS ADCs typically do not have the input buffer.

With no internal buffer, an external driver may be required to isolate the signal source from the ADC input. When subjected to the transient currents, the driver must settle to the required accuracy in an interval that is less than approximately equal to one-half the sampling clock period.

The external ADC driver may be required to perform other functions such as gain, level shifting, single-ended to differential conversion, as well as isolating the signal source from the ADC input.

The external driver should be selected so that it does not degrade either the ac or dc performance of the ADC.

The bandwidth of the driver is generally very high in order to keep distortion low; therefore, some noise filtering between the driver and the ADC is generally desirable.

However, one should not automatically assume that an external driver is required because a few ADCs are designed to interface directly with transducers (primarily some sigma-delta and SAR ADCs).

In any event, the ADC data sheet must be carefully studied to understand what type of analog input driver is suitable if one is required.

Single-Ended DC-Coupled Amplifier Drivers for ADCs

Op Amp Gain and Level Shifting Circuits

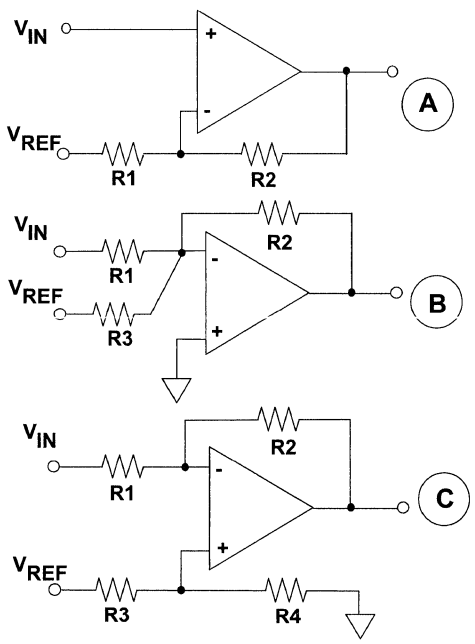

$$V_{out} = \left(1 + \frac{R2}{R1}\right) V_{in} - \frac{R2}{R1} V_{ref}$$

$$\text{NOISE GAIN} = 1 + \frac{R2}{R1}$$

$$V_{out} = -\frac{R2}{R1} V_{in} - \frac{R2}{R3} V_{ref}$$

$$\text{NOISE GAIN} = 1 + \frac{R2}{R1||R3}$$

$$V_{out} = -\frac{R2}{R1} V_{in} + \left(\frac{R4}{R3 + R4}\right)\left(1 + \frac{R2}{R1}\right) V_{ref}$$

$$\text{NOISE GAIN} = 1 + \frac{R2}{R1}$$

In dc-coupled applications, the drive amplifier may be required to provide gain and offset voltage, to match the signal to the input voltage range of the ADC. This figure summarizes various op amp gain and level shifting options. The circuit of A operates in the non-inverting mode, and uses a low impedance reference voltage, V_{REF}, to offset the output. Gain and offset interact according to the equation:

$$V_{OUT} = [1 + (R2/R1)] \cdot V_{IN} - [(R2/R1) \cdot V_{REF}].$$

The circuit in B operates in the inverting mode, and the signal gain is independent of the offset. The disadvantage of this circuit is that the addition of R3 increases the noise gain, and hence the sensitivity to the op amp input offset voltage and noise. The input/output equation is given by:

$$V_{OUT} = -(R2/R1) \cdot V_{IN} - (R2/R3) \cdot V_{REF}.$$

The circuit in C also operates in the inverting mode, and the offset voltage V_{REF} is applied to the non-inverting input without noise gain penalty. This circuit is also attractive for single-supply applications ($V_{REF} > 0$). The input/output equation is given by:

$$V_{OUT} = -(R2/R1) \cdot V_{IN} + [R4/(R3+R4)][\,1 + (R2/R1)] \cdot V_{REF}.$$

Note that the circuit of A is sensitive to the impedance of V_{REF}, unlike the counterparts in B and C. This is because the signal current flows into/from V_{REF}, due to V_{IN} operating the op amp over its common-mode range. In the other two circuits the common-mode voltages are fixed, and no signal current flows in V_{REF}. However, a dc current flows from the reference in B and C, so the output impedance of the reference must be added to R3 in performing the calculations.

Single-Ended Level Shifter with Gain Requires Rail-to-Rail Op Amp

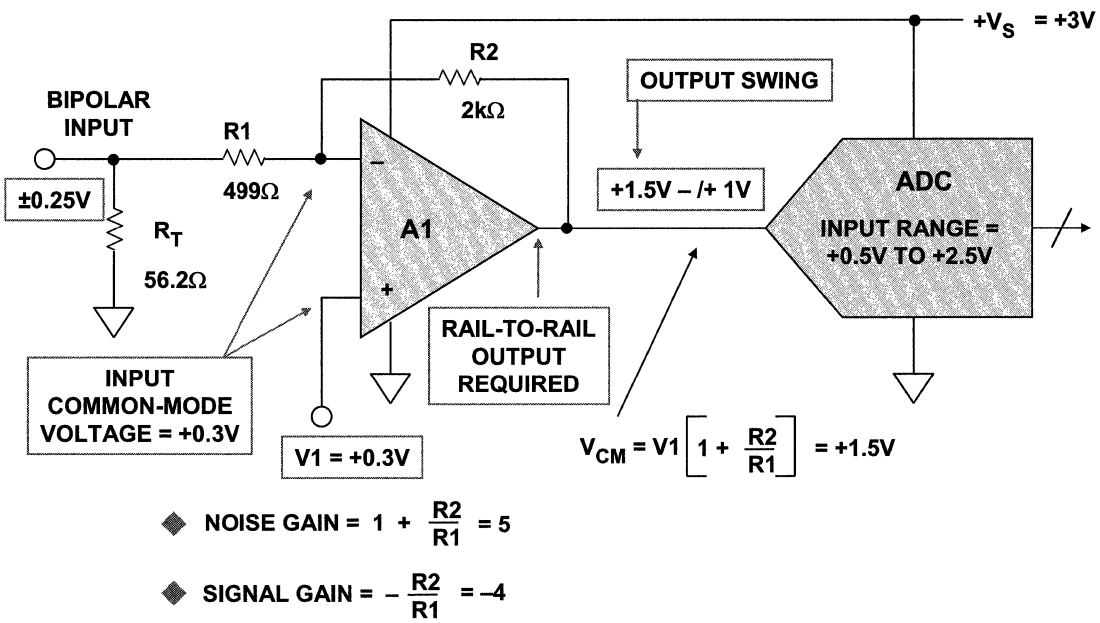

This circuit represents a typical single-supply ADC driver interface example. The is ideally suited to a single-supply level shifter and is similar to C shown in the previous figure. It will now be examined further in light of single-supply and common-mode issues. This figure shows the amplifier driving an ADC with an input range of +0.5V to +2.5V. Note that the entire circuit must operate on a single +3V supply.

The input range of the ADC (+0.5V to +2.5V) determines the output range of the A1 op amp. In order to drive the signal to within 0.5V of each rail, a rail-to-rail output stage is required. The signal gain of the op amp is set to 4, thereby amplifying the 0.5V p-p input signal to 2V p-p.

The input common-mode voltage of A1 is set at +0.3V which generates the required output offset of +1.5V. This produces the +1.5V offset when the bipolar input signal is at 0V. Note that some non rail-to-rail single-supply op amps can accommodate this input common-mode voltage when operating on a single +3V supply; however, the amplifier data sheet must be consulted.

This relatively simple circuit is an excellent example of where careful analysis of dc voltages is invaluable to the amplifier selection process.

Note that there will usually be some noise filtering between the amplifier output and the ADC input, but this will be discussed in more detail later.

An understanding of rail-to-rail input and output structures is needed to select the proper drive amplifier, and this discussion follows.

A True Rail-to-Rail Input Stage

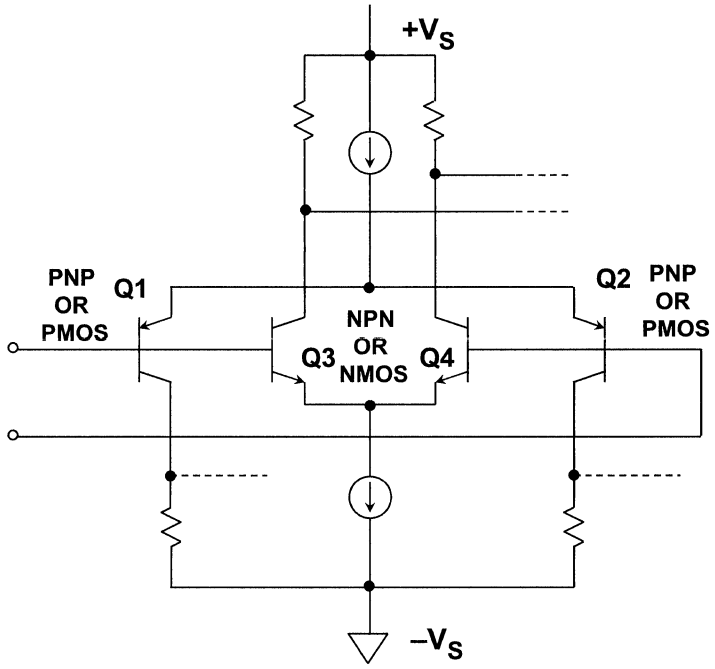

A simplified diagram of what has become known as a true rail-to-rail input stage is shown in this figure. Note that this requires use of two long tailed pairs, one of PNP bipolar transistors Q1-Q2, the other of NPN transistors Q3-Q4. Similar input stages can also be made with CMOS or JFET differential pairs.

The NPN pair is operational when the input common-mode voltage is at or near the positive rail, and the PNP pair is operational when the input common-mode voltage is at or near the negative rail.

Rail-to-rail amplifier input stage designs must transition from one differential pair to the other differential pair, somewhere along the input common-mode voltage range. Where this transition region occurs depends upon the particular amplifier design under consideration. Some have it set near the positive rail, some near mid-supply, and some near the negative rail. Some have an externally programmable transition region.

When the common-mode voltage is in the transition region, the input bias current will most likely change value and direction, and the common-mode rejection may be degraded. Other specifications may also be affected (such as the offset voltage), so the data sheet of the amplifier must be carefully studied to ensure that this is not a problem for the required system common-mode operating voltage.

At this point it should be noted that not all single-supply op amps are rail-to-rail.

Popular Op Amp Output Stage

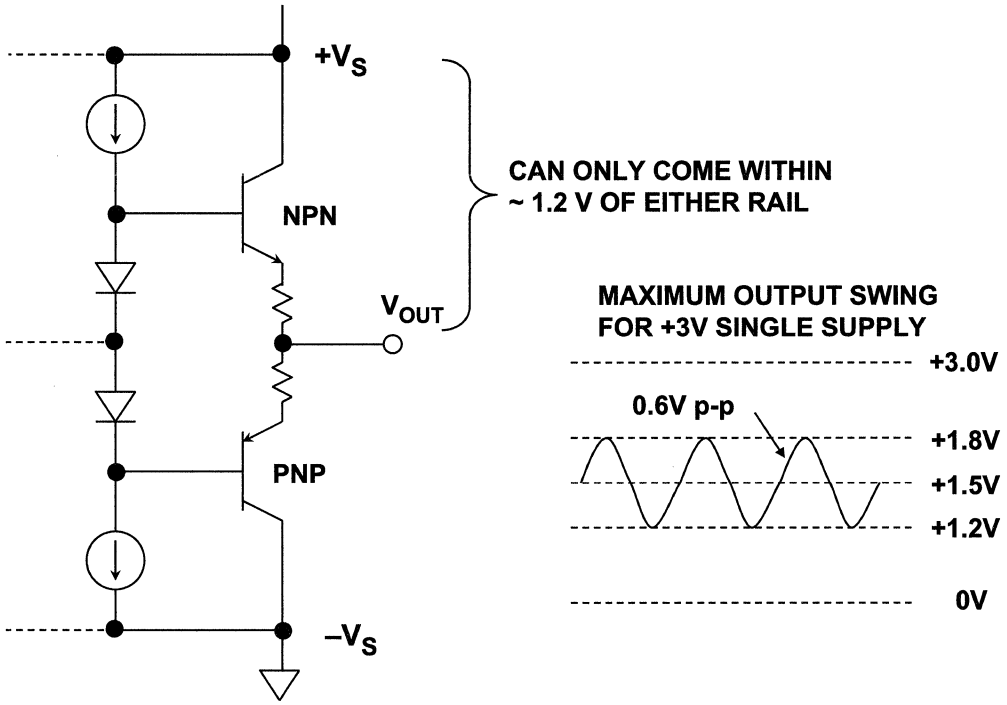

We will now examine output stages of op amps. This is a standard emitter-follower (common collector) output used in complementary bipolar processes. It has low output impedance and is relatively insensitive to capacitive loading.

However, the output can go no closer than about 1.2V to each supply rail. The headroom requirement can be even greater than 1.2V for some op amps which use this output structure, depending on the design.

On low supply voltages, such as 3V, this stage has only 0.6V peak-to-peak output voltage swing, centered on a common-mode voltage of +1.5V. In a very few applications (especially in differential output amplifiers) this swing may be adequate. In most single-ended applications, however, more signal swing is required.

"Almost" Rail-to-Rail Output Stages

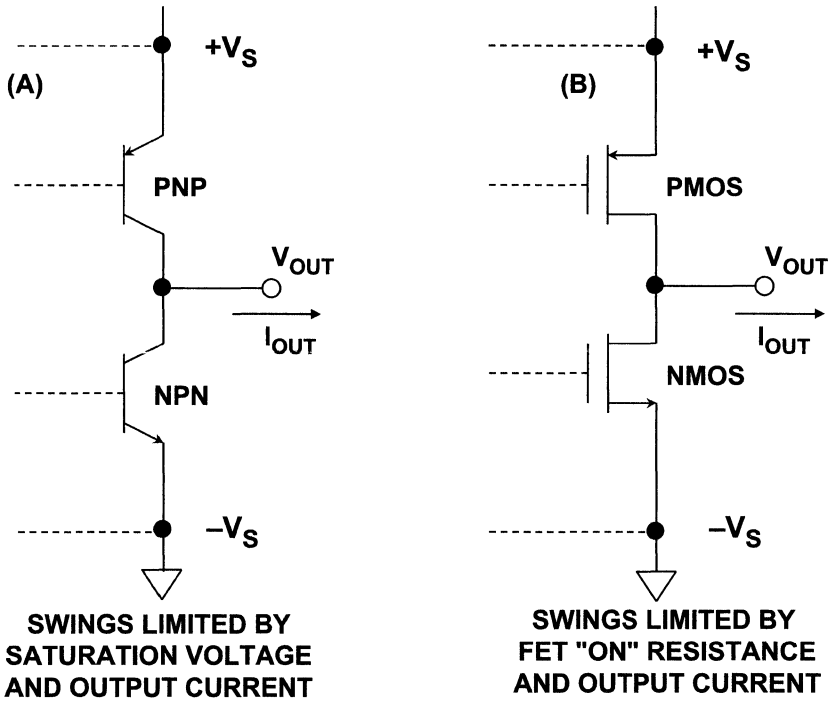

The complementary common-emitter/common-source output stages shown in A and B allow the op amp output voltage to swing much closer to the rails, but these stages have much higher open-loop output impedance than do the emitter follower-based stages previously discussed.

In practice, however, the amplifier's high open-loop gain and the applied feedback can still produce an application with low output impedance (particularly at frequencies below 10Hz). What should be carefully evaluated with this type of output stage is the loop gain within the application, with the load in place. Typically, the op amp will be specified for a minimum gain with a load resistance of 10kΩ (or more). Care should be taken that the application loading doesn't drop lower than the rated load, or gain accuracy may be lost.

It should also be noted that these output stages will cause the op amp to be more sensitive to capacitive loading than the emitter-follower type. Again, this will be noted on the device data sheet, which will indicate a maximum of capacitive loading before overshoot or instability will be noted.

The complementary common emitter output stage using BJTs in A cannot swing completely to the rails, but only to within the transistor saturation voltage (VCESAT) of the rails. For small amounts of load current (less than 100μA), the saturation voltage may be as low as 5 to 10mV, but for higher load currents, the saturation voltage can increase to several hundred mV (for example, 500mV at 50mA).

On the other hand, an output stage constructed of CMOS FETs as in B can provide nearly true rail-to-rail performance, but only under no-load conditions. If the op amp output must source or sink substantial current, the output voltage swing will be reduced by the I·R drop across the FET's internal "on" resistance. Typically this resistance will be on the order of 100Ω for precision amplifiers, but it can be less than 10Ω for high current drive CMOS amplifiers.

For the above basic reasons, it should be apparent that there is no such thing as a true rail-to-rail output stage, hence the title "Almost" Rail-to-Rail Output Stages. The best any op amp output stage can do is an "almost" rail-to-rail swing, when it is lightly loaded.

Input Circuit of AD7276 2.35V-3.6V, 12-Bit, 3MSPS 6-Lead TSOT ADC

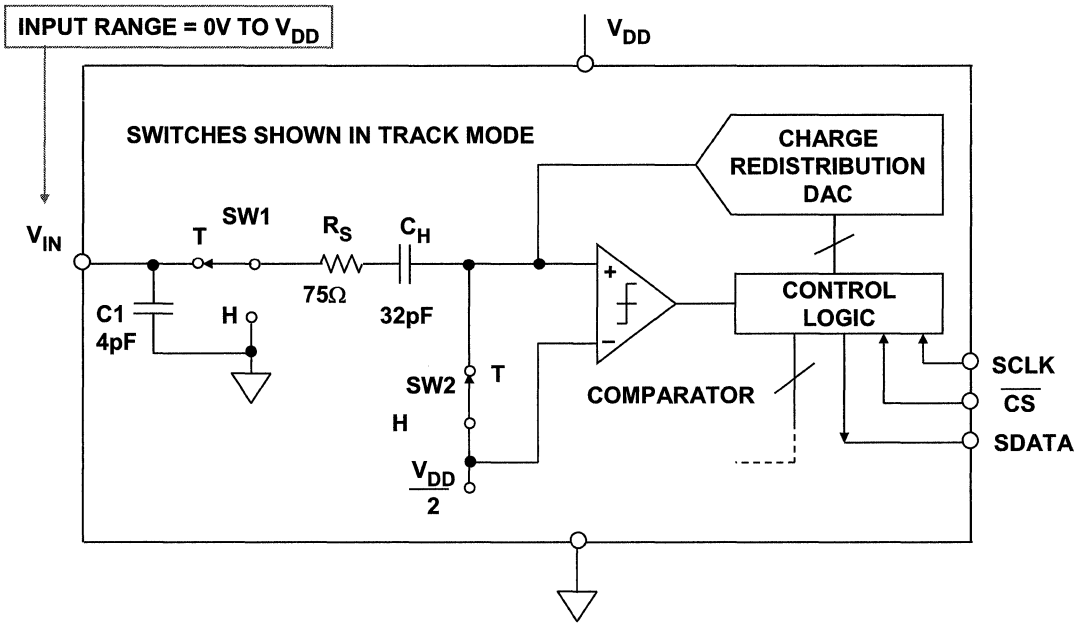

The AD7276/AD7277/AD7278 are 12-/10-/8-bit, low power (12.6mW), successive approximation ADCs, respectively. The parts operate from a single 2.35V to 3.6V power supply and feature throughput rates of up to 3MSPS. The parts contain a low noise, wide bandwidth track-and-hold circuit that can handle input frequencies in excess of 55MHz. The conversion process and data acquisition are controlled using the serial clock, allowing the devices to interface with microprocessors or DSPs. The input signal is sampled on the falling edge of CS (bar), and the conversion is also initiated at this point. There are no pipeline delays associated with the part. The reference for the part is taken internally from V_{DD}. This allows the widest dynamic input range to the ADC; therefore, the analog input range for the part is 0 to V_{DD}. The conversion rate is determined by the SCLK. For 3MSPS operation, SCLK is 48MHz. The CS (bar) signal does not have to be synchronized to SCLK.

A simplified block diagram of the series is shown in this figure. This ADC utilizes a standard successive approximation architecture based on a switched capacitor CMOS charge redistribution DAC. The input CMOS switches, SW1 and SW2, comprise the sample-and-hold function, and are shown in the track mode in the diagram. Capacitor C1 represents the equivalent parasitic input capacitance, C_H is the hold capacitor, and R_S is the equivalent on-resistance of SW2. In the track mode, SW1 is connected to the input, and SW2 is closed. In this condition, the comparator is balanced, and the hold capacitor C_H is charged to the value of the input signal. Note that the drive circuit must be capable of driving this capacitance, and the series resistance must not be high enough to limit the bandwidth. Assertion of convert start CS (bar) starts the conversion process: SW2 opens, and SW1 is connected to ground, causing the comparator to become unbalanced. The control logic and the charge redistribution DAC are used to add and subtract fixed amounts of charge from the hold capacitor to bring the comparator back into balance. At the end of the appropriate number of clock pulses, the conversion is complete.

Under certain conditions, the AD7276-family can be directly connected to the source as described in the next figure.

Low Source Resistances Can Drive the AD7276 Input Directly

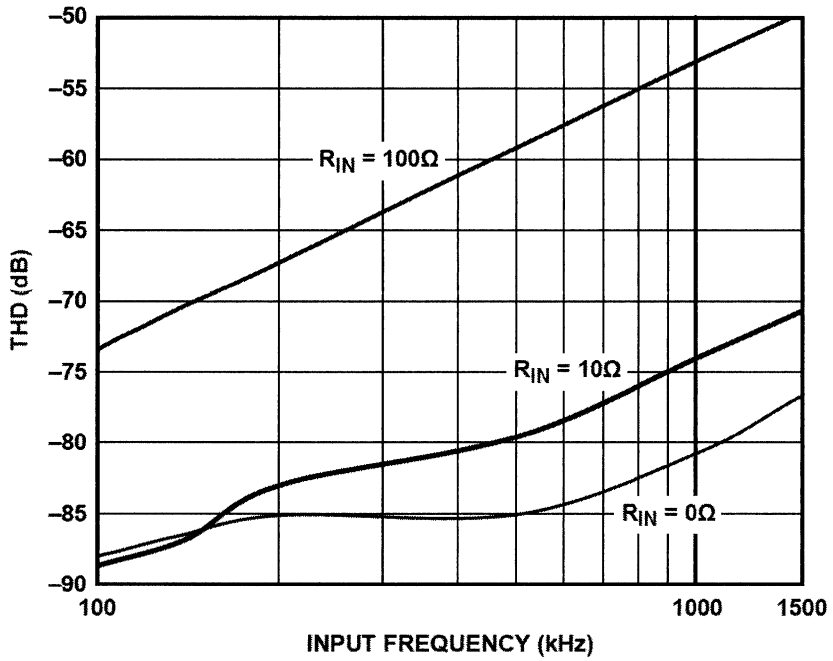

This figure shows the AD7276 THD as a function of the analog input frequency and the source resistance. This allows the user to determine if an external buffer amplifier is required, based on the system THD requirement and the source resistance.

The reason for the strong dependence of THD on source resistance is because of the highly nonlinear nature of the input circuit. Larger source resistances increase the effects of the nonlinearity caused by the changing capacitance as the sample-and-hold switches.

Op Amp Driver for AD7276 Requires Dual Supply Op Amp

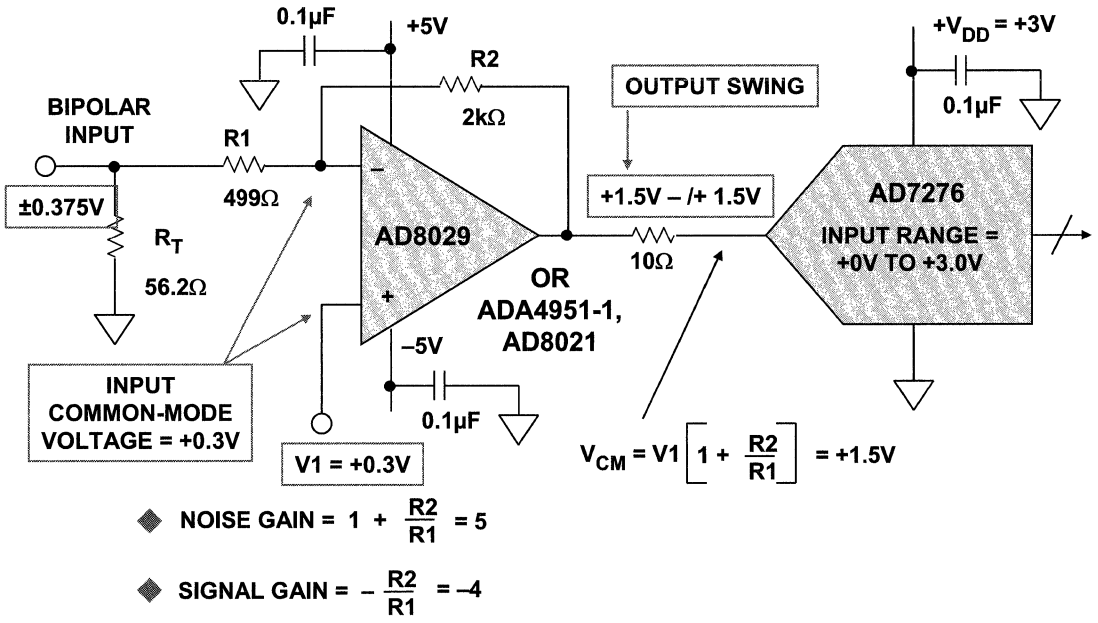

If an external buffer amplifier is needed to drive the AD7276, it must operate on separate supplies, because the output stage must drive the signal between 0V and +3V. A rail-to-rail output amplifier operating on a single +3V supply would cause the signal to be clipped when it approached either 0V or +3V.

This shows the AD8029 operating on dual 5V supplies acting as a level shifter and a negative gain-of-4. The +0.3V dc level on the non-inverting input is amplified by the noise gain (5) to provide the +1.5V common-mode output level.

The design procedure starts by determining the signal gain required. The output swing is 3V p-p, and the input swing is 0.75V p-p, so the signal gain must be –4. This sets the ratio of R2 to R1.

The noise gain is 5, therefore a common-mode voltage of +0.3V is required to develop the output offset of +1.5V.

Note that since the drive amplifier operates on ±5V supplies, and the ADC on a +3V supply, care must be taken that the amplifier does not overdrive the ADC, especially under power-up conditions. A suitable clamping network may therefore be required to protect the ADC from overdrive.

Differential Amplifier Drivers for ADCs

Simplified Input Circuit for a Typical Unbuffered Switched Capacitor CMOS Sample-and-Hold

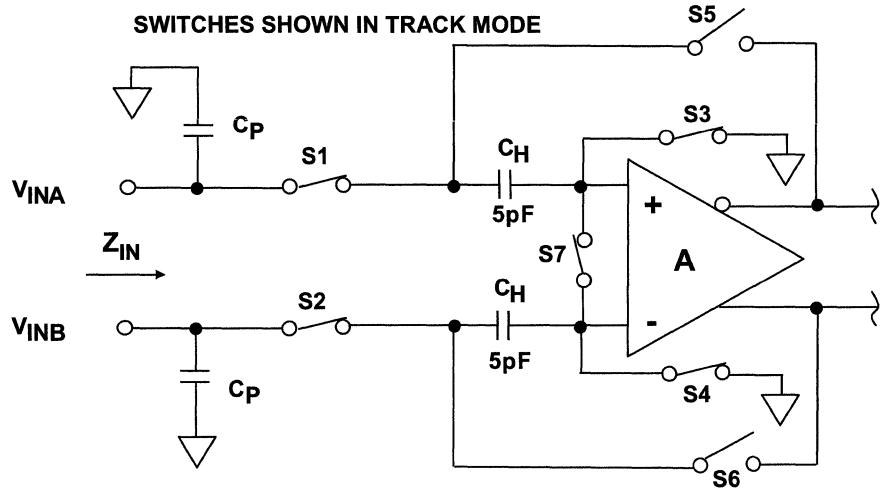

Z_{IN} IS A FUNCTION OF:
- TRACK MODE VS. HOLD MODE
- INPUT FREQUENCY

This figure shows a simplified input circuit for an unbuffered CMOS switched capacitor ADC. Most high performance CMOS switched capacitor pipelined ADCs have differential inputs. The differential structure is typically carried through most of the ADC. This makes matching requirements easier as well as reduces second-order products. In addition, the differential structure helps in common-mode noise rejection.

Note that the SHA switches are connected directly to each of the inputs. Switching transients can be significant, because there is no isolation buffer. The drive amplifier settling time to the transients must be fast enough so that the amplifier settles to the required accuracy in less than one-half the sampling period (this settling time must include the effects of any external series resistance).

The differential input impedance of this structure is dynamic and changes when the SHA switches between the sample mode and the hold mode. In addition, the impedance is a function of the analog input frequency.

In the track mode (shown in the figure), the input signal charges and discharges the hold capacitors, C_H. When the circuit switches to the hold mode the switches reverse their positions, and the voltage across the hold capacitors is transferred to the outputs.

It is highly recommended that this type of input be driven differentially for common-mode rejection of the switching transients. While it is possible to drive them single-ended (with one input connected to the appropriate common-mode voltage), degradation in SFDR will occur because the even-order distortion products are no longer rejected.

Typical Single-Ended (A) and Differential (B) Input Transients of CMOS Switched Capacitor ADC

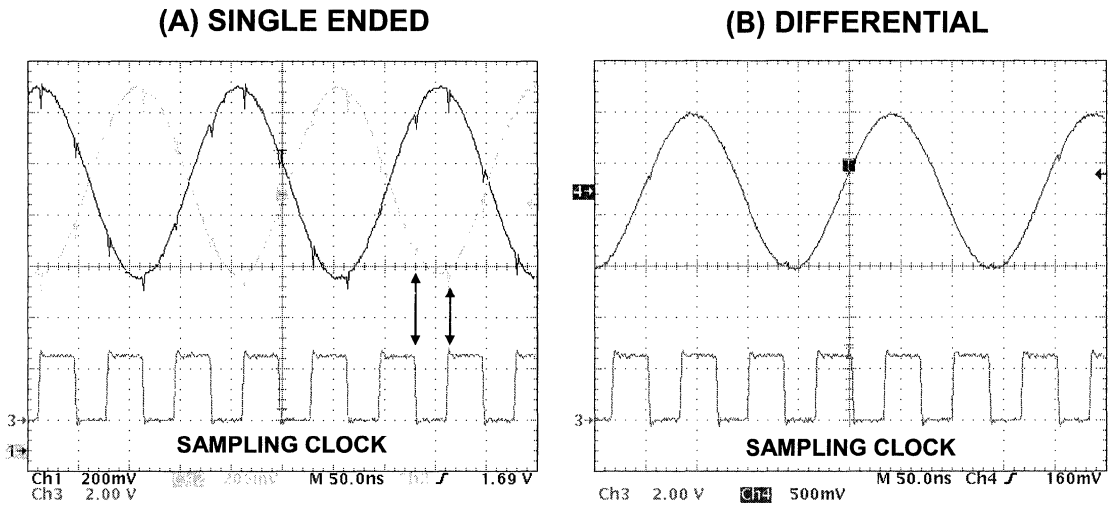

- **Differential charge transient is symmetrical around mid-scale and dominated by linear component**
- **Common-mode transients cancel with equal source impedance**

Note: Data Taken with 50Ω Source Resistances

Figure (A) shows each of the differential inputs of a typical unbuffered CMOS ADC as well as the sampling clock. The inputs were driven with a 50Ω source resistance. Note that a transient occurs on each edge of the sampling clock because of the switching action previously described.

Figure (B) shows the differential input signal to the ADC under the same conditions as (A). Note that most of the transient current glitches are cancelled because they are common-mode signals.

Note that for cancellation to be optimum the two inputs must be driven from a balanced source impedance (the real and reactive components of the impedance must be matched).

Advantages of Differential Analog Input Interfaces for Data Converters

- **Differential inputs give twice the signal swing vs. single-ended (Especially important for low voltage single-supply operation)**
- **Differential inputs help suppress even order distortion products**
- **Many IF/RF components such as SAW filters and mixers are differential**
- **Differential inputs suppress common-mode ADC switching noise including LO feed-through from mixer and filter stages**
- **Differential ADC designs allow better internal component matching and tracking than single-ended. Less need for trimming**

- **If you drive them single-ended, you will have degradation in distortion and noise performance**
- **However, many signal sources are single-ended, so the differential amplifier is useful as a single-ended to differential converter**

This list summarizes the advantages of using differential analog inputs for ADCs. In the real world, however, many signals are single-ended, and a convenient method is required to convert them to differential signals with minimum degradation in noise and distortion.

A family of differential amplifiers has been developed specifically for this purpose and are described in the next few pages.

The first two differential amplifiers discussed are the ADA4941 and the ADA4922. These amplifiers are optimum drivers for the 16- and 18-bit family of PulSAR successive approximation ADCs.

Another class of differential amplifiers is designed specifically for higher speed ADCs.

ADA4941 Driving AD7690 18-Bit PulSAR® ADC in +5V Application

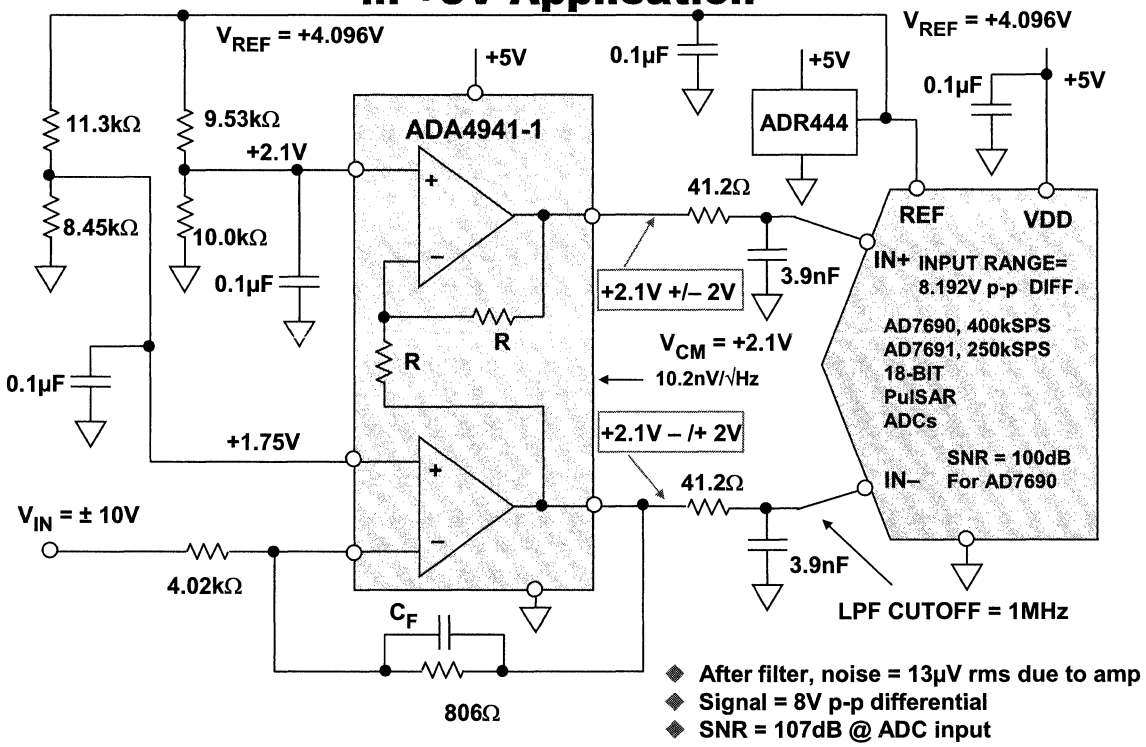

This figure shows the ADA4941-1 driving the 18-bit PulSAR family of ADCs which have switched capacitor inputs. The ADA4941-1 is a low power, low noise differential driver for ADCs up to 18 bits in systems that are sensitive to power. Small signal bandwidth is 31MHz. A resistive feedback network can be added to achieve gains greater or less than 2. The ADA4941-1 provides essential benefits, such as low distortion and high SNR that are required for driving high resolution ADCs. With a wide input voltage range (0V to 3.9V on a single 5V supply), rail-to-rail output, high input impedance, and a user-adjustable gain, the ADA4941-1 is designed to drive single-supply ADCs with differential inputs found in a variety of low power applications, including battery-operated devices and single-supply data acquisition systems.

In this application, the two resistor dividers set the output common-mode voltage of the ADA4941-1 to +2.1V so that the output only has to go to within 100mV of ground. This allows sufficient headroom for the rail-to-rail output stages of the amplifier and allows the entire circuit to operate on a single +5V supply.

The input range of the AD7690 and AD7691 is $2V_{REF}$ p-p differential. The reference used is the ADR444 which is a 4.096V reference. The 41.2Ω resistors and the 3.9nF capacitors for a lowpass filter with a cutoff frequency of 1MHz, suitable for use with the AD7690 which has an input bandwidth of 9MHz. A lower frequency cutoff frequency would be used with the 250kSPS AD7692 PulSAR ADC.

The output noise spectral density of the ADA4941-1 is 10.2nV/√Hz. Integration over the noise bandwidth of the filter yields:

$$v_{rms} = v_n\sqrt{BW} = 10.2\times10^{-9}\sqrt{(1.57\times1\times10^{6})} = 13\mu V.$$

The peak-to-peak signal is 8V, and the rms value of the signal is therefore 2.83V.

The SNR due to the op amp is therefore SNR = $20\log(2.83 \div 13\times10^{-6})$ = 107dB, which is 7dB better than the 100dB SNR of the ADC.

Positioning the Noise Reduction Filter to Reduce the Effects of the Op Amp Noise

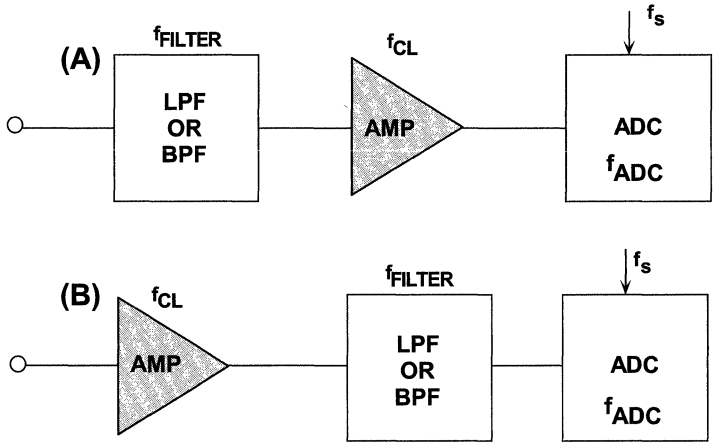

- **ADCs typically have very high input bandwidths, usually much greater than $f_S/2$**
- **Low distortion drive amplifiers typically have high bandwidths**
- **Placing a simple LPF or BPF placed between the AMP and the ADC is an excellent noise reduction technique**
- **The output capacitor of the filter absorbs some of the ADC input transient currents.**

ADCs typically have input bandwidths much greater than their maximum sampling rates. For instance, a 100MSPS ADC may have an input bandwidth of 700MHz.

A good drive amplifier also has a bandwidth which is much greater than the sampling rate in order to give good distortion performance over the bandwidth of interest.

The wideband noise from an op amp will therefore be integrated over the full input bandwidth of the ADC if the filter is placed ahead of the op amp as in (A).

The most desirable location for the noise reduction filter is between the amplifier and the ADC as shown in (B). However, the amplifier must be capable of driving the net impedance presented by the filter and the ADC.

The rms noise at the output of the filter is easily calculated from the following equation:

$$v_{rms} = v_n\sqrt{BW},$$

where v_n is the wideband voltage noise spectral density of the op amp (expressed in nV/√Hz), and BW is the equivalent noise bandwidth of the filter (see next slide). The SNR of the output of the filter due to the op amp noise can then be calculated knowing the peak-to-peak value of the input signal to the ADC. This SNR can then be compared to the SNR of the ADC.

The input noise of the ADC (not including the op amp) can be calculated based on the ADC SNR by using the equation SNR = 20log(Vsignal/Vnoise) and solving for Vnoise. Vsignal is simply the rms value of the ADC fullscale input signal.

The total input noise from both the ADC and the op amp can be calculated by combining the two noise sources on a root-sum-square basis because they are uncorrelated.

Good reductions in noise can be achieved with just a simple 1- or 2-pole filter as will be shown in the next figure.

Relationship Between Equivalent Noise Bandwidth and 3-dB Bandwidth for Butterworth Filter

NUMBER OF POLES	EQ. NOISE BW / 3dB BW
1	1.57
2	1.11
3	1.05
4	1.03
5	1.02

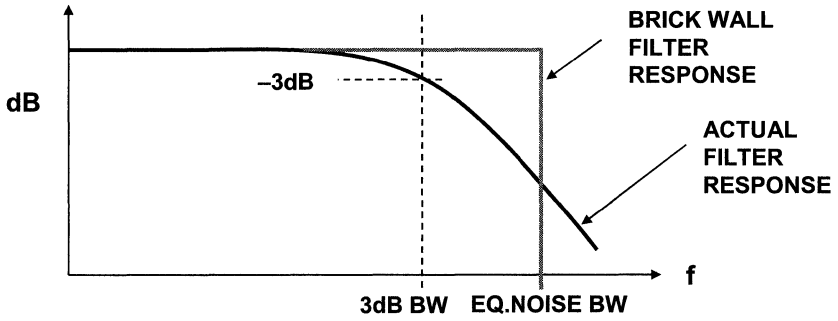

The equivalent noise bandwidth of a filter is the bandwidth of a "brick wall" filter which has the same effect on broadband Gaussian noise as an actual filter which has a finite transition region.

This figure shows the ratio of the equivalent noise bandwidth to the 3dB bandwidth for Butterworth filters with one to five poles. For a single-pole filter, the ratio is $\pi/2 = 1.57$. Notice that since the ratio is only 1.11 for a 2-pole filter, adding additional poles offers very little improvement in noise reduction.

Most of the driver circuits shown in the following figures in this section contain at least a single-pole RC noise reduction filter between the drive amplifier and the ADC. In many cases, the R and C values are optimized based on empirical data because of the transient nature of the CMOS ADC inputs.

ADA4922-1 Driving AD7634 18-Bit *i*CMOS PulSAR ADC in ±12V Industrial Application

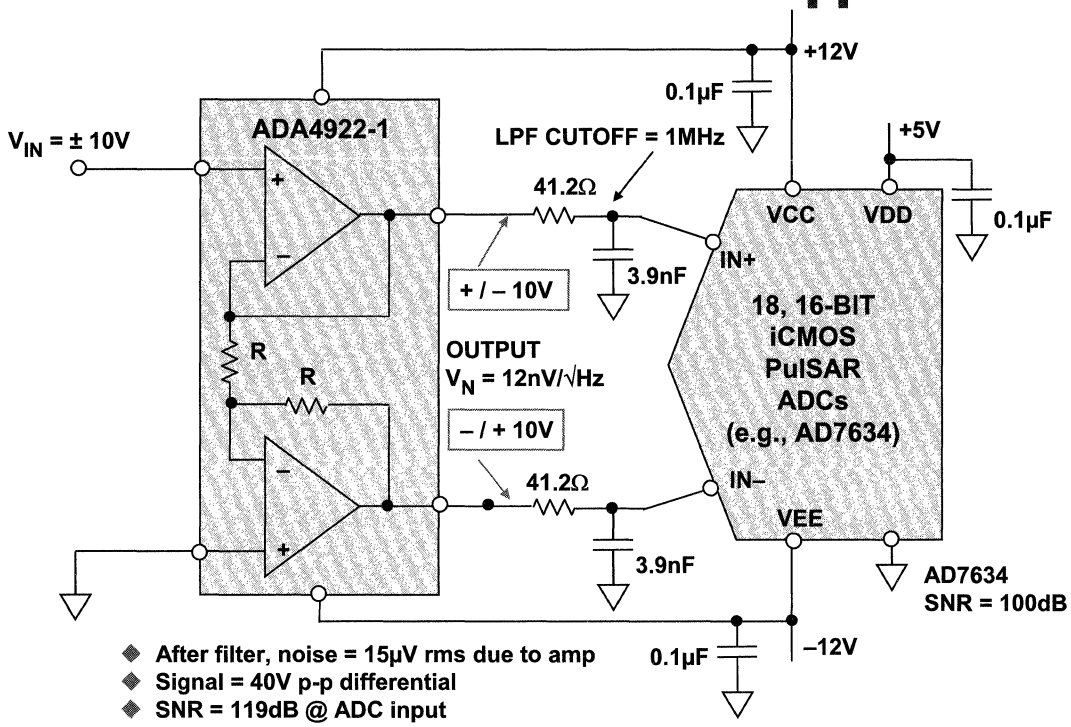

There are many industrial applications where signals as great as ±10V are standard. This figure shows a simple method for performing a single-ended to differential conversion using the ADA4922-1 driving a 16-bit or 18-bit *i*CMOS PulSAR ADC. The *i*CMOS family of PulSAR ADCs has a low power front end which operates high voltage supplies up to ±12V. The rest of the ADC operates on a low voltage power supply which is typically 5V.

The ADA4922-1 is a differential driver for 16-bit to 18-bit ADCs that have differential input ranges up to 40V p-p. Small signal bandwidth is 38MHz. Configured as an easy-to-use, single-ended-to-differential amplifier, the ADA4922-1 requires no external components to drive ADCs. The ADA4922-1 provides essential benefits such as low distortion and high SNR that are required for driving ADCs with resolutions up to 18 bits. With a wide supply voltage range (5V to 26V), high input impedance, and fixed differential gain of 2, the ADA4922-1 is designed to drive ADCs found in a variety of applications, including industrial instrumentation.

The ADA4922-1 is manufactured on ADI's proprietary second-generation XFCB process that enables the amplifier to achieve excellent noise and distortion performance on high supply voltages. The ADA4922-1 is available in an 8-lead 3mm × 3mm LFCSP as well as an 8-lead SOIC package. Both packages are equipped with an exposed paddle for more efficient heat transfer. The ADA4922-1 is rated to work over the extended industrial temperature range, −40°C to +85°C.

Noise calculations using the 1MHz lowpass filter yield 15µV rms for the op amp. The signal range of the ADC is 40V p-p, which is 14.14V rms. This yields an SNR of 119dB due to the op amp alone. Using the AD7634 SNR of 100dB, the rms ADC input noise contribution is calculated to be 141µV rms. The combined input ADC noise is therefore 142µV rms, and the contribution due to the op amp is almost negligible.

DC-Coupled Single-Supply Level Shifter for Driving AD922x ADC Input

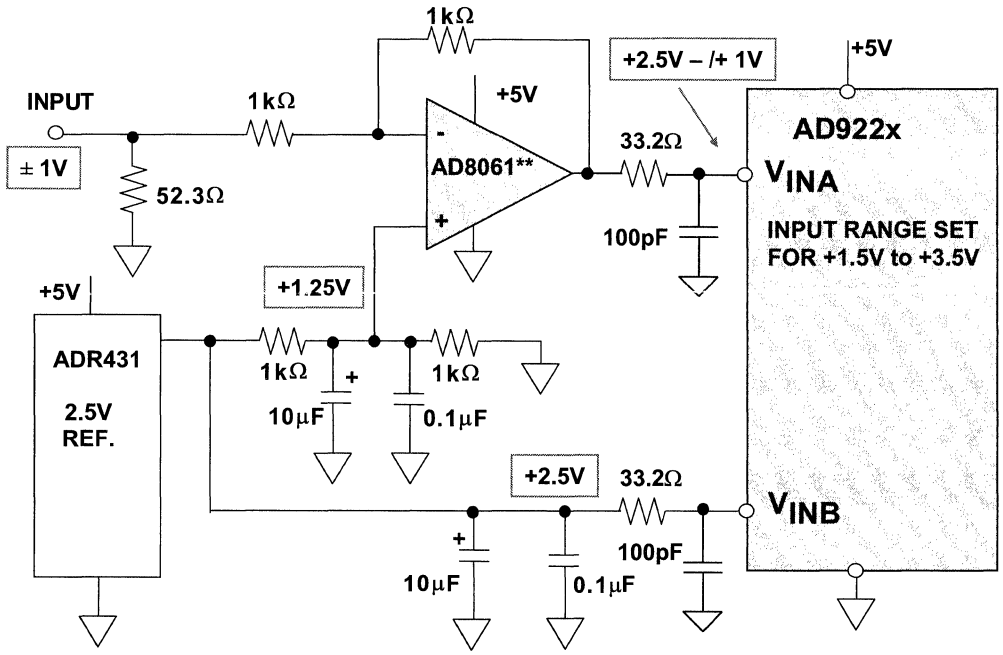

Distortion performance will most likely be compromised if a differential input ADC is driven single-ended. However, if single-ended drivers must be used, care should be taken that the source impedance is balanced as shown here. Balancing the source impedance (both R and C) allows some cancellation of the common-mode current transients produced by the ADC input SHA.

This circuit is designed to operate on a single +5V supply. It accepts a bipolar ±1V input signal and interfaces it to the input of the ADC whose range is set for 2V p-p with a 2.5V common-mode voltage. The AD8061 rail-to-rail output op amp is used, although others are suitable depending upon bandwidth and distortion requirements (for example, the AD8027, AD8031, or AD8091). The +1.25V input common-mode voltage for the AD8061 is developed by a voltage divider from the external ADR431 2.5V reference.

AD813x and ADA493x Differential ADC Drivers Functional Diagram and Equivalent Circuit

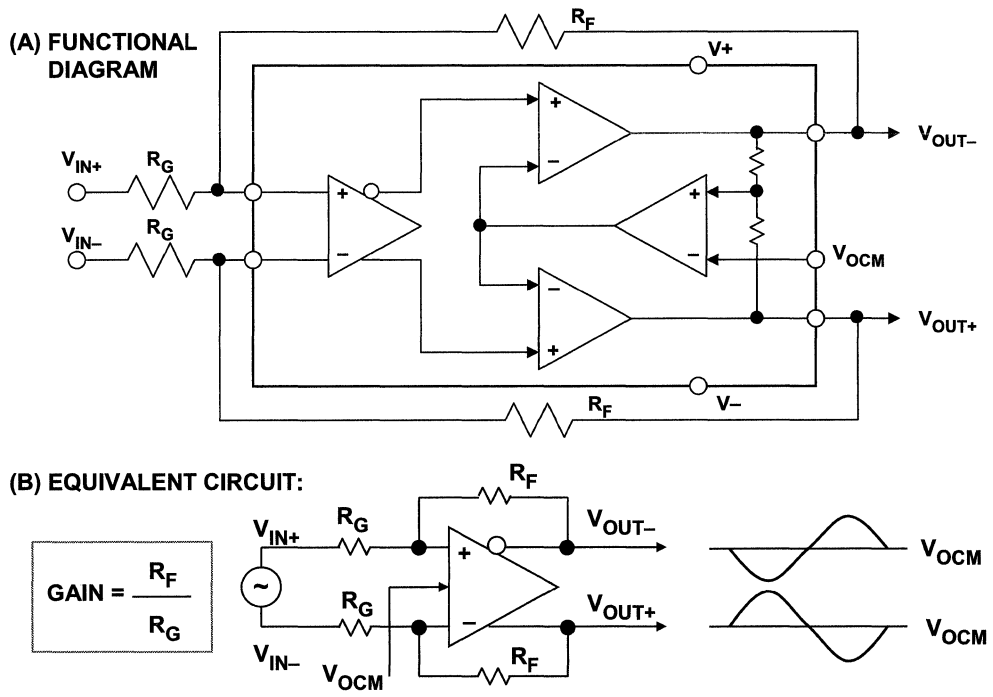

A block diagram of the AD813x and ADA493x family of fully differential amplifiers optimized for higher speed ADC driving is shown in this figure. The (A) diagram shows the details of the internal circuit, and (B) shows the equivalent circuit. The gain is set by the external resistors R_F and R_G, and the common-mode voltage is set by the voltage on the V_{OCM} pin. The internal common-mode feedback forces the V_{OUT+} and V_{OUT-} outputs to be balanced, i.e., the signals at the two outputs are always equal in amplitude but 180° out of phase per the equation,

$$V_{OCM} = (V_{OUT+} + V_{OUT-}) / 2.$$

The amplifier uses two feedback loops to separately control the differential and common-mode output voltages. The differential feedback, set with external resistors, controls only the differential output voltage. The common-mode feedback controls only the common-mode output voltage. This architecture makes it easy to arbitrarily set the output common-mode level in level shifting applications. It is forced, by internal common-mode feedback, to be equal to the voltage applied to the V_{OCM} input, without affecting the differential output voltage. The result is nearly perfectly balanced differential outputs of identical amplitude and exactly 180° apart in phase over a wide frequency range. The circuit can be used with either a differential or a single-ended input, and the voltage gain is equal to the ratio of R_F to R_G.

The V_{OCM} pin is internally biased at a voltage approximately equal to the midsupply point (average value of the voltages on V+ and V−). Relying on this internal bias results in an output common-mode voltage that is within about 100mV of the expected value. In cases where more accurate control of the output common-mode level is required, it is recommended that an external source, or resistor divider (made up of 10kΩ resistors or less), be used. In addition, the V_{OCM} pin should be decoupled to ground using a ceramic capacitor (0.01µF to 0.1µF).

DC-Coupled AD8138 Driving AD9235 12-Bit, 20/40/65MSPS CMOS ADC, Baseband Signal

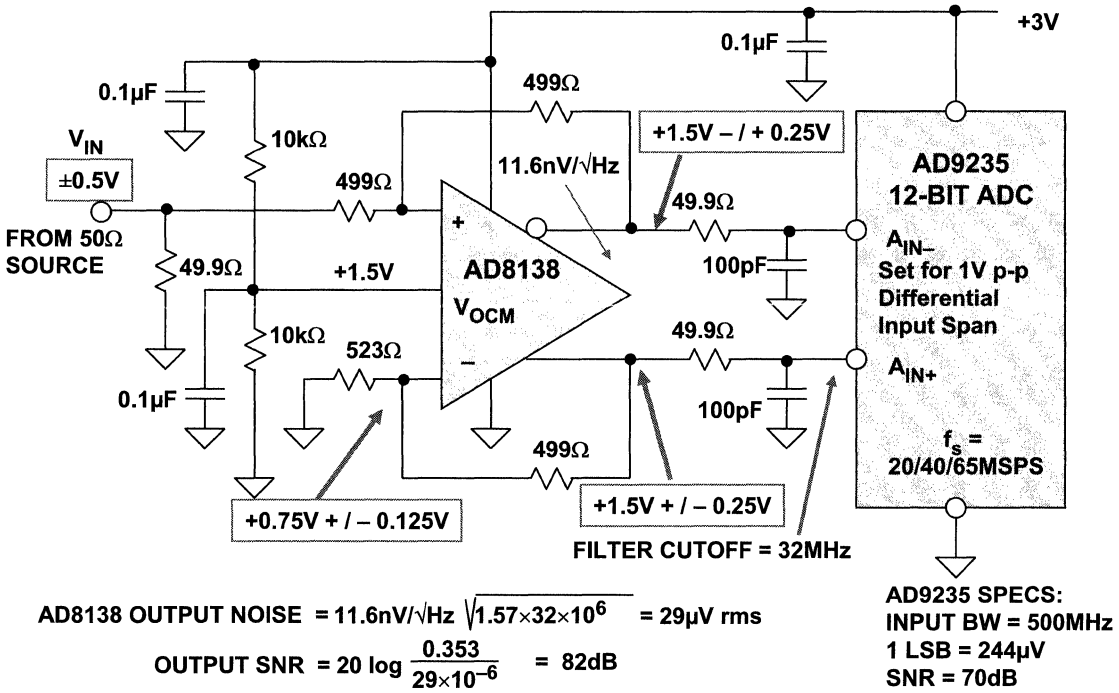

AD8138 OUTPUT NOISE = $11.6nV/\sqrt{Hz}\ \sqrt{1.57\times32\times10^6}$ = 29µV rms

OUTPUT SNR = $20\log\frac{0.353}{29\times10^{-6}}$ = 82dB

This figure represents a typical application of the AD8138 differential amplifier as a single-ended to differential ADC driver for the AD9235 12-bit, 65MSPS CMOS ADC.

The AD8138 has a 3dB bandwidth of 320MHz and delivers a differential signal with 85dBc SFDR for a 20MHz signal.

The input range of the AD9235 is set for 1V p-p differential; therefore, each input of the AD8138 only swings between +1.25V and +1.75V. This is within the output drive capability of the AD8138 even though it does not have a "rail-to-rail" output stage. This allows the AD8138 to be operated on the same +3V supply as the AD9235 ADC.

The 523Ω resistor matches the net drive impedance seen by the noninverting input (499Ω + 50Ω||49.9Ω ≈ 523Ω).

Note the simple RC input filtering circuit which reduces the effects of the transient currents as well as the amplifier noise. The cutoff frequency of the RC combination is 32MHz. The output voltage noise spectral density of the AD8138 is 11.6nV/√Hz, resulting in 29µV rms noise in the 32MHz bandwidth. The corresponding SNR is 82dB, which is 12dB better than the 70dB SNR of the AD9235, so the amplifier has only a minimal contribution to the overall system noise.

A Differential Amplifier Gain Calculator design tool is available at www.analog.com/DesignCenter, which can be used to check the input and output common-mode voltage ranges of the differential amplifier series for different power supplies, gain settings, and input signal ranges.

ADA4937-1 Driving AD6645 in +5V DC-Coupled Application

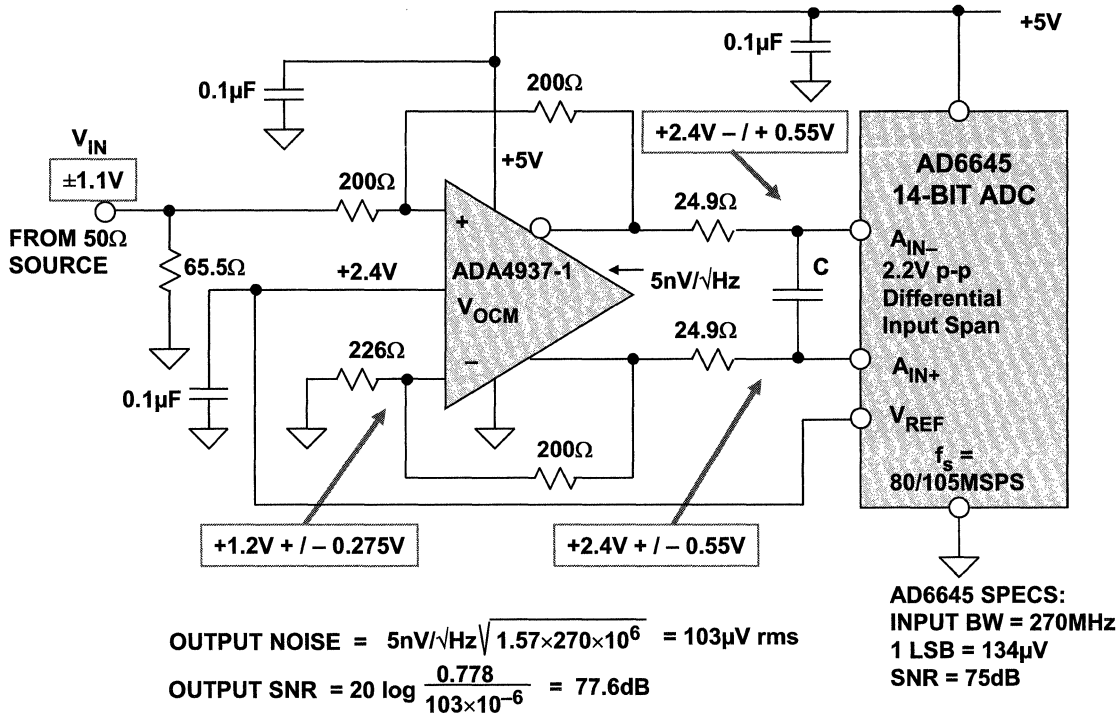

This figure as well as the following three figures illustrate very similar circuits. However, each circuit illustrates how subtle differences in ADC common-mode voltage, signal swing, and supply voltage affect the choice of differential drive amplifier.

The ADA4937-1 is one of the latest in the series of differential amplifiers and is optimized for operation on a single +5V supply. In this figure, it is used as a level shifter to drive the AD6645 14-bit 80/105MSPS ADC. The AD6645 operates on a 2.2V p-p differential signal with a common-mode voltage of +2.4V. This means that each output of the ADA4937 must swing between 1.85V and 2.95V which is within the output drive capability of the ADA4937-1 operating on a single +5V supply.

The input signals must swing between 0.925V and 1.475V which falls within the allowable input range of the ADA4937-1 operating on a single +5V supply.

The 65.5Ω input termination resistor in parallel with the 200Ω gain setting resistor makes the overall impedance approximately 50Ω. Note that a 226Ω resistor is inserted in series with the inverting input. This is to match the net impedance seen by the noninverting input (200Ω + 65.5Ω||50Ω ≈ 226Ω).

The output noise voltage spectral density of the ADA4937-1 is only 5nV/√Hz. This value includes the contributions of the feedback and gain resistors and is for G = 1. Integrated over the input bandwidth of the AD6645 (270MHz), this yields an output noise of 103µV rms. This corresponds to an SNR of 77.6dB due to the amplifier. Note that the integration must be over the full input bandwidth of the ADC since there is no external noise filter.

The SNR of the AD6645 is 75dB which corresponds to an input noise of 138µV rms. The combined noise due to the op amp (103µV) and the ADC (138µV) is 172µV, yielding an overall SNR of 73dB.

If the full bandwidth of the AD6645 is not required, a single-pole noise reduction filter can be added by selecting an appropriate value for C.

ADA4938-1 (0V, 10V) Driving AD9446 in DC-Coupled Application

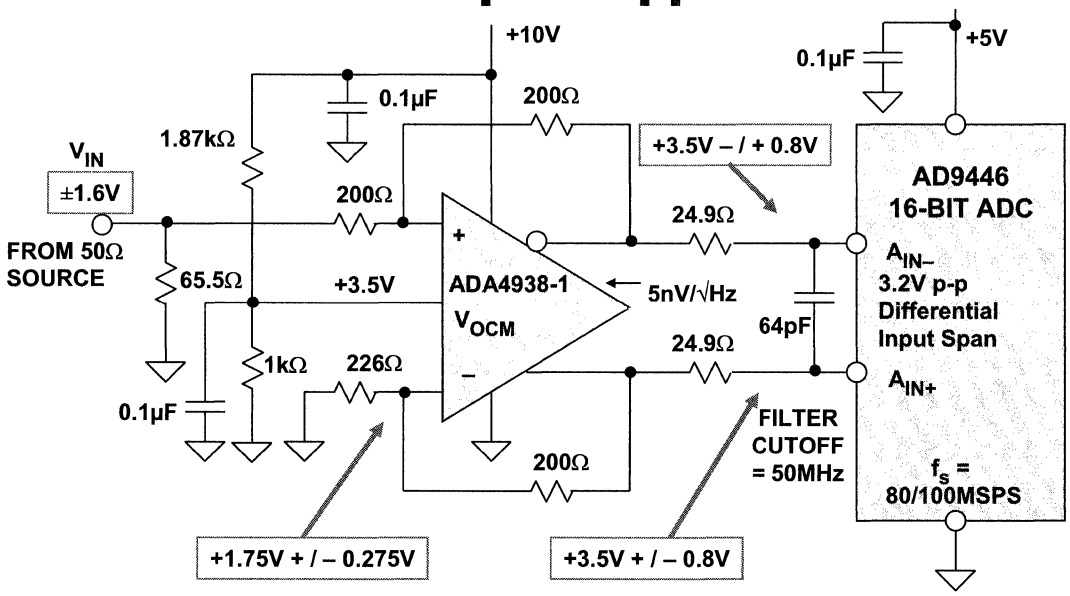

$$\text{OUTPUT NOISE} = 5\text{nV}/\sqrt{\text{Hz}}\sqrt{1.57\times50\times10^{6}} = 44\mu\text{V rms}$$

$$\text{OUTPUT SNR} = 20\log\frac{1.13}{44\times10^{-6}} = 88.2\text{dB}$$

This circuit is very similar to the previous figure, but the AD9446 ADC requires an input voltage of +3.5V ±0.8V on each differential input. This means that each output of the differential amplifier must swing between +2.7V and +4.3V. The +4.3V requirement is outside the capability of the AD4937-1 driver operating on a +5V supply, so the ADA4938-1 driver operating on a +10V supply must be used.

The noise filter has a cutoff of 50MHz. The output noise of the driver (5nV/√Hz) integrated over this bandwidth is 44µV rms. This yields an SNR of 88.2dB for the driver which is 8.2dB better than the SNR of the AD9446.

Note that since the drive amplifier operates on +10V and the ADC on +5V, care must be taken that the amplifier does not overdrive the ADC input and cause damage. Power supply sequencing can also be an issue if power is applied to the amplifier before power is applied to the ADC. Suitable protection circuitry may therefore be required.

ADA4938-1 (0V,+10V) Driving AD9445 in DC-Coupled Application

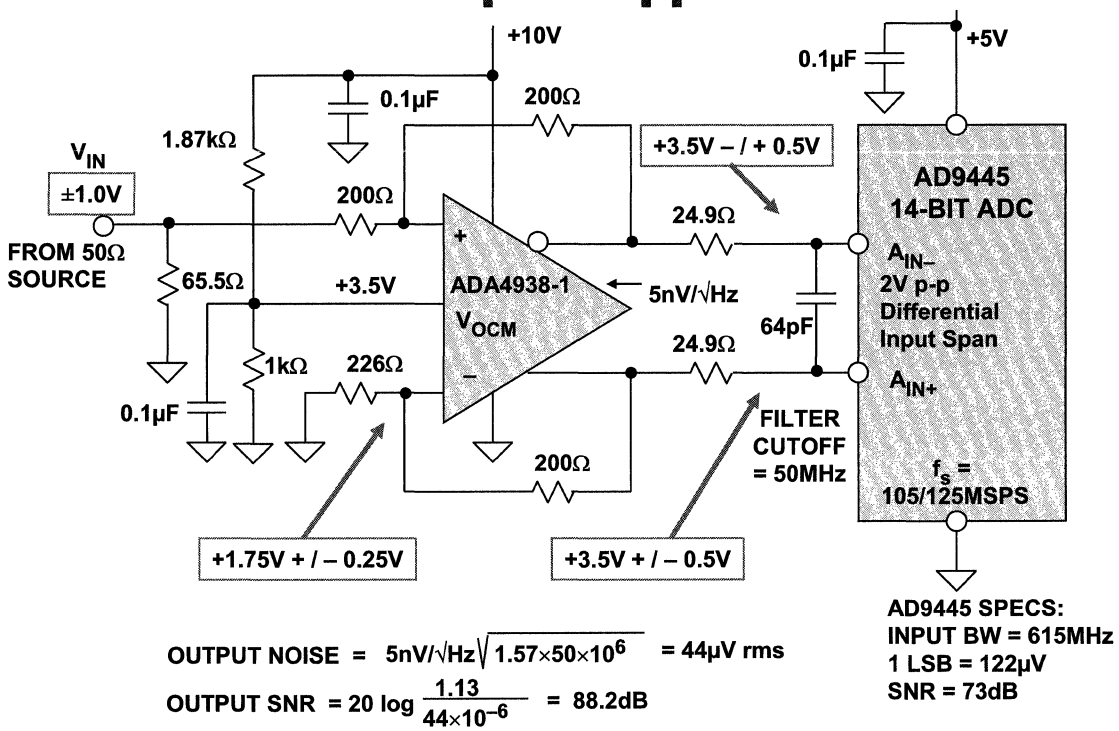

$$\text{OUTPUT NOISE} = 5\text{nV}/\sqrt{\text{Hz}}\sqrt{1.57\times50\times10^{6}} = 44\mu\text{V rms}$$

$$\text{OUTPUT SNR} = 20\log\frac{1.13}{44\times10^{-6}} = 88.2\text{dB}$$

This figure shows the ADA4938-1 driving the AD9445 14-bit 105/125MSPS ADC. The circuit is very similar to the previous circuit with the exception of signal amplitude. Again, the ADA4938-1 amplifier operating on a +10V supply is required because each input of the AD9445 must be driven to +4V, which would not be possible with a drive amplifier operating on a +5V supply.

As in the previous example, suitable precautions against ADC overdrive must be taken since the amplifier operates on a +10V supply and the ADC on a +5V supply.

ADA4938-1 (±5V) Driving AD9246 1.8V ADC in DC-Coupled Application

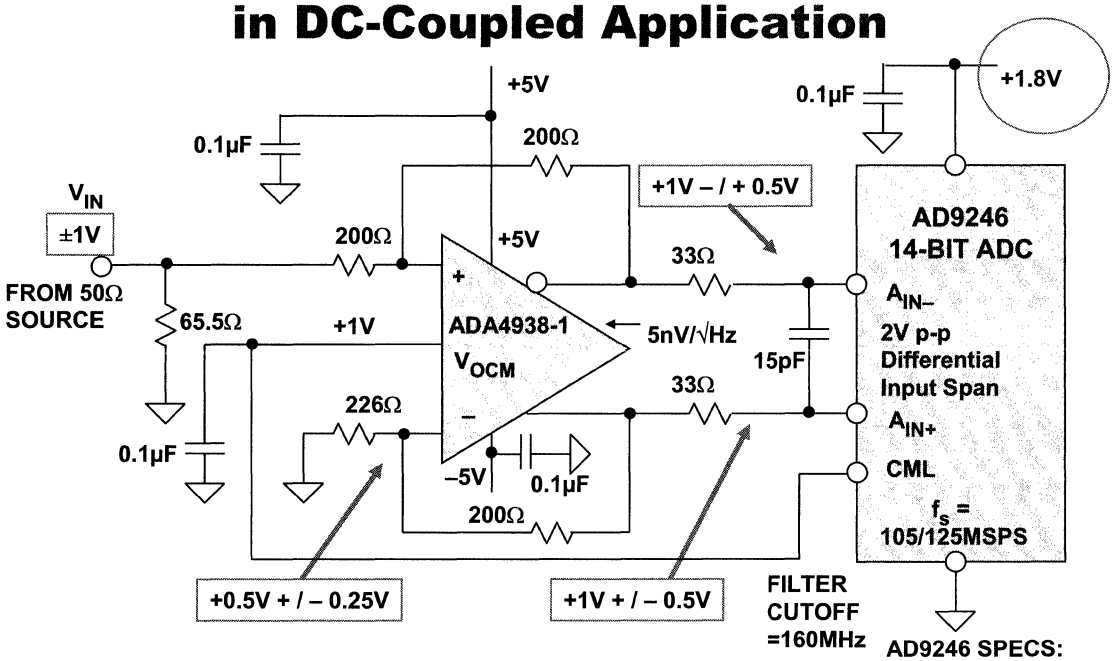

The AD9246 is a low power (395mW) 14-bit, 105/125MSPS ADC which operates on an analog supply of +1.8V.

The input signal to the AD9246 is 2V p-p differential with a common-mode voltage of +1V; therefore, each output of the drive amplifier must swing between +0.5V and +1.5V. This requires a dual supply differential driver such as the ADA4938-1 operating on ±5V supplies as shown in the figure.

As in the previous circuits, suitable precautions must be taken against ADC overdrive because the drive amplifier operates on ±5V supplies, while the ADC operates on a single +1.8V supply.

Equivalent Input Circuit Models for Buffered (BiCMOS) and Unbuffered (CMOS) Pipelined ADCs

Buffered and Unbuffered Differential ADC Inputs Structures

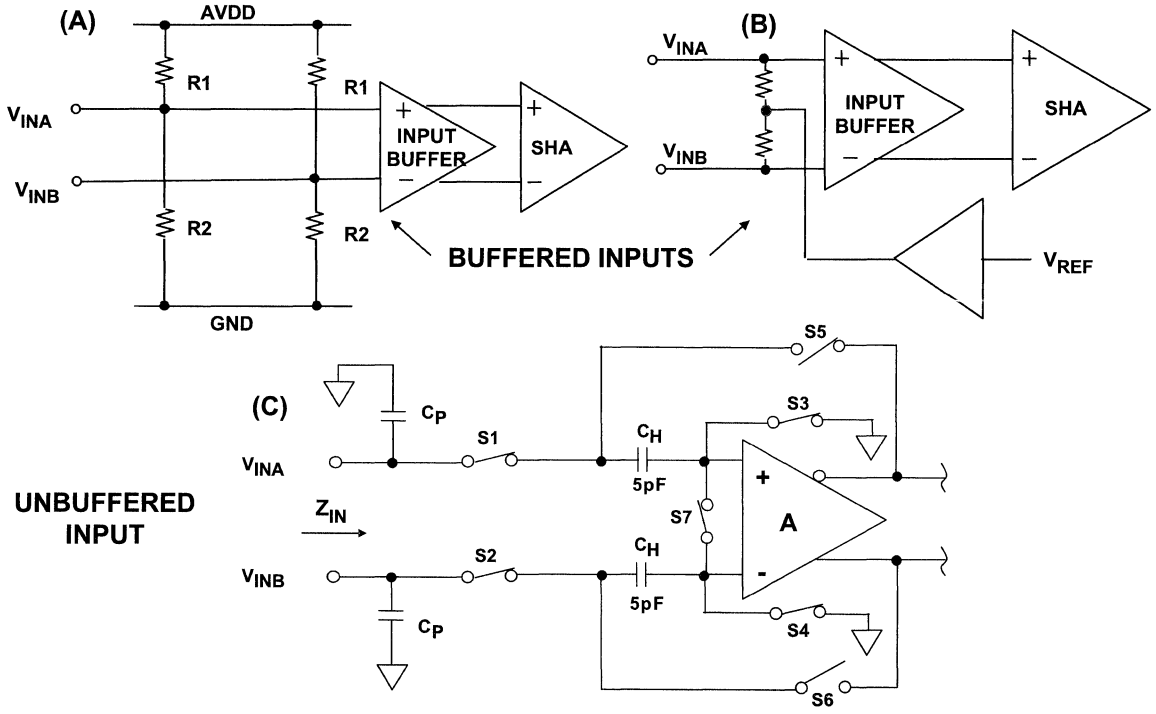

In designing the input interface circuit, it is important to know whether or not there is an input buffer present to isolate the input from the sample-and-hold circuit.

High performance pipelined BiCMOS ADCs (such as the AD6645, AD9445, and AD9446 generally have the input buffer on-chip. Two popular input structures for buffered input BiCMOS ADCs are shown in (A) and (B) in the figure.

CMOS pipelined ADCs typically dissipate less power and have slightly lower performance than BiCMOS ones, and generally don't have the input buffer, and as shown in (C), the input is connected directly to the sample-and-hold switches. The unbuffered input structure typically generates more transient currents than the buffered structure and is more difficult to drive.

When choosing an ADC, it is important to know if the input structure is buffered or unbuffered. In many cases, the input buffer is actually shown in the functional block diagram of the ADC which appears on the first page of the data sheet. In some cases, the input structure can be determined from the applications section of the data sheet where the description of the converter's operation is contained. As mentioned above, most CMOS ADCs have an unbuffered input, while BiCMOS ADCs have a buffered input—however, there can be exceptions.

Input Impedance Model for Buffered and Unbuffered Input ADCs

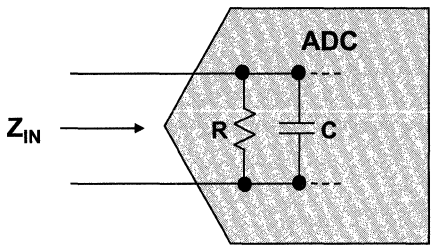

BUFFERED INPUT

- **R and C are constant over frequency**
- **Typically:**

  R: 1kΩ - 2kΩ
  C: 1.5pF - 3pF

UNBUFFERED INPUT

- **R and C vary with both frequency and mode (track/hold)**
- **Use Track mode R and C at the input frequency of interest**

We can model the input impedance of both the buffered and unbuffered input structures as a resistance in parallel with a capacitance. In the case of the buffered input ADC, the R and C values are constant over input frequency. Typical values of R range from 1kΩ to 2kΩ, and typical values of C range from 1.5pF to 3pF, depending on the particular part. However, the data sheet for the ADC should always be consulted because there are some exceptions.

The unbuffered input structure is much more difficult to model because the R and C values change dynamically with both analog input frequency and whether the ADC is in the track mode or the hold mode.

For purposes of modeling the unbuffered input, it is the track mode impedance that is of most interest in designing the interface.

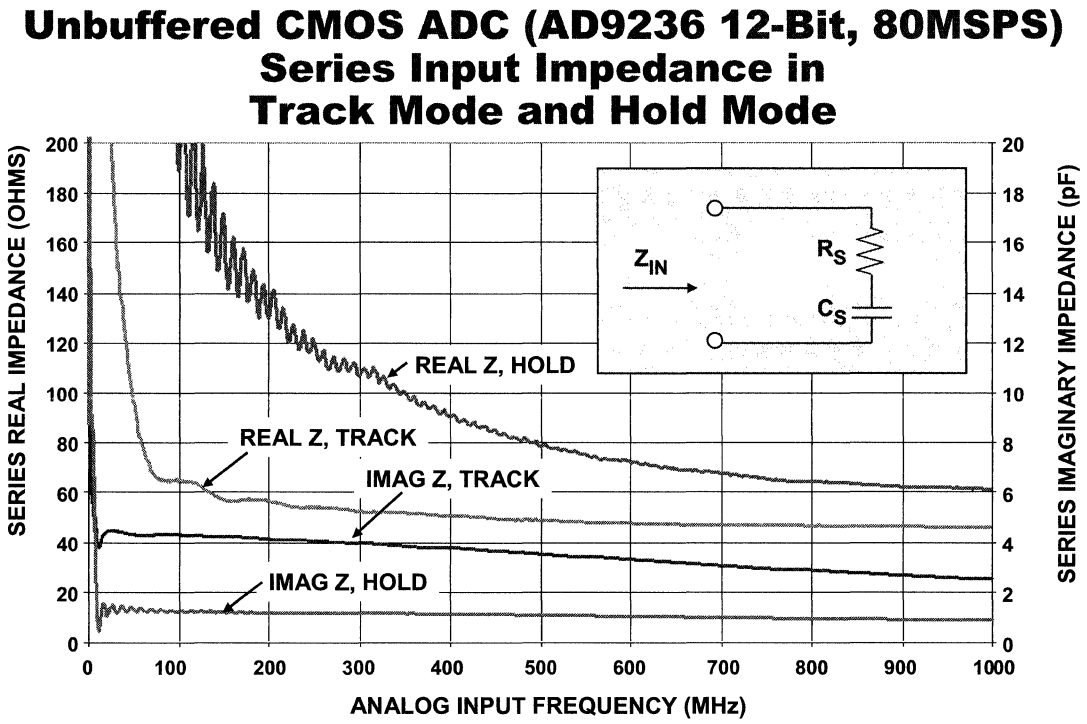

A sampling network analyzer can be used to measure the track mode and hold mode input impedance as a function of analog input frequency. See reference below.

This figure shows the real and imaginary part of the input impedance for the AD9236 12-bit, 80MSPS CMOS ADC, which is typical of other members of the family.

The figure shows the series mode input impedance, however, the parallel mode input impedance is of more interest.

Rob Reeder, "Frequency Domain Response of Switched-Capacitor ADCs," Application Note *AN-742*, Analog Devices., www.analog.com/DesignCenter.

Converting Between Series and Parallel Equivalent Circuits

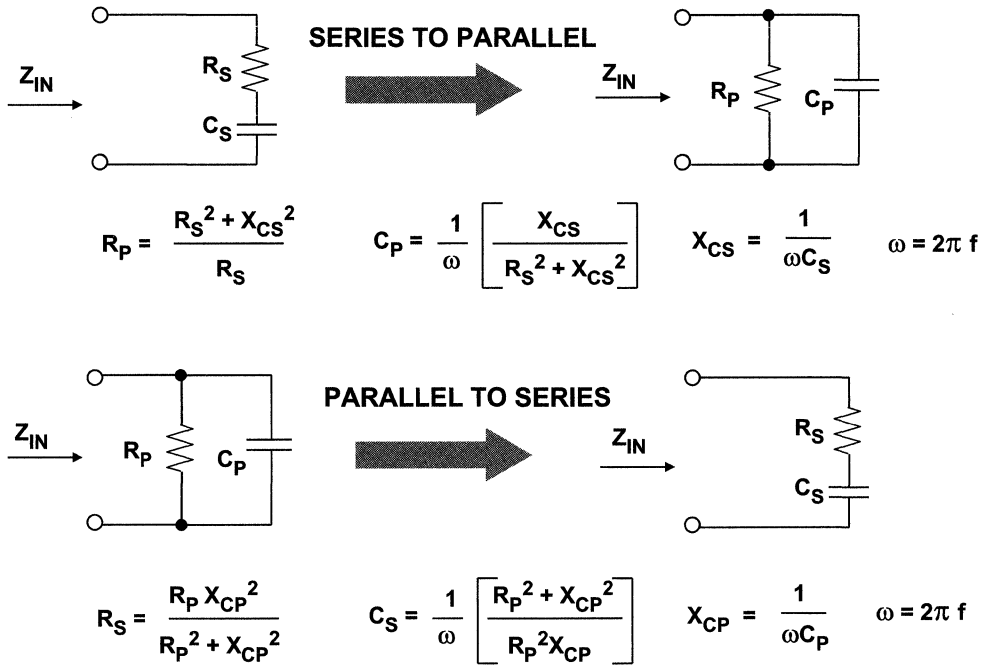

The formulas shown in this figure can be used to convert between the series and parallel equivalent circuits if required.

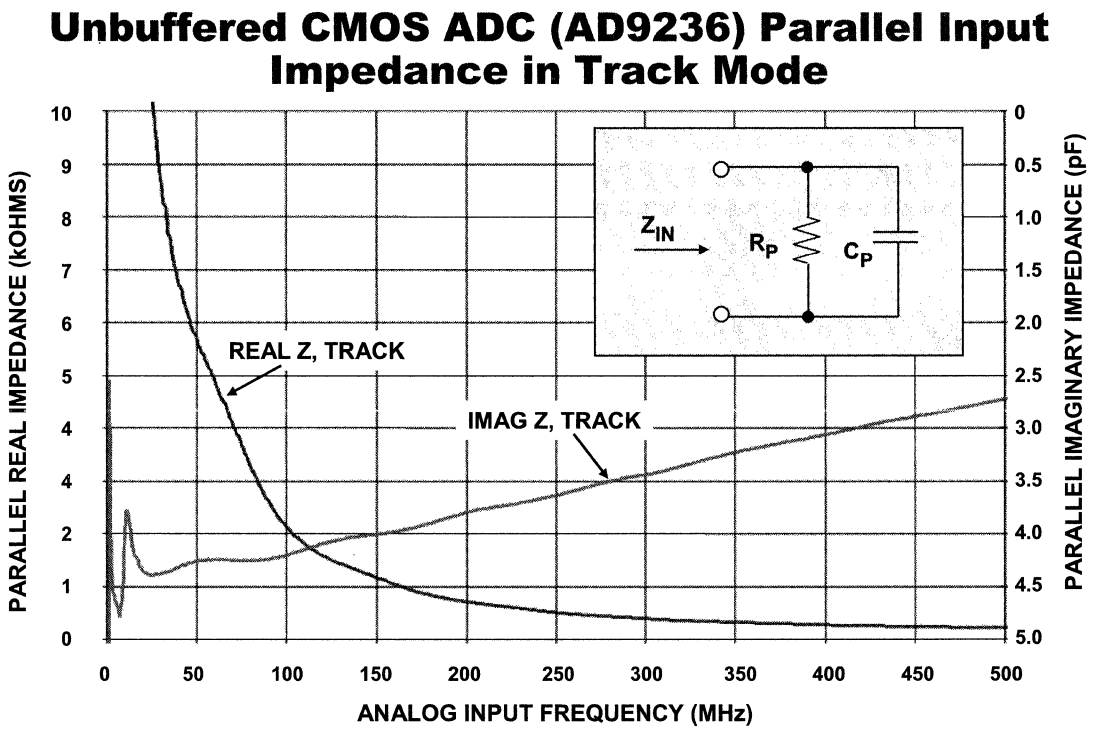

This shows only the parallel equivalent circuit for the track mode impedance as a function of analog input frequency. In general, parallel equivalent circuit is more useful.

The real (resistive) part of the input impedance is very high at lower frequencies (baseband) and to less than 2kΩ above 100MHz.

The imaginary (capacitive) part of the input impedance starts out at approximately 4.5pF at low frequencies and drops gradually as the frequency is increased.

Real and imaginary impedances are available as S-parameter in spreadsheet format on the Analog Devices' website in the evaluation board section of the product information.

Basic Principles of Resonant Matching

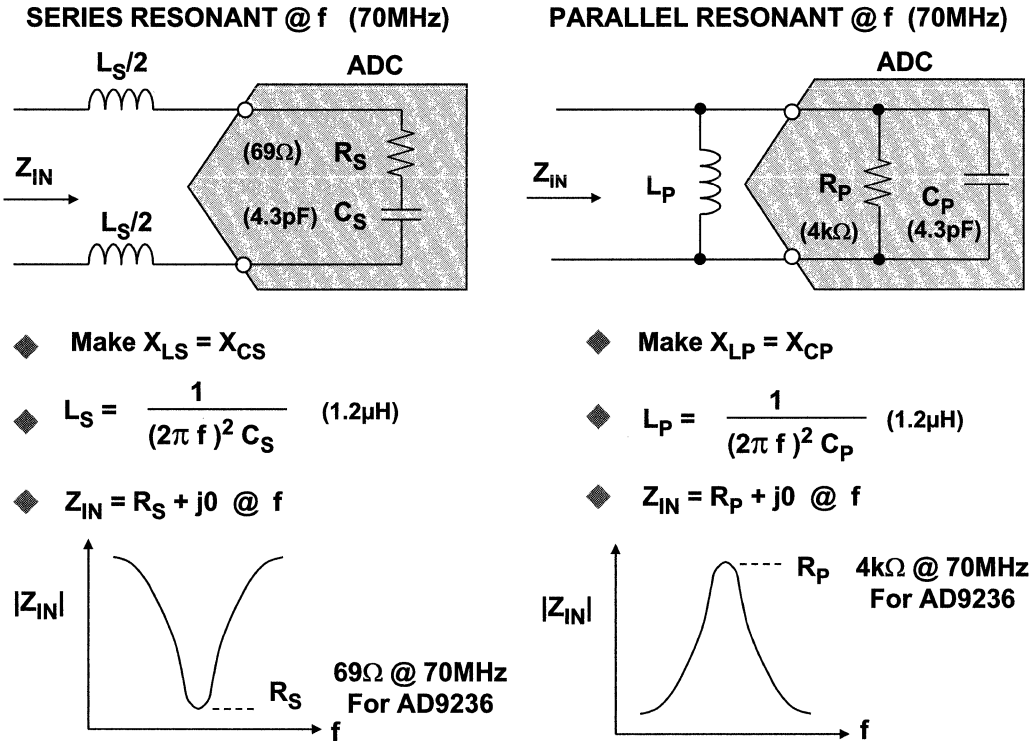

Now that we have a method of determining the impedance of the input of the ADCs, it is useful to consider the principle of resonant matching when the input signal is an IF signal with a limited bandwidth.

Resonant matching is achieved by simply adding the appropriate external series or parallel inductance to cancel the effects of the series or parallel capacitance of the input circuit of the ADC. This presents a resistive load to the driver at the IF frequency of interest.

The first diagram shows the series approach. At a 70MHz IF frequency, the reactance of the 1.2μH inductor matches that of the 4.3pF capacitor, and the net impedance Z_{IN} is resistive and equal to 69Ω. Note that the inductor value in each leg is equal to the calculated value divided by two.

The second diagram shows the parallel equivalent circuit for the same ADC. At the resonant IF frequency (70MHz), the input impedance is resistive and equal to 4kΩ.

The parallel resonant approach is preferable because at resonance it yields a much higher input impedance than the series approach. The high impedance is easier to drive with low distortion amplifiers. The reference below describes the technique of resonant matching in more detail.

Chris Bowick, *RF Circuit Design*, Newnes/Elsevier, 1982, ISBN 0-7506-9946-9.

Resonant Matched Design Example:

AD8370 Variable Gain Amplifier Driving AD9236 12-Bit, 80MSPS ADC with 70MHz IF

This design example is for a 70MHz IF frequency sampled at 76.8MHz (customer specified frequency). The sampling process downconverts the IF frequency to 6.8MHz. The AD9236 12-bit, 80MSPS ADC is driven with the AD8370 variable gain amplifier (VGA) followed by a 4-pole noise reduction lowpass filter.

The input impedance of the ADC at 70MHz is used to determine the value of the appropriate parallel resonant inductor, L_P. This transforms the input impedance of the ADC to a resistance at the IF frequency.

The lowpass filter is then designed based on the resistive source impedance of the amplifier, the equivalent resistive load of the ADC, and the desired filter transfer function.

The AD8370 is a programmable gain amplifier that has two ranges: −11dB to +17dB and +6dB to +34dB. The differential input impedance is 200Ω, and the differential output impedance is 100Ω. The AD8370 has 7dB noise figure at maximum gain, two-tone IP3 of +35dBm at 70MHz, 3dB bandwidth of 750MHz, and a 40dB precision gain range. It is controlled via a serial 8-bit digital interface and operates on a +3V to +5V power supply.

AD8370/AD9236 Matching and Antialias Filter Interface Design for 70MHz IF

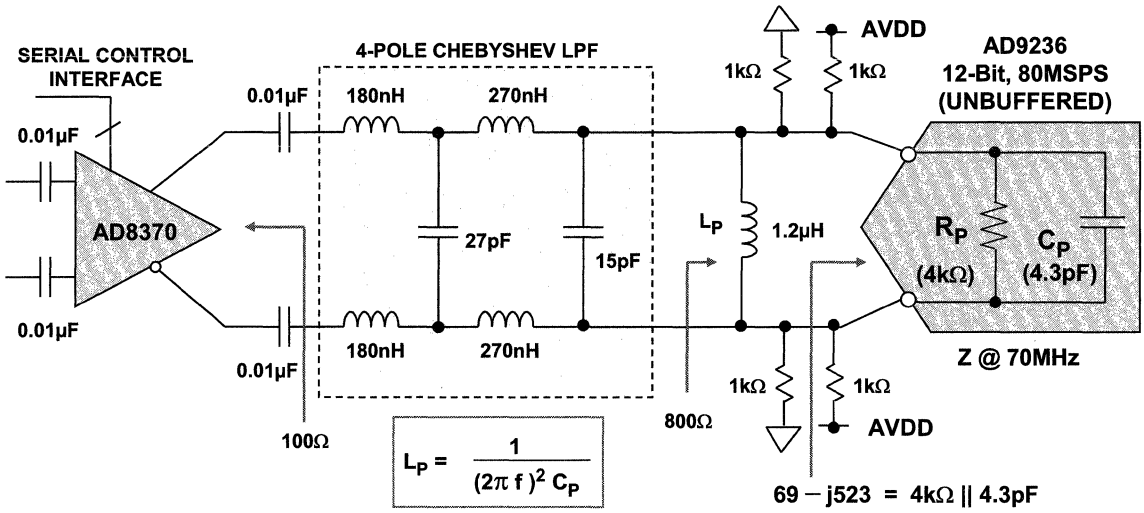

- **The resonant shunt inductor in combination with the 1kΩ bias resistors presents an ~ 800Ω differential load.**
- **The anti-aliasing filter is designed to be sourced from 100Ω and loaded into 800Ω and was optimized using a filter design program to use standard values and includes the effects of inductor parasitics.**

The final design for the 70MHz IF digitizer is shown in this figure. The first step is to determine the input impedance of the AD9236 at 70MHz. This data is available from an on-line spreadsheet and is 4kΩ in parallel with 4.3pF. The 1.2µH external parallel inductor resonates with the 4.3pF ADC capacitance at 70MHz and is calculated using the formula in the figure.

The four 1kΩ resistors are required to set the common-mode voltage for the differential inputs to the ADC. The four resistors form a 1kΩ differential resistance. The equivalent resistive load to the filter output is therefore 1kΩ||4kΩ = 800Ω.

The filter design is a fourth-order Chebyshev with 0.5dB ripple, 100Ω source, and 800Ω load. The cutoff frequency is set to 83MHz, and the filter is designed to give greater than 10dB attenuation at 100MHz. The values were adjusted to obtain standard value components and account for inductor parasitics. Other filter response types could be used here if desired, such as Butterworth.

Details of this type of resonant design using amplifier drivers can be found in Reference 1. The filters can be easily designed using a number of filter design programs as in Reference 2.

---

1. Eric Newman and Rob Reeder, "A Resonant Approach to Interfacing Amplifiers to Switch-Capacitor ADCs," *Application Note AN-827*, Analog Devices, www.analog.com.

2. Filter design programs such as Filter Lite 5.0 from Nuhertz Technologies, or Agilent Technologies Advanced Design System (ADS).

Simulated Response of Interface

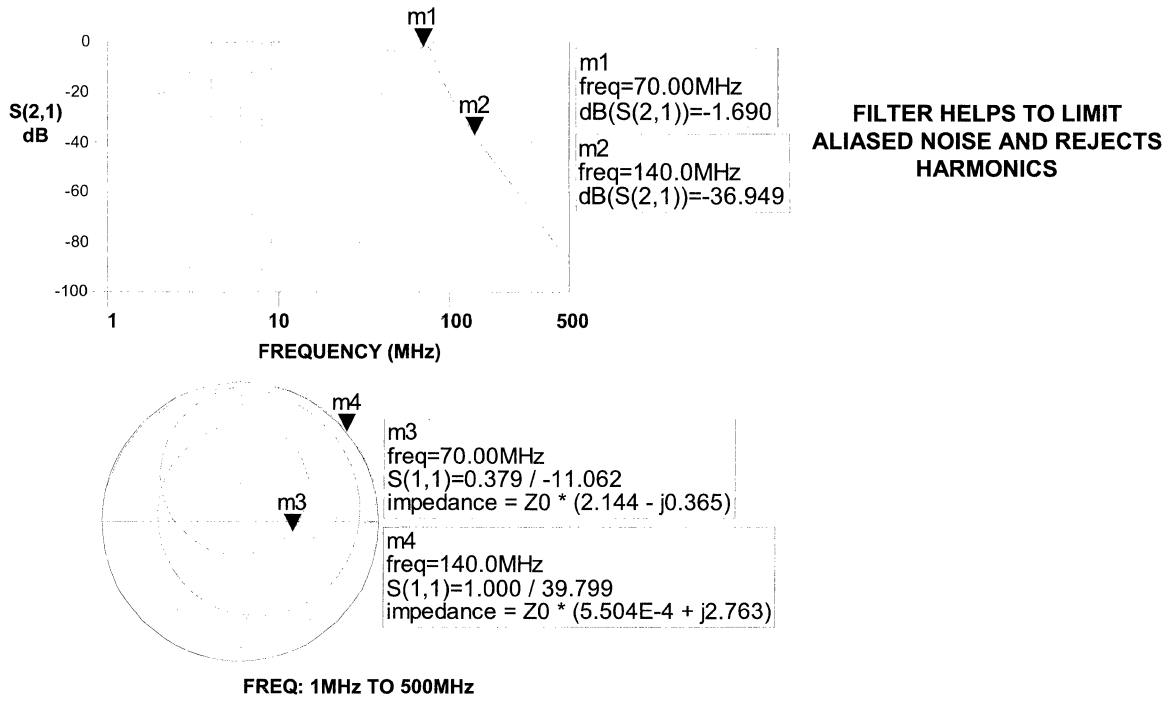

This figure shows the simulated response of the entire interface, including the resonant inductor and the ADC input impedance. Also shown is a Smith chart of the input impedance to the filter.

Before and After Adding Matching Analog Antialiasing Filter Network

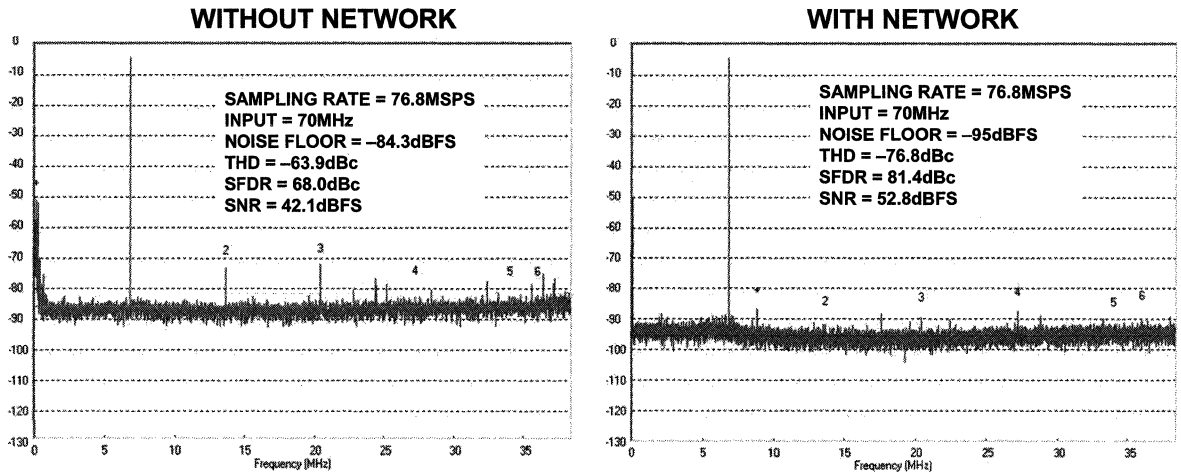

- **SFDR Improved by 13.4dB, SNR improved by 10.7dB**
- **Note: Measured at maximum gain of 35dB (gain code 255, high gain mode) using 76.8MHz sampling clock**

This figure shows that adding the matching filter network improves the SFDR of the system by 13.4dB and the SNR by 10.7dB.

The improvement in SFDR is primarily because the load presented to the driving amplifier at 70MHz IF is resistive due to the resonant matching and is also a higher impedance than would be the case without the matching circuit. (The 4pF input capacitance has a reactance of 568Ω at 70MHz).

The improvement in SNR is primarily due to the filter bandlimiting the broadband noise to 80MHz. Otherwise the amplifier output noise would be integrated over the full 500MHz input bandwidth of the ADC.

Wideband Design Example:

AD8352 Variable Gain Amplifier Driving AD9445 14-Bit, 125MSPS Buffered Input ADC

In this example, we are digitizing a wideband signal with the AD9445 14-bit, 125MSPS ADC and desire to preserve as much of the ADC input bandwidth as possible.

The AD9445 has 615MHz input bandwidth and an SFDR of 95dBc for a 100MHz input.

For the driver, we have chosen the AD8352 2GHz bandwidth differential amplifier because it has a resistor programmable gain range of 3db to 21dB. The amplifier also has low noise (2.7nV/$\sqrt{Hz}$ referred to the input for a gain setting of 10dB) and low distortion (82dBc HD3 at 100MHz).

The lower end of the bandwidth requirement is approximately 10MHz.

AD8352 2GHz Differential Amplifier Driving AD9445 14-Bit, 125MSPS ADC

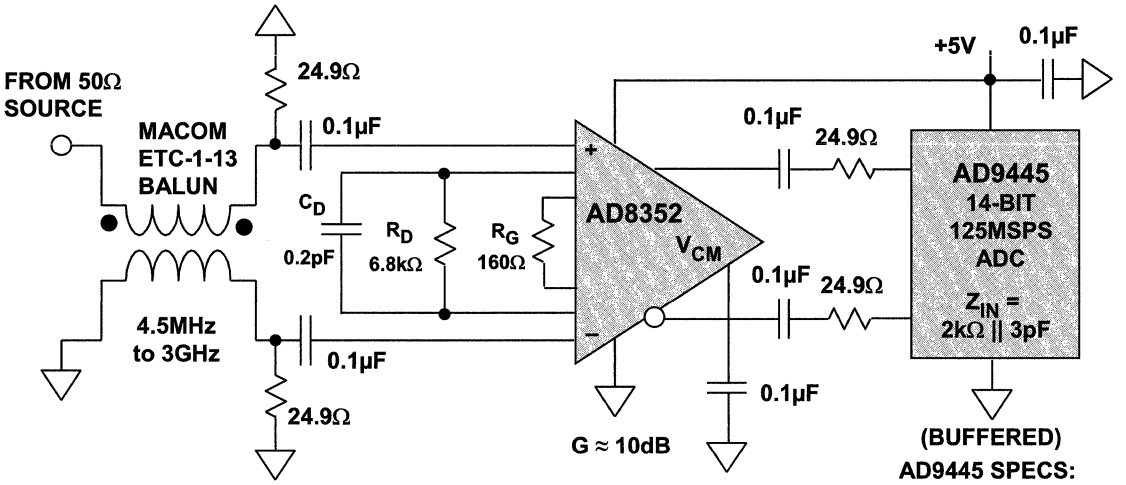

- OUTPUT NOISE OF AD8352 FOR 10dB GAIN = 8.5nV/√Hz
- INTEGRATED OVER 615MHz INPUT BW OF AD9445 = 264µV rms
- INPUT REFERRED NOISE OF AD9445 = 158µV rms
- TOTAL NOISE = 307µV rms
- SNR = 67dB FOR 2V P-P INPUT

AD9445 SPECS:
2V p-p FS Diff.
INPUT BW = 615MHz
1 LSB = 122µV
SNR = 73dB

This shows the optimum circuit configuration for driving the AD9445 with the 2GHz AD8352 in a wideband application. The balun converts the single-ended input to differential to drive the AD8352. Although it is possible to configure the AD8352 to accept a single-ended input (see AD8352 data sheet), optimum distortion performance is obtained if it is driven differentially as shown.

The C_D/R_D network is chosen to optimize the third-order intermodulation performance of the AD8352. The values are selected based on the desired gain and are given in the data sheet.

The performance of the circuit is shown in the next figure. Note that SFDR is 83dBc for a 98.9MHz input signal sampled at 105MSPS.

Noise performance can be predicted as follows:

The output noise spectral density of the AD8352 for G = 10 is 8.5nV/√Hz. Since there is no input filter, this must be integrated over the entire input bandwidth of the AD9445, 615MHz:

$$V_{NAMP} = 8.5nV/Hz\sqrt{(1.57\times615\times10^{6})} = 264\mu V \text{ rms.}$$

The input noise of the ADC is calculated from the SNR or 73dB:

$$V_{NADC} = 0.707\div10^{SNR/20} = 158\mu V \text{ rms.}$$

The total noise is calculated by:

$$V_{TOTAL} = \sqrt{(V_{NAMP}^{2} + V_{NADC}^{2})} = 307\mu V \text{ rms.}$$

This results in a combined system SNR of:

$$SNR = 20\log(0.707\div307\times10^{-6}) = 67dB$$

FFT Data for AD8352 Driving AD9445
Input = 98.9MHz, Sampling Rate = 105MSPS

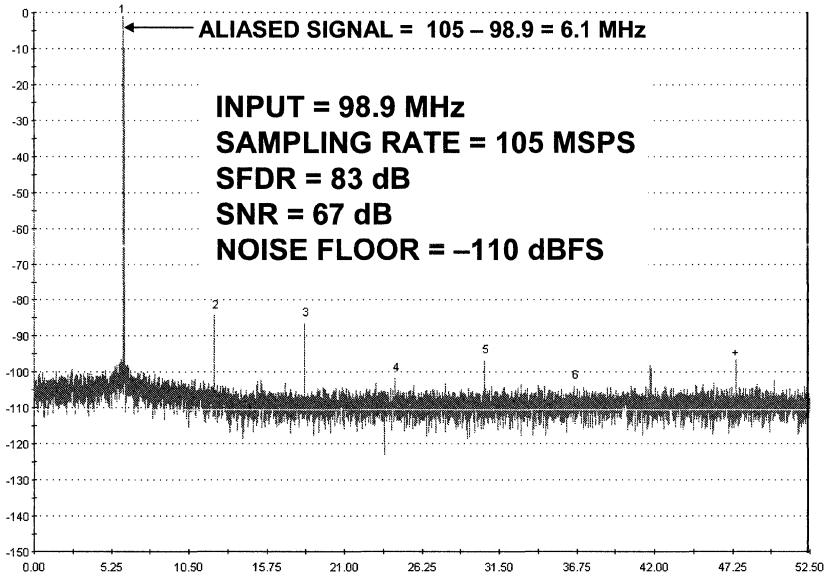

This figure shows the FFT output for the AD8352 driving the AD9445 with an input signal of 98.9MHz and a sample rate of 105MSPS. The circuit is identical to the one in the previous figure.

SFDR is 83dBc, and SNR is 67dB, representing exceptional performance for this IF input frequency.

Transformer Drivers

Transformers are popular ADC drivers in IF/RF applications. However, they are somewhat more difficult than amplifiers to analyze and apply as ADC drivers, and cannot pass dc or low frequency signals.

Transformers do not add additional noise to the system, but cannot generally be used for voltage gains more than two. In addition, they are somewhat difficult to analyze, and manufacturers do not typically provide detailed models on the data sheets.

The entire load on the transformer's secondary winding (including a noise filter if used) is reflected back to the primary winding. The resulting input load is therefore more difficult to control than that of an amplifier driver because output load of an amplifier is well isolated from its input.

The transformer is used as a single-ended to differential converter in many applications. At frequencies above about 100MHz, the parasitic capacitance between the primary and secondary windings can cause phase and amplitude unbalance which, in turn, can increase the distortion of the ADC. This may be remedied by either using a two-transformer configuration or using a higher performance transformer.

Transformer Coupling into the AD9446 16-Bit, 80/100MSPS BiCMOS Buffered Input ADC, Baseband Signal

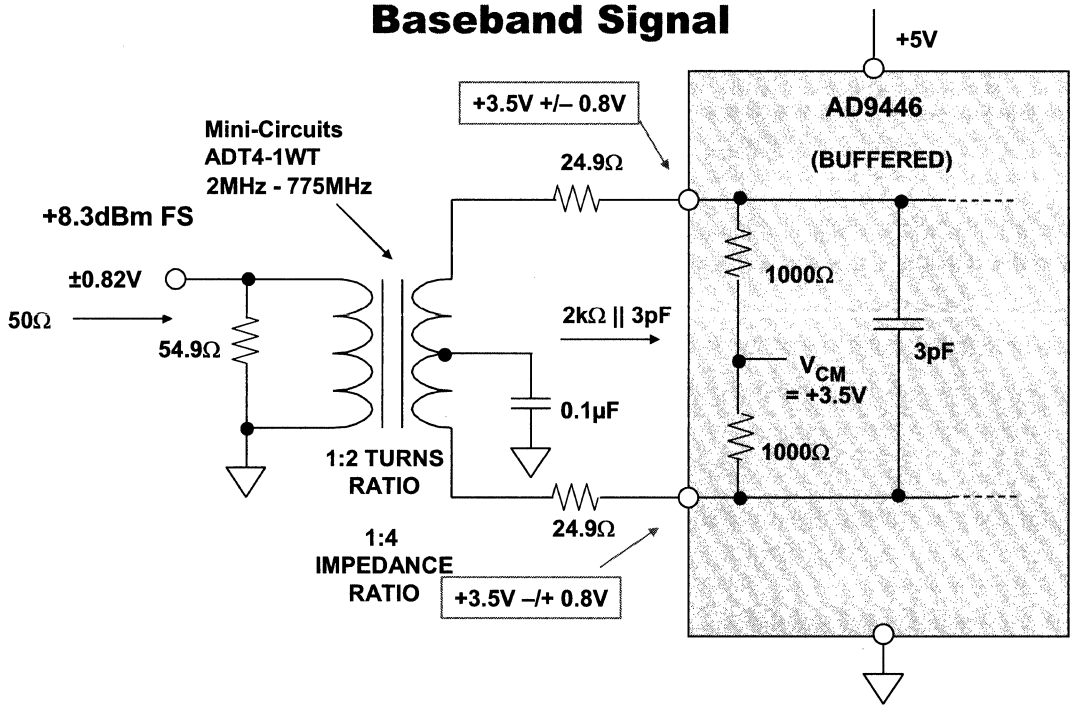

This figure shows a typical baseband transformer driver used as a single-ended to differential converter. The AD9446 16-bit, 80/100MSPS ADC is fabricated on a BiCMOS process, and has buffered inputs. Therefore, input transient currents are minimal.

The 24.9Ω resistors in series with the input isolate the transformer from the input capacitance of the ADC. The common-mode voltage of +3.5V is generated internally in the ADC.

This circuit uses a 1:2 turns ratio transformer for voltage gain. The total resistive load seen by the secondary of the transformer is 2000Ω + 24.9Ω + 24.9Ω = 2050Ω. This resistance is divided by 4 to reflect it to the primary winding as 512Ω. Therefore, in order for the input of the transformer to be a 50Ω termination for the signal, a 54.9Ω resistor is added in parallel with the equivalent 512Ω resistance.

For additional information regarding optimizing a transformer-coupled ADC input interface, refer to the following reference:

Rob Reeder,"Transformer-Coupled Front-End for Wideband A/D Converters," *Analog Dialogue*, Vol. 39, Number 2, 2005, www.analog.com.

Differential Amplifiers vs. RF Transformer/Balun Drivers

DIFFERENTIAL AMPLIFIERS	RF TRANSFORMERS/BALUNS
◆ Allow dc coupling, gain, offset adj.	◆ Only work in ac-coupled apps.
◆ Provide single-ended to differential conversion	◆ Provide single-ended to differential conversion
◆ Add noise	◆ No additional noise added
◆ Add distortion	◆ Add less distortion
◆ I/O Impedances well defined	◆ I/O Impedance analysis difficult
◆ Input isolated from output	◆ Input and output interact
◆ Add power to system	◆ No additional power added
◆ Can provide voltage gain	◆ Gain > 2 difficult at IF frequencies
◆ Good for baseband and medium IF frequency (up to 100MHz)	◆ Good for baseband, medium, and high IF frequency

This figure compares differential amplifier drivers with transformer, or balun drivers. Note that a *balun* is a transformer with a bifilar winding and stands for "balanced-unbalanced." Baluns typically have higher bandwidth and lower parasitics than traditional RF transformers.

Transformers become difficult to use in low distortion applications requiring voltage gains greater than 2. On the other hand, differential amplifiers are capable of providing 10dB to 30dB voltage gain, depending on the device. The downside of differential amplifiers, of course, is the additional noise they add to the system.

Transformer Insersion Loss and Return Loss

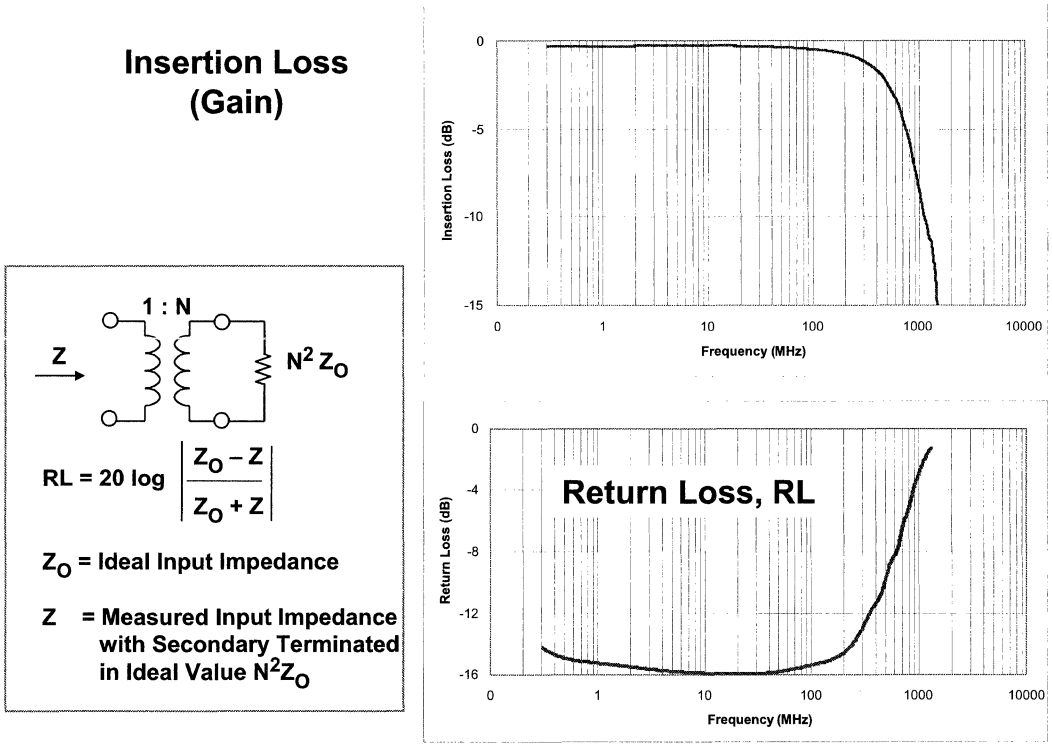

A transformer can be viewed simplistically as a bandpass filter. The transformer *insertion loss* is essentially a bandwidth specification, but it should not be the only consideration when designing with a transformer.

Return Loss is the effective impedance as seen by the primary when the secondary is terminated. For example, if you have an ideal 1:2 turns ratio (1:4 impedance ratio), transformer you would expect a 50Ω impedance reflected to the primary when the secondary is terminated with 200Ω. However, this is not always true. The impedance reflected to the primary varies with frequency. In general, as the turns ratio increases, so does the variability of the return loss. An example is shown on the next page.

The figure shows how to calculate the return loss. The secondary is terminated in a resistor equal to N^2Z_O, where N is the turns ratio, and Z_O is the ideal transformer impedance. The actual input impedance is then measured, and this value is Z. The return loss, RL, is then calculated from the formula $RL = 20\log|[(Z_O - Z) / (Z_O + Z)]|$. The bottom graph in the figure shows the return loss plotted as a function of input frequency for a typical RF transformer.

It is important to know the return loss at the IF frequency of interest so that adjustments can be made in the terminations to make the actual input impedance reflected to the primary the correct value. This ensures a good match to the source and minimizes reflections due to impedance mismatch.

Data for Mini-Circuits TC1-1-13M Balun

Data Used by Permission of Mini-Circuits,
P.O. Box 350166, Brooklyn, New York 11235-0003
http://www.minicircuits.com

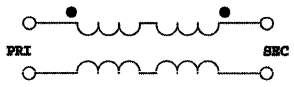

Transformer Electrical Specifications (T_{AMB}=25°C)

Ω RATIO	FREQUENCY (MHz)	INSERTION LOSS*			PHASE UNBALANCE (Deg.) Typ.		AMPLITUDE UNBALANCE (dB) Typ.	
		3 dB MHz	2 dB MHz	1 dB MHz	1 dB bandwidth	2 dB bandwidth	1 dB bandwidth	2 dB bandwidth
1	4.5-3000	2000-3000	1000-2000	4.5-1000	2	3	0.5	0.5

* Insertion Loss is referenced to mid-band loss, 0.5 dB typ.

Typical Performance Data

FREQUENCY (MHz)	INSERTION LOSS (dB)	INPUT R. LOSS (dB)	AMPLITUDE UNBALANCE (dB)	PHASE UNBALANCE (Deg.)
4.50	0.18	31.52	0.69	176.19
10.00	0.18	34.60	0.56	178.22
50.00	0.19	33.50	0.56	180.11
100.00	0.24	29.68	0.55	179.81
500.00	0.46	19.52	0.45	179.19
1000.00	0.68	16.22	0.14	178.41
1500.00	0.90	15.89	0.29	179.11
2000.00	1.11	16.97	0.71	181.28
2500.00	1.62	12.88	0.78	185.79
3000.00	3.02	6.79	0.49	167.68

Turns ratio, insertion loss (gain), and return loss are three commonly specified transformer parameters as shown in this figure. Amplitude and phase unbalance may also be specified.

However, the parasitic inductances and capacitances associated with the transformer are rarely specified on the manufacturer's data sheet. In some cases these parameters can be obtained by contacting the manufacturer directly. Otherwise, estimates or actual measured values can be used. In most cases some optimization in the actual circuit is required regardless of the simulation results.

Baseband Sampling Applications for Buffered Input ADCs

- The input interface for buffered ADCs is fairly simple to design
- When IFs are ≥ 100MHz, 2nd-order distortions begin to rise because the of transformer's amplitude and phase imbalance
- To solve this problem two transformers may be needed (see next page)

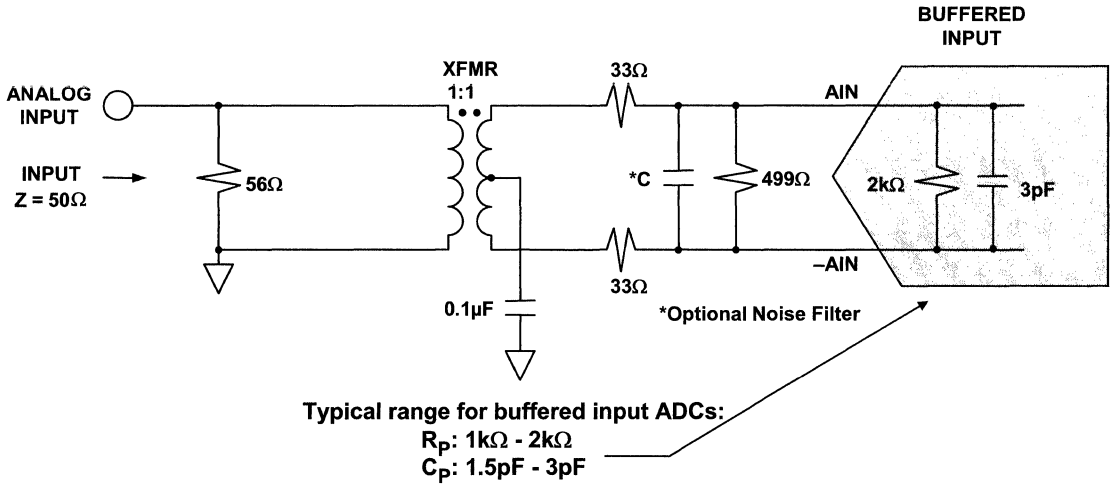

This figure shows a generic buffered input ADC driven by a transformer in a baseband application, where the input signal has a bandwidth up to $f_S/2$. The buffered ADC input impedance is 2kΩ in parallel with 3pF. The secondary termination is split between the 499Ω parallel resistor and the two 33Ω series resistors. The 33Ω series resistors isolate the secondary winding from the ADC input capacitance and the 499Ω resistor reduces the effects of the 3pF input ADC input capacitance.

The 499Ω resistor in parallel with the 2kΩ ADC resistor forms a 400Ω equivalent resistor. This makes the net secondary resistive load equal to 33Ω + 33Ω + 400Ω = 466Ω. This is reflected to the primary as 466Ω, since the transformer is 1:1. The 56Ω input resistor in parallel with 466Ω gives a net input termination resistance of 50Ω.

Most buffered input ADCs provide an internal dc common-mode bias voltage on the two inputs, so there is no need for an external bias network, as in the case of unbuffered input ADCs.

The center tap of the secondary winding is decoupled to ground to ensure that the input to the ADC is balanced.

A small capacitor, C, can be added to filter high frequency noise if desired.

Double Transformers/Baluns May Improve Performance of Buffered ADCs for IF Frequencies > 100MHz

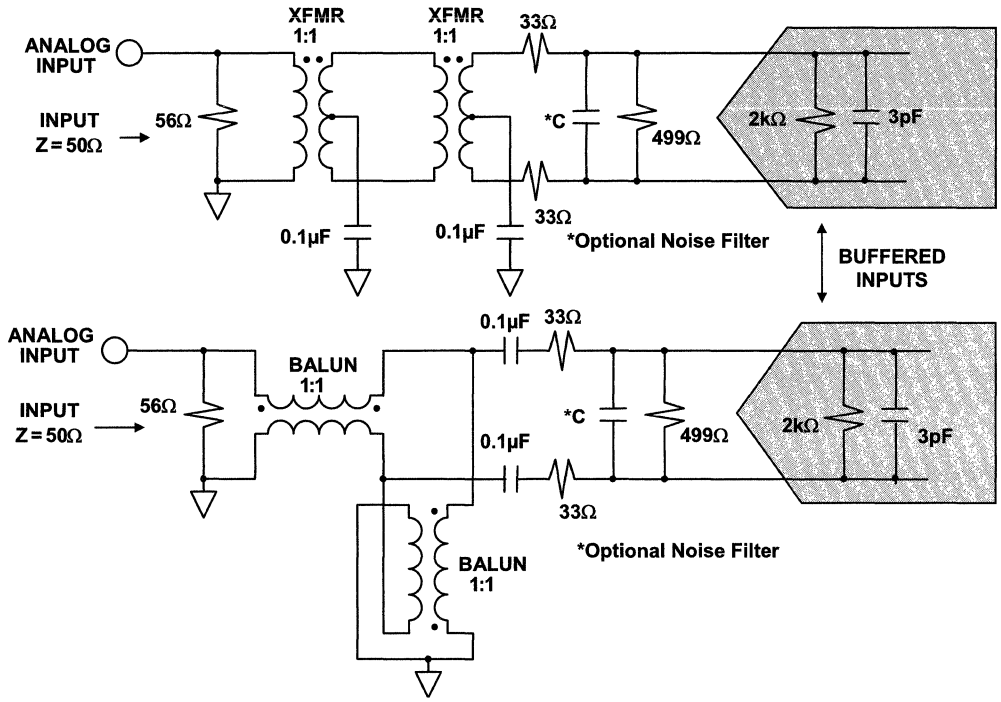

For IF frequencies above about 100MHz, the parasitic capacitance between the primary and secondary windings of the transformer may cause enough amplitude and phase unbalance to affect second-order distortion performance. The major reason for this unbalance is because one side of the primary winding is grounded and has no signal while the other side of the winding is driven by the signal and can couple into the secondary through the parasitic capacitance.

This problem can be addressed in two ways. One solution is to select a higher performance transformer (at additional cost) with better phase and amplitude balance.

Another solution is to add a second transformer (or balun) as shown in the figure. The second transformer serves to "distribute" the unbalance between the two transformers, thereby reducing the overall unbalance.

Regardless of the approach selected, some experimentation is required in order to achieve the optimum performance.

Details of the double transformer approach can be found in the reference.

Rob Reeder and Ramya Ramachandran, "Wideband A/D Converter Front-End Design Considerations – When to Use a Double Transformer Configuration," *Analog Dialogue* Vol. 40, Number 3, 2006, Analog Devices, www.analog.com.

Double Transformer Improves 2nd Harmonic Distortion by 10.5dB at 290MHz IF

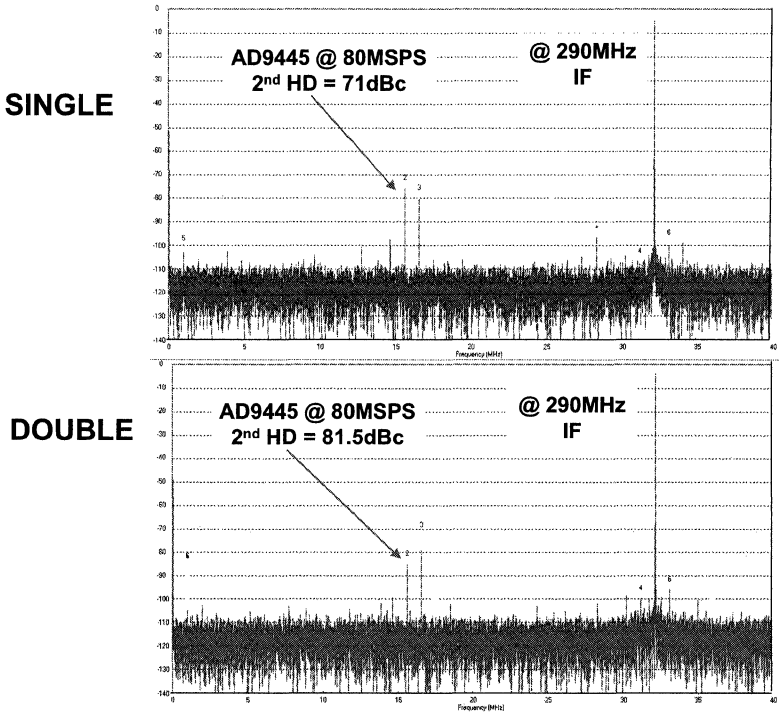

Here is an example using the AD9445-125MSPS ADC sampling at 80MSPS with an input frequency of 290MHz. Note the second harmonic is 71dBc. This data was gathered using a single 1:1 impedance ratio transformer on the front end.

Adding a second transformer reduces the second harmonic to 81.5dBc, an improvement of 10.5dB.

Switched-Capacitor ADC Input Configurations for Wideband Signals

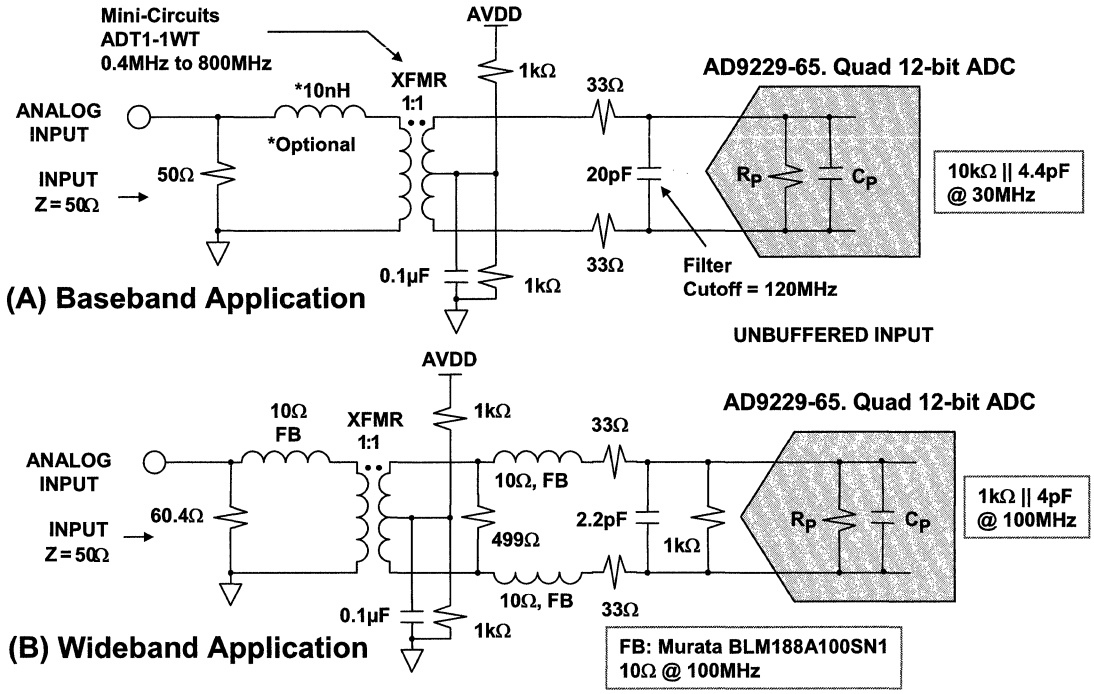

This figure shows two examples of unbuffered ADC input configurations: baseband and wideband.

In baseband applications (A), the analog input frequency to the ADC is less than $f_S/2$, and the resistive component of the ADC input impedance is high. This makes the input circuit easier to design. The two 33Ω resistors isolate the transformer from the ADC input transient currents as well as form a simple noise filter with the 20pF differential capacitor. The cutoff frequency is approximately 120MHz.

For baseband applications an inductor in series with the transformer's primary can be used to alter the bandwidth response of the transformer by peaking the gain in the passband and providing a steeper rolloff outside the passband. The inductor has the effect of adding a pole in the transfer function. The value of inductance depends on the desired amount of peaking and bandwidth requirement. However, the designer should note that this peaking can be undesirable where flatness of response and well-behaved phase response are important criteria.

For wideband signals which extend well beyond $f_S/2$, the design becomes more critical. In (B) the secondary resistive termination is split between R_p, the 1kΩ resistor, the 499Ω resistor, and the 33Ω series resistors. The net secondary termination is 265Ω. The 60.4Ω input resistor in parallel with the 265Ω equivalent resistance yields the desired 50Ω input termination.

The two series ferrite beads have been added to minimize gain peaking at the higher IF input frequencies. The part selected has a resistance of 10Ω at 100MHz. Determining the optimum ferrite bead generally involves some empirical work.

Rob Reeder, "Transformer-Coupled Front-End for Wideband A/D Converters," *Analog Dialogue*, Vol. 39, Number 2, 2005, Analog Devices, www.analog.com.

Transformer Driver Resonant Matched Design Example: Unbuffered CMOS ADC, IF = 170MHz, Sampling Rate = 65MSPS

This design example is typical of the process required to optimize a transformer driven unbuffered CMOS ADC operating on an IF frequency. Resonant matching is used to optimize the response at the desired IF frequency. The system specifications represent those of an actual customer.

In this design, a 170MHz IF is digitized at a 65MSPS sampling rate. The bandwidth of the IF signal is 20MHz.

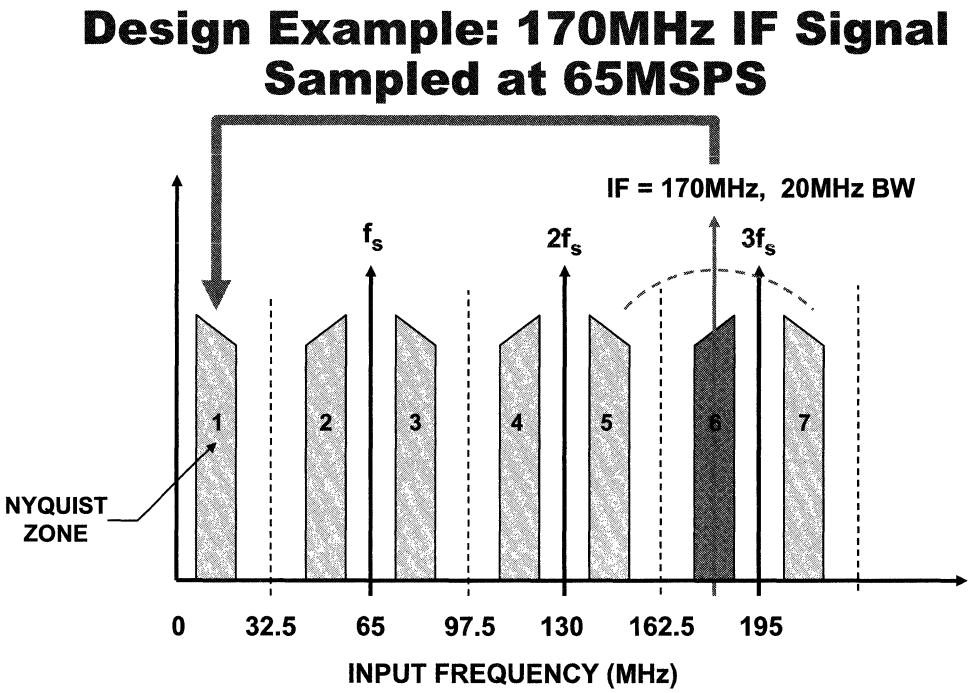

This figure shows the frequency spectrum of the 170MHz IF signal sampled at 65MSPS. The IF signal lies in the sixth Nyquist zone. The sampling process downconverts it to the first Nyquist zone as shown.

The actual bandwidth of the signal is 20MHz and, therefore, a sampling rate of at least 40MSPS is required. In this example, the sampling rate was selected to be 65MSPS.

Design Example: Design Requirements and ADC Requirements

Design Requirements

	Input Impedance (Ohm)	VSWR	Passband Flatness (dB)	–3dB IF BW (MHz)	SNR (dBc)	SFDR (dBc)	Input Drive Level (dBm)
Ideal Value	50	1	0.5	100	74	90	5
Design Limit	20	2	3	200	65	70	12

ADC Requirements

Sample Rate = 65MSPS **SNR = ≥ 65dB** **IF = 170MHz** **Band = 20MHz (160-180MHz)**	**AD9238, 12-bit, 65MSPS, 3V, Dual ADC** **Input Bandwidth = 500MHz** **Unbuffered CMOS Switched Capacitor Input Structure**

The basic design requirements and ADC requirements are presented in these tables. The IF frequency is 170MHz, and the bandwidth is 20MHz (160-180MHz). Input impedance is important in IF systems because mismatches reduce the amount of power transferred. Attention to input impedance is especially important in transformer drivers, because all of the load on the secondary is reflected back to the primary. The input impedance should not only be the proper value but should be primarily resistive at the IF frequencies of interest. This is quite different from an amplifier driver, where the output is relatively isolated from the input.

The desired passband flatness in the IF bandwidth (160-180MHz) is 0.5dB with a 3dB upper limit. The desired 3dB IF bandwidth is 100MHz, but 200MHz is acceptable. SNR should be between 65dB and 74dB, and SFDR between 70dB and 90dB. Input drive level should be no greater than +12dBm.

The system sampling rate was selected to be 65MSPS based on the 20MHz IF bandwidth. In order to achieve greater than 65dB SNR, a 12-bit ADC is required. The AD9238 12-bit, 65MSPS CMOS ADC was selected. This ADC has an unbuffered switched capacitor input structure as previously discussed. The input bandwidth of the AD9238 is 500MHz which is sufficient to handle the 170MHz IF input.

The AD9238 is a dual ADC and operates on a single +3V power supply.

Design Example: AD9238 Switched Cap ADC Baseline Performance Without R and L

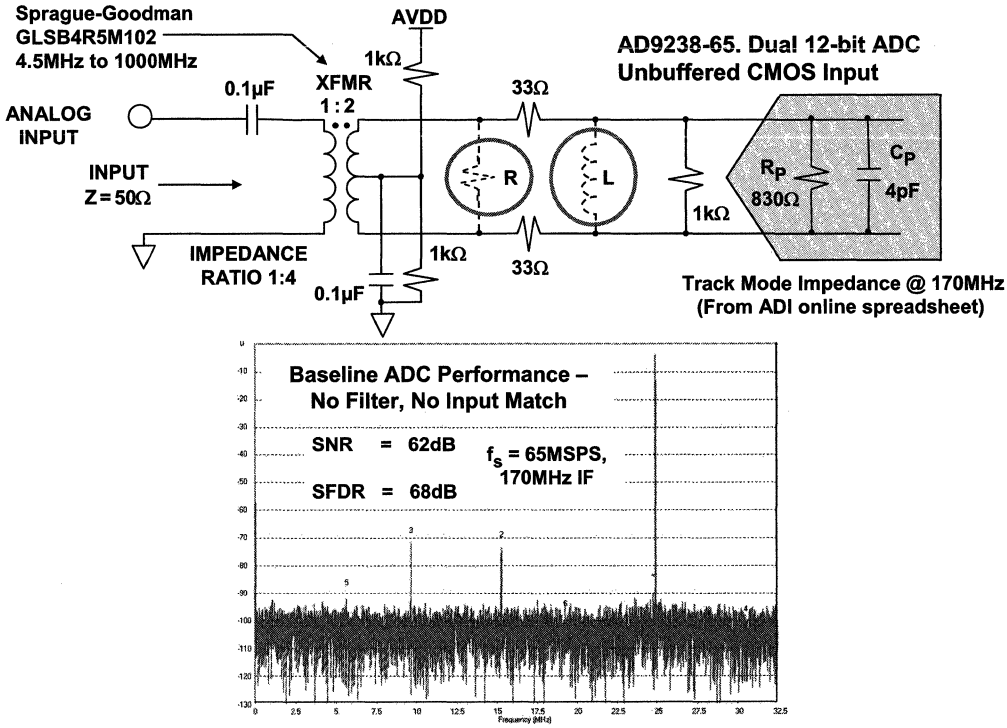

This figure shows the basic circuit without resonant matching as well as its performance when sampling a 170MHz IF signal at 65MSPS. SNR is 62dB, and SFDR is 68dB.

From the Analog Devices' online spreadsheets, the track mode impedance of the AD9238 is $R_P = 830\Omega$ in parallel with $C_P = 4pF$ for an IF input frequency of 170MHz.

The 1kΩ differential resistor is to de-Q the parasitic C_P on the ADC. The 33Ω series resistors inserted in series with each analog input aid in reducing the charge injection kickback into the transformer.

The transformer selected is a Sprague-Goodman, part number GLSB4R5M102 which has a turns ratio of 1:2 and an impedance ratio of 1:4.

The objective is to select the values of the parallel L to resonate with the 4pF capacitor at 170MHz and to select the value of R to make the overall input impedance 50Ω.

It should be emphasized that the process requires some experimental optimization, primarily because an exact model for the transformer parasitic inductances, resistances, and capacitances is not available. PC board parasitics also affect the optimum values.

Design Example: Determining Initial Values of Shunt R and L

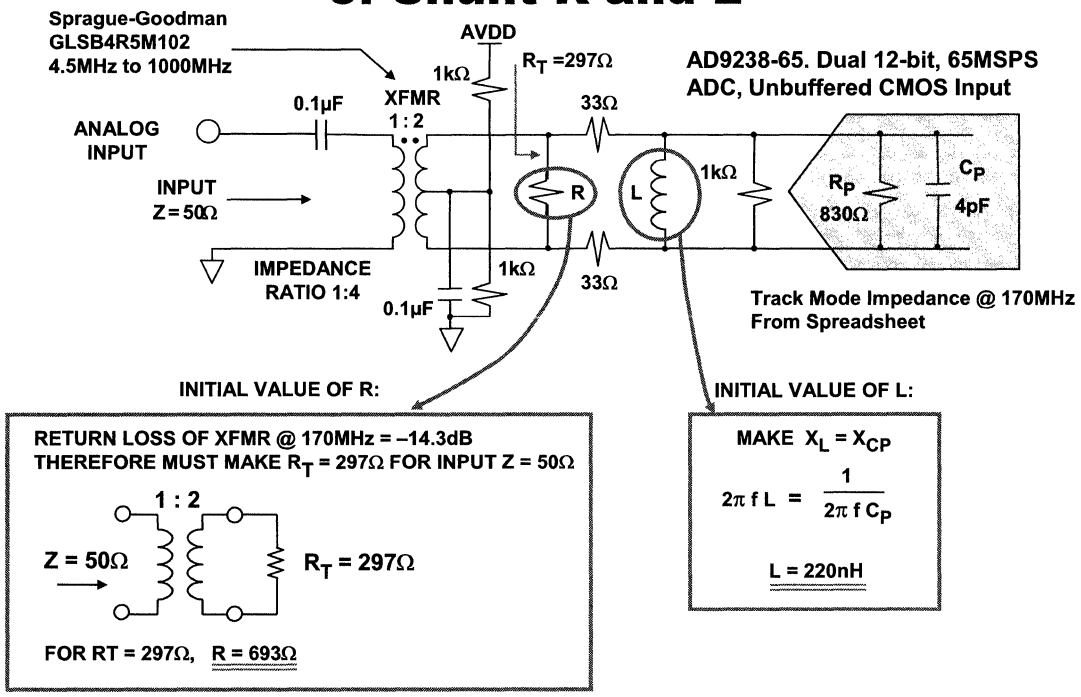

Selecting the initial value of L is a simple process of setting $X_L = X_{CP}$ and solving for L with f = 170MHz as shown in the figure. The value obtained is L = 220nH.

The return loss of the transformer at 170MHz is –14.3dB. This means that a net termination resistance of 297Ω (rather than the expected value of 200Ω) is required to make the effective input impedance 50Ω. The proper value for R_T is calculated as follows.

The equation for return loss is: $RL = 20\log|(Z_O - Z)/(Z_O + Z)|$. Substitute RL = –14.3dB, and solve the equation for Z. The result is $Z = 0.673Z_O$. This says that if the secondary is terminated in the theoretically ideal value of 200Ω, the input impedance is actually 0.673×50Ω = 33.65Ω. Therefore, the secondary must be terminated in 297Ω in order to make the actual input impedance 50Ω.

Without the resistor R, the net resistive termination on the secondary of the transformer is equal to:

$$830\Omega||1000\Omega + 33\Omega + 33\Omega + = 519.5\Omega.$$

Therefore, R must be 693Ω in order to make the net termination resistance 297Ω. This is the starting point of the design.

Design Example: AD9238 Switched Cap ADC – Iterate to Achieve Performance

IF = 170MHz, f_s = 65MSPS

DESIGN REQUIREMENTS	REV 1	REV 2	REV 3	REV 4 - Final
R || L	693Ω || 220nH	400Ω || 120nH	560Ω || 240nH	620Ω || 120nH
Input Impedance (50Ω)	47	23	32	48
VSWR (1-2)	1.06	2.17	1.56	1.04
Passband Flatness (0-3dB)	4.04	1.5	2.8	0.17
IF –3dB Bandwidth (100-200MHz)	174	174	214	150
SNR (65-74dB)	69	69	69	69
SFDR (70-90dB)	73.5	77	78	80
Input Drive Level (5-12dBm)	9.06	7.6	5.8	5.9

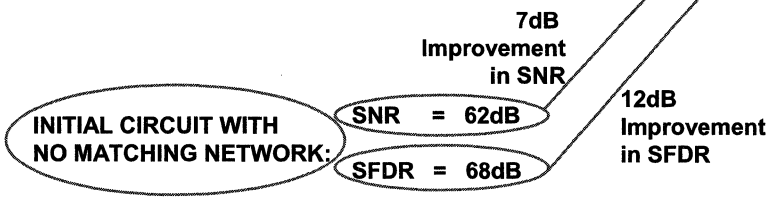

This shows the results of four interations of the design, starting with the calculated values of R and L. The final values (REV 4) yield performance within the required limits.

The primary reason for the interative design process is the lack of a good model for the transformer.

Note that the addition of the proper resonant matching network improves the SNR by 7dB and the SFDR by 12dB.

Design Example: Final Configuration for Input Resonant Match

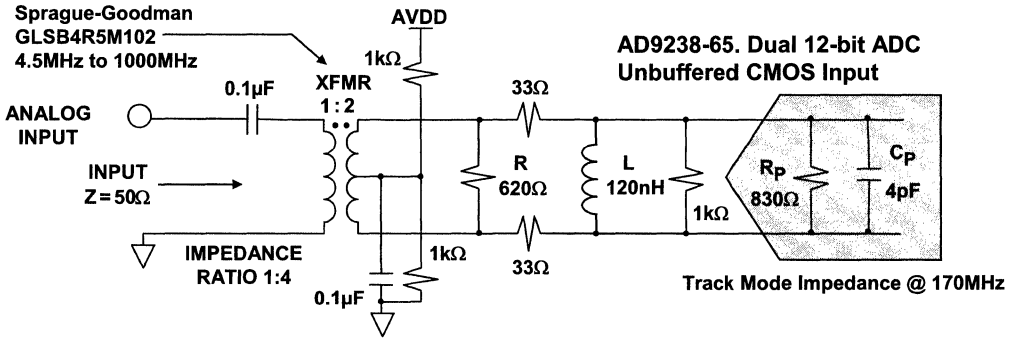

	BEFORE MATCH	AFTER MATCH	CHANGE
SNR	62dB	69dB	+7dB
SFDR	68dB	80dB	+12dB

This figure shows the final design and the improvement achieved by resonant matching.

Design Example: AD9238 Switched Cap ADC - Final Results – Bandwidth & Passband Flatness

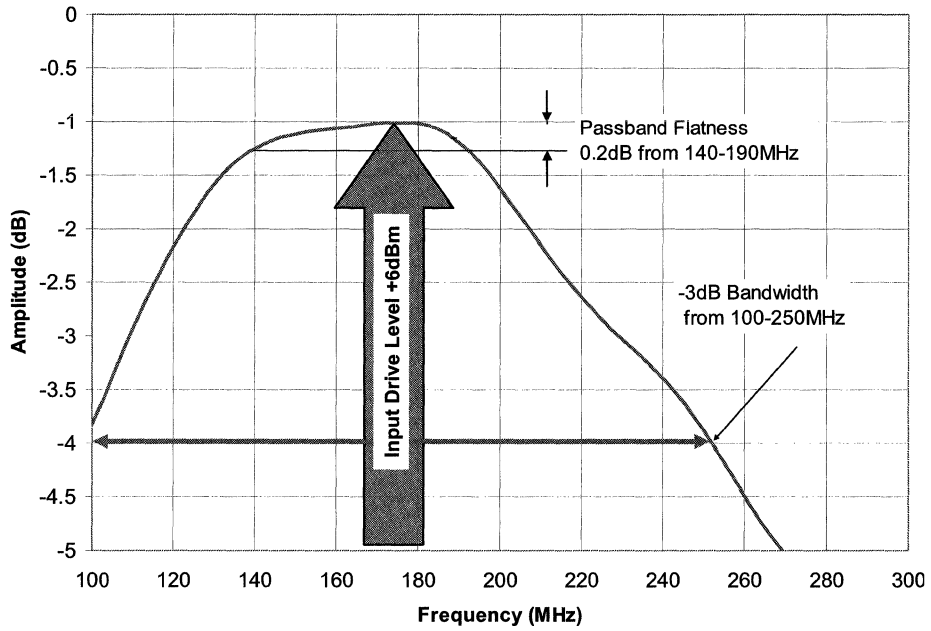

This figure shows the response of the interface. Note that the passband flatness is 0.2dB from 140MHz to 190MHz, and the 3dB bandwidth is 150MHz.

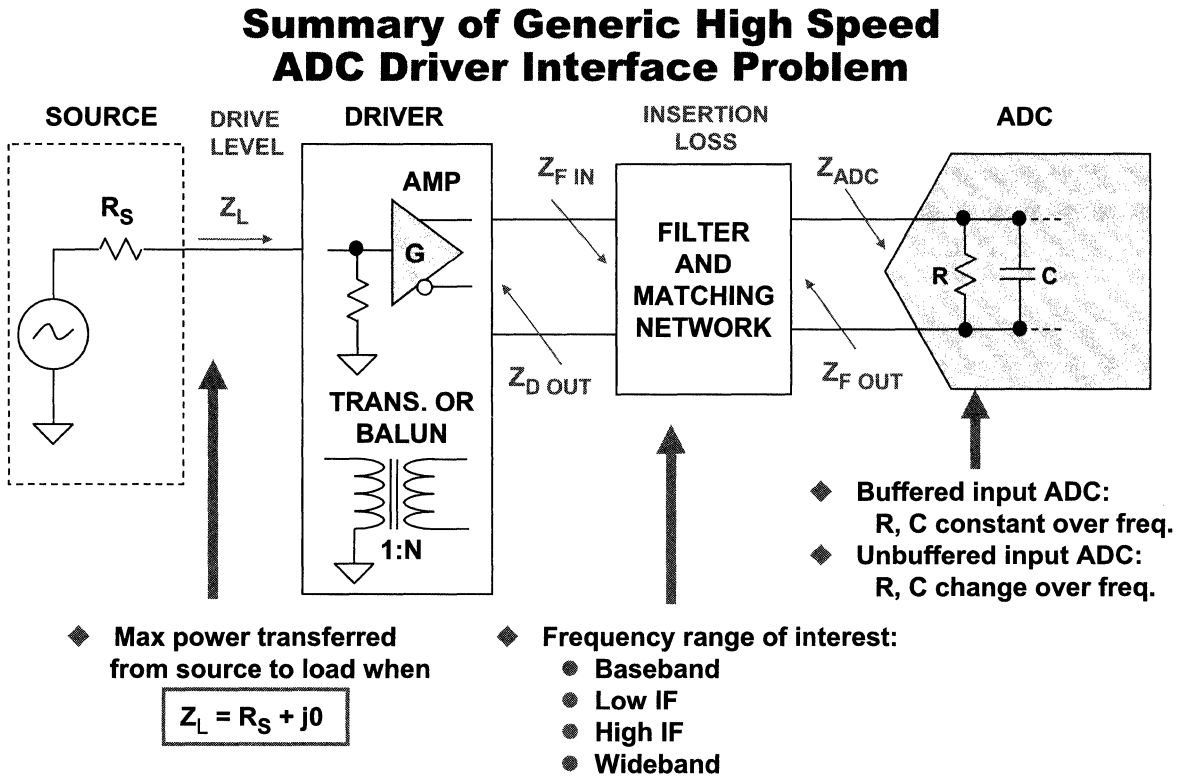

This figure summarizes the various considerations involved in designing the ADC analog input interface.

Amplifier drivers provide good isolation between the ADC input circuit and the signal source, and also provide signal gain if required. Current state-of-the-art differential amplifier output noise performance is approximately 5nV/√Hz.

Transformer drivers present more of a design problem because the secondary load is always reflected to the primary. In addition, good models containing transformer parasitics are not always available. However, transformers do not add additional noise to the system.

In this section we have also shown how resonant matching improves the dynamic performance in IF sampling applications. Dynamic ADC input impedance values are available which allow the design to be optimized at the desired IF frequency.

The filter and matching network requirements for baseband and broadband applications have also been discussed in this section.

References

1. Rob Reeder, "Frequency Domain Response of Switched-Capacitor ADCs," *Application Note AN-742*, Analog Devices, www.analog.com.
2. Eric Newman and Rob Reeder, "A Resonant Approach to Interfacing Amplifiers to Switch-Capacitor ADCs," *Application Note AN-827*, Analog Devices, www.analog.com.
3. Rob Reeder, "Transformer-Coupled Front-End for Wideband A/D Converters," *Analog Dialogue*, Vol. 39, Number 2, 2005, Analog Devices, www.analog.com.
4. Rob Reeder, Mark Looney, and Jim Hand, "Pushing the State of the Art with Multichannel A/D Converters," *Analog Dialogue,* Vol. 39, Number 2, 2005, Analog Devices, www.analog.com.
5. Rob Reeder and Ramya Ramachandran, "Wideband A/D Converter Front-EndDesign Considerations – When to Use a Double Transformer Configuration," July Issue, *Analog Dialogue*, Vol. 40, Number 3, 2006, Analog Devices, www.analog.com.
6. ADC Input S-parameter data:
   Go to product webpage, click on "Evaluation Boards," upload
   S-parameter data in a spreadsheet (where available).

Sampling Clock Drivers

Effect of Sampling Clock Phase Noise on Ideal Digitized Sinewave

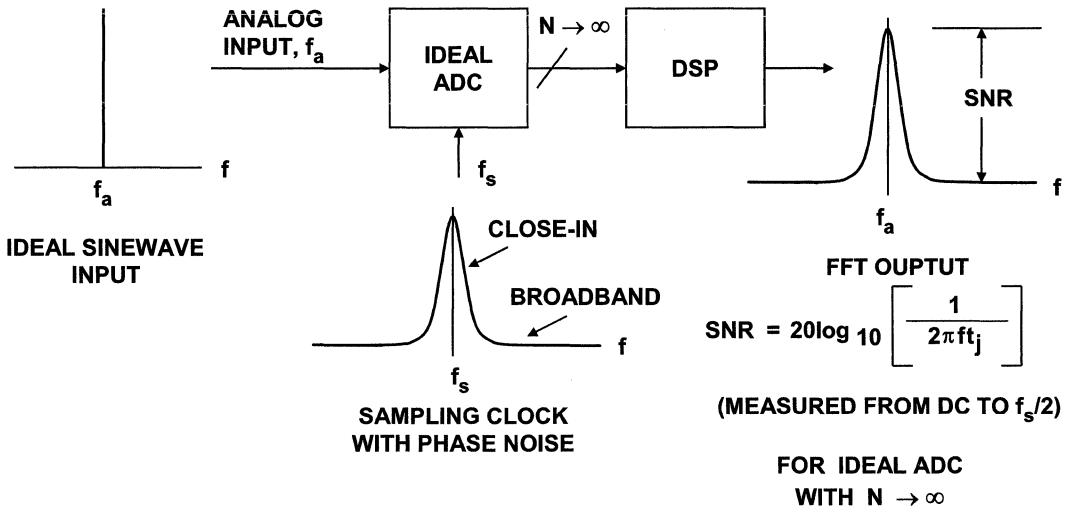

This figure shows an ideal ADC digitizing an ideal sinewave where the only error is phase noise on the sampling clock. Note that the sampling clock phase noise appears on the reconstructed sinewave output of the FFT. This is because of the inherent mixing that occurs between the sampling clock and the input signal.

The figure also shows that the overall SNR (measured from dc to $f_S/2$) due to broadband sampling clock jitter for an ideal ADC ($N \rightarrow \infty$) is given by the familiar equation,

$$SNR = 20\log[1/2\pi ft_j],$$

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

The total broadband jitter is made up of the external clock jitter and the aperture jitter of the ADC itself. Although the two components combine on an rss basis, the external clock jitter is typically the dominant component and the most often neglected.

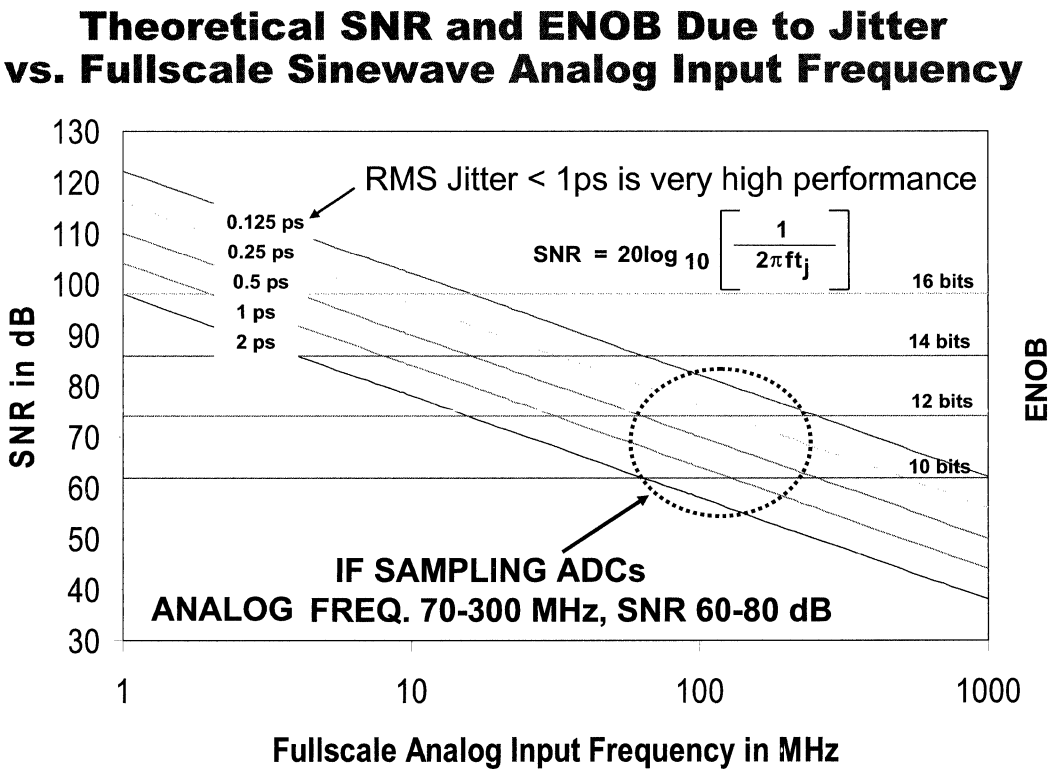

This figure plots the equation for jitter-based SNR and shows SNR on the left and ENOB (effective number of bits) on the right. They are related by the simple equation SNR = 6.02N + 1.76dB, where N = ENOB.

Note that for a fullscale 30MHz input, 14-bit ENOB requires that the rms jitter be no more than 0.3ps, or 300fs.

The equation assumes an ADC of infinite resolution, where the only error is the noise produced by the clock jitter.

IF sampling ADCs typically operate on IF frequencies between 70MHz 300MHz. The required SNR is typically between 60dB and 80dB.

Clock jitter less than 1ps rms represents a high performance clock.

Additive RMS Jitter of Logic Gates/Drivers

◆ FPGA (driver gates only, not including internal gates or dll/pll)	33-50 ps**
◆ 74LS00	4.94 ps *
◆ 74HCT00	2.20 ps *
◆ 74ACT00	0.99 ps *
◆ MC100EL16 PECL	0.7 ps **
◆ AD951x family	0.22 ps **
◆ NBSG16, Reduced Swing ECL (0.4V)	0.2 ps **
◆ ADCLK9xx, ECL Clock Driver Family	<0.1 ps**

*** Calculated values based on degradation in ADC SNR**
**** Manufacturers' specification**

To put jitter in perspective, this shows the typical rms jitter of several types of standard logic gates.

If several gates are put in series, the total rms jitter is the root-sum-square (rss) of the individual gate jitter. Four similar gates in series yields a jitter of $\sqrt{4} = 2$ times the jitter of each individual gate.

The data for the 74-series was obtained by measuring the degradation in an ADC SNR due to the addition of the gates in the sampling clock path. The 74-series is typically not specified for jitter.

The MC100EL16 and NBSG16 gates shown represent manufacturer's specifications.

The new Analog Devices' ADCLK9xx-series of drivers has less than 0.1ps rms jitter.

A high performance low jitter ADC test set can be used to measure the external clock jitter. This measurement technique is described further in the following reference:

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

Calculating Jitter from Phase Noise

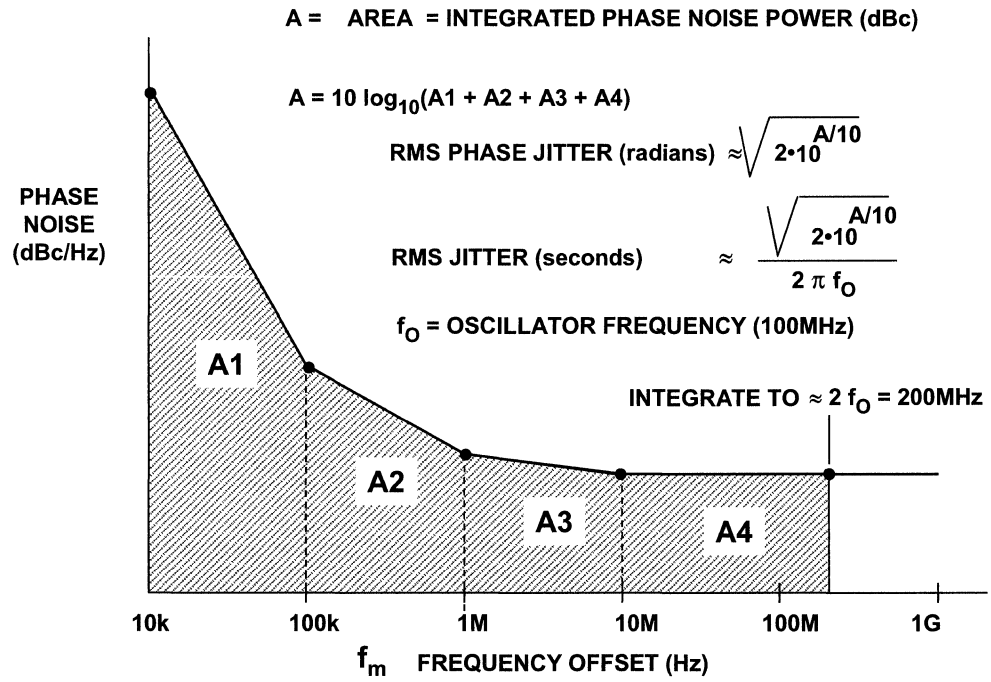

Most oscillators specify phase noise rather than jitter. Phase noise can be converted into jitter by integrating the phase noise over frequency as shown in this figure.

The process is quite similar to converting the voltage noise spectral density of an op amp into an rms voltage.

Note that the individual areas A1, A2, A3, and A4 are each expressed as ratios of phase noise power/carrier power. Therefore, they can be added directly.

The value for integrated phase noise power in dBc is obtained by A = 10log10(A1+A2+A3+A4).

The lower bandwidth of integration depends on the specified low frequency resolution of the system.

The upper limit for the integration should be approximately twice the oscillator frequency (or sampling clock frequency in the case of sampled data systems). This is a reasonable approximation to the bandwidth of the ADC clock input, and the actual number used is not that critical to the final result. This also assumes that there is no filter between the oscillator and the ADC sampling clock input.

Spreadsheet programs are available to make these calculations by simply inputting the various data points. The ADIsimPLL program will also convert phase noise into jitter.

More details on converting phase noise into jitter can be found in the following reference:

Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN-0916550273 Chapter 6. Also available as *Data Conversion Handbook*, Elsevier-Newnes, 2005, ISBN: 0750678410, Chapter 6.

Jitter Calculations for Low Noise 100MHz Crystal Oscillators
(Phase Noise Data Used with Permission of Wenzel Associates)

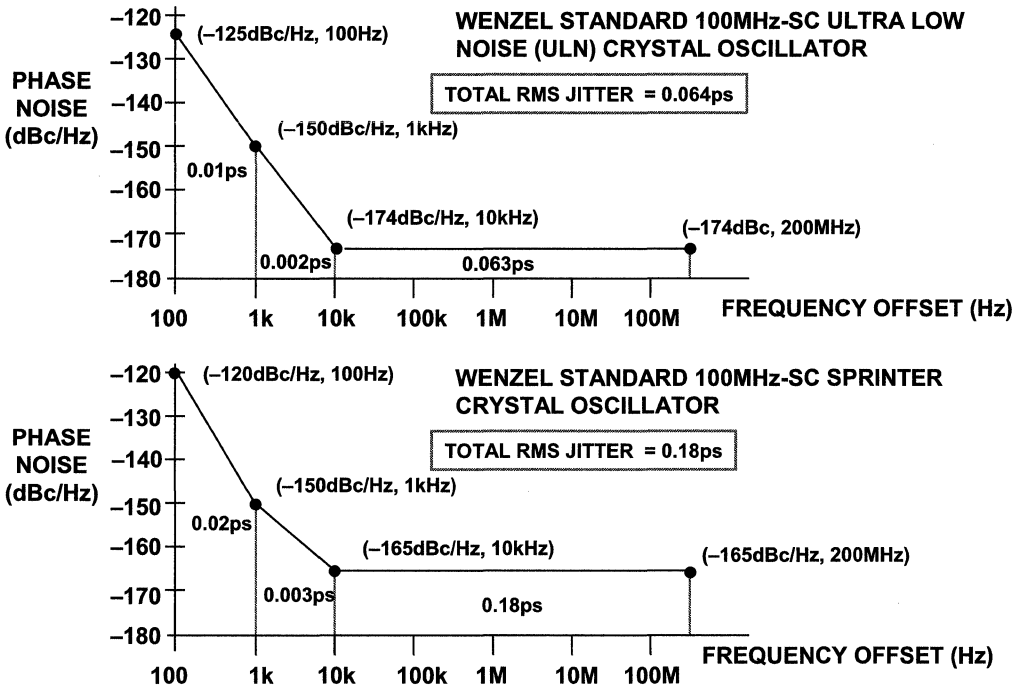

This figure shows the phase noise of two high quality crystal oscillators from Wenzel Associates. The graph is broken up into regions, and the jitter in each region is calculated along with the total jitter. The overall jitter is primarily determined by the "broadband" region.

The top graph is for the ULN-series and the calculations yield 64fs rms jitter.

The bottom graph is for the Sprinter-series oscillator and is about three times worse at 180fs.

In most cases the broadband component dominates the total jitter, and a reasonable approximation can be made by simply integrating the rectangular portion of the curve.

Ultra Low Jitter (<100 fs) Differential Clock Source

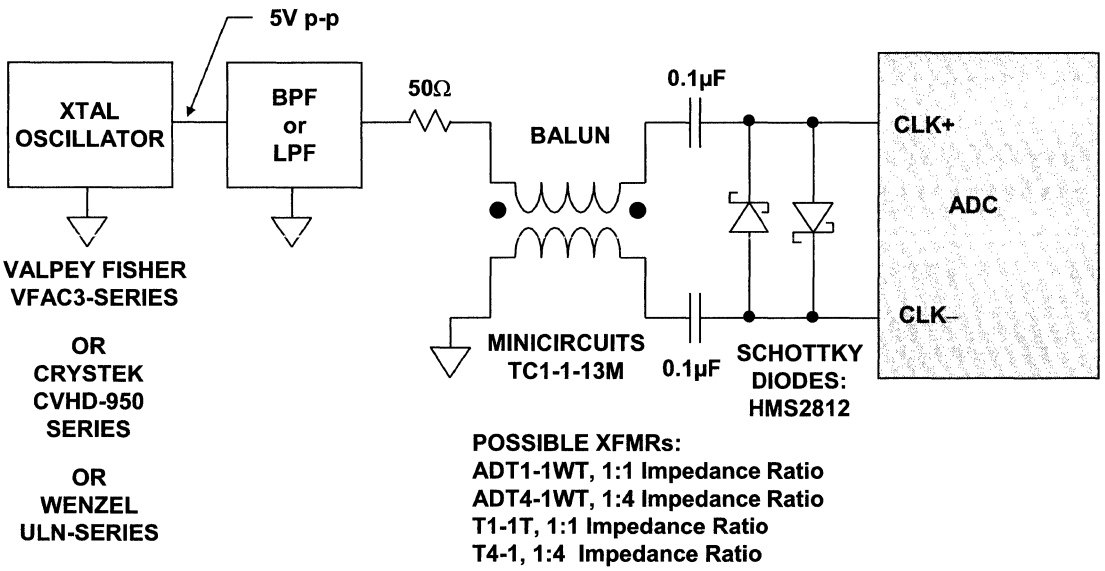

This figure shows a method for generating a sampling clock with less than 100fs jitter. The crystal oscillator is either a Valpey Fisher VFAC3-series (www.valpeyfisher.com), Crystek CVHD-950-series (www.crystek.com) or a Wenzel ULN-series (www.wenzel.com).

The output of the oscillator is filtered with either a lowpass or bandpass filter to remove harmonics and noise. A bandpass filter should be used for fixed frequency applications, but a lowpass filter allows some adjustment of the sampling clock frequency.

The single-ended filtered output must then be converted into a differential signal to drive the sampling clock input of the high speed ADC. This conversion can be performed using a balun or a transformer.

The output of the transformer is clamped with fast recovery Schottky diodes in order to present a high slew-rate differential square wave to the ADC sampling clock input. The diodes clamp at about 0.4V forward voltage, so the amplitude of the input square wave is about 0.8V p-p differential. This represents an optimum drive level for the CLK inputs.

Using a Phase-Locked Loop (PLL) and Bandpass Filter to Condition a Noisy Clock Source

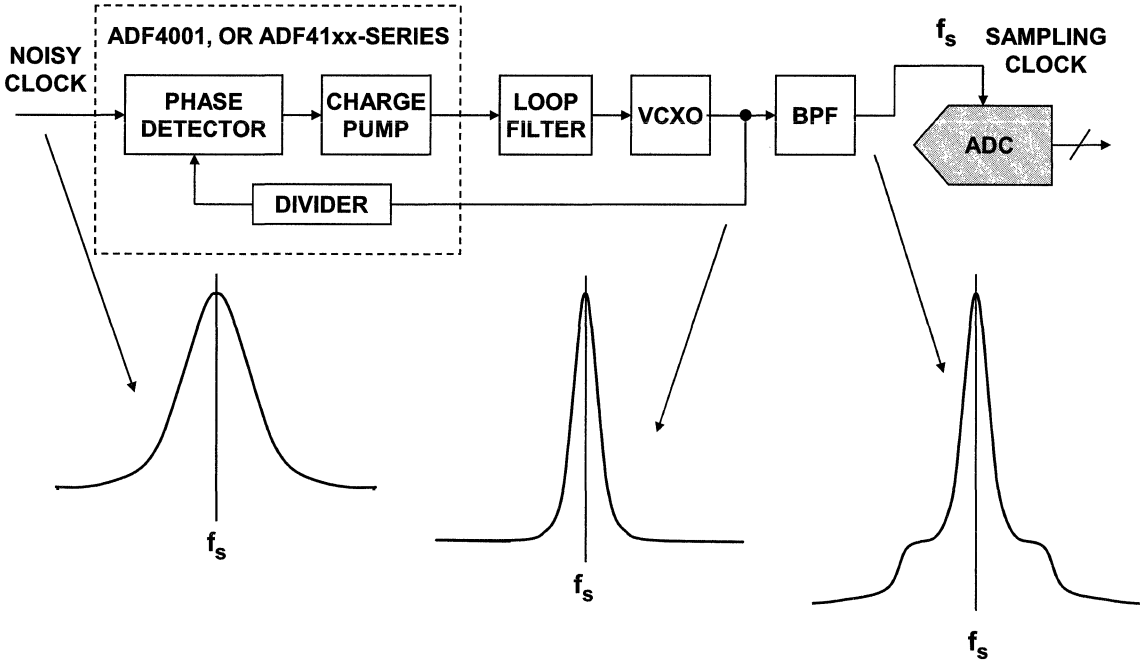

Only the highest end systems can afford the dedicated crystal oscillators of the type previously discussed.

Many systems therefore use a low-cost combination of a phase-locked loop (PLL) and a VCXO (voltage-controlled crystal oscillator) followed by a bandpass filter to generate a low jitter sampling clock from a noisy system clock.

Analog Devices makes a number of PLLs and clock generation and distribution ICs suitable for this function, and these will be discussed in Section 3.

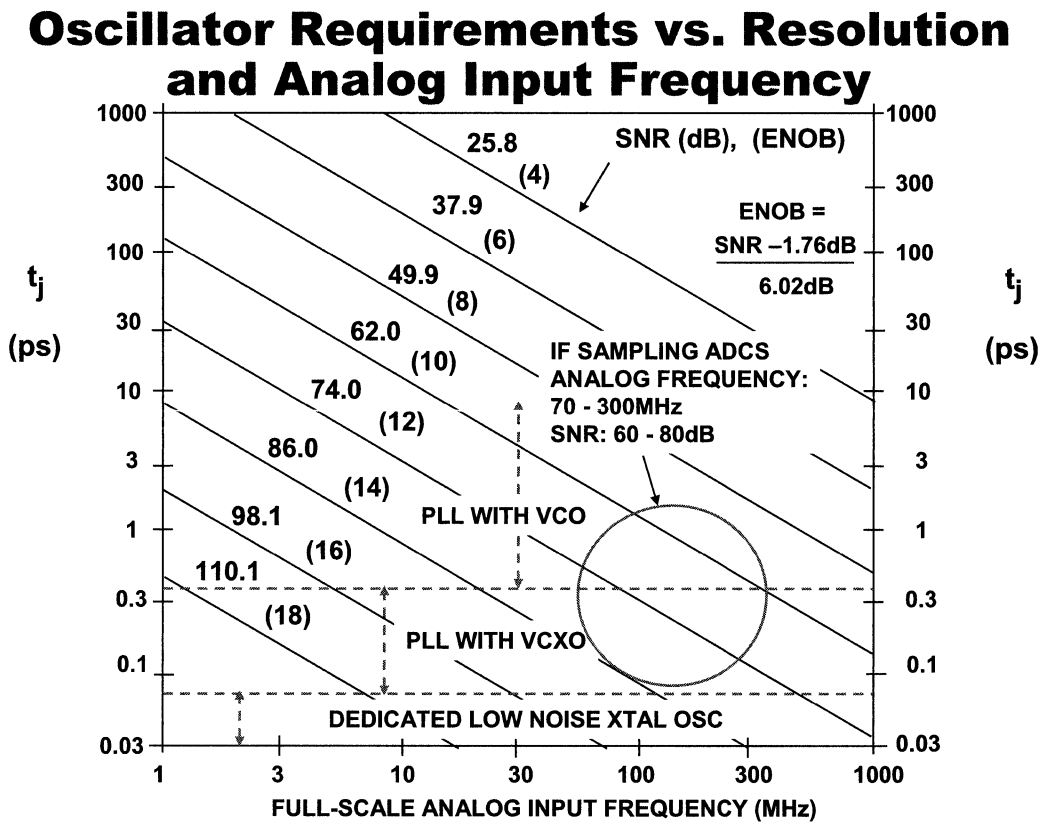

This figure shows that standard PLLs with VCOs can be used in many applications for generating clean sampling clocks.

PLLs with VCXOs can provide lower jitter solutions which are suitable for all but the most demanding applications where dedicated low-noise XTAL oscillators are required.

Square or Sine Source Input using a 1.8V or 3V CMOS Single-Ended Clock Driver

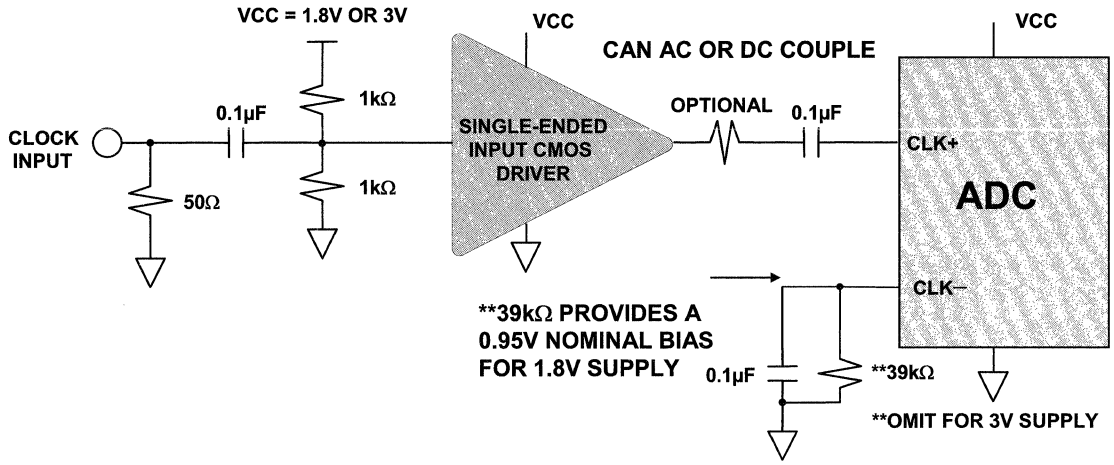

POSSIBLE 1.8V DRIVERS:
74VCX86 – LOW VOLTAGE VCX LOGIC
NC7NP04 – ULP TINY LOGIC SERIES
(FAIRCHILD)

POSSIBLE 3V DRIVERS:
74VCX86 – LOW VOLTAGE VCX LOGIC
NC75VC04 – TINY LOGIC SERIES
(FAIRCHILD)

In order to achieve high performance, most pipelined ADCs are designed for differential sampling clock inputs. This provides good common-mode rejection and optimizes the signal for driving the internal sample-and-hold and minimizes the additive jitter at the sampling clock input.

In some applications where optimum jitter performance is not required, differential clock inputs can be driven single-ended.

In this figure, a low jitter CMOS gate is used to drive the CLK+ input of the ADC, while the CLK– input is biased to the appropriate value of 0.95V for a 1.8V supply. Note that this bias level is highly product specific, so the data sheet must always be consulted to determine the optimum bias circuit and level.

A simple resistor divider is used at the input of the gate driver to set the proper dc bias level.

Several possible low jitter 1.8V and 3V CMOS drivers are listed in the figure. In all cases, the jitter specification on the CMOS driver should be compatible with the overall system requirements.

Low Jitter Single-Ended to Differential PECL Driver

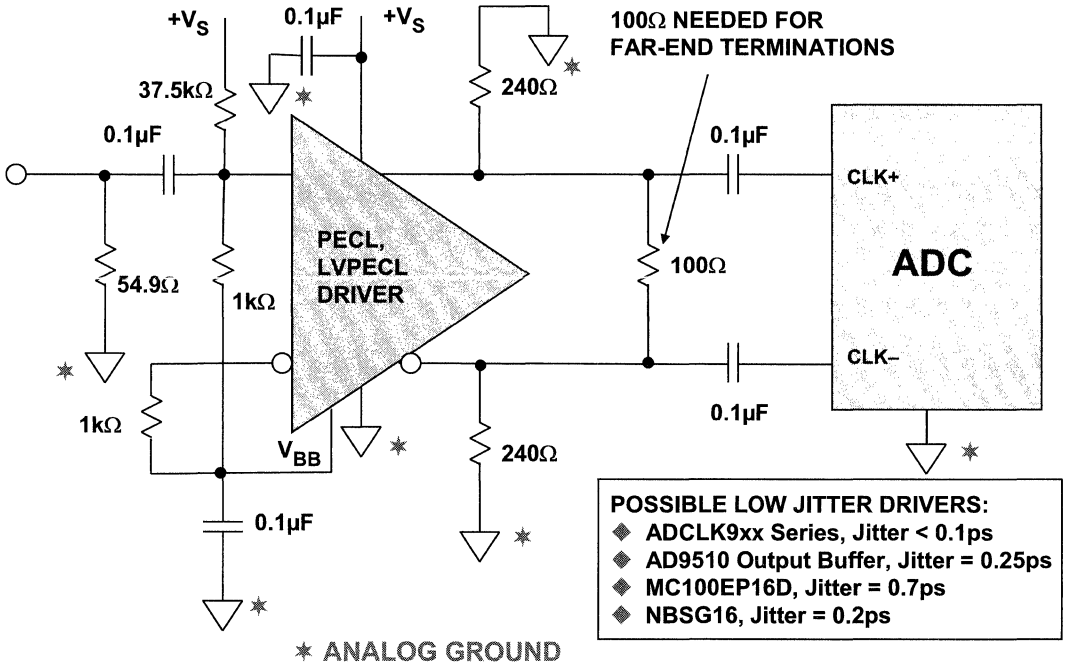

Most high-performance ADCs have differential clock inputs for lowest jitter.

Low jitter PECL or reduced signal PECL drivers can be used to convert a single-ended signal to a differential one. These are made by ON-SEMI, and have specified jitter as low as 0.2ps.

The Analog Devices' new ADCKL9xx-series of low jitter clock drivers have additive jitter less than 0.1ps.

Note that the clock driver circuit as well as the ADC should be referenced to the same ground plane. More discussion on grounding and decoupling can be found in Section 4.

ADC Data Outputs

Typical CMOS Digital Output Drivers

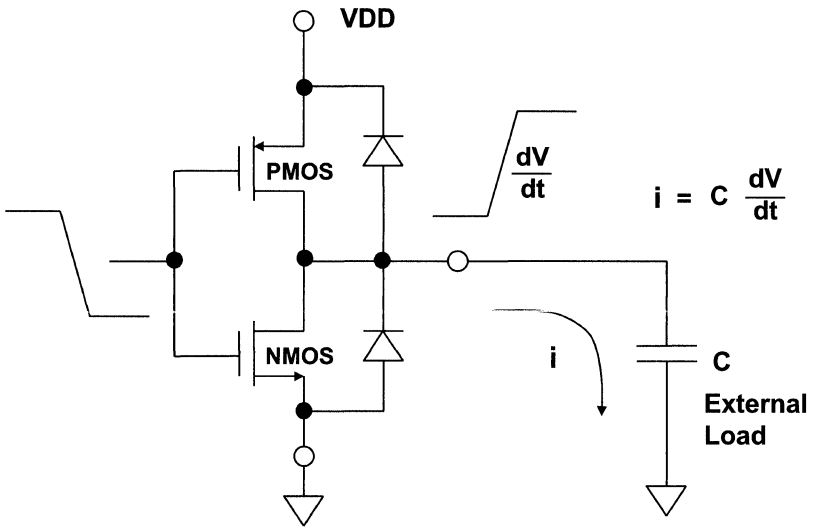

The final interface to be considered is the data output. Although these are digital outputs, they should be treated with care, because transient currents can increase the noise and distortion of the ADC by coupling back into the analog input.

Typical CMOS drivers shown in this figure are capable of generating large transient currents, especially when driving capacitive loads.

Particular care must be taken with CMOS data output ADCs so that these currents are minimized and do not generate additional noise and distortion in the ADC.

Use Series Resistance to Minimize Charging Current of CMOS Digital Outputs

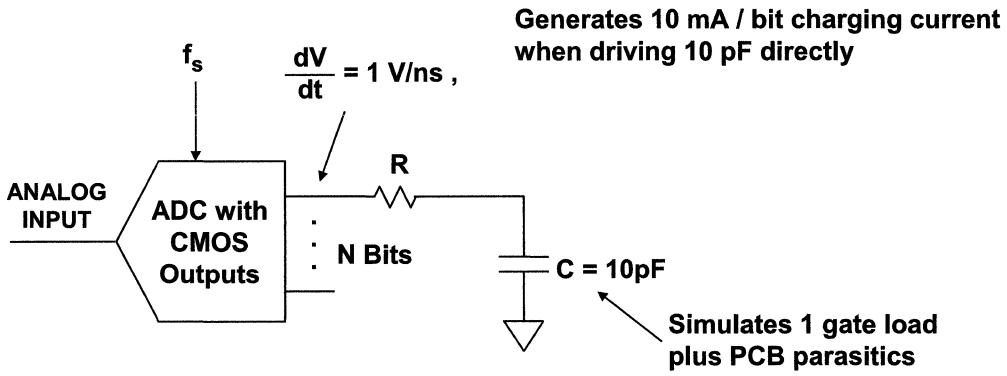

- Make RC < 0.1 $\left[\frac{1}{f_s}\right]$
- For f_s = 100 MSPS, RC < 1 ns
- If C = 10 pF, R = 100 Ω

Consider the case of a 16-bit parallel CMOS output ADC. With a 10pF load on each output (simulates one gate load plus PCB parasitics), each driver generates a charging current of 10mA when driving a 10pF load. The total transient current for the 16-bit ADC can therefore be as high as 16×10mA = 160mA.

These transients currents can be suppressed by adding a small resistor, R, in series with each data output. The value should be chosen so that the RC time constant is less than 10% of the total sampling period. For f_S = 100MSPS, RC should be less than 1ns. With C = 10pF, an R of about 100Ω is optimum. Choosing larger values of R can degrade output data settling time and interfere with proper data capture.

Capacitive loading on CMOS ADC outputs should be limited to a single gate load (usually an external data capture register). Under no circumstances should the data output be connected directly to a noisy data bus. An intermediate buffer register must be used to minimize direct loading of the ADC outputs.

LVDS Driver Designed in CMOS

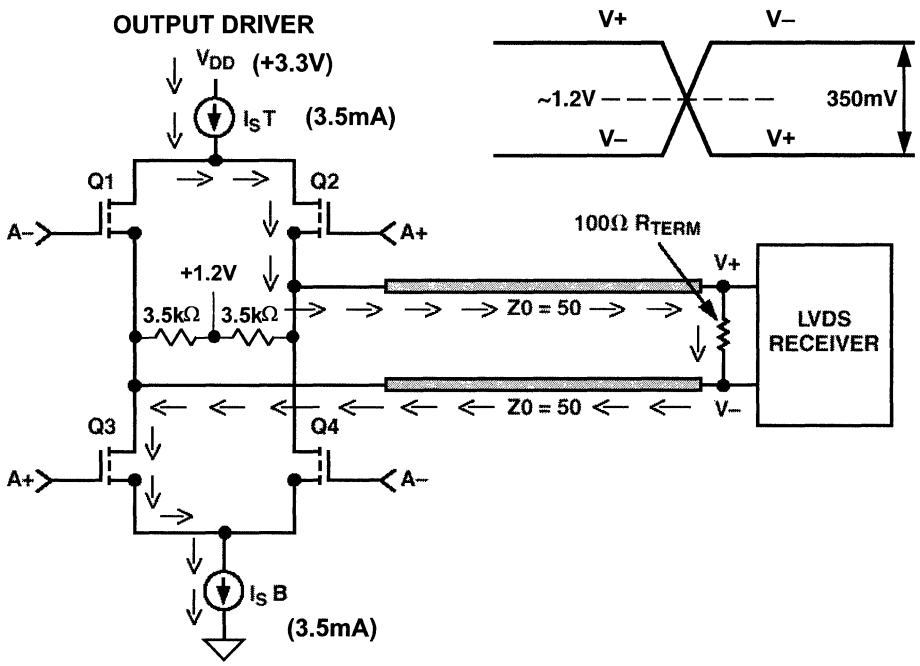

Emitter-coupled-logic (ECL) has long been known for low noise and its ability to drive terminated transmission lines with rise/fall times less than 2ns. The family presents a constant load to the power supply, is non-saturating, and the low-level differential outputs provide a high degree of common-mode rejection. However, ECL dissipates lots of power.

Recently, low-voltage-differential-signaling (LVDS) logic has attained widespread popularity because of similar characteristics, but with lower amplitudes and lower power dissipation than ECL. The LVDS logic swing is typically 350mV peak-to-peak centered about a common-mode voltage of +1.2V. A typical driver and receiver configuration is shown in this figure. The driver consists of a nominal 3.5mA current source with polarity switching provided by PMOS and NMOS transistors. The output voltage of the driver is nominally 350mV peak-to-peak at each output, and can vary between 247mV and 454 mV. The output current can vary between 2.47mA and 4.54mA. The LVDS receiver is terminated in a 100Ω line-to-line. According to the LVDS specification, the receiver must respond to signals as small as 100mV, over a common-mode voltage range of 50mV to +2.35V. The wide common-mode receiver voltage range is to accommodate ground voltage differences up to ±1V between the driver and receiver.

LVDS outputs for high-performance ADCs should be treated differently than standard LVDS outputs used in digital logic. While standard LVDS can drive 1 to 10 meters in high-speed digital applications (dependent on data rate), it is not recommended to let a high-performance ADC drive that distance. It is recommended to keep the output trace lengths short (< 2 in.), minimizing the opportunity for any noise coupling onto the outputs from the adjacent circuitry, which may get back to the analog inputs. The differential output traces should be routed close together, maximizing common-mode rejection, with the 100Ω termination resistor close to the receiver. Users should pay attention to PCB trace lengths to minimize any delay skew.

LVDS also offers some benefits in reduced EMI, because the EMI fields generated by the opposing LVDS currents tend to cancel each other.

AD9228 Quad 12-Bit, 40/60MSPS, 1.8V LVDS Output ADC

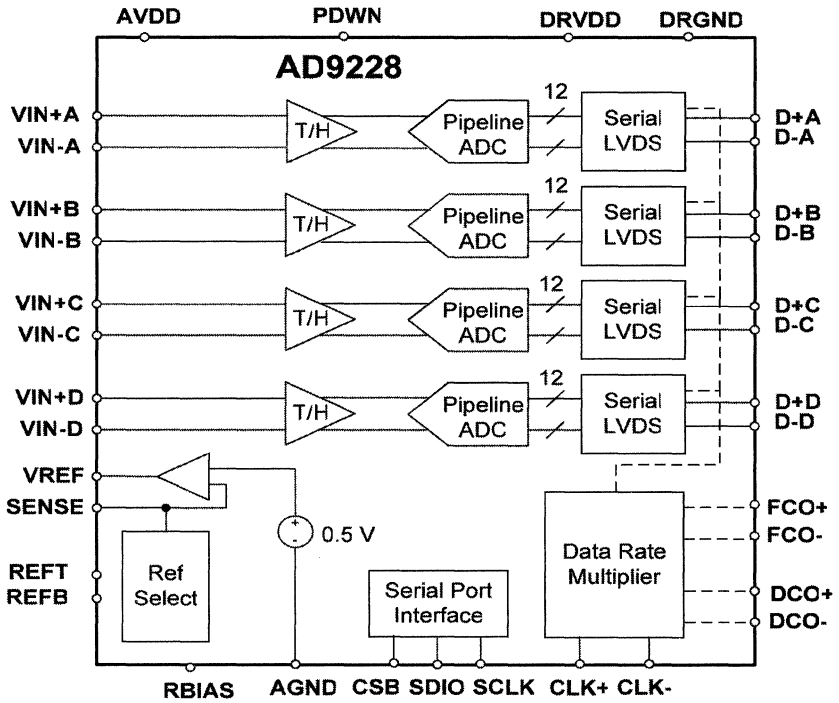

LVDS interface data rates can be as high as 800Mbits/s, thereby making serial data transfer practical even for some high speed ADCs. For instance, the AD9228 quad 12-bit, 65-MSPS ADC uses four serial LVDS outputs, each operating at 780Mbits/s. A functional block diagram of the quad ADC is shown in this figure.

The AD9228 has an on-chip track-and-hold circuit and is designed for low cost, low power, small size, and ease of use. The converter operates up to 65MSPS conversion rate and is optimized for outstanding dynamic performance where a small package size is critical. The ADC requires a single +1.8V power supply and CMOS/TTL sample rate clock for full performance operation. Total power dissipation is 477mW.

No external reference or driver components are required for many applications. A separate output power supply pin supports LVDS-compatible serial digital output levels. The ADC automatically multiplies the sample rate clock for the appropriate LVDS serial data rate. An MSB trigger is provided to signal a new output byte. In addition, a power down mode is supported.

AD9228 LVDS Output Data Timing Diagram

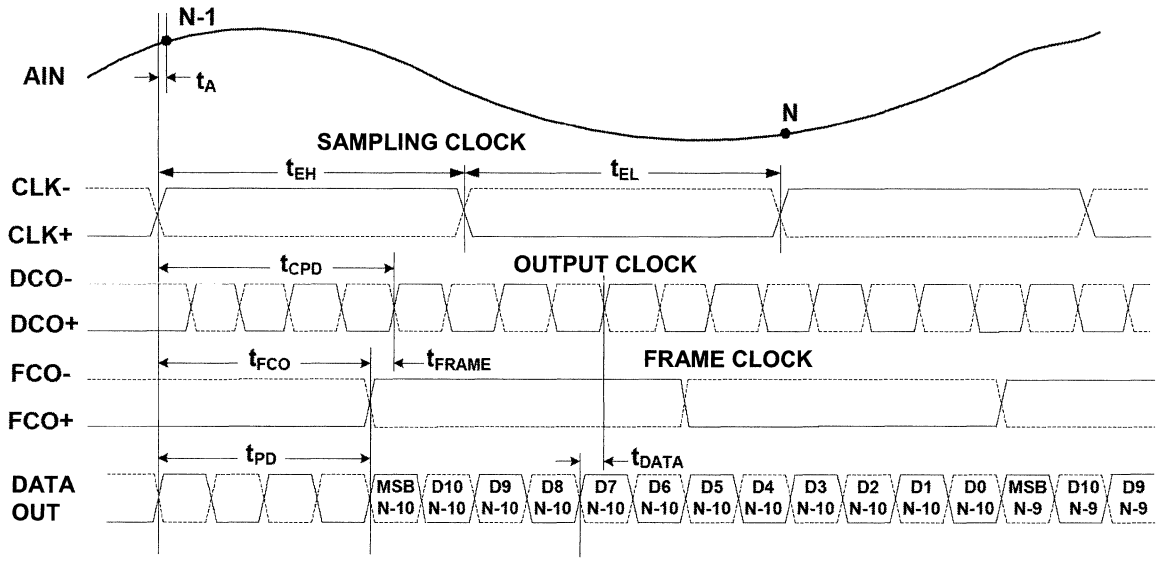

This figure shows the data output timing diagram for the AD9228. Data from each ADC is serialized and provided on a separate channel. The data rate for each serial stream is equal to 12-bits times the sample clock rate, with a maximum of 780MHz (12-bits × 65MSPS = 780MHz).

Data is clocked out of the AD9228 on the rising and falling edges of DCO. The MSB clock (FCO) is used to signal the MSB of a new output byte and is equal to the sampling clock rate.

The use of high-speed serial LVDS data outputs in the AD9228 results in a huge savings in the pin count, compared with parallel outputs. A total of 48 data pins would be required to provide four individual parallel 12-bit single-ended CMOS outputs. Using serial LVDS, the AD9228 requires only four differential LVDS data outputs, or eight pins, thereby saving a total 40 pins. In addition, the use of LVDS rather than CMOS reduces digital output transient currents and the overall ADC noise.

In most applications, the ADCs drive an FPGA, where the serial-to-parallel conversion can be easily performed.

Output Driving Summary

- **CMOS Outputs**
  - **More common output logic standard interface**
  - **Limited speed capabilities due to board parasitics, etc. (≈ 150MSPS)**
  - **More noise due to larger signal swing causing transient currents and digital ground bounce**
  - **Series resistor in each data output recommended to reduce transient current**
- **LVDS Outputs**
  - **Becoming more popular**
  - **Much faster than single-ended CMOS**
  - **Less transient noise because LVDS uses current-mode switching and small signal swing**
  - **Good common-mode rejection**
  - **Differential resistor termination required**
  - **Twice the number of traces to route if using parallel outputs**
  - **Simpler data capture compared to single-ended CMOS demuxed solution**

CMOS output ADCs retain their popularity for applications where the optimum noise performance is not mandatory. The noise produced by the high output transient currents can be reduced by placing an appropriate resistor in series with each data output. CMOS data outputs greater than about 150MSPS require demuxed outputs.

LVDS is becoming more popular because it is much faster than single-ended CMOS and generates less noise. Although LVDS requires two pins per bit output, the pin count is no greater than single-ended demuxed CMOS which also requires two pins per bit output.

High Speed System Applications

SECTION 3
DACs, DDSs, PLLs, AND CLOCK DISTRIBUTION

High Speed CMOS DACs

www.analog.com/txdacs

PMOS Transistor Current Switches

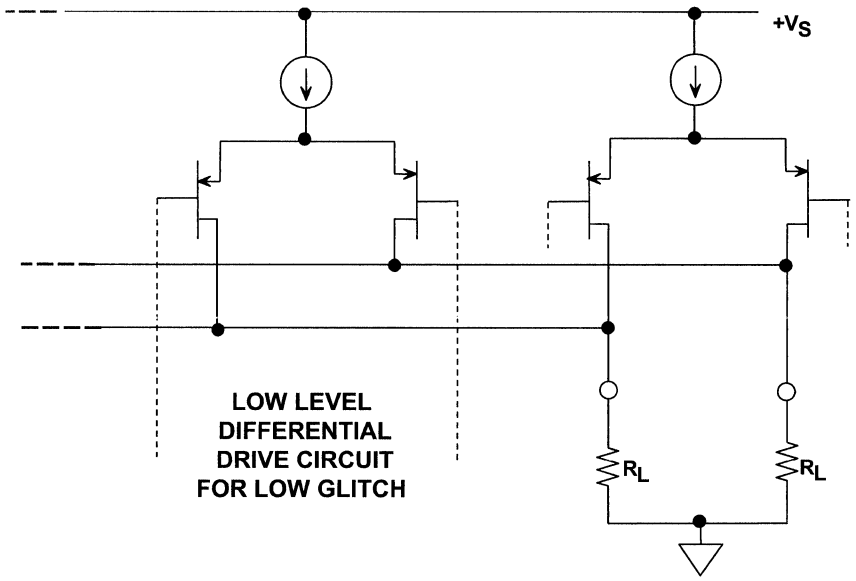

This is a typical CMOS transistor current steering switch based on PMOS transistors. It is used in the TxDAC family.

The advantage of PMOS is that the output sources current and provides a signal which is above ground in a single-supply system. Both true and complementary current outputs are supplied in low distortion DACs allowing differential operation to minimize second-order distortion products.

An NMOS output stage would sink current, and provide a negative output voltage in a single-supply system. NPN or NMOS transistor-based DACs operate in this manner.

The outputs of the PMOS switches generally have 1V or 2V of compliance, and can usually go about 1V below ground. (Compliance is the allowable output voltage range over which the DAC meets its linearity specifications).

The basic PMOS current switch is driven from a low-level differential signal whose level and amplitude has been optimized to produce a minimum switching glitch at the DAC output. The following references describe the basic method used in optimizing the current switching action in the TxDAC family as well as other details of their design:

---

1. Doug Mercer, "Low Power Approaches to High Speed CMOS Current Steering DACs," *CCIC 2006 Conference Proceedings*.

2. Doug Mercer, "A Low Power Current Steering Digital-to-Analog Converter in 0.28 micron CMOS," *ISLPED 2005 Digest of Technical Papers*.

3. Doug Mercer, "A Study of Error Sources in Current Steering Digital-to-Analog Converters," *CCIC 2004 Conference Proceedings*.

4. W. Schofield, Doug Mercer, and L. St. Onge, "A 16b 400MS/s DAC with <–80dBc IMD to 300MHz and <–160dBm/Hz Noise Power Spectral Density," *ISSCC 2003 Digest of Technical Papers*.

5. Doug Mercer, "A 16b D/A Converter with Increased Spurious Free Dynamic Range," *IEEE Journal of Solid State Circuits*, vol. 29, no. 10, pp. 1180-1185, October 1994.

High Speed 3-Bit Binary DAC with Complementary Current Outputs

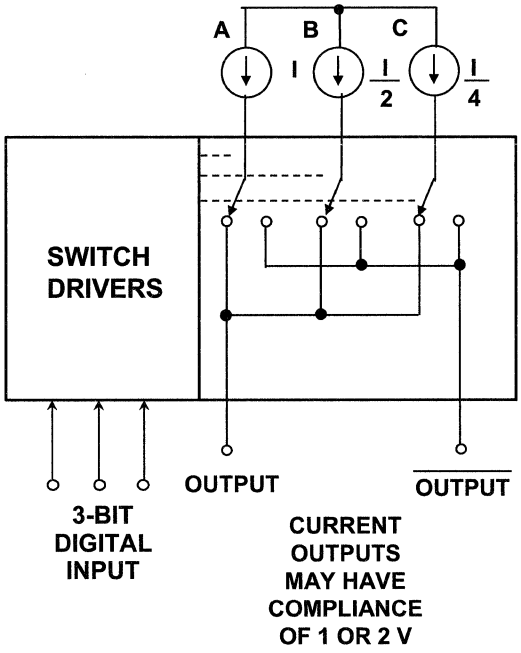

PMOS current switches available in digital CMOS processes can be used to construct simple binary-weighted DACs as shown here for a 3-bit DAC. There are two basic problems in using this architecture for high resolution fast DACs.

The first problem is maintaining the correct binary ratio between the currents for high resolutions. For resolutions greater than about eight bits, some type of trimming is required in order to maintain the required accuracy.

The second problem is the switching glitches of the binary DAC tend to be code dependent. That is, the switching transient generated at the bit-1 transition, 011 to 100, tends to be larger than the switching transient generated at the bit-2 transitions, 101 to 110, or 001 to 010.

The code-dependent glitches produce various unwanted harmonics of the fundamental output frequency. Basic sampling theory tells us that the higher order products which fall outside the Nyquist bandwidth (dc to $f_s/2$) will alias back into the Nyquist bandwidth where they may add distortion to the desired output frequency.

High Speed 3-Bit Thermometer (Fully Decoded) DAC with Complementary Current Outputs

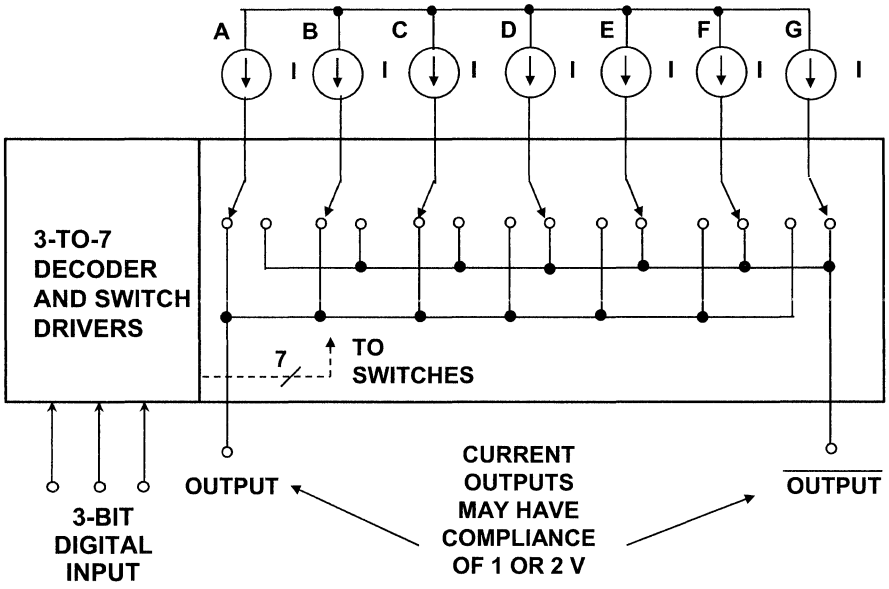

While it is impossible to totally eliminate DAC switching glitches, it is possible to use architectures which produce glitches that are code-independent and occur at the DAC update rate.

This figure shows a "thermometer" or "fully-decoded" DAC based on equal value current sources and switches. This architecture minimizes the code dependent glitches at the expense of an increase in the number of current sources and switches to obtain the same 3-bit resolution.

The thermometer DAC also makes it easier to match each stage, since the currents are equal.

The problem with this architecture is the physical size and power required to implement high resolutions. For instance, an 8-bit thermometer DAC requires 255 current sources and switches. (Note that the concept of a thermometer DAC is similar to that of a flash ADC).

The thermometer DAC and binary DAC architectures can be combined to produce a "segmented" architecture which makes an excellent compromise between performance and chip area.

Typical TxDAC® 14-Bit CMOS Segmented DAC Core

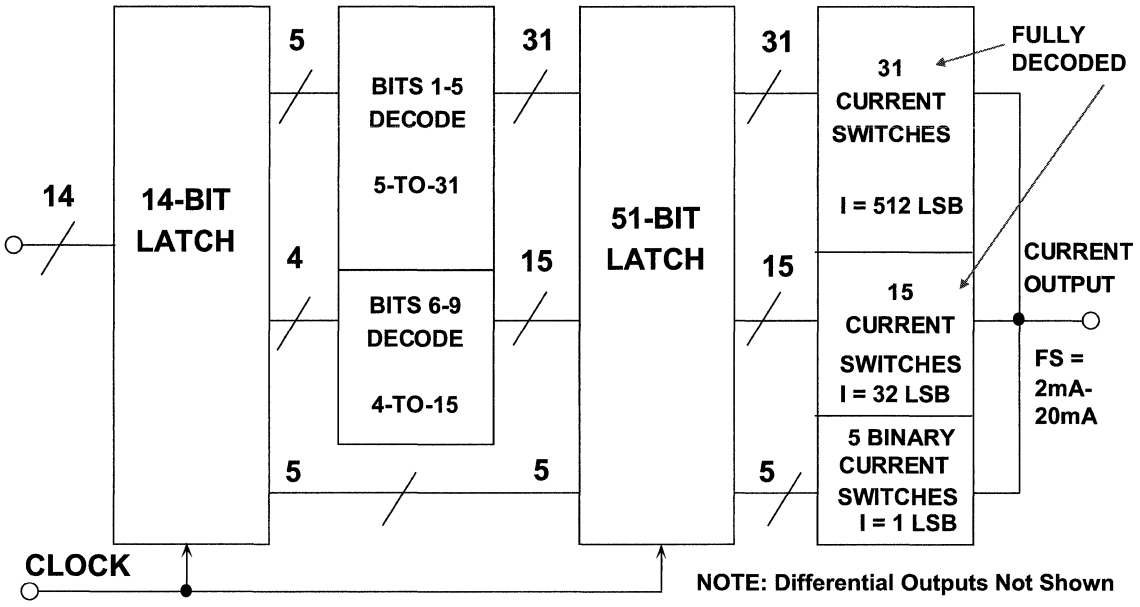

This figure shows a typical example of a 14-bit core of the TxDAC family. This is an excellent illustration of the "segmented" DAC architecture. The five MSBs are fully decoded into 31 equal current switches. The next four bits are fully decoded into 15 equal current switches. The five LSBs are binarily decoded.

Note that a second latch (51 bits) is used to drive the output current switches. This eliminates the delay "skew" added by the decoding logic which follows the input 14-bit latch.

All the currents are combined to produce the final output current.

This architecture yields excellent low-distortion performance.

In practice, the DAC is fully differential for better linearity and to minimize second-order distortion.

Location of First Nine Harmonic Products: Output Signal = 7MHz, DAC Update Rate = 20MSPS

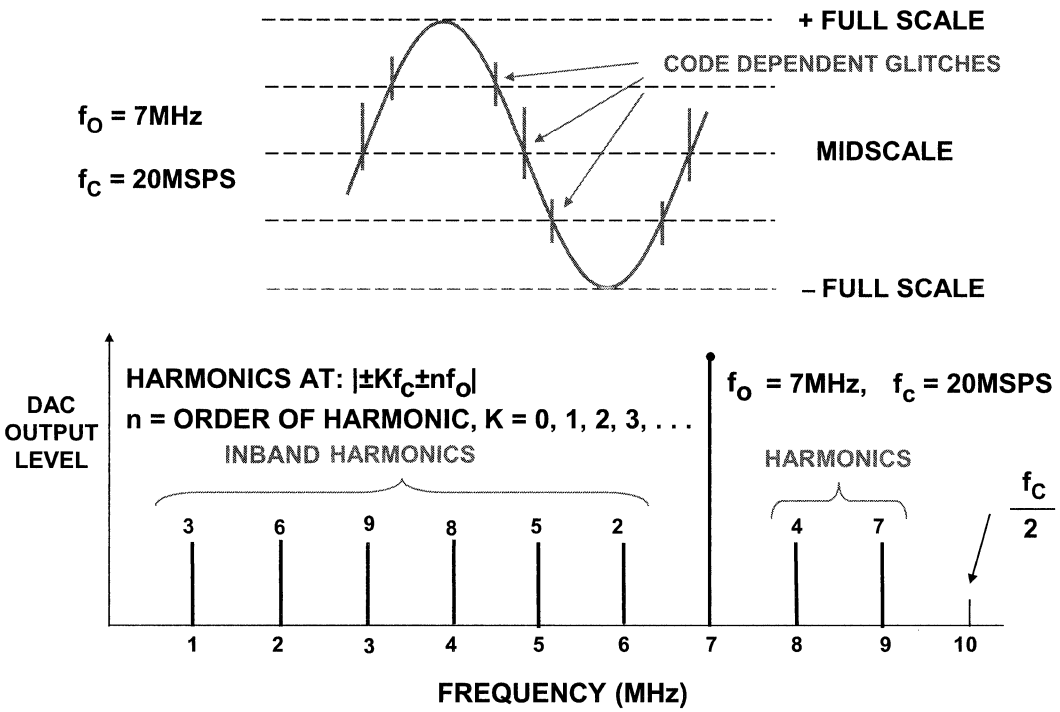

This figure shows how code dependent glitches produce harmonics which alias back into the Nyquist bandwidth.

In this case, the DAC is updated at a 20MSPS rate and produces a 7-MHz output signal. The figure shows the location of the first nine harmonics.

Aliased harmonics of f_O fall at frequencies equal to $|\pm Kf_C \pm nf_O|$, where n is the order of the harmonic, and K = 0, 1, 2, 3, ...

It is often desirable to carefully select the DAC update rate and the output frequency such that the worst harmonics do not interfere with the signal band of interest. Locating the position of the higher order harmonics is facilitated by an on-line design tool from Analog Devices, the "DAC Harmonic Image" tool. This tool is described in more detail later in this section.

Measuring DAC Distortion and SNR with an Analog Spectrum Analyzer

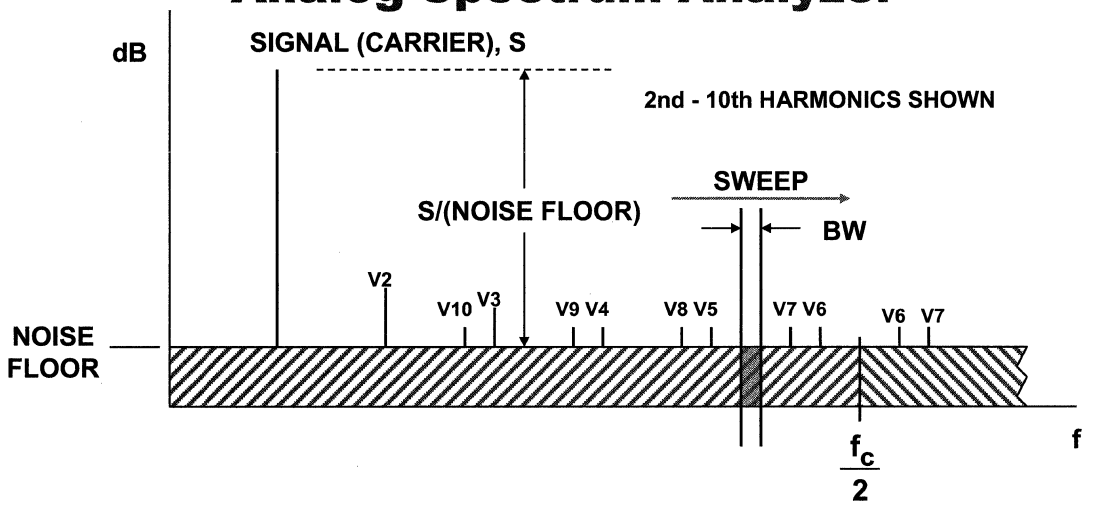

- BW = ANALYZER RESOLUTION BANDWIDTH
- $SNR = S/(NOISE\ FLOOR) - 10\log_{10}\left[\frac{f_c/2}{BW}\right]$

Analog spectrum analyzers are used to measure the distortion and SFDR of high performance DACs. Care must be taken such that the front end of the analyzer is not overdriven by the fundamental signal. If overdrive is a problem, a bandstop (notch) filter can be used to filter out the fundamental signal such that the spurious components can be observed.

Spectrum analyzers can also be used to measure the SNR of a DAC provided attention is given to bandwidth considerations. SNR of an ADC is normally defined as the signal-to-noise ratio measured over the Nyquist bandwidth dc to $f_c/2$. However, spectrum analyzers have a resolution bandwidth which is less than $f_c/2$—this therefore lowers the analyzer noise floor by the *process gain* which is equal to $10 \cdot \log10[f_c/(2 \cdot BW)]$, where BW is the resolution noise bandwidth of the analyzer.

It is important that the noise bandwidth (not the 3-dB bandwidth) be used in the calculation, however the error is small assuming that the analyzer narrowband filter is at least two poles. The ratio of the noise bandwidth to the 3-dB bandwidth of a one-pole Butterworth filter is 1.57 (causing an error of 1.96 dB in the process gain calculation). For a two-pole Butterworth filter, the ratio is 1.11 (causing an error of only 0.45 dB in the process gain calculation). Using more than two poles makes essentially no difference in the effectiveness of the noise filter.

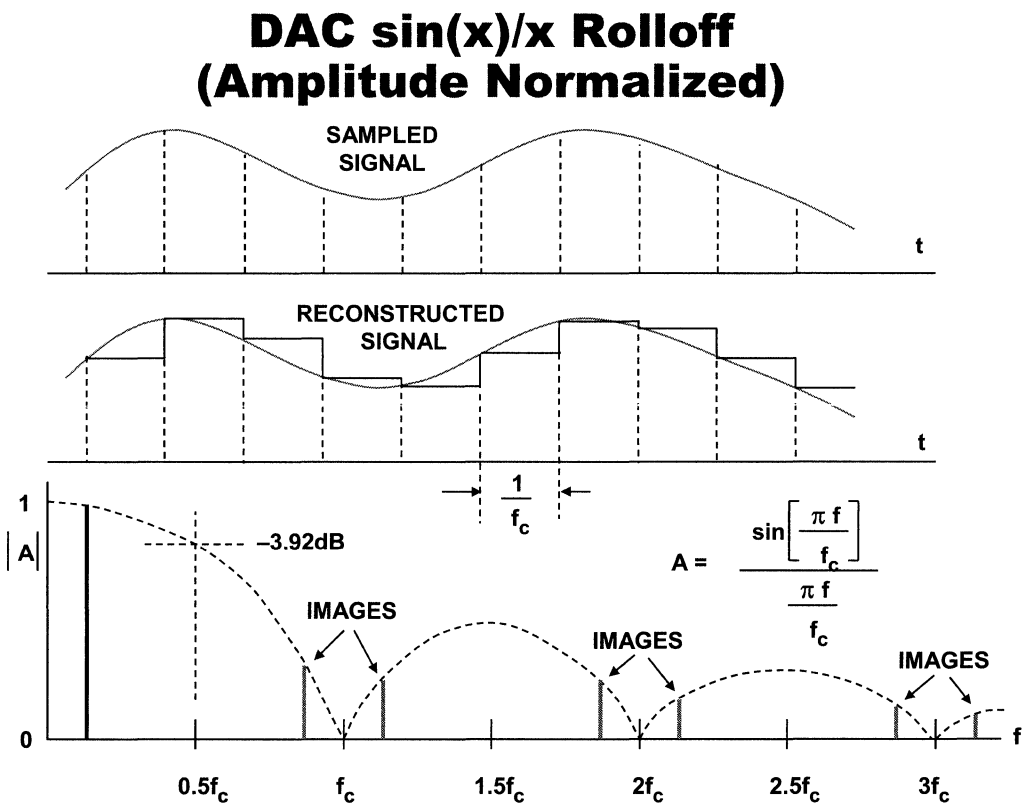

Here we see the effects of the sin(x)/x rolloff which occurs when the reconstructed signal is not composed of ideal impulses, but NRZ analog "boxcar" data as shown. This is often referred to as a "zero-order hold" output and affects all DACs.

The NRZ boxcar output is used because it is easier to generate with more energy and less distortion than data composed of narrow pulses.

At $f_C/2$, the output is attenuated by 3.92dB.

Theoretically, digital upconversion could be achieved by using a bandpass filter to pass one of the higher order images rather than the baseband signal. This method is limited, however, because of the reduced amplitude caused by the sinx/x rolloff. Details of the technique can be found in the following references:

---

1. Ken Gentile, "Digital Upconverter IC Tames Complex Modulation," *Microwaves and RF*, August 2000.

2. Allen Hill and Jim Surber, "Using Aliased-Imaging Techniques in DDS to Generate RF Signals, *RF Design*, September 1993, pp. 31-36.

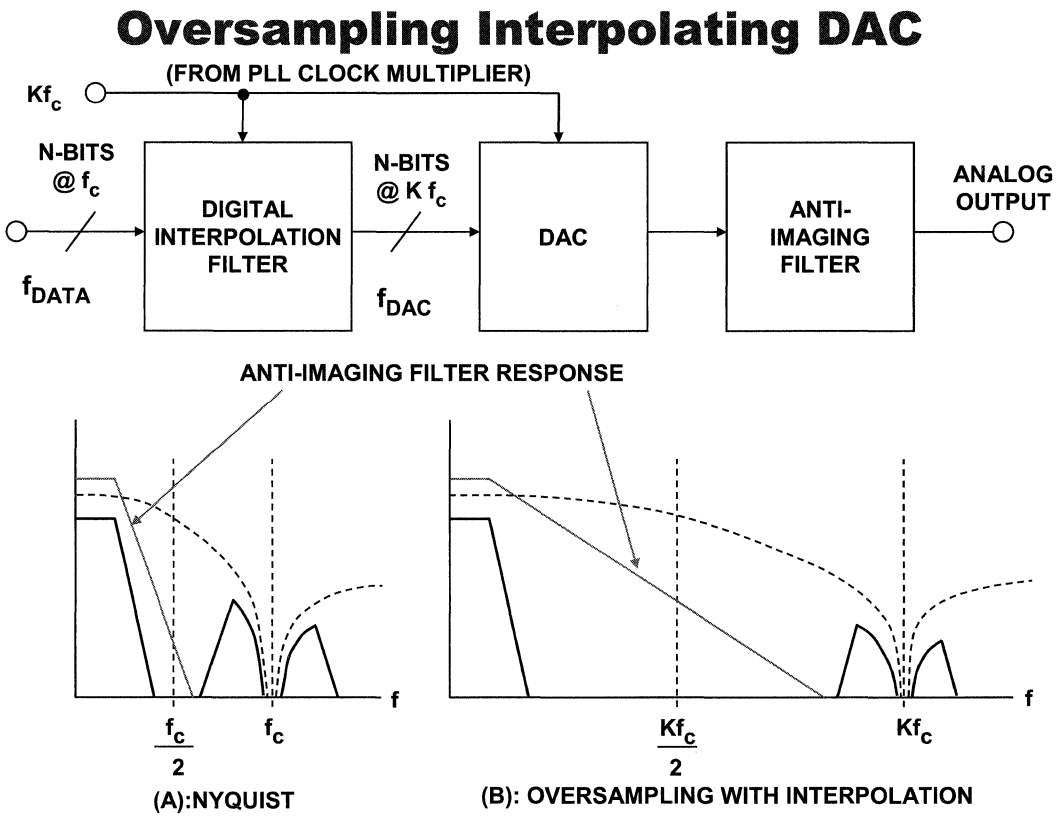

In ADC-based systems, oversampling can ease the requirements on the antialiasing filter. In a DAC-based system (such as DDS), the concept of oversampling and interpolation can be used in a similar manner. This concept is common in digital audio CD players, where the basic update rate of the data from the CD is 44.1kSPS. Early CD players used traditional binary DACs and inserted "Zeros" into the parallel data, thereby increasing the effective update rate to 4-times, 8-times, or 16-times (sometimes more) the fundamental throughput rate. The 4×, 8×, or 16× data stream is passed through a digital interpolation filter which generates the extra data points. The high oversampling rate moves the image frequencies higher, thereby allowing a less complex reconstruction filter with a wider transition band and less phase shift. The sigma-delta 1-bit DAC architecture uses a much higher oversampling rate and represents the ultimate extension of this concept and has become popular in modern CD players.

The same concept of interpolation can be applied to high speed DACs used in communications applications as shown here, relaxing the requirements on the output filter as well as increasing the SNR due to process gain.

Note that the traditional Nyquist condition is shown in (A), where the maximum output frequency of the DAC is generally no more than $f_C/3$. Increasing the DAC output data rate, f_{DAC}, by a factor of K allows a much less complex anti-imaging filter as shown in (B). The digital interpolation filter generates the extra data points as shown in the next figure.

Oversampling and Interpolation in the Time Domain

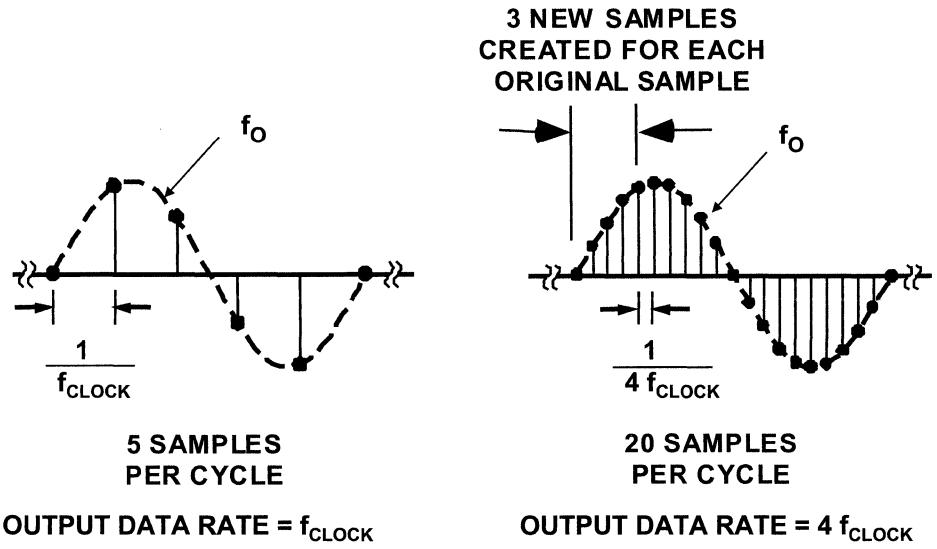

This example shows how oversampling and digital interpolation works in the time domain. The figure on the left shows the basic input data to the interpolating DAC before interpolation. In this example, there are five samples per cycle of the output sinewave. The output data rate is equal to f_{CLOCK}.

The right-hand diagram shows what happens when the interpolation filter is used to increase the sample rate by a factor of 4. Now, there are 20 samples per cycle, and the digital interpolation filter creates three "new" samples for each original sample. The new output data rate is now equal to $4f_{CLOCK}$.

It is important to note that the bandwidth of the output signal is always limited by the Nyquist criteria, which is determined by the initial output data rate, f_{CLOCK}. The interpolation process does not increase the information contained in the original input data stream, it simply adds the extra data points, increases the output data rate, and makes filtering the images an easier problem to solve.

Using the DAC Harmonic Image Tool

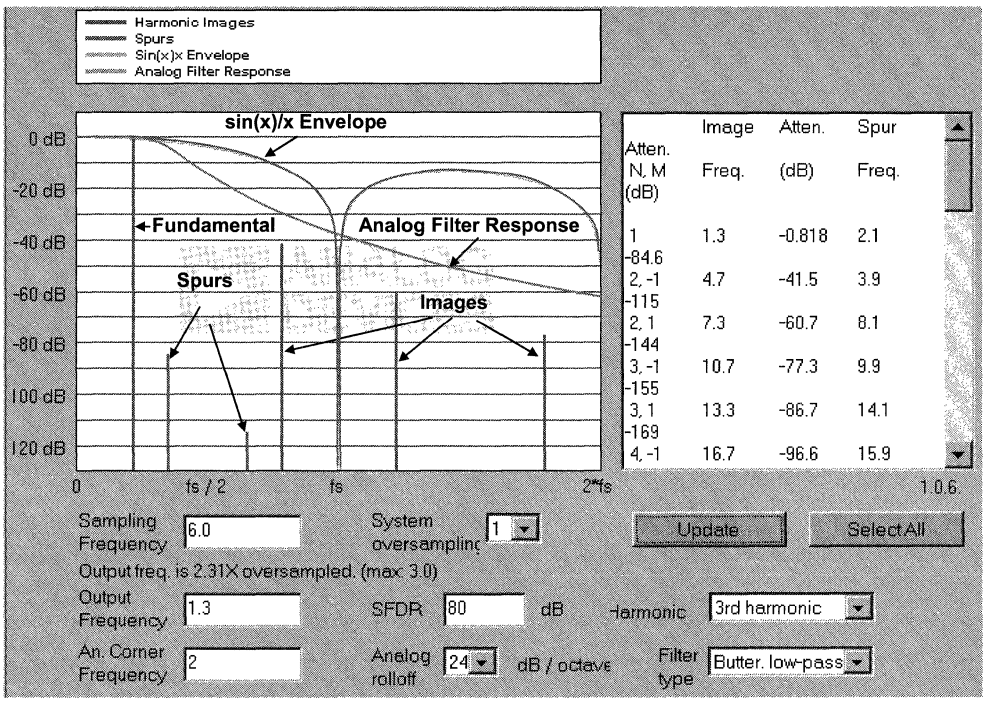

The DAC Harmonic Image tool demonstrates harmonic images and spurs in an idealized DAC output. Second and third harmonics are supported by the tool.

Images are located at N·Fs + M·Fo, where M = ±1. Without external filtering image amplitudes roll off as sin(x)/x where x = π·Fo/Fs. Spurious 2nd or 3rd harmonics of the fundamental are assumed to result from DAC nonlinearities and so remain folded within the baseband. These spurs in turn have their own harmonic images that roll off as sin(x)/x.

To illustrate suppression of these images, the tool can apply a simulated post-DAC analog filter. The output signal is assumed to have been generated at the DAC update rate (or perfectly reconstructed from a lower data rate). Harmonics must be supressed by analog filtering.

The total attenuation of a given harmonic is the sum of the sin(x)/x attenuation and the analog filter attenuation.

Harmonic Images Design Tool (Ideal DAC)
Placement of 2nd Harmonic Spurs

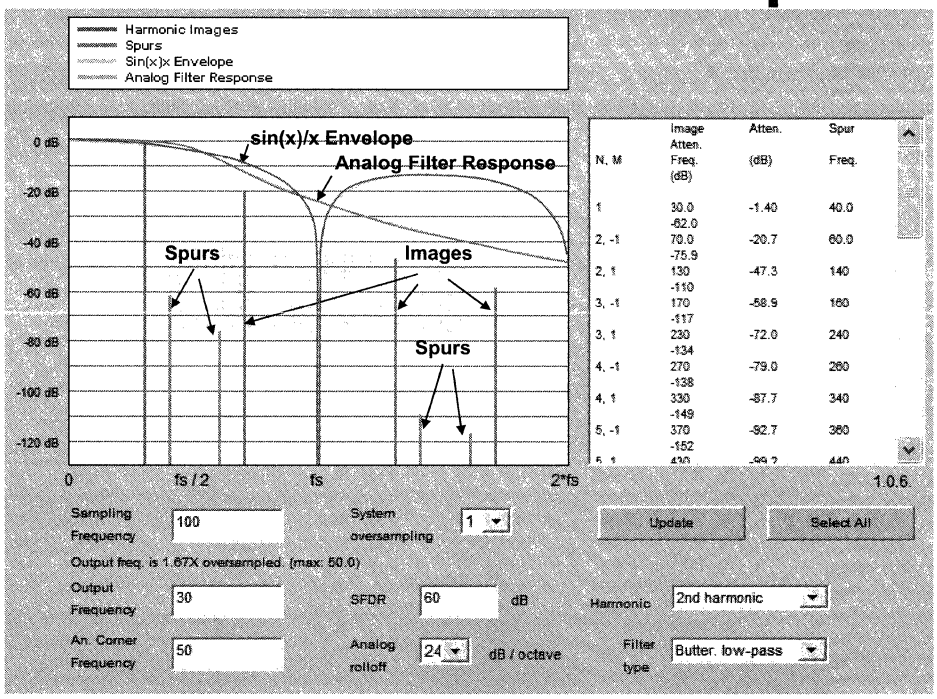

In order to use the tool, enter the Sampling Frequency or Oversampling Ratio in the fields provided.

Hit "Enter" or click "Update" to recompute the display. This slide shows the locations of the second harmonics.

Select the Analog Corner Frequency. Behavior of this hypothetical DAC is assumed frequency-independent, so only frequency ratios matter—the units must be uniform, but are otherwise irrelevant.

Select filter rolloff and type. Butterworth, Chebychev, lowpass and highpass filters are supported with rolloffs up to 48dB/octave. The analog filter response is shown in green.

Image frequencies and amplitudes are shown in a table at top right. The first column gives the N value and whether it's ± Fo (M = ±1). The table data is selectable and can be copied and pasted into a spreadsheet. Use "Select all" to conveniently select the entire contents before copying.

Experiment by changing the analog filter parameters and comparing the results to changing the Oversampling Ratio. Changing the oversampling ratio changes the Sampling Frequency but the reverse is not true—it's sampling frequency that's important and the oversampling ratio menu is just a convenience.

SFDR is used to set the relative level of spurs, which are assumed here to be the result of DAC nonlinearity. Harmonic selects whether these distortion spurs are most prominent at 2× or 3× (default) the Output Frequency. By varying the output frequency in small increments, you can see how the harmonics move and fold.

SFDR is a DAC data sheet parameter and, in general, depends on both the sample rate and output frequency, among other variables. However, a single compromise number is often specified.

Harmonic Images Design Tool (Ideal DAC)
Placement of 3rd Harmonic Spurs

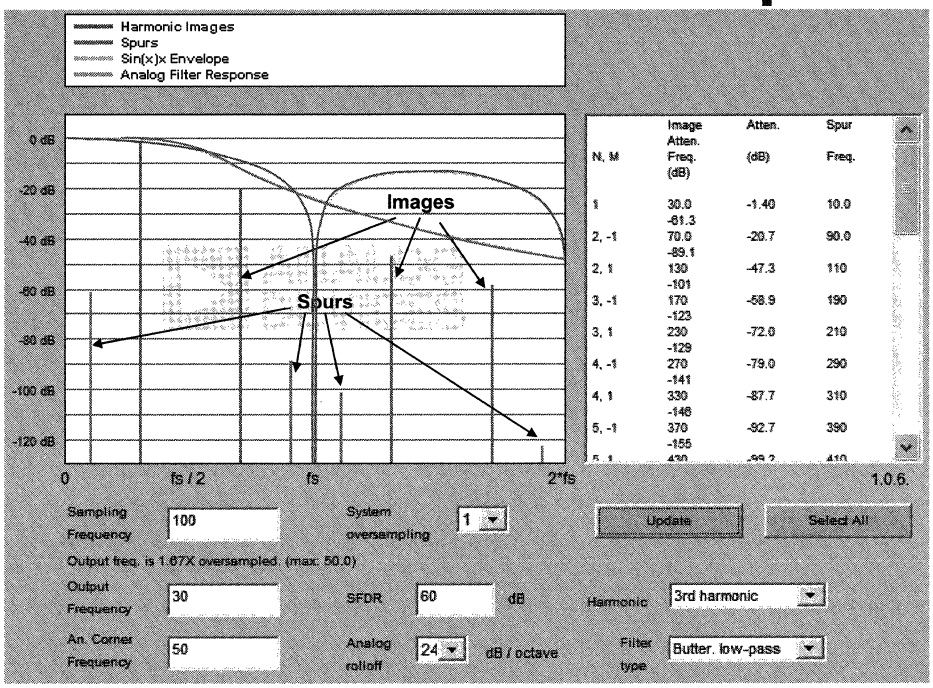

This figure is for the same conditions, but shows the placement of the third harmonics.

Again, the program assumes that the level of the harmonics is equal to the SFDR of 60dB.

For instance, the harmonic at 90MHz is attenuated by the analog filter attenuation (approximately 20dB) and the sin(x)/x rolloff (approximately 20dB) for a total attenuation of 40dB, placing the spur level at 60 + 40 = 100dB below fullscale.

Harmonic Images Tool (AD9777 Interpolating TxDAC)

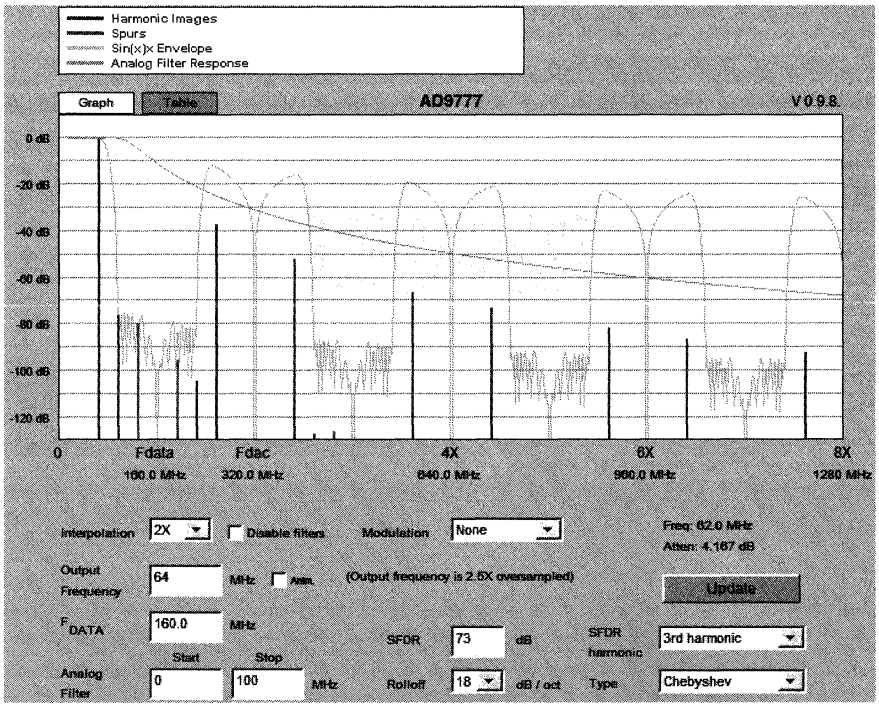

The tool above shows the harmonic images and spurs for single frequency output from an AD9777. The model of the AD9777 is simplified and idealized—only SFDR is modeled, and it is assumed frequency-independent. See the data sheet for actual performance data.

For an ordinary DAC (see previous DAC Harmonic Images Calculator), images are located at $N \cdot F_{DAC} \pm F_{OUT}$ and follow a sin(x)/x envelope (shown by default). The AD9777 contains an integral interpolator which doubles, quadruples, or octuples the input data rate by stuffing zeros between successive samples and then (optionally) filtering the result. Both the filtered and unfiltered images then create further images and spurs at the DAC data rate, according to the $N \cdot F_{DAC} \pm F_{OUT}$ rule. The magnitude response of the AD9777 combining its internal interpolation with the sin(x)/x envelope is shown.

The AD9777 can optionally modulate the filtered data stream by F_{DAC} / 2, 4, or 8, shifting the output spectrum and creating mirror images around the modulating frequency. These images follow separate magnitude envelopes so two envelopes are shown in different colors in this case. This release only models real modulation. Complex modulation is planned for a future release.

Spurious second or third harmonics of each image are assumed to result from D/A conversion nonlinearities and so are folded within the first Nyquist Zone (NZ) of F_{DAC}. These spurs then have their own harmonic images that roll off as sin(x)/x (where $x = \pi \cdot F_{SPUR}/F_{DAC}$).

To show external selection/suppression of desired/undesired images and spurs, the tool can apply a simulated post-DAC analog filter.

DAC Applications in Transmitters

www.analog.com/txdacs

Applying Nyquist's Criteria to Interpolating and Modulating DACs

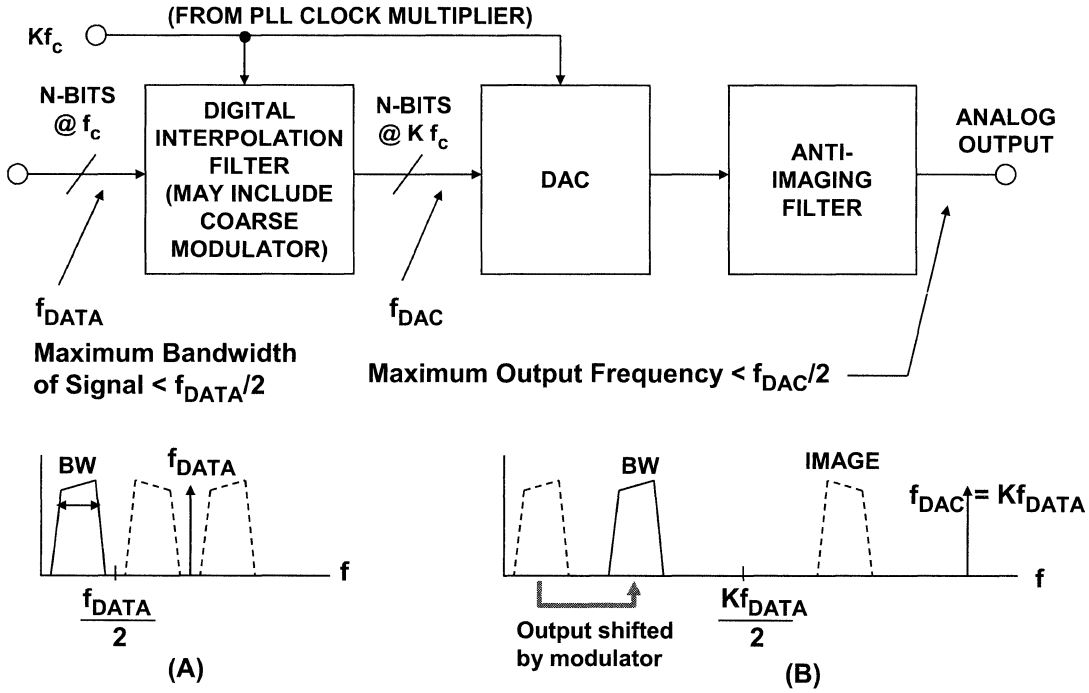

Before starting the discussion on DACs in transmitters, it is important to expand upon a thought mentioned in a previous figure. Nyquist's criterion still applies regardless of interpolation or modulation!

The diagram labeled (A) shows the traditional case of data generated at a rate equal to f_{DATA} generating a signal having a bandwidth BW. Nyquist's criteria says that BW must be less than $f_{DATA}/2$ as shown.

The effects of digital interpolation and modulation are shown in (B). In addition to oversampling and interpolation, the spectrum of the signal has been shifted using digital modulation. Even so, the bandwidth of the signal is still BW—limited by the original Nyquist bandwidth, $f_{DATA}/2$.

It is also important to note that the maximum output frequency is limited by the rate at which the DAC is updated, f_{DAC}. Typically the maximum usable output frequency is approximately $f_{DAC}/3$, because of the non-ideal anti-imaging filter.

Although it is possible to use the higher images of the DAC output as the actual signal (the main baseband signal as well as the other images must be remove by filtering), the images suffer from attenuation due to the sin(x)/x rolloff, and may not provide sufficient amplitude to be usable. Therefore, this is not a popular approach.

In summary, the maximum bandwidth of the final DAC output signal is determined by one-half the input data rate, f_{DATA}, and the maximum IF output frequency by $f_{DAC}/2$.

For example, a 4× interpolating DAC which accepts input data at 250MSPS (f_{DATA} = 250MSPS) and outputs data at 1GSPS (f_{DAC} = 1GSPS) could reasonably reconstruct a signal having a bandwidth of 83MHz centered about an IF frequency of 300MHz.

Categories of DACs in the TxDAC® Family

- **Fast LVDS DACs: eg., AD9736 1.2GSPS, 14-bits, 2× interpolation**
  - **f_{DATA} (max) = f_{DAC} (max) = 1.2GSPS**
- **Interpolating DACs**
- **Dual DACs**
- **Dual Interpolating DAC with coarse digital modulation**
  - **f_{DATA} (max) = 250MSPS (CMOS), f_{DAC} (max) = 1GSPS, 2×, 4×, 8× interpolation**
- **Mixer DAC with on-chip digital quadrature modulator: AD9957**
  - **f_{DATA} (max) = 250MHz (CMOS), f_{DAC} (max) = 1.2GSPS, 4× to 252× interpolation (factors of 4 only)**
- ***http://www.analog.com/txdacs***

The TxDAC® family of CMOS DACs from Analog Devices has grown rapidly in the last 10 years. Early products in the family provided the basic DAC with very little added digital functionality. Smaller geometry CMOS processes have allowed the addition of a large number of digital features to the family.

Modern members of the TxDAC family offer such important features as digital interpolation, modulation, dual DACs, and "mixer" DACs containing on-chip digital quadrature modulators.

The product family is currently too broad to do more than describe a few key new technology leading products and refer the reader to

http://www.analog.com/txdacs

For input data rates greater than about 250MSPS, LVDS receivers are utilized rather than traditional single-ended CMOS. For instance, the AD9736 accepts LVDS input data at a rate up to 1.2 GSPS. DACs such as these provide on-chip data alignment circuitry to ensure that input data is clocked at the proper time.

Two Popular Methods for RF Upconversion

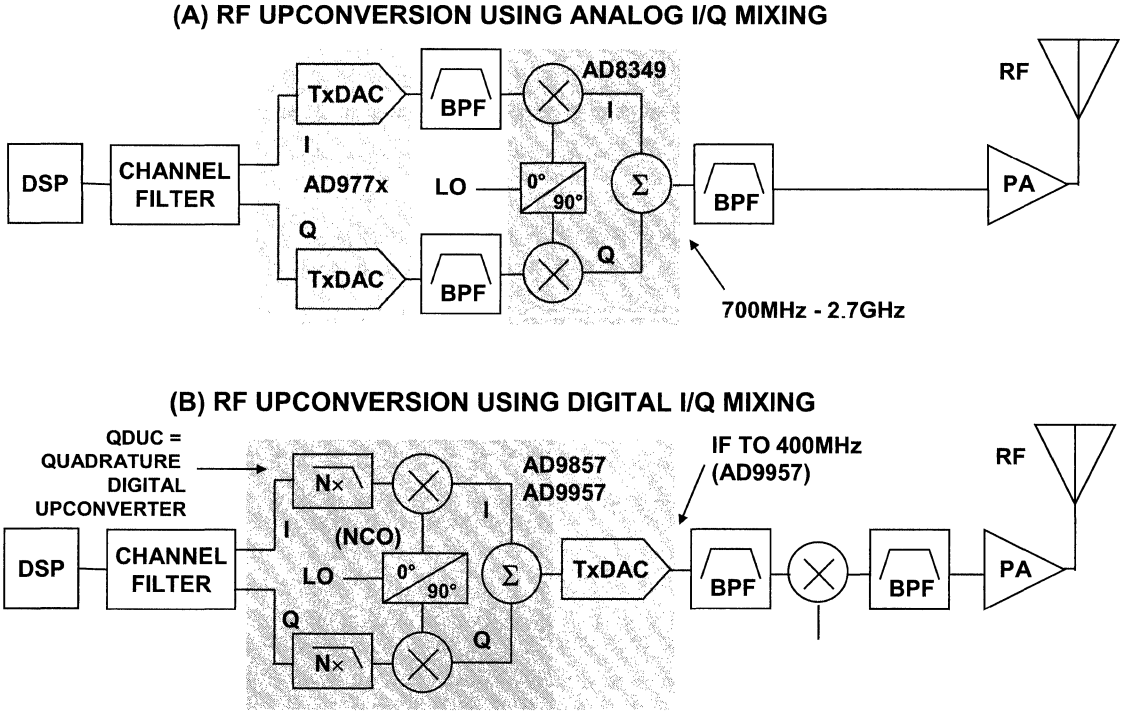

This figure shows two popular methods for performing RF upconversion in a digital transmitter. Traditional analog I/Q mixing is utilized in (A) using a dual DAC (AD977x) and an analog I/Q modulator (AD8349). The output of the quadrature modulator is at the RF frequency and is passed on to the power amplifier.

In (B) the digital I/Q data is passed through a quadrature digital upconverter (QDUC) such as the AD9857 (200MHz) or AD9957 (400MHz) and then to a single on-chip DAC which provides an IF frequency output (up to 400MHz for the AD9957). The IF frequency then passes through a final stage of upconversion using traditional analog techniques. Devices such as the AD9857/AD9957 are commonly referred to as "mixer" DACs.

Both techniques are currently being used in transmitters. Analog I/Q mixing produces the RF signal directly. Digital I/Q mixing requires an additional upconversion stage, but eliminates the need for an analog I/Q mixer.

AD9779 Multicarrier WCDMA Signal, 4× Interpolation, f_{DATA} = 122.88 MSPS, f_{DAC}/4 Modulation

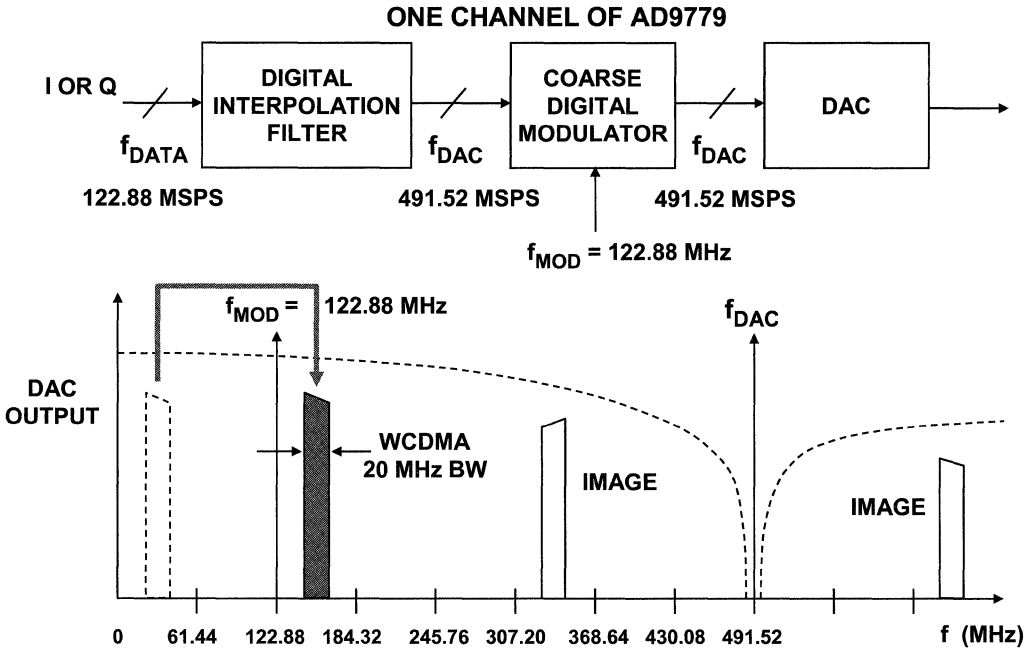

The addition of a digital modulator to an interpolating DAC allows the baseband output IF signal to be positioned in the desired Nyquist zone. In this case, the AD9779 16-bit DAC input data rate is 122.88MSPS. The signal bandwidth is 20MHz which corresponds to a wideband CDMA multicarrier signal (WCDMA) The digital interpolation filter then increases the data rate by 4× to 491.52MSPS.

The on-chip coarse digital modulator in the AD9779 then serves to place the WCDMA signal in the third Nyquist zone referenced to the input data rate.

The AD9776/AD9778/AD9779 are dual 12/14/16-bit DACs capable of output data rates up to 1GSPS with input data rates to 250MSPS.

AD9779 Multicarrier WCDMA Signal, 4× Interpolation, f_{DATA} = 122.88 MSPS, f_{DAC}/4 Modulation

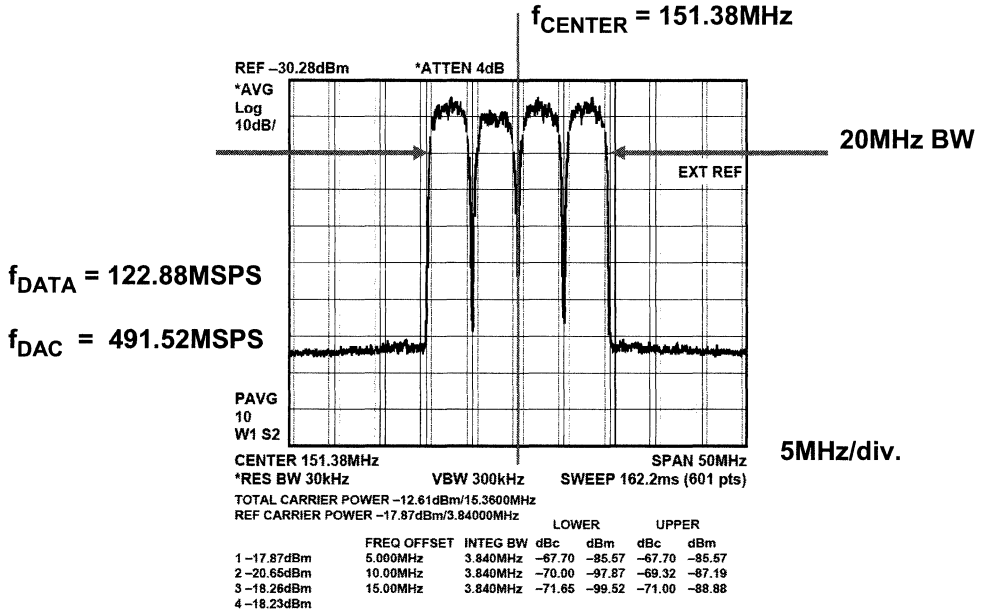

This figure shows the actual analog output of the AD9779 reconstructing a WCDMA signal with an input data rate of 122.88MSPS, and an output data rate of 491.52MSPS. The digital modulator is used to position the signal in the third Nyquist zone referenced to the input data rate.

A functional diagram of the AD9779 family is shown in the next figure.

AD9776/AD9778/AD9779 12/14/16-Bit Dual 1GSPS DACs

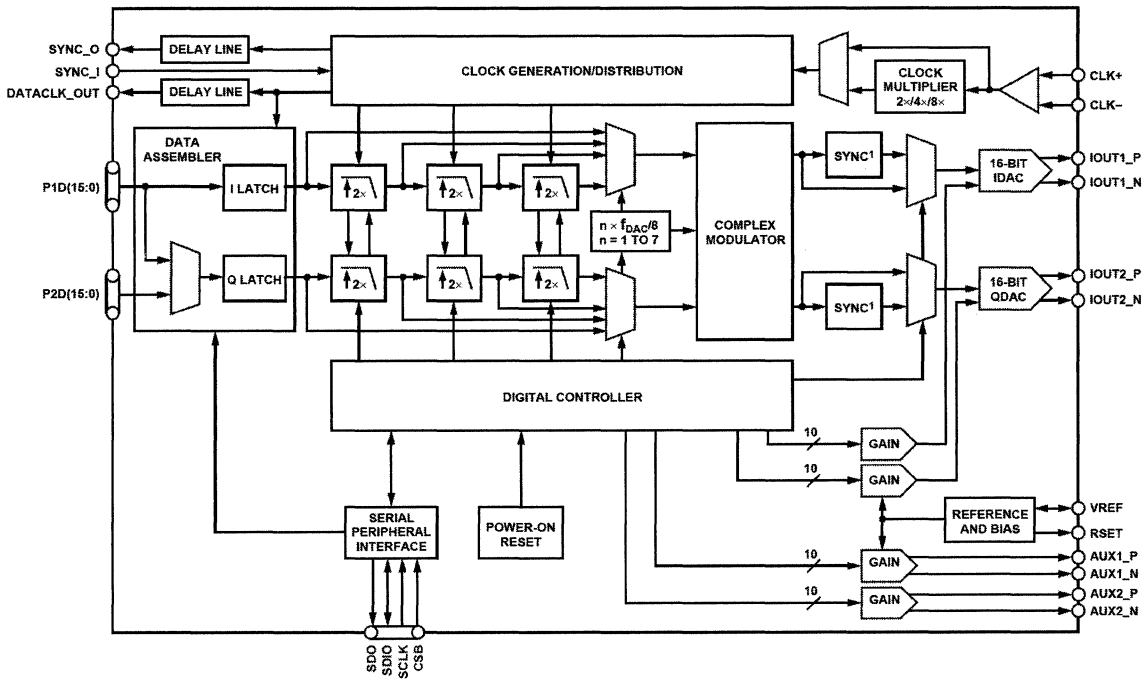

The AD9776/AD9778/AD9779 are dual, 12-/14-/16-bit, high dynamic range DACs that provide an interpolated sample rate of 1 GSPS, thus permitting multicarrier signal generation up to high IF frequencies. As previously discussed, they include features optimized for direct conversion transmit applications, including complex digital modulation, and gain and offset compensation. The DAC outputs are optimized to interface seamlessly with analog quadrature modulators, such as the AD8349.

A serial peripheral interface (SPI) provides for programming/readback of many internal parameters. The output current can be programmed over a range of 10mA to 30mA. The devices are manufactured on an advanced 0.18µm CMOS process and operate from 1.8V and 3.3V supplies for a total power consumption of 1.0W. They are enclosed in 100-lead TQFP packages.

The devices provide an SFDR of 78dBc for output frequencies up to 100MHz and have a single carrier WCDMA adjacent channel leakage ratio (ACLR) = 79 dBc @ 80MHz IF.

AD9957 1 GSPS Quadrature Digital Upconverter (QDUC) "Mixer" DAC

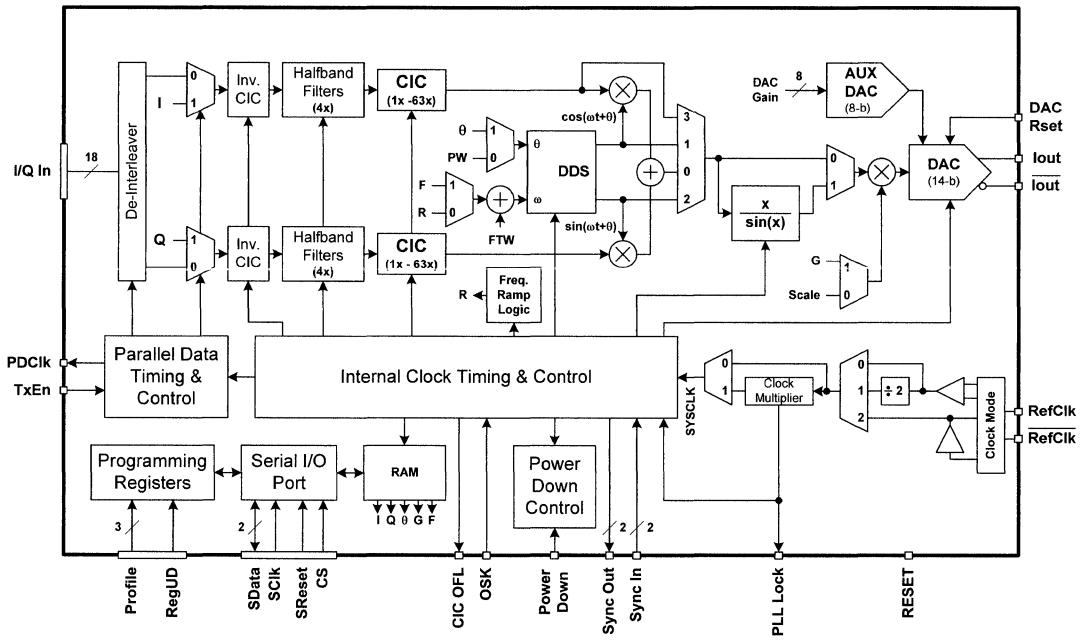

The AD9957 is a 1GSPS quadrature digital upconverter capable of generating IF frequencies up to 400MHz. Input data can be accepted on an 18-bit parallel I/Q port at a rate up to 250MHz. The AD9957 allows interpolation rates from 4× to 252×.

The on-chip numerically controlled oscillator (NCO) uses a 32-bit frequency tuning word.

An auxiliary DAC is also included which can be used to provide some on-ramping to the output power of the DAC without sacrificing any of the dynamic range.

The AD9957 is packaged in a 100-lead TQFP.

The following reference describes the details of the operation of the AD9857, the 200MHz predecessor of the AD9957. The parts share many similar operational characteristics.

---

1. Ken Gentile, "Digital Upconverter IC Tames Complex Modulation," *Microwaves and RF*, August 2000.

Buffering DAC Outputs

www.analog.com/txdacs
www.analog.com/amps
www.analog.com/diffamps

Generalized Model of a High Speed CMOS DAC Output such as the AD978x and AD977x Series

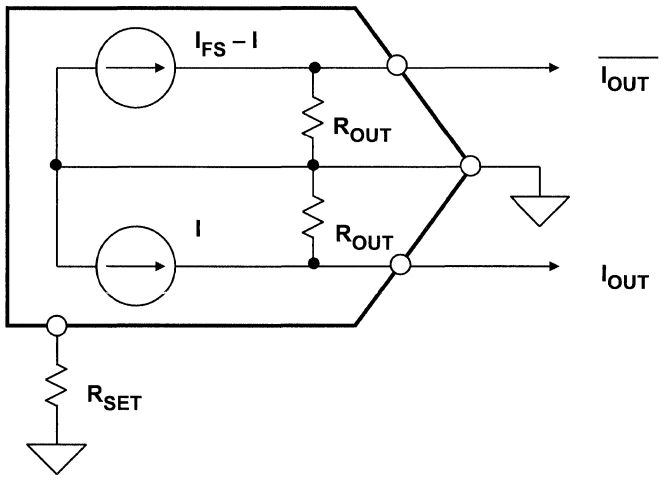

- I_{FS} 2 - 20mA typical
- R_{OUT} > 100kΩ
- Output compliance voltage < ±1V for best performance

Now, we will look at some issues relating to high-speed DAC outputs. This figure shows the equivalent output circuit for the Analog Devices' CMOS TxDAC family.

Outputs are differential currents which can be set from 2 to 20mA FS by an external RSET resistor.

Output impedance is greater than 100kΩ, and the output compliance voltage is ±1V. Best performance in terms of SNR and SFDR is generally obtained with an output voltage between 0.5V p-p and 1V p-p (10mA to 20mA into 50Ω. The next few figures show typical output circuits.

Differential Transformer Coupling

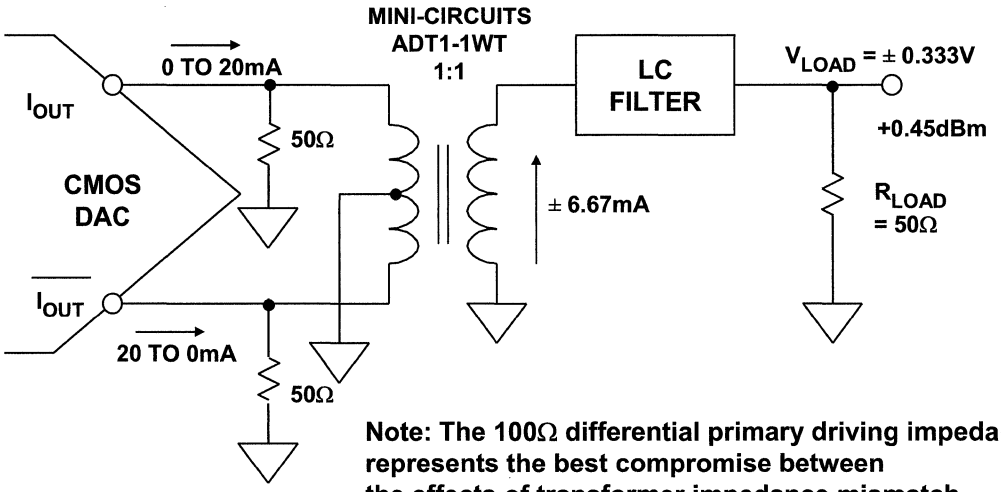

Note: The 100Ω differential primary driving impedance represents the best compromise between the effects of transformer impedance mismatch and DAC SNR performance.

In this application, the differential outputs are configured to drive the primary winding of a transformer. This converts the differential output signal into a single-ended signal.

Note that the 100Ω primary driving impedance was chosen in order to optimize the DAC SNR and SFDR. The resulting impedance mismatch does not significantly affect the performance of the transformer.

The secondary winding single-ended load is 50Ω. This is reflected back to the primary as a 25Ω load on each of the differential DAC outputs. The effective dc load on each DAC output is therefore 50Ω||25Ω = 16.7Ω.

The resulting single-ended output signal level is ±333mV p-p into a 50Ω load, which is +0.45dBm.

Transformer Coupling out of the AD9786 on Evaluation Board

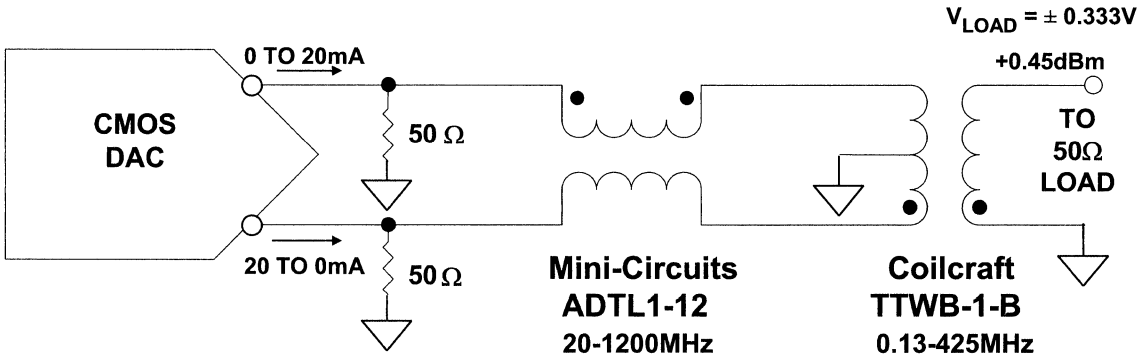

- **Transmission Line Transformer in series with outputs to help cancel HD2**
- **RF Transformer from Coilcraft (TTWB-1-B) shows better performance for IFs at 200-300 MHz**

In order to obtain the best possible second harmonic distortion performance at high IF frequencies (generally > 100MHz), a double transformer output can be used as shown in this figure. The double transformer configuration helps eliminate the unbalance caused by the parasitic capacitance between the primary and secondary windings of the transformers.

In some cases, equivalent performance can be achieved with a more expensive single high performance transformer. In any case, some experimentation may be required to achieve optimum results in demanding applications.

Differential DC Coupling Using a Dual-Supply Op Amp

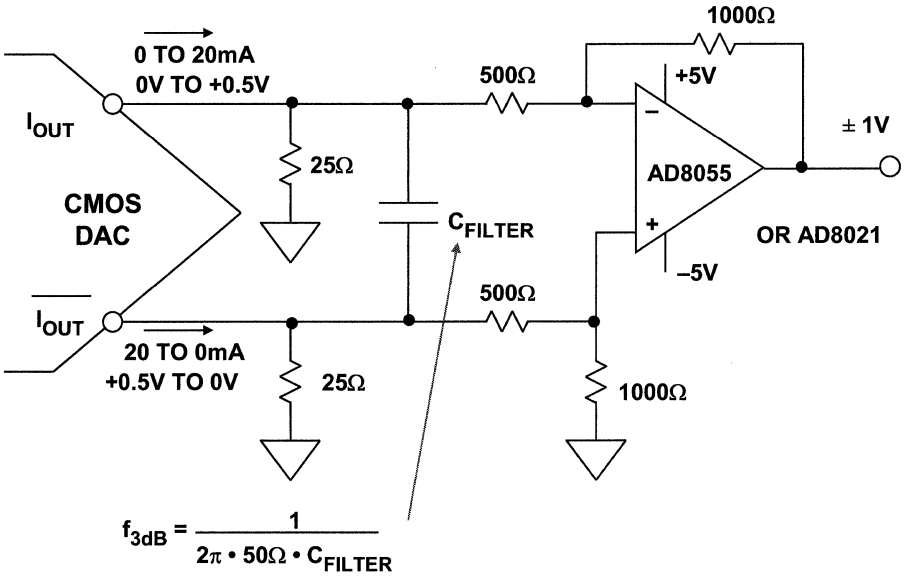

The output current drive for most members of the TxDAC family can be set for up to 20mA with an external resistor. In most applications this results in sufficient load power so that additional output buffering is not required.

For applications which do require an output buffer, a high speed op amp can be used to perform the differential to single-ended conversion as shown in this figure.

Note that the op amp is not used directly as an I/V converter, but is configured as a "difference" amplifier to amplify the voltage developed across the 25Ω load resistors. The amplifier gain is set for a factor of 2 which develops a final output voltage of 2V p-p. Since the output swings above and below ground, a dual-supply op amp is required.

The C_{FILTER} capacitor forms a differential filter with the equivalent 50Ω differential output load. This filter reduces any slew-induced distortion of the op amp as well as broadband noise, and the optimum cutoff frequency of the filter is determined empirically to give the best tradeoff between SFDR and SNR while maintaining the required bandwidth.

The op amp must be carefully selected such that it meets the overall system noise and distortion requirements. In some cases at high IF frequencies, the optimum performance can only be achieved with a transformer output driver as previously discussed.

Differential DC Coupling with a Single-Supply Op Amp

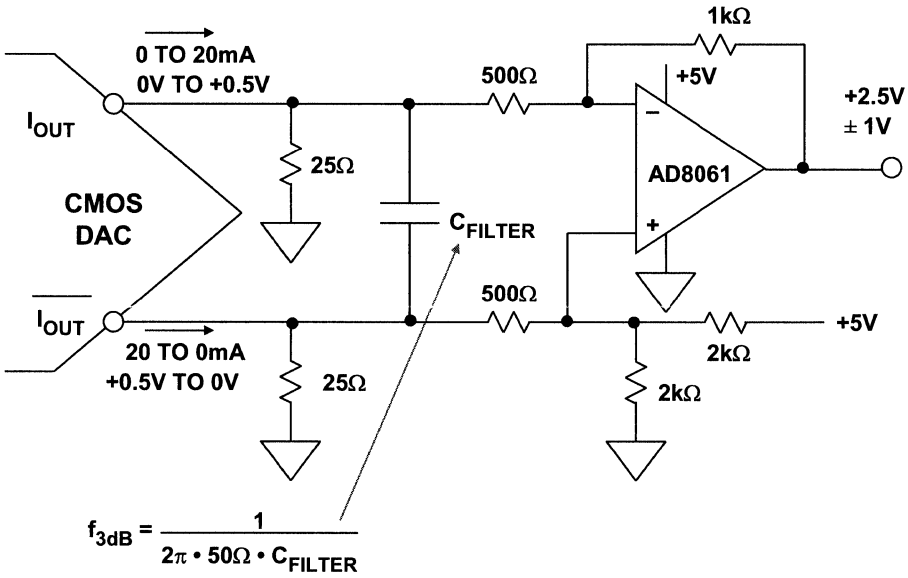

A modified form of the the circuit in the previous figure can be operated on a single supply, provided the common-mode voltage of the op amp is set to mid-supply (+2.5V). This is shown here, where the AD8061 op amp is used.

The input common-mode voltage of the AD8061 is –0.2V to +3.2V when operating on a +5V supply. The output common-mode range is 0.3V to +4.5V when the output is terminated with a 150Ω resistor connected to mid-supply.

The output voltage of the circuit is 2V p-p centered around a common-mode voltage of +2.5V. This common-mode voltage is developed from the +5V supply using a resistor divider, and the supply must be heavily decoupled to prevent amplification of the power supply noise.

An alternative is to eliminate the 2kΩ resistor voltage divider completely and connect the non-inverting input of the op amp to a 2.5V voltage reference through a 1kΩ resistor.

Buffering High-Speed DAC Outputs Using the AD813x or ADA493x Differential Op Amps

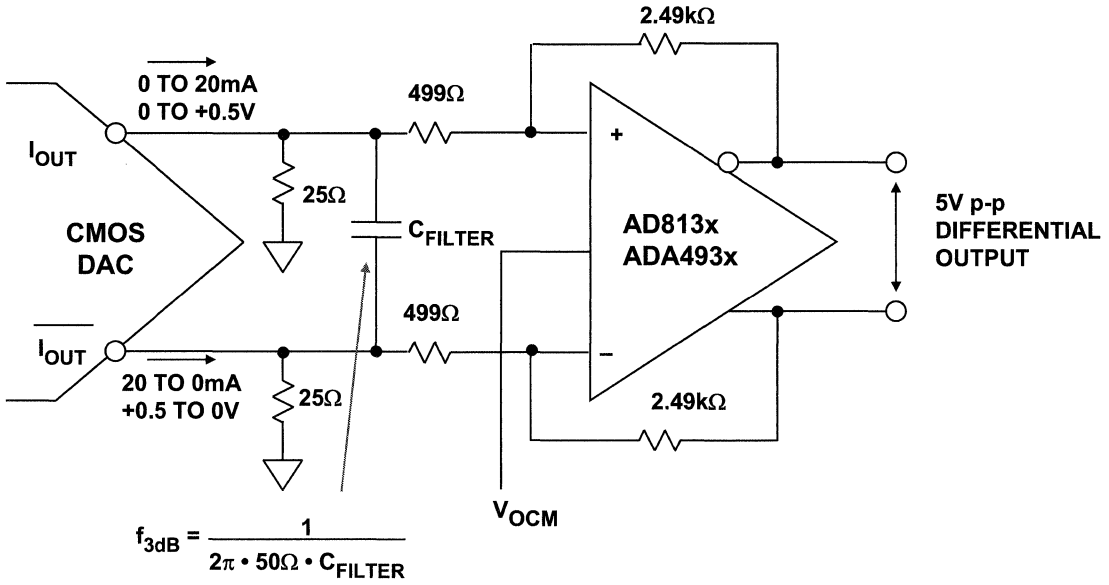

If a buffered differential voltage output is required from a current output TxDAC, the AD813x- or ADA493x-series of differential amplifiers can be used as shown here.

The DAC output current is first converted into a voltage that is developed across the 25Ω resistors. The voltage is amplified by a factor of 5 using the AD813x. This technique is used in lieu of a direct I/V conversion to prevent fast slewing DAC currents from overloading the amplifier and introducing distortion.

The V_{OCM} input on the AD813x can be used to set a final output common-mode voltage within the range of the AD813x. Adding a pair of series output resistors will allow transmission lines to be driven.

As a final note of caution regarding DAC output amplifier buffers, it is extremely important to carefully select the amplifier based on the bandwidth, noise, and distortion requirements of the system. At high IF frequencies, the only acceptable solution may be to either use the DAC current outputs directly or use a suitable RF transformer.

DAC Evaluation Hardware and Software

www.analog.com/txdacs

High Speed Converter Group DAC Bench Testing System

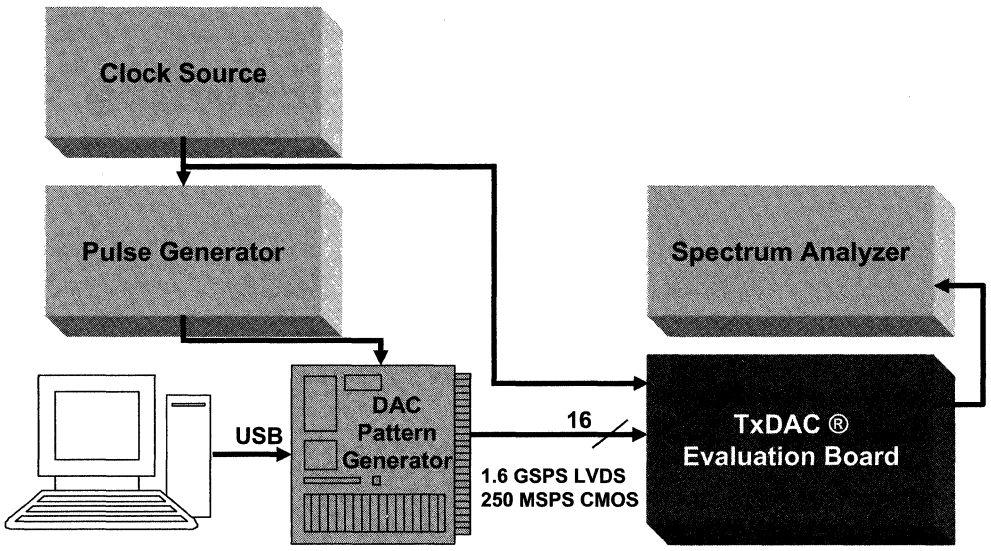

This figure shows a typical test setup for measuring the distortion and noise of a DAC. The first consideration, of course, is the generation of the digital signal to drive the DAC. To achieve this, modern arbitrary waveform generators (for example, Tektronix AWG2021 with Option 4) or word generators (Tektronix DG2020) allow almost any waveform to be synthesized digitally in software, and are mandatory in serious frequency domain testing of DACs. In most cases, these generators have standard waveforms pre-programmed, such as sinewaves and triangle waves, for example. In many communications applications, however, more complex digital waveforms are required, such as two-tone or multi-tone sinewaves, QAM, GSM, and CDMA test signals, etc. In many cases, application-specific hardware and software exists for generating these types of signals and can greatly speed up the evaluation process. Analog Devices offers a DAC Pattern Generator which will be described shortly.

The spectrum analyzer chosen to measure the distortion and noise performance of the DAC should have at least 10dB more dynamic range than the DAC being tested. The "maximum intermodulation-free range" specification of the spectrum analyzer is an excellent indicator of distortion performance. However, spectrum analyzer manufacturers may specify distortion performance in other ways. Modern communications DACs such as the TxDAC®-series require high performance spectrum analyzers such as the Rhode and Schwartz FSEA30. As in the case of oscilloscopes, the spectrum analyzer must not be sensitive to overdrive. This can be easily verified by applying a signal corresponding to the full-scale DAC output, measuring the level of the harmonic distortion products, and then attenuating the signal by 6dB or so and verifying that both the signal and the harmonics drop by the same amount. If the harmonics drop more than the fundamental signal drops, then the analyzer is distorting the signal. In some cases, an analyzer with less than optimum overdrive performance can still be used by placing a bandstop filter in series with the analyzer input to remove the frequency of the fundamental signal being measured. The analyzer looks only at the remaining distortion products. This technique will generally work satisfactorily, provided the attenuation of the bandstop filter is taken into account when making the distortion measurements. Obviously, a separate bandstop filter is required for each individual output frequency tested, and therefore multi-tone testing is cumbersome.

High Speed DAC Pattern Generator (DPG)

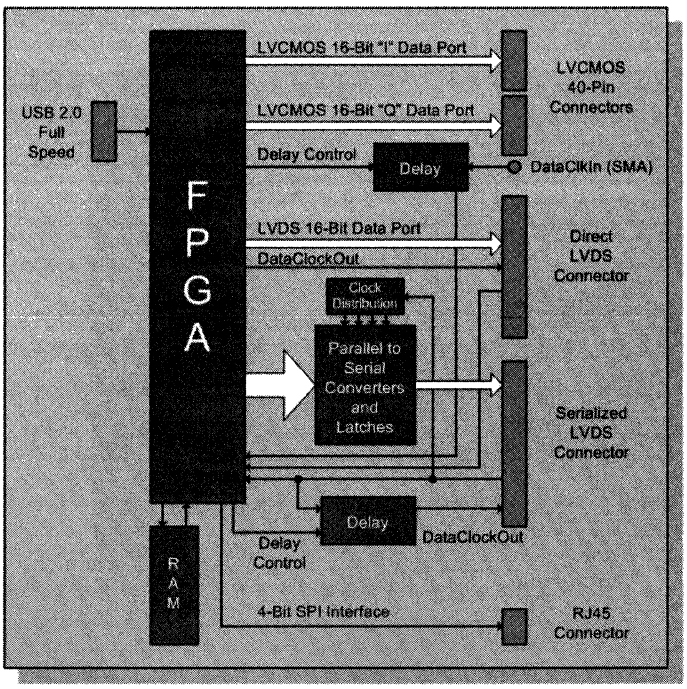

- Provides high-speed digital data to ADI's DAC Evaluation Boards
- CMOS and LVDS data formats supported
- Up to 1.6 GSPS in LVDS mode, 250 MSPS in CMOS mode
- 512 MB RAM for complex waveform generation

The Analog Devices high speed DAC Pattern Generator (DPG) provides high speed digital data to ADI's DAC evaluation boards. Both CMOS and LVDS data formats are supported.

The board can output data at a rate up to 1.6GSPS in the LVDS mode and 250MSPS in the CMOS mode.

There is 512MB of on-board RAM for complex waveform generation.

High Speed DAC Pattern Generator (DPG)

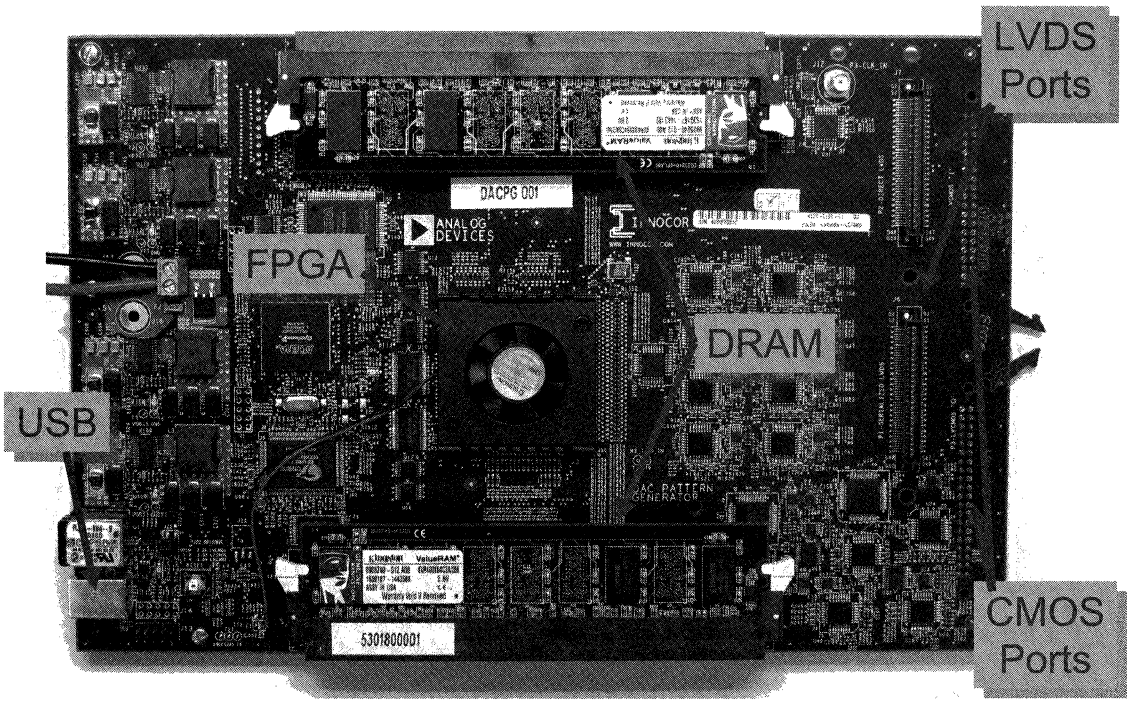

This figure shows a photograph of the DAC Pattern Generator (DPG) board. Note the locations of the USB, LVDS, and CMOS data ports. The DPG interfaces to the PC via a standard USB port.

High Speed DAC Pattern Generator (DPG)

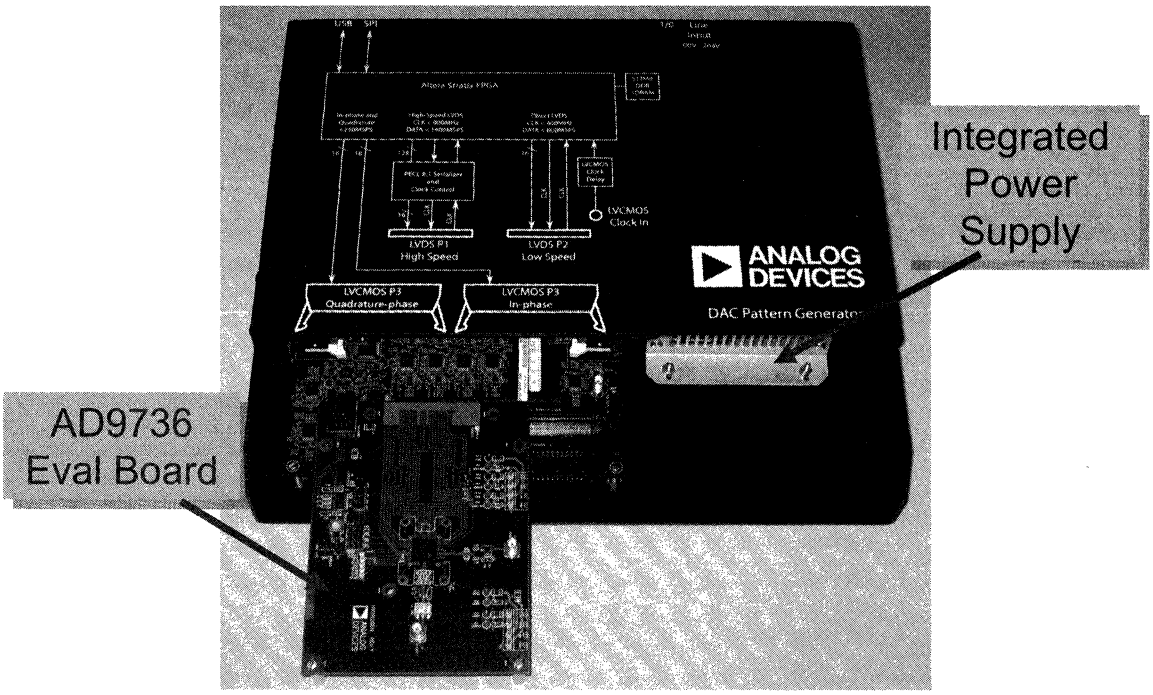

The DAC Pattern Generator board is designed to interface with the individual TxDAC evaluation boards as shown here for the AD9736. The DPG is shown here with an integrated power supply.

VisualDAC™

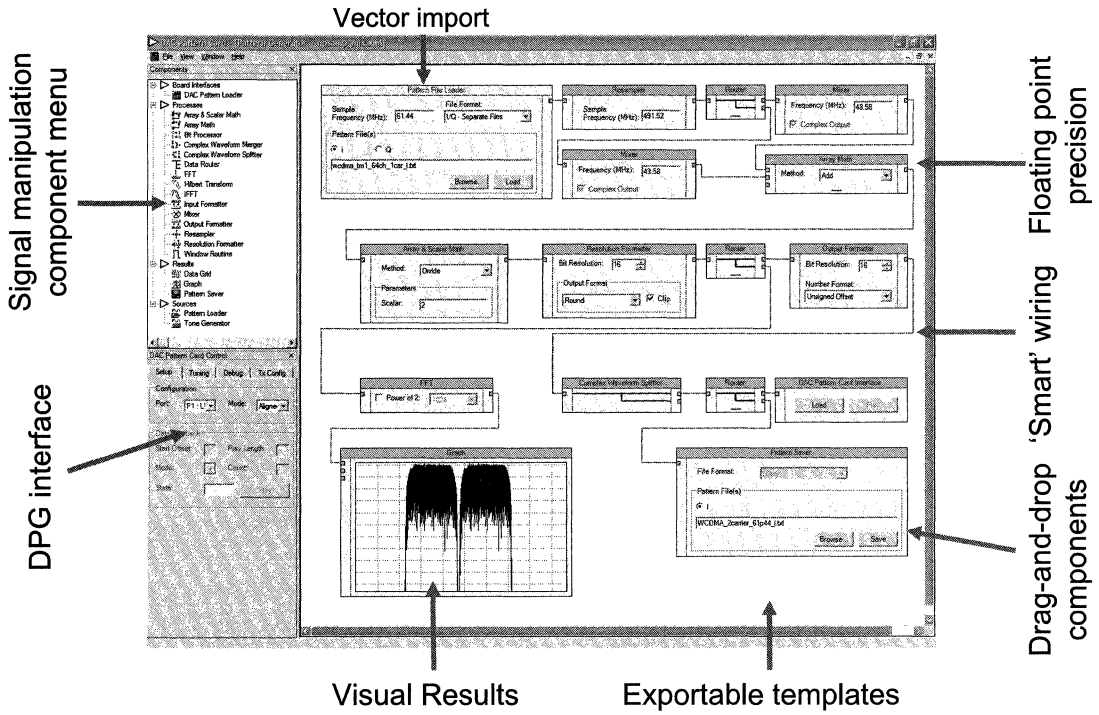

The DPG board is operated using VisualDAC™ software which is a graphical user interface as shown in this figure. The software is used to generate the various complex waveforms required to test the TxDAC series.

Direct Digital Synthesis

www.analog.com/dds

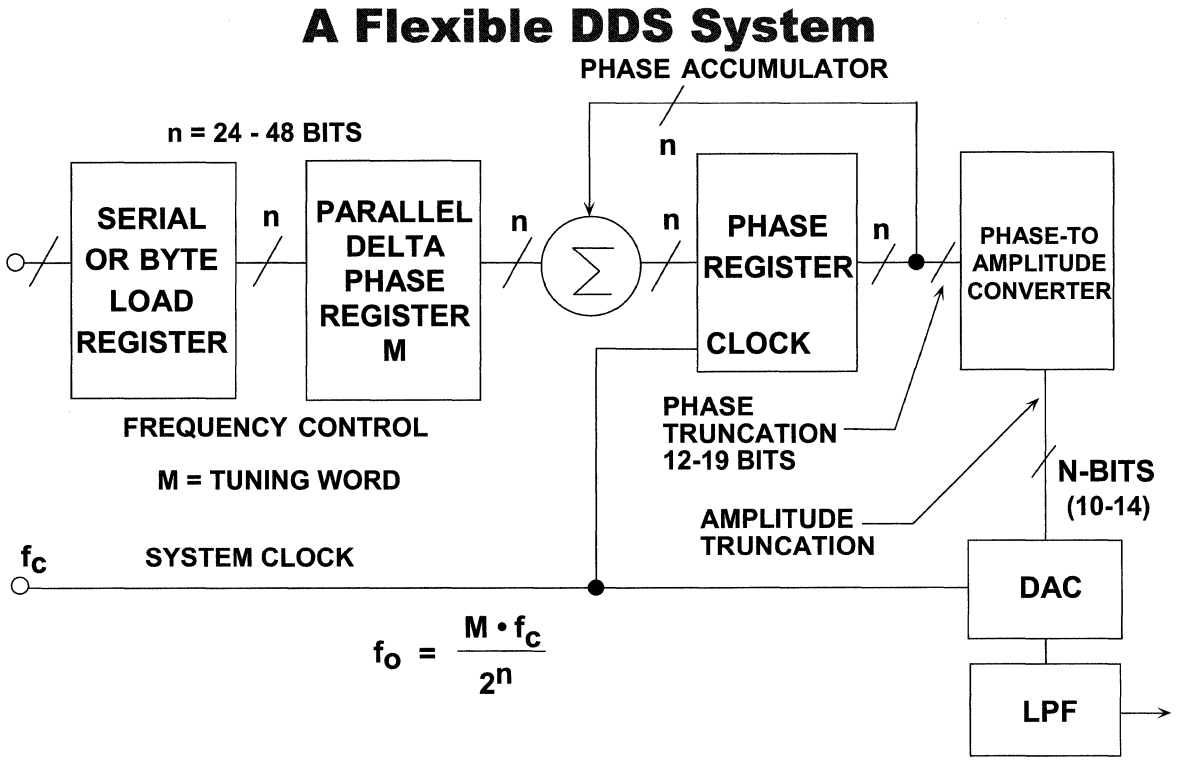

The heart of the DDS system is the phase accumulator whose contents are updated once each clock cycle. Each time the phase accumulator is updated, the digital number, M, stored in the delta phase register is added to the number in the phase accumulator register. Assume that the number in the delta phase register is 00...01 and that the initial contents of the phase accumulator is 00...00. The phase accumulator is updated by 00...01 on each clock cycle. If the accumulator is 32-bits wide, 232 clock cycles (over 4 billion) are required before the phase accumulator returns to 00...00, and the cycle repeats.

The truncated output of the phase accumulator serves as the address to a sine (or cosine) lookup table. Each address in the lookup table corresponds to a phase point on the sinewave from 0° to 360°. The lookup table contains the corresponding digital amplitude information for one complete cycle of a sinewave. The lookup table therefore maps the phase information from the phase accumulator into a digital amplitude word, which in turn drives the DAC. In practice, only data for 90° is required because the quadrature data is contained in the two MSBs. In order to further reduce the size of the lookup tables, various proprietary algorithms have been developed to compute the sine values, however, the fundamental concept is still the same.

For an n-bit phase accumulator (n generally ranges from 24 to 32 in most DDS systems), there are 2^n possible phase points. The digital word in the delta phase register, M, represents the amount the phase accumulator is incremented each clock cycle. If f_c is the clock frequency, then the frequency of the output sinewave is equal to $Mf_c/2^n$. This equation is known as the DDS "tuning equation." Note that the frequency resolution of the system is equal to $f_c/2^n$. For n = 32, the resolution is greater than one part in four billion! In a practical DDS system, all the bits out of the phase accumulator are not passed on to the lookup table, but are truncated, thereby reducing the size of the lookup table without affecting frequency resolution. The amount of truncation depends upon the resolution and performance of the output DAC. In general, the phase address information should have two to four bits more resolution than the DAC, but this can vary some from product to product. The objective is to use enough resolution in the lookup table address so that the overall noise and distortion of the analog output signal is limited by the DAC and not the effects of phase truncation.

Signal Flow Through the DDS Architecture

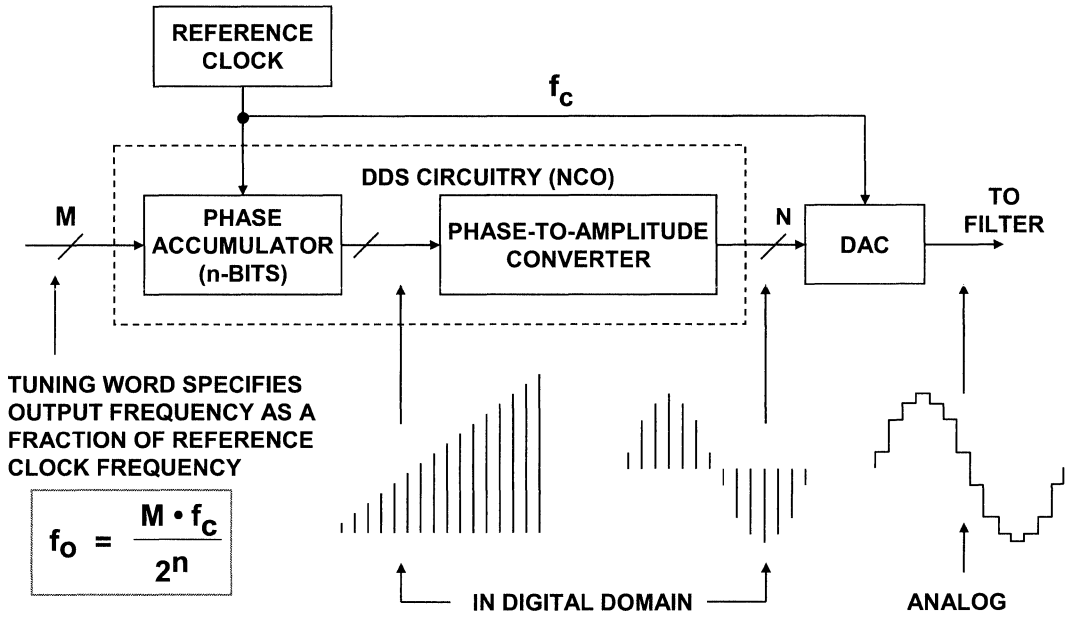

This figure shows the signal flow through the DDS architecture. The phase accumulator is actually a modulus M counter that increments its stored number each time it receives a clock pulse. The magnitude of the increment is determined by the binary input number or word (M) contained in the delta phase register that is summed with the overflow of the counter. The digital phase information from the phase accumulator is converted into a corresponding digital amplitude by the phase-to-amplitude converter. Finally, the DAC converts the digital amplitude into a corresponding analog signal.

The DDS architecture allows the frequency to be changed instantaneously by simply changing the tuning word M. There is no phase hit when the frequency is changed, thereby making the DDS approach ideal for frequency hopping applications.

When IC DDS systems became popular in the mid 1980s, the digital NCO was generally fabricated on a CMOS process, and the high speed DAC on a bipolar process, thereby yielding a two-chip solution. Today, however, modern CMOS processes are suitable for not only the digital circuits but for the high performance DAC as well (as illustrated by the many TxDACs currently offered by Analog Devices). Modern DDS systems therefore are fully integrated and include many additional options as well.

The following article is an excellent reference on DDS operation, especially the analysis of output spurious components.

David Brandon, "DDS Design," *EDN*, May 13, 2004, pp. 71-84.

Effect of Ratio of Sampling Clock to Output Frequency on SFDR for Ideal 12-bit DAC

(A) f_{OUT} = 2.0000 MHz, f_S = 80.0000 MHz

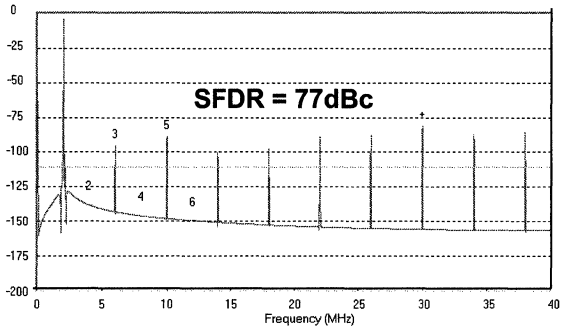

(B) f_{OUT} = 2.0111 MHz, f_S = 80.0000 MHz

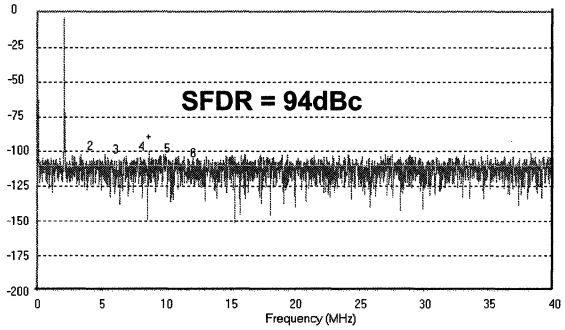

FFT SIZE	= 8192
THEORETICAL 12-BIT SNR	= 74dB
FFT PROCESS GAIN	= 36dB
FFT NOISE FLOOR	= 110dBFS

In DDS systems it is important to utilize frequency planning so that the output frequency is not an exact submultiple of the clock frequency. This can be demonstrated using the ADIsimADC program and the ideal 12-bit ADC model. Note that the same effect occurs in ADCs as well as DACs. In ADC applications, however, the noise on the input signal tends to mask the effect somewhat, although it can still be observed.

In (A), the clock frequency is 80.0000MHz, and the output frequency is 2.0000MHz (an exact ratio of 40). Note that the quantization noise is concentrated at harmonics of the fundamental frequency. This limits the SFDR to approximately 77dBc, even though the calculated SNR is still 74dB.

In (B), the output frequency is changed to 2.0111MHz. Now, the quantization noise is spread more or less uniformly over the entire Nyquist bandwidth. The SFDR is approximately 94dBc (just slightly above the FFT noise floor of 110dBFS—averaging a number of FFTs will cause the SFDR to approach the noise floor even closer).

Note that for each bit added to the ideal DAC, the above numbers increase by approximately 6dB.

This effect can be largely avoided by careful selection of clock and output frequencies in systems using DDS.

Recent DDS technology (SpurKiller) which will be discussed later helps ease the frequency planning imposed constraints by the targeted reduction of specific spurs.

AD9858 1GSPS DDS with Phase Detector and Analog Multiplier

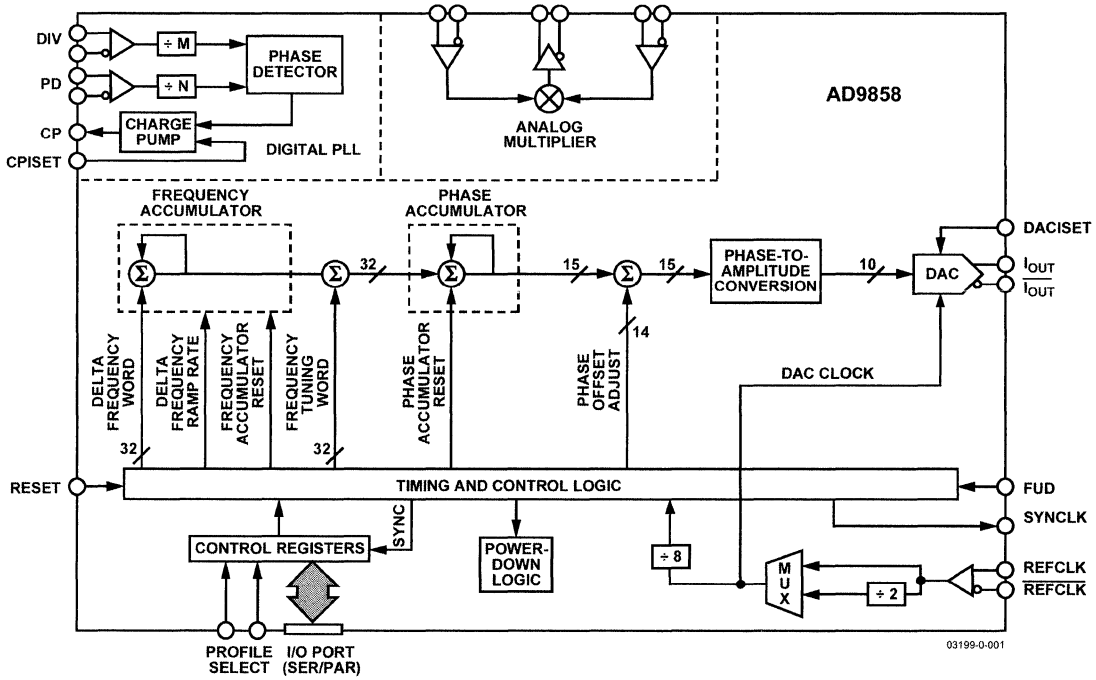

This figure shows a block diagram of one of Analog Devices' recent DDS ICs. The AD9858 is a direct digital synthesizer (DDS) featuring a 10-bit DAC operating up to 1GSPS. The AD9858 uses advanced DDS technology, coupled with an internal high speed, high performance DAC to form a digitally programmable, complete high frequency synthesizer capable of generating a frequency-agile analog output sine wave at up to 400MHz.

The AD9858 is designed to provide fast frequency hopping and fine tuning resolution (32-bit frequency tuning word). The frequency tuning and control words are loaded into the AD9858 via parallel (8-bit) or serial loading formats. The AD9858 contains an integrated charge pump (CP) and phase frequency detector (PFD) for synthesis applications requiring the combination of a high speed DDS along with phase-locked loop (PLL) functions.

Another benefit of the DDS approach to frequency synthesis is the ability to very accurately and reliably inject phase offsets. 14 bits of phase control allow fine adjustments to ~0.022 degrees as well as supporting phase hopping in the same way that frequency hopping is supported.

Note that a second accumulator has been included in the DDS architecture. Shown here specifically driving the frequency tuning word, this enables a very easy and well controlled method for sweeping across a range of frequencies rather than hold at, or jumping between specific frequencies. The newest DDS chips allow the second accumulator to drive the phase and amplitude control functions of the DDS as well, so that either of these may be swept instead of the frequency.

An analog mixer is also provided on-chip for applications requiring the combination of a DDS, PLL, and mixer, such as frequency translation loops, tuners, and so on. The AD9858 also features a divide-by-two on the clock input, allowing the external clock to be as high as 2 GHz.

The AD9858 is specified to operate over the extended industrial temperature range of –40°C to +85°C.

DDS Single Loop Upconversion Using the AD9858

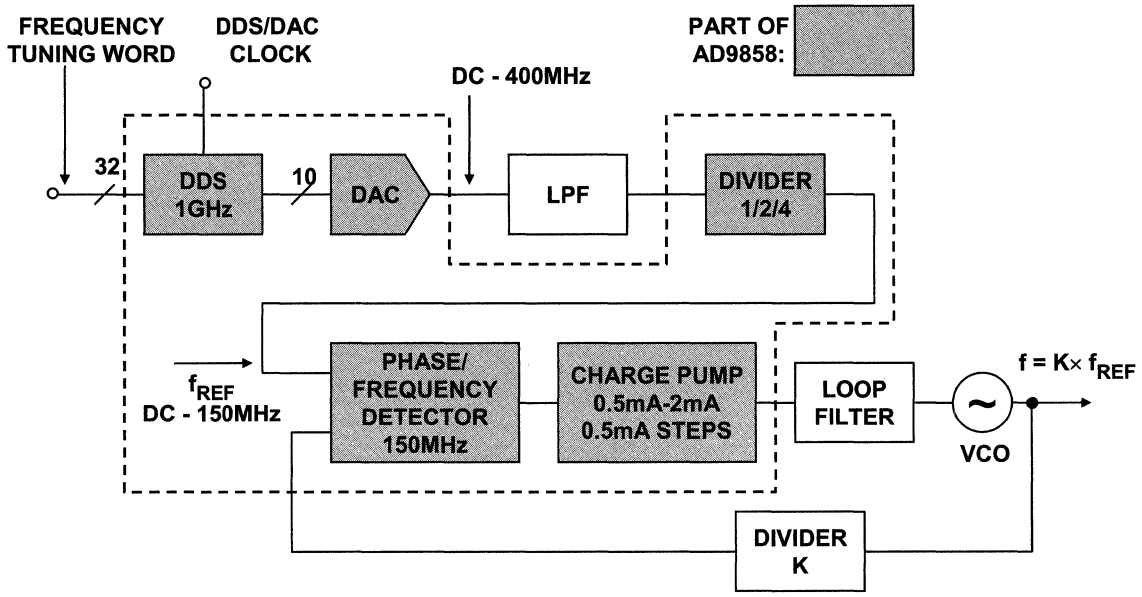

This figure shows the AD9858 configured as an upconverter using the internal phase detector and charge pump along with external filtering and a VCO to form a high-speed PLL.

The basic AD9858 DDS can generate a frequency up to 400MHz. The PLL circuitry, in conjunction with a high frequency VCO and divider, is capable of multiplying the reference frequency well up into the GHz region. The reference frequency into the phase detector can be as high as 150MHz.

SpurKiller Technology

- **Use an auxiliary DDS channel to add in a signal at the same frequency and amplitude as the spur, but 180° out of phase with the highest spur...**

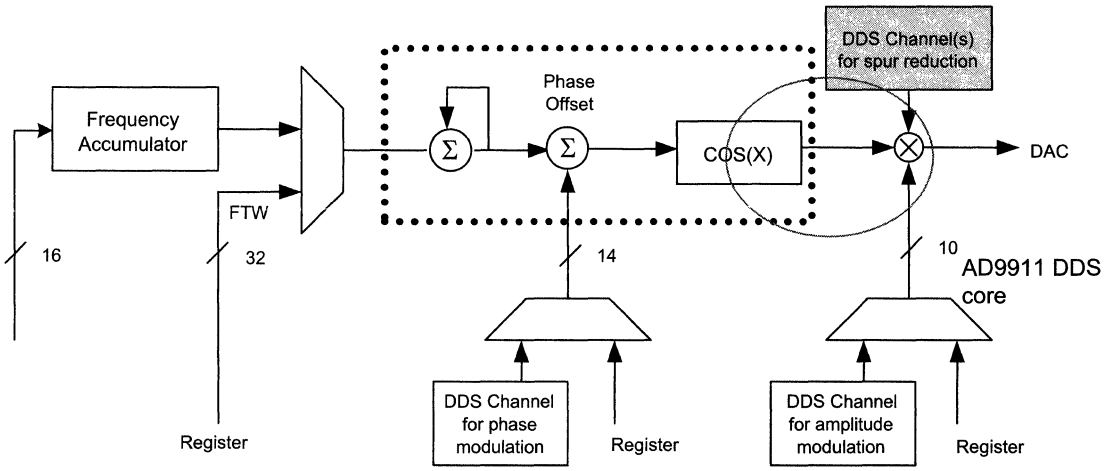

The frequencies at which spurs appear in a DDS output spectrum are simple functions of the sampling rate and the programmed output frequency and are therefore predictable. In addition, they are fairly repeatable from device to device (assuming the same sampling rate and output frequency). In addition, the relative phase of each spur does not change.

The amplitude of the spur does change somewhat over temperature and with supply voltage as well as process variation.

The Analog Devices' SpurKiller technology uses an auxilary DDS channel to add in a signal at the same frequency and amplitude as the worst spur, but 180° out of phase. In effect, this removes most of the spur from the output spectrum.

This figure shows the basic concept used in the first SpurKiller DDS, the AD9911. The auxiliary DDS spur reduction channel is shaded in the figure.

The AD9911 is described in more detail in the next figure.

AD9911 SpurKiller 500MHz DDS

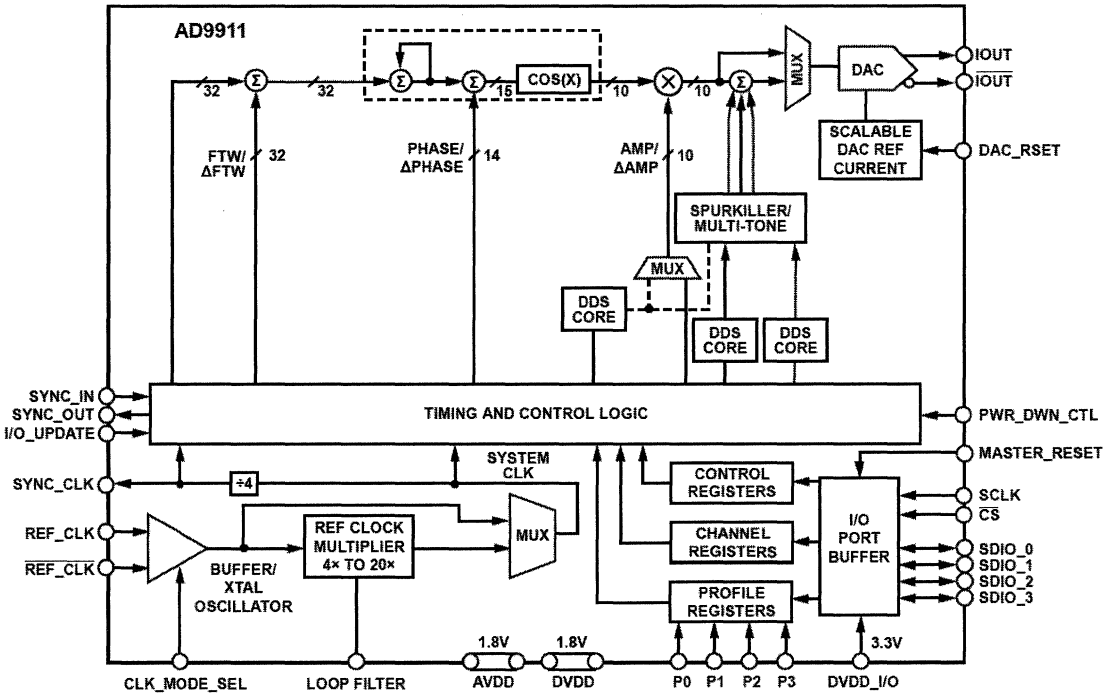

The AD9911 is a complete direct digital synthesizer (DDS) that operates up to 500MSPS. This device includes a high speed DAC with excellent wideband and narrowband spurious-free dynamic range (SFDR) as well as three auxiliary DDS cores without assigned digital-to-analog converters (DACs). These auxiliary channels are used for spur reduction, multitone generation, or test-tone modulation.

The AD9911 is the first DDS to incorporate SpurKiller technology and multitone generation capability. Multitone mode enables the generation up to four concurrent carriers; frequency, phase, and amplitude can be independently programmed. Multitone generation can be used for system tests, such as inter-modulation distortion and receiver blocker sensitivity. SpurKilling enables customers to improve SFDR performance by reducing the magnitude of harmonic components and/or the aliases of those harmonic components.

Test-tone modulation efficiently enables sine wave modulation of amplitude on the output signal using one of the auxiliary DDS cores.

The AD9911 can perform modulation of frequency, phase, or amplitude (FSK, PSK, ASK). Modulation is implemented by storing profiles in the register bank and applying data to the profile pins. In addition, the AD9911 supports linear sweep of frequency, phase, or amplitude for applications such as radar and instrumentation.

The Results of Using SpurKiller Technology on a DDS Output Spur

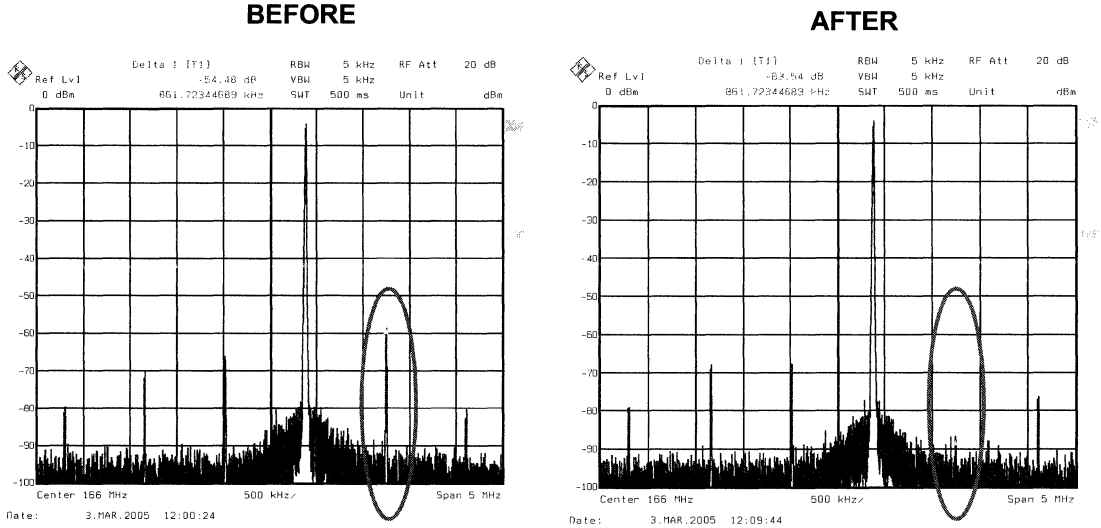

OUTPUT FREQUENCY = 166MHz

This figure shows the AD9911 DDS output spectrum before and after applying the SpurKiller technique. The output frequency is 166MHz, and the update rate is 500MHz. The worst spur is highlighted and appears approximately 60dB below fullscale.

The basic SpurKiller procedure is as follows:

1. Locate the worst spur.
2. Program the SpurKiller DDS channel to produce a signal at a frequency slightly higher than the spur frequency.
3. Program the SpurKiller signal so that its amplitude matches that of the worst spur.
4. Program the SpurKiller signal to match the exact frequency of the worst spur. The frequency of the spur can be determined mathematically based on the update rate and the output frequency. The interactive DDS Interactive On-Line Design Tool can be of great assistance here.
5. Adjust the phase of the SpurKiller signal so that the worst spur is minimized.
6. Depending upon how accurately the SpurKiller frequency matches that of the spur, the SpurKiller frequency may require a slight adjustment for optimum results. In most cases, this won't be necessary if the spur frequency is calculated using the design tool.

The figure shows that after applying the SpurKiller technique, the worst spur is reduced in amplitude by nearly 30dB.

The current design of the AD9911 has a considerable amount of part to part variation in terms of spur magnitude. The next generation SpurKiller parts should be released in 2007. Initial silicon on these parts shows the ability to reduce spurs by at least 10dB using a single programming word across multiple parts.

DDS On-Line Interactive Design Tool (ADIsimDDS)

DDS Design Tool Main Screen

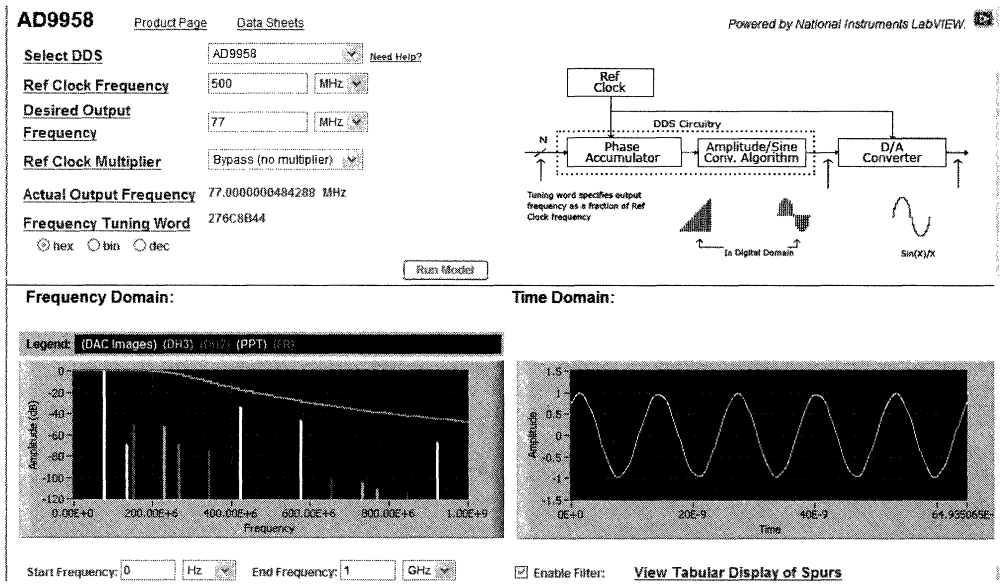

The purpose of this on-line interactive design tool is to assist a user in selecting and evaluating Analog Devices' DDS ICs. It allows a user to select a device, enter the desired operating conditions, and evaluate its general performance. It should be noted that the tool uses mathematical equations to approximate the overall performance of the selected device and does not calculate all possible errors. Therefore, the tool should be used as a design aid only and is not intended to be used as a replacement for actual hardware testing and evaluation. The basic procedure for operating the DDS tool is as follows:

1. Select a DDS device from the pulldown menu at the upper left. The "Need Help?" link (at right) can be used to obtain additional information and to make selections from a sortable DDS Interactive Selection Table (IST)

2. Input the Reference Clock Frequency

3. Input the Desired Output Frequency. Due to Nyquist requirements, the output frequency is generally limited to approximately 40% of the Reference Clock Frequency

4. Input the Ref. Clock Multiplier, if applicable. Some DDS devices offer an internal Ref. Clock Multiplier Circuit that can be used to increase the Clock Frequency

5. The Actual Output Frequency will be displayed. The actual output frequency is determined by the value of the digital tuning word and, therefore, will be slightly different than the desired output frequency.

6. The Frequency Tuning Word is displayed. This is the required digital word to be written to the DDS to obtain the desired output frequency. The tuning word value can be displayed in decimal, hexadecimal or binary by selecting the corresponding radio button.

7. Run the model. Observe the locations of the spurs as well as the time domain representation of the output signal.

DDS Design Tool: Tabular Display of Spurs

Tabular Display of Spurs (hide) (top) KEY: f1 = Ref Clock, f2 = Desired Output Frequency, PPT Spur = Primary Phase Truncation Spur

DAC Images			DAC Harmonics 2nd Order			DAC Harmonics 3rd Order			Other Spurs		
Formula	Frequency (Hz)	Attenuation (dBc)	Formula	Frequency (Hz)	Attenuation (dBc)	Formula	Frequency (Hz)	Attenuation (dBc)	Name	Frequency (Hz)	Attenuation (dBc)
f1	7.700000E+7	-0.341531	2f1	1.540000E+8	-50.341531	3f1	2.310000E+8	-50.341531	PPT Spur	1.360016E+8	-70.014874
f2-f1	4.230000E+8	-15.138524	f2-2f1	3.460000E+8	-65.138524	f2-3f1	2.690000E+8	-65.138524			
f2+f1	5.770000E+8	-17.835233	f2+2f1	6.540000E+8	-67.835233	f2+3f1	7.310000E+8	-67.835233			
2f2-f1	9.230000E+8	-21.915751	2f2-2f1	8.460000E+8	-71.915751	2f2-3f1	7.690000E+8	-71.915751			
2f2+f1	1.077000E+9	-23.256031	2f2+2f1	1.154000E+9	-73.256031	2f2+3f1	1.231000E+9	-73.256031			
3f2-f1	1.423000E+9	-25.675815	3f2-2f1	1.346000E+9	-75.675815	3f2-3f1	1.269000E+9	-75.675815			
3f2+f1	1.577000E+9	-26.568351	3f2+2f1	1.654000E+9	-76.568351	3f2+3f1	1.731000E+9	-76.568351			
4f2-f1	1.923000E+9	-28.291302	4f2-2f1	1.846000E+9	-78.291302	4f2-3f1	1.769000E+9	-78.291302			
4f2+f1	2.077000E+9	-28.960447	4f2+2f1	2.154000E+9	-78.960447	4f2+3f1	2.231000E+9	-78.960447			
5f2-f1	2.423000E+9	-30.298785	5f2-2f1	2.346000E+9	-80.298785	5f2-3f1	2.269000E+9	-80.298785			
5f2+f1	2.577000E+9	-30.834005	5f2+2f1	2.654000E+9	-80.834005	5f2+3f1	2.731000E+9	-80.834005			
6f2-f1	2.923000E+9	-31.928293	6f2-2f1	2.846000E+9	-81.928293	6f2-3f1	2.769000E+9	-81.928293			
6f2+f1	3.077000E+9	-32.374267	6f2+2f1	3.154000E+9	-82.374267	6f2+3f1	3.231000E+9	-82.374267			
7f2-f1	3.423000E+9	-33.299855	7f2-2f1	3.346000E+9	-83.299855	7f2-3f1	3.269000E+9	-83.299855			
7f2+f1	3.577000E+9	-33.682096	7f2+2f1	3.654000E+9	-83.682096	7f2+3f1	3.731000E+9	-83.682096			
8f2-f1	3.923000E+9	-34.484083	8f2-2f1	3.846000E+9	-84.484083	8f2-3f1	3.769000E+9	-84.484083			
8f2+f1	4.077000E+9	-34.818531	8f2+2f1	4.154000E+9	-84.818531	8f2+3f1	4.231000E+9	-84.818531			
9f2-f1	4.423000E+9	-35.526056	9f2-2f1	4.346000E+9	-85.526056	9f2-3f1	4.269000E+9	-85.526056			
9f2+f1	4.577000E+9	-35.823335	9f2+2f1	4.654000E+9	-85.823335	9f2+3f1	4.731000E+9	-85.823335			
10f2-f1	4.923000E+9	-36.456313	10f2-2f1	4.846000E+9	-86.456313	10f2-3f1	4.769000E+9	-86.456313			

8. View Tabular Display of Spurs. Clicking the View Tabular Display of Spurs Link located below the Time Domain Plot allows the user to view Sur Data in a Tabular Format.

9. The log box shown below the time domain plot is used to display and flag errors that may occur. Additionally, it may contain suggestions for resolving the issue.

DDS Design Tool: Display Options and Filter Selection

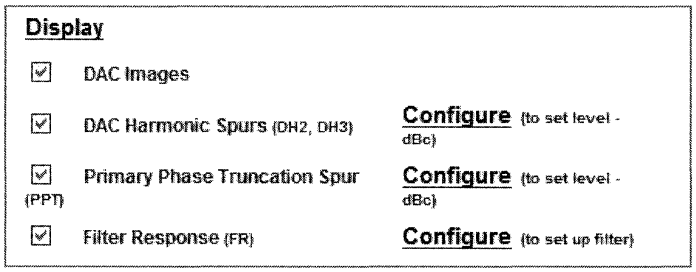

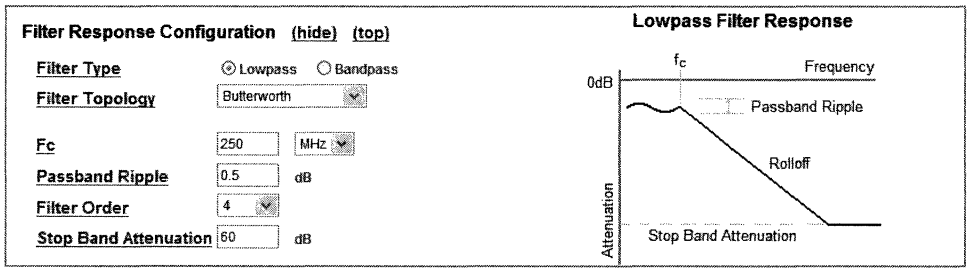

10. Modify the Display Response if desired. The tool provides check boxes and cofiguration capabilities just below the frequency domain plot. These allow a user to configure and to select to include or exclude images and spurs in the frequency domain plot.

11. Configure the Filter Response if desired. This part of the tool allows selection of the Filter Type, Filter Topology, Cutoff Frequency, Passband Ripple, Filter Order, and Stopband Attenuation.

Phase Locked Loops

www.analog.com/pll

Basic Phase Locked Loop (PLL) Model

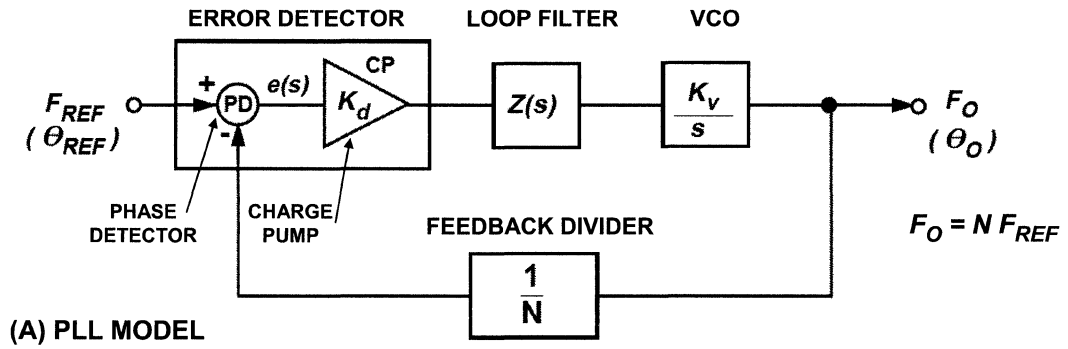

(A) PLL MODEL

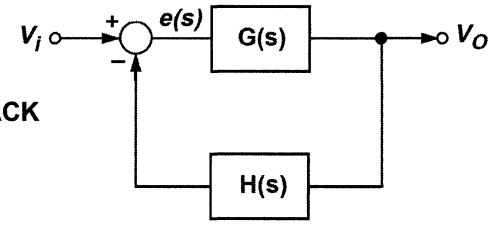

(B) STANDARD NEGATIVE FEEDBACK CONTROL SYSTEM MODEL

The top diagram (A) shows the basic model for a PLL. A PLL is a feedback system combining a voltage controlled oscillator (VCO) and a phase detector connected so that the VCO maintains a constant phase angle relative to a reference signal. The PLL can be analyzed as a negative feedback system using Laplace Transform theory with a forward gain term, G(s), and a feedback term, H(s), as shown in diagram (B). The usual equations for a negative feedback system apply:

$$Forward\ Gain = G(s),\ [s = j\omega = j2\pi f]$$

$$Loop\ Gain = G(s) \times H(s)$$

$$Closed\text{-}Loop\ Gain = \frac{G(s)}{1 + G(s)H(s)}$$

The basic blocks of the PLL are the Error Detector (composed of a phase detector and a charge pump), Loop Filter, VCO, and a Feedback Divider.

Negative feedback forces the error signal, e(s), to approach zero at which point the feedback divider output and the reference frequency are in phase and frequency lock, and $F_O = NF_{REF}$.

Referring to the above diagram, a system for using a PLL to generate higher frequencies than the input, the VCO oscillates at an angular frequency of ω_O. A portion of this signal is fed back to the error detector, via a frequency divider with a ratio 1/N. This divided down frequency is fed to one input of the error detector. The other input in this example is a fixed reference signal. The error detector compares the signals at both inputs. When the two signal inputs are equal in phase and frequency, the error will be constant and the loop is said to be in a "locked" condition.

Phase/Frequency Detector (PFD) Driving Charge Pump (CP)

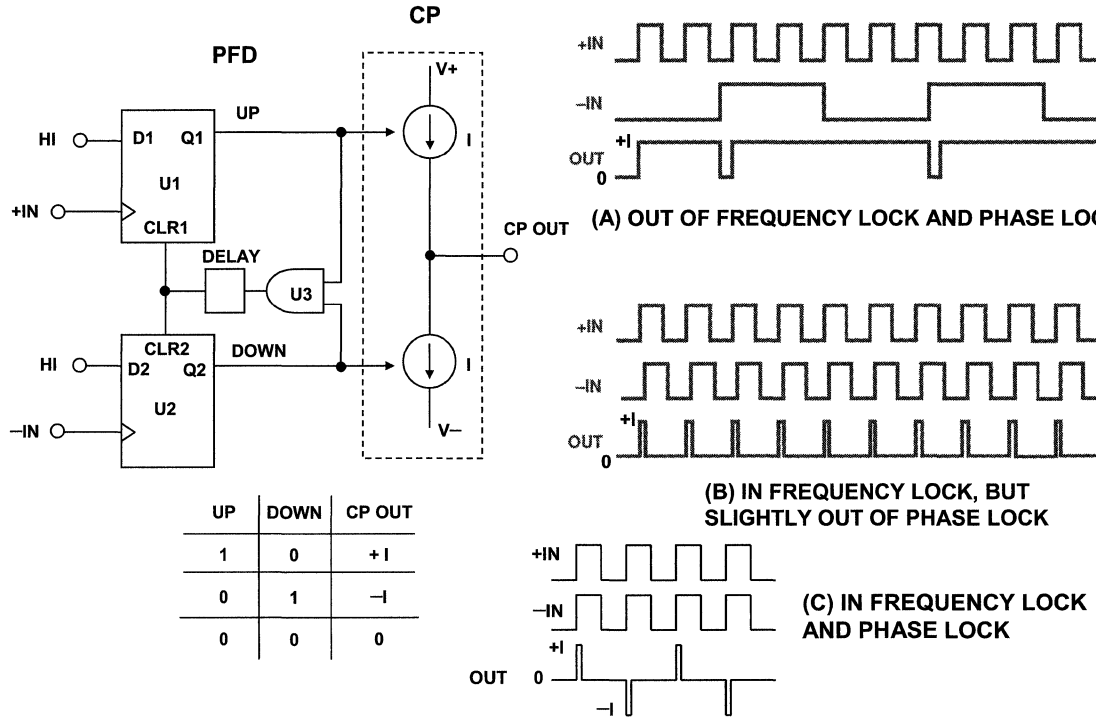

UP	DOWN	CP OUT
1	0	+I
0	1	–I
0	0	0

This shows a popular implementation of a PFD, basically consisting of two D-type flip flops. One Q output enables a positive current source; and the other Q output enables a negative current source. Assuming that, in this design, the D-type flip flop is positive-edge triggered, the possible states are shown in the logic table.

Consider now how the circuit behaves if the system is out of lock and the frequency at +IN is much higher than the frequency at –IN, as shown in (A).

Since the frequency at +IN is much higher than that at –IN, the UP output spends most of its time in the high state. The first rising edge on +IN sends the output high and this is maintained until the first rising edge occurs on –IN. In a practical system this means that the output, and thus the input to the VCO, is driven higher, resulting in an increase in frequency at –IN. This is exactly what is desired. If the frequency on +IN were much lower than on –IN, the opposite effect would occur. The output at OUT would spend most of its time in the low condition. This would have the effect of driving the VCO in the negative direction and again bring the frequency at –IN much closer to that at +IN, to approach the locked condition. Figure (B) shows the waveforms when the inputs are frequency-locked and close to phase-lock. Since +IN is leading –IN, the output is a series of positive current pulses. These pulses will tend to drive the VCO so that the –IN signal become phase-aligned with that on +IN. When this occurs, if there were no delay element between U3 and the CLR inputs of U1 and U2, it would be possible for the output to be in high-impedance mode, producing neither positive nor negative current pulses. This would not be a good situation. The VCO would drift until a significant phase error developed and started producing either positive or negative current pulses once again. Over a relatively long period of time, the effect of this cycling would be for the output of the charge pump to be modulated by a signal that is a subharmonic of the PFD input reference frequency. Since this could be a low frequency signal, it would not be attenuated by the loop filter and would result in very significant spurs in the VCO output spectrum, a phenomenon known as the "backlash" effect. The delay element between the output of U3 and the CLR inputs of U1 and U2 ensures that it does not happen. With the delay element, even when the +IN and –IN are perfectly phase-aligned, there will still be a current pulse generated at the charge pump output as shown in (C). The duration of this delay is equal to the delay inserted at the output of U3 and is known as the anti-backlash pulse width.

Note that if the +IN frequency is lower than the –IN frequency and/or the +IN phase lags the –IN phase, then the output of the charge pump will be a series of negative current pulses—the reverse of the condition shown in (A) and (B) above.

Adding an Input Reference Divider and a Prescaler to the Basic PLL

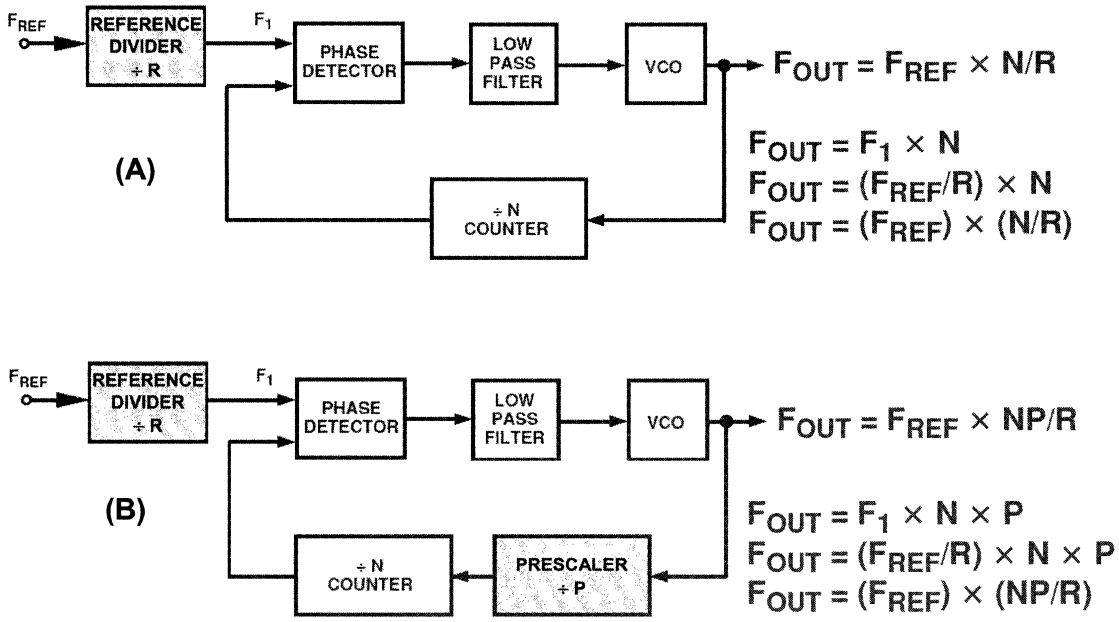

In the classical Integer-N synthesizer, the resolution of the output frequency is determined by the reference frequency applied to the phase detector. So, for example, if 200kHz spacing is required (as in GSM phones), then the reference frequency must be 200kHz. However, getting a stable 200kHz frequency source is not easy. A sensible approach is to take a good crystal-based high frequency source and divide it down. For example, the desired frequency spacing could be achieved by starting with a 10MHz frequency reference and dividing it down by 50. This approach is shown in the diagram (A).

The "N counter," also known as the N divider, is the programmable element that sets the relationship between the input and output frequencies in the PLL. The complexity of the N counter has grown over the years. In addition to a straightforward N counter, it has evolved to include a prescaler, which can have a dual modulus. This structure has grown as a solution to the problems inherent in using the basic divide-by-N structure to feed back to the phase detector when very high-frequency outputs are required. For example, let's assume that a 900MHz output is required with 10Hz spacing. A 10MHz reference frequency might be used, with the R-Divider set at 1000. Then, the N-value in the feedback would need to be of the order of 90,000. This would mean at least a 17- bit counter capable of dealing with an input frequency of 900MHz. To handle this range, it makes sense to precede the programmable counter with a fixed counter element to bring the very high input frequency down to a range at which standard CMOS will operate. This counter, called a prescaler, is shown in diagram (B).

However, note that using a standard prescaler as shown reduces the system resolution to F1×P. This issue can be addressed by using a dual-modulus prescaler which has the advantages of a standard prescaler, but without loss of resolution. A dual-modulus prescaler is a counter whose division ratio can be switched from one value to another by an external control signal. It's use is described in the next figure.

Adding a Dual Modulus Prescaler to the PLL

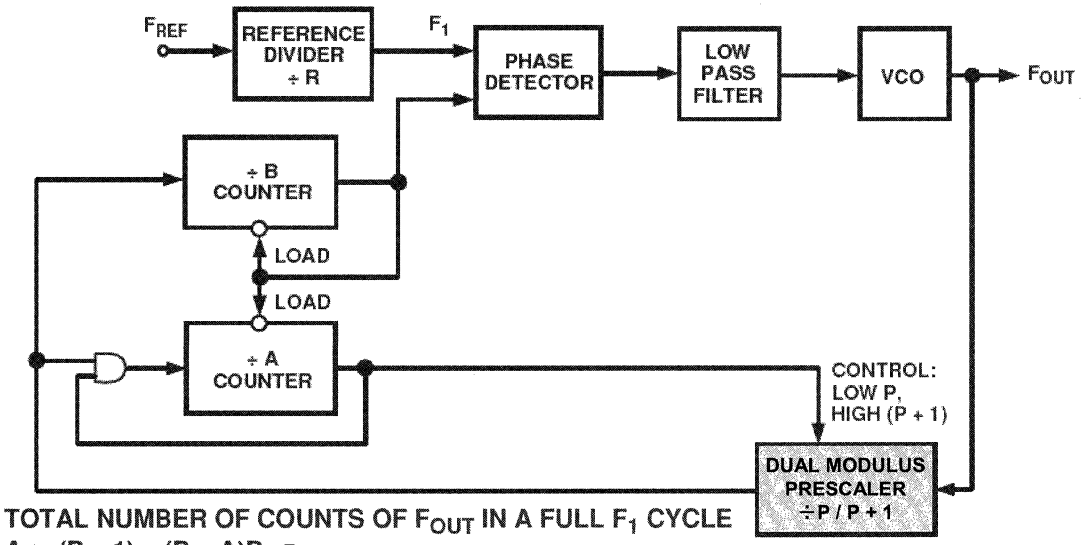

TOTAL NUMBER OF COUNTS OF F_{OUT} IN A FULL F_1 CYCLE

$A \times (P + 1) + (B - A)P$ =

$AP + A + BP - AP$ =

$BP + A$, THEREFORE

$F_{OUT} = F_1 \times (BP + A)$

$F_{OUT} = (F_{REF}/R) \times (BP + A)$

By using the dual-modulus prescaler with an A and B counter, one can still maintain output resolution of F1. However, the following conditions must be met:

1. The output signals of both counters are High if the counters have not timed out.
2. When the B counter times out, its output goes Low, and it immediately loads both counters to their preset values.
3. The value loaded to the B counter must always be greater than that loaded to the A counter.

Assume that the B counter has just timed out and both counters have been reloaded with the values A and B. Let's find the number of VCO cycles necessary to get to the same state again.

As long as the A counter has not timed out, the prescaler is dividing down by P + 1. So, both the A and B counters will count down by 1 every time the prescaler counts (P + 1) VCO cycles. This means the A counter will time out after ((P + 1) × A) VCO cycles. At this point the prescaler is switched to divide-by-P. It is also possible to say that at this time the B counter still has (B – A) cycles to go before it times out. How long will it take to do this: ((B – A) × P).

The system is now back to the initial condition where we started.

The total number of VCO cycles needed for this to happen is :

$$N = [A \times (P + 1)] + [(B - A) \times P]$$
$$= AP + A + BP - AP$$
$$= BP + A.$$

Therefore, $F_{OUT} = (F_{REF}/R) \times (BP + A)$.

Key PLL Specifications

- Input RF Frequency (Minimum/Maximum)
- Channel Spacing
- Loop Bandwidth and Phase Margin
- Frequency Lock Time
- Phase Lock Time
- Output Frequency Error
- Output Phase Error
- Phase Noise and Phase Jitter
- Reference Spurs

There are many specifications to consider when designing a PLL. The input RF frequency range and the channel spacing determine the value of the R and N counter and the prescaler parameters.

The loop bandwidth determines the frequency and phase lock time. Since the PLL is a negative feedback system, phase margin and stability issues must be considered.

Spectral purity of the PLL output is specified by the phase noise and the level of the reference-related spurs.

Many of these parameters are interactive; for instance, lower values of loop bandwidth lead to reduced levels of phase noise and reference spurs, but at the expense of longer lock times and less phase margin.

Because of the many tradeoffs involved, the use of a PLL design program such as the Analog Devices' ADIsimPLL allows these tradeoffs to be evaluated and the various parameters adjusted to fit the required specifications. The program not only assists in the theoretical design, but also aids in parts selection and determines component values.

Oscillator Phase Noise and Spurs

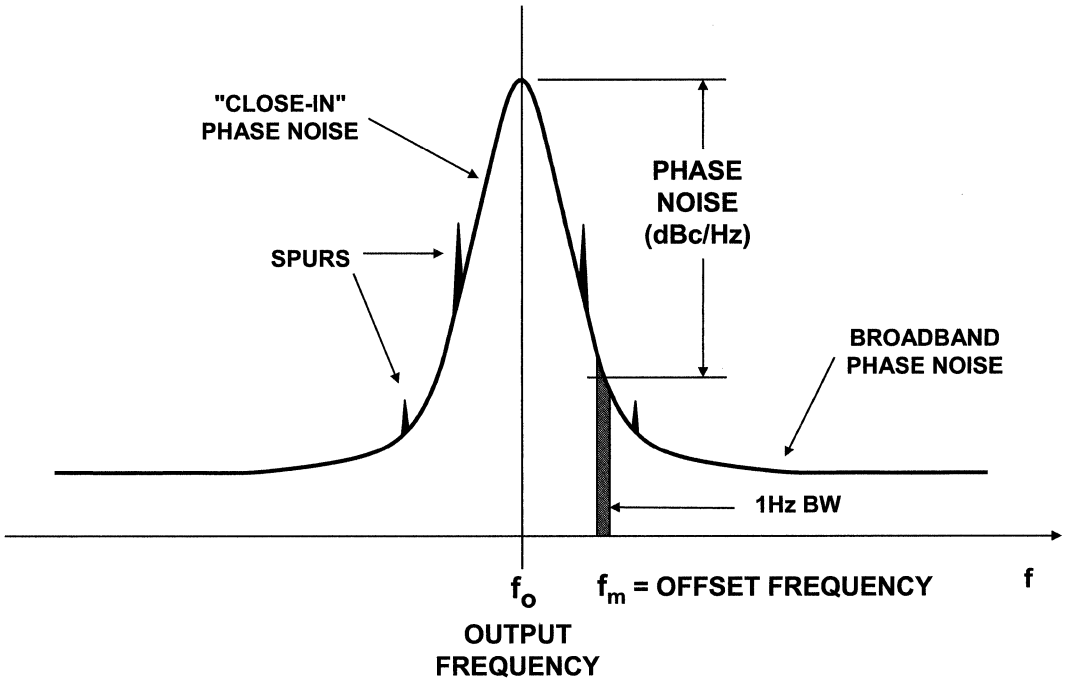

A PLL is a type of oscillator, and in any oscillator design, frequency stability is of critical importance. We are interested in both long-term and short-term stability. Long-term frequency stability is concerned with how the output signal varies over a long period of time (hours, days, or months). It is usually specified as the ratio, $\Delta f/f$ for a given period of time, expressed as a percentage or in dB.

Short-term stability, on the other hand, is concerned with variations that occur over a period of seconds or less. These variations can be random or periodic. A spectrum analyzer can be used to examine the short-term stability of a signal. This diagram shows a typical spectrum, with random and discrete frequency components causing a broad skirt and spurious peaks.

The discrete spurious components could be caused by known clock frequencies in the signal source, power line interference, and mixer products. The broadening caused by random noise fluctuation is due to phase noise. It can be the result of thermal noise, shot noise, and/or flicker noise in active and passive devices.

The phase noise spectrum of an oscillator shows the noise power in a 1Hz bandwidth as a function of frequency. Phase noise is defined as the ratio of the noise in a 1Hz bandwidth at a specified frequency offset, f_m, to the oscillator signal amplitude at frequency f_o.

Phase Noise in dBc/Hz Versus Frequency Offset from Output Frequency

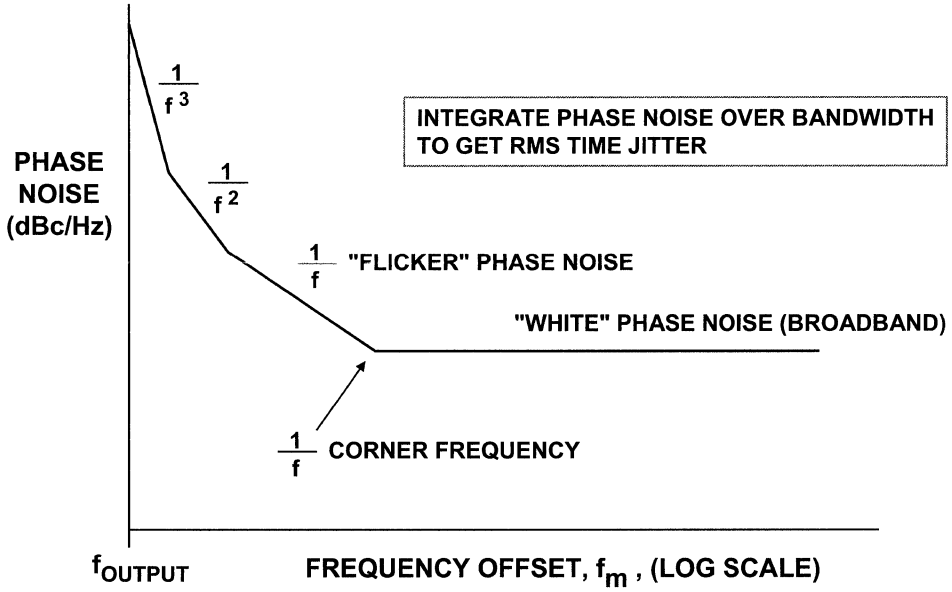

It is customary to characterize an oscillator in terms of its single-sideband phase noise as shown here, where the phase noise in dBc/Hz is plotted as a function of frequency offset, f_m, with the frequency axis on a log scale. Note the actual curve is approximated by a number of regions, each having a slope of $1/f^x$, where x = 0 corresponds to the "white" phase noise region (slope = 0dB/decade), and x = 1 corresponds to the "flicker" phase noise region (slope = –20dB/decade). There are also regions where x = 2, 3, 4, and these regions occur progressively closer to the carrier frequency.

Note that the phase noise curve is somewhat analogous to the input voltage noise spectral density of an amplifier. Like amplifier voltage noise, low 1/f corner frequencies are highly desirable in an oscillator.

In some cases, it is useful to convert phase noise into time jitter. This can be done by basically integrating the phase noise plot over the desired frequency range. The ADIsimPLL program performs this calculation if needed. Details of how to convert phase noise into time jitter can be found in the following reference:

Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN-0916550273 Chapter 6. Also available as *Data Conversion Handbook*, Elsevier-Newnes, 2005, ISBN: 0750678410, Chapter 6.

The ability to perform this conversion between phase noise and time jitter is especially useful when using the PLL output to drive an ADC sampling clock. Once the time jitter is known, its effect on the overall ADC SNR can be evaluated.

The ADIsimPLL program (to be discussed shortly) performs the conversion between phase noise and time jitter.

ADF4360 Family of Integrated PLLs

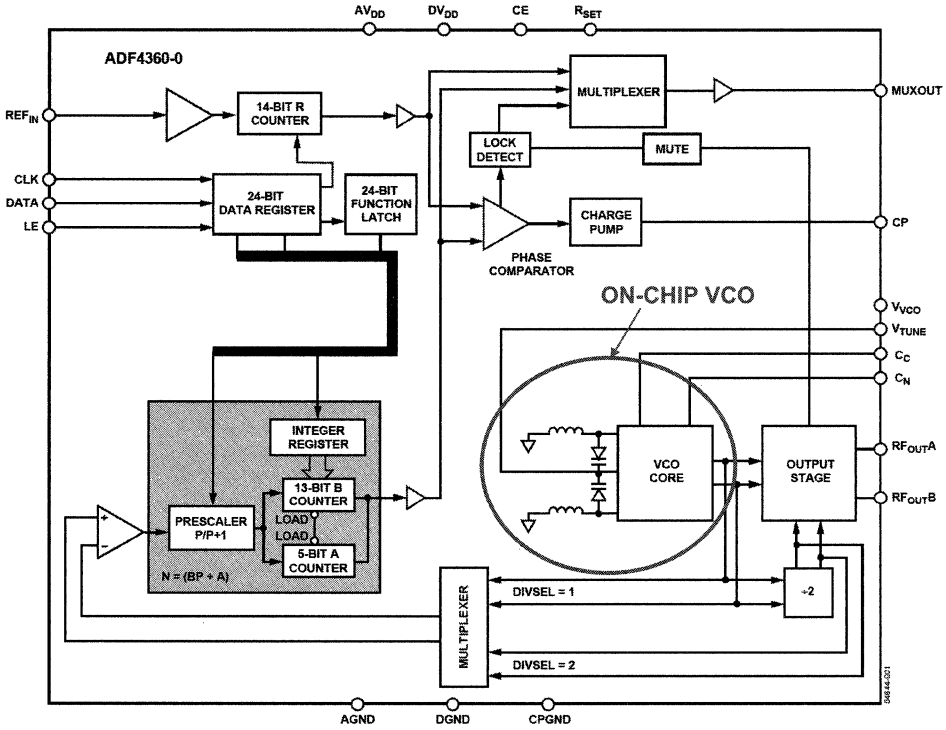

We will now examine a couple of examples of PLL ICs from Analog Devices. It would be impossible to discuss all of them, so for complete coverage refer to www.analog.com/PLL.

The ADF4360 family of integrated PLLs include the integer-N synthesizer with prescalars and VCO. They are optimized for the output frequency ranges given in parentheses: ADF4360-0 (2400MHz to 2725MHz), ADF4360-1 (2050MHz to 2450MHz), ADF4360-2 (1850MHz to 2150MHz), ADF4360-3 (1600MHz to 1950MHz), ADF4360-4 (1450MHz to 1750MHz), ADF4360-5 (1200MHz to 1400MHz), ADF4360-6 (1050MHz to 1250MHz), ADF4360-7 (350MHz to 1800MHz). The reference input frequency can range from 10MHz to 250MHz.

The ADF4360-7 has a prescaler, but the center frequency is set by external inductors.

The ADF4360-8 also has the center frequency set by external inductors. This allows an output frequency range of between 65MHz to 400MHz. The ADF4360-8 has no RF prescaler.

All members of the ADF 4360 family operate on +3.0V to +3.6V and have 1.8V logic interface compatibility. The output power level can be programmed from −13dBm to −4dBm, and the VCO current is programmable in 5mA steps between 5mA and 20mA. The total power supply current ranges between 25mA to 50mA (or 45mA), depending on the particular part and the programmed settings.

All members of the family are packaged in a 24-lead CSP.

Block Diagram of ADF4110 Family

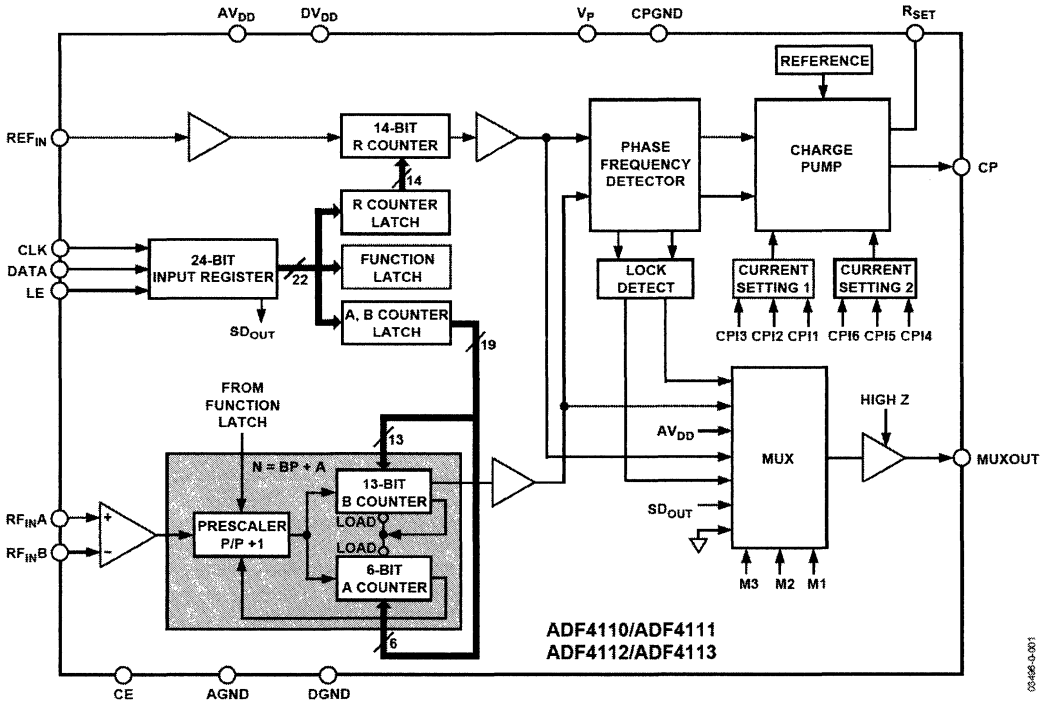

The ADF4110 family of synthesizers consists of single devices, and the ADF4210 family consists of dual versions. The block diagram for the ADF4110 is shown here. It contains the reference counter, the dual-modulus prescaler, the N counter and the PFD blocks previously described. The VCO must be supplied externally, and for low phase noise applications, crystal VCOs (VCXOs) are often used.

The ADF4110 family of frequency synthesizers can be used to implement local oscillators in the upconversion and downconversion sections of wireless receivers and transmitters. They consist of a low noise digital PFD (phase frequency detector), a precision charge pump, a programmable reference divider, programmable A and B counters, and a dual-modulus prescaler (P/P + 1). The A (6-bit) and B (13-bit) counters, in conjunction with the dual-modulus prescaler (P/P + 1), implement an N divider (N = BP + A). In addition, the 14-bit reference counter (R counter) allows selectable REF_{IN} frequencies at the PFD input.

A complete phase-locked loop (PLL) can be implemented if the synthesizer is used with an external loop filter and voltage controlled oscillator (VCO). Control of all the on-chip registers is via a simple 3-wire interface. The devices operate with a power supply ranging from 2.7V to 5.5V and can be powered down when not in use.

RF input frequencies for the various parts in the family are as follows (+5V supply).

ADF4110: 80/550 MHz (min/max)

ADF4111: 0.08/1.4 GHz (min/max)

ADF4112: 0.1/3.0 GHz (min/max)

ADF4113: 0.2/4.0 GHz (min/max)

Integer-N Compared to Fractional-N Synthesizer

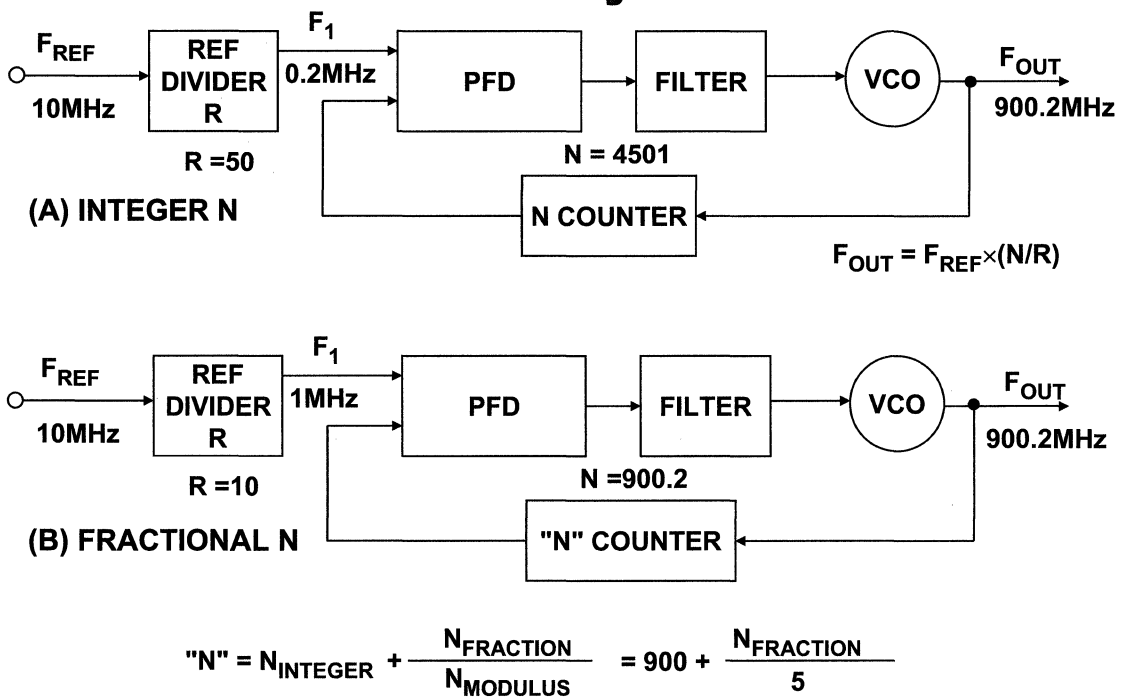

$$\text{"N"} = N_{INTEGER} + \frac{N_{FRACTION}}{N_{MODULUS}} = 900 + \frac{N_{FRACTION}}{5}$$

Fractional-N PLLs have been utilized since the 1970s. As has been discussed, the resolution at the output of an integer-N PLL is limited to steps of the PFD input frequency as shown in (A), where the PFD input is 0.2MHz. Fractional-N allows the resolution at the PLL output to be reduced to small fractions of the PFD frequency as shown in (B), where the PFD input frequency is 1MHz. It is possible to generate output frequencies with resolutions of 100s of Hz, while maintaining a high PFD frequency. As a result the N-value is significantly less than for integer-N. Since noise at the charge pump is multiplied up to the output at a rate of 20logN, significant improvements in phase noise are possible. For a GSM900 system, the fractional-N ADF4252 offers phase noise performance of –103dBc/Hz, compared with –93dBc/Hz for the ADF4106 integer-N PLL. Also offering a significant advantage is the lock-time improvement made possible by fractional-N. The PFD frequency set to 20MHz and loop bandwidth of 150kHz will allow the synthesizer to jump 30MHz in less than 30µs. Current base stations require two PLL blocks to ensure that LOs can meet the timing requirements for transmissions. With the super-fast lock times of fractional-N, future synthesizers will have lock time specs that allow the two "ping-pong" PLLs to be replaced with a single fractional-N PLL block.

The downside of fractional-N PLLs is higher spurious levels. A fractional-N divide by 900.2 consists of the N-divider dividing by 900 80% of the time, and by 901 20% of the time. The average division is correct, but the instantaneous division is incorrect. Because of this, the PFD and charge pump are constantly trying to correct for instantaneous phase errors. The heavy digital activity of the sigma-delta modulator, which provides the averaging function, creates spurious components at the output. The digital noise, combined with inaccuracies in matching the hard-working charge pump, results in spurious levels greater than those allowable by most communications standards. Only recently have fractional-N parts, such as the ADF4252, made the necessary improvements in spurious performance to allow designers to consider their use in traditional integer-N markets.

ADF4153 Fractional-N Synthesizer

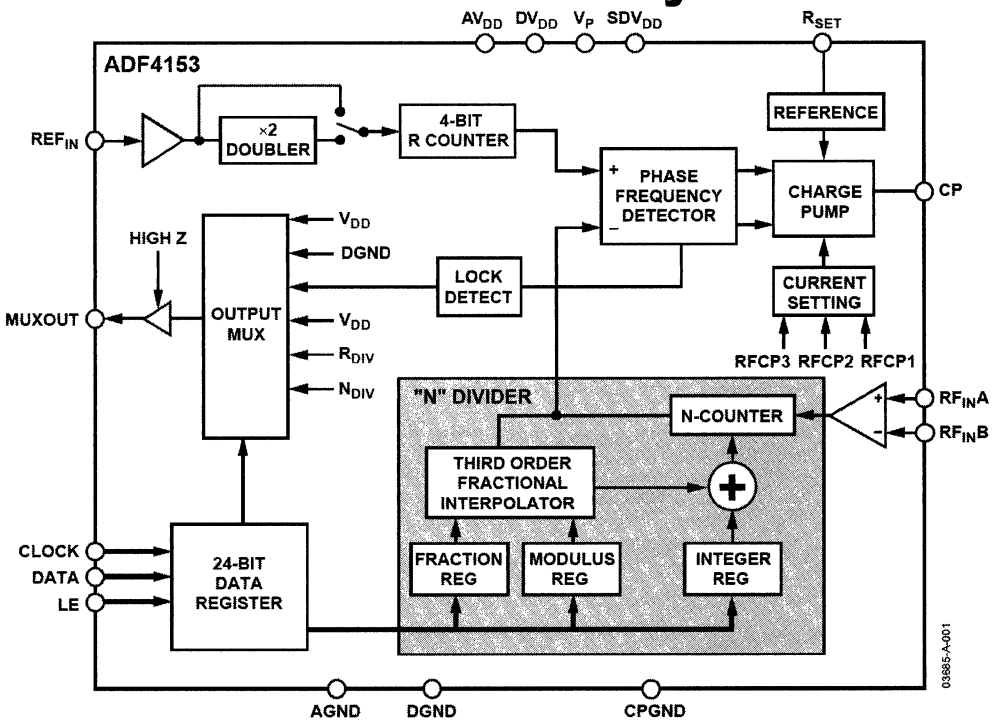

"N" = $N_{INTEGER}$ + $N_{FRACTION}$ / $N_{MODULUS}$

The ADF4153 is a fractional-N frequency synthesizer that implements local oscillators in the upconversion and downconversion sections of wireless receivers and transmitters. It consists of a low noise digital phase frequency detector (PFD), a precision charge pump, and a programmable reference divider. There is a third order Σ-Δ-based fractional interpolator to allow programmable fractional-N division. The INTEGER, FRACTION, and MODULUS registers define an overall N divider, where "N" = NINTEGER + NFRACTION/NMODULUS.

In addition, the 4-bit reference counter (R counter) allows selectable REFIN frequencies at the PFD input. A complete phase-locked loop (PLL) can be implemented if the synthesizer is used with an external loop filter and a voltage controlled oscillator (VCO). Control of all on-chip registers is via a simple 3-wire interface. The device operates with a power supply ranging from 2.7V to 3.3V and can be powered down when not in use.

"Noise" and "spur" mode allows the user to optimize a design either for improved spurious performance or for improved phase noise performance. When the lowest spur setting is chosen, dither is enabled. This randomizes the fractional quantization noise so that it looks more like white noise rather than spurious noise. This means that the part is optimized for improved spurious performance. This operation would normally be used when the PLL closed-loop bandwidth is wide, for fast-locking applications. (Wide-loop bandwidth is seen as a loop bandwidth greater than 1/10 of the RFOUT channel step resolution (fRES)). A wide-loop filter does not attenuate the spurs to a level that a narrow-loop bandwidth would. When the low noise and spur setting is enabled, dither is disabled. This optimizes the synthesizer to operate with improved noise performance. However, the spurious performance is degraded in this mode compared to the lowest spurs setting. To further improve noise performance, the lowest noise setting option can be used, which reduces the phase noise. As well as disabling the dither, it also ensures that the charge pump is operating in an optimum region for noise performance. This setting is extremely useful where a narrow-loop filter bandwidth is available. The synthesizer ensures extremely low noise, and the filter attenuates the spurs.

ADF4153 Operating in Lowest Noise Mode for N = 8581/130

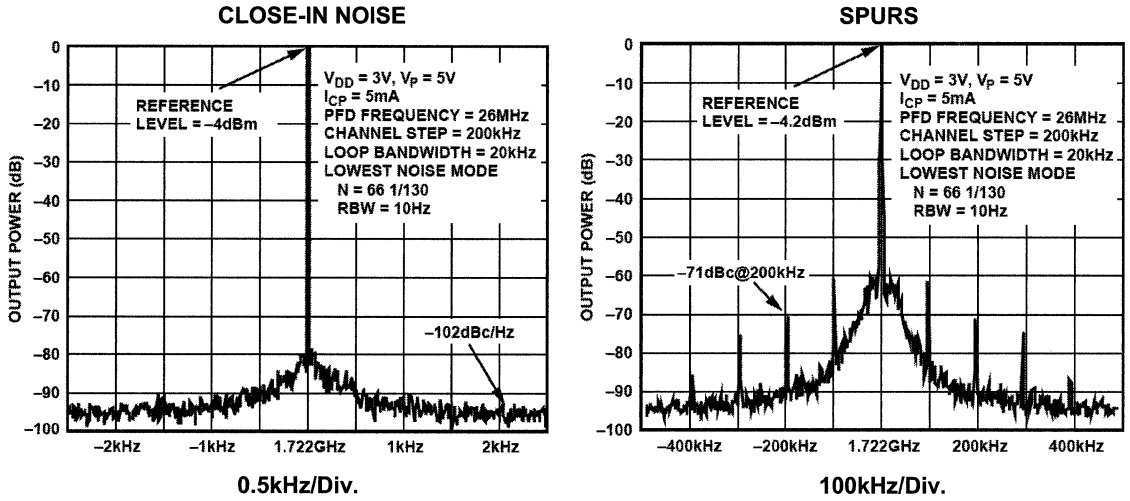

This figure shows the noise and spur performance for the ADF4153 operating in the lowest noise mode. The PFD input frequency is 26MHz, channel step size is 200kHz, loop bandwidth is 20kHz, N = 66 1/130, and the receiver bandwidth is set for 10Hz. The output frequency is 1.7162GHz.

Note that the noise is approximately –102dBc/Hz, and the spur at 200kHz offset is –71dBc.

ADF4153 Operating in Lowest Spur Mode for N = 8581/130

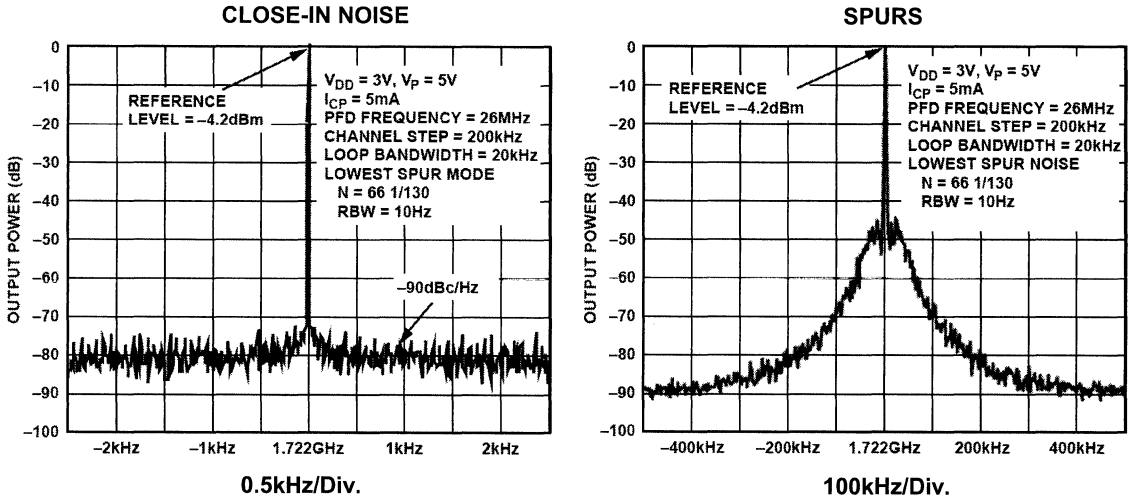

This figure shows the noise and spur performance for the ADF4153 operating in the lowest spur mode. The PFD input frequency is 26MHz, channel step size is 200kHz, loop bandwidth is 20kHz, N = 66 1/130, and the receiver bandwidth is set for 10Hz. The output frequency is 1.7162GHz.

Note that the noise is approximately –90dBc/Hz, and the spur at 200kHz offset indistinguishable from the noise level which is –80dBc.

The ADF4153 also has a "lowest noise and spur" mode. In this mode, the noise is approximately –95dBc/Hz, and the spur at 200kHz offset is –74dBc. These values are measured under the same conditions as the data presented in the two figures.

PLL Design Software

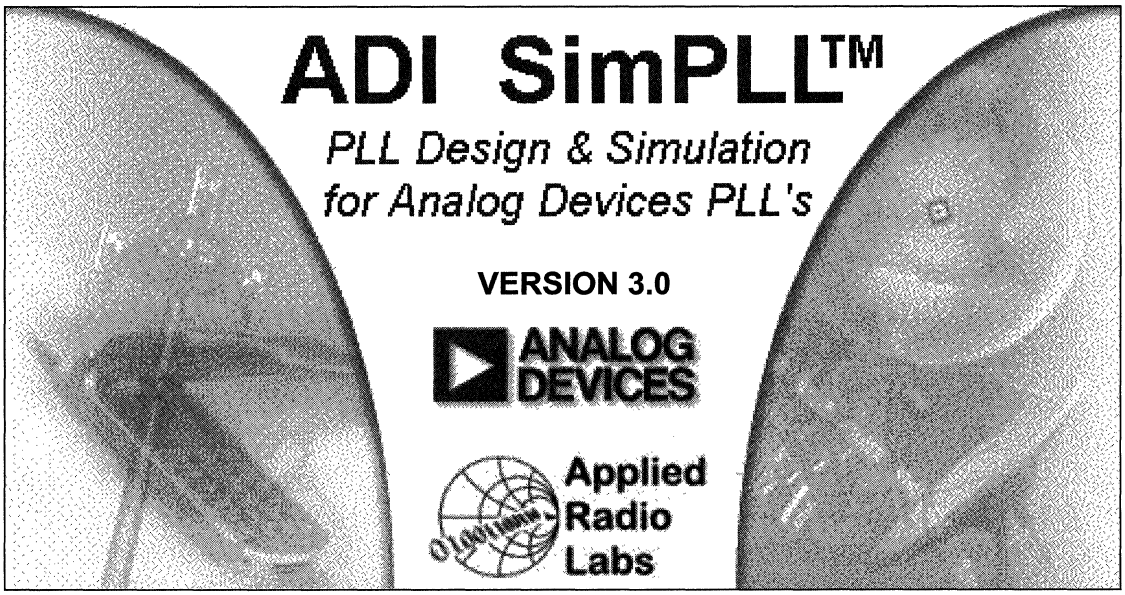

www.analog.com/adisimpll

The ADIsimPLL™ software is a complete PLL design package which can be downloaded from the Analog Devices' website at www.analog.com/adisimpll. The software has a user-friendly graphical interface, and a complete comprehensive tutorial for first-time users.

Traditionally, PLL Synthesizer design relied on published application notes to assist in the design of the PLL loop filter. It was necessary to build prototype circuits to determine key performance parameters such as lock time, phase noise, and reference spurious levels. Optimization was accomplished by "tweaking" component values on the bench and repeating lengthy measurements.

ADIsimPLL both streamlines and improves the traditional design process. Starting with the "new PLL wizard," a designer constructs a PLL by specifying the frequency requirements of the PLL, selecting an integer-N or fractional-N implementation, and then choosing from a library of PLL chips, library or custom VCO, and a loop filter from a range of topologies. The program designs a loop filter and displays key parameters including phase noise, reference spurs, lock time, lock detect performance, and others.

ADIsimPLL operates with spreadsheet-like simplicity and interactivity. The full range of design parameters such as loop bandwidth, phase margin, VCO sensitivity, and component values can be altered with real-time updates of the simulation results. This allows the user to easily optimize the design for specific requirements. Varying the bandwidth, for example, enables the user to observe the tradeoff between lock time and phase noise in real-time, and with bench-measurement accuracy.

ADIsimPLL includes accurate models for phase noise, enabling reliable prediction of the synthesizer closed-loop phase noise. Users report excellent correlation between simulation and measurement. If required, the designer can work directly at the component level and observe the effects of varying individual component values.

ADIsimPLL Design Process

- **Choose reference frequency, output frequency range, and channel spacing**
- **Select PLL chip from list**
- **Select VCO**
- **Select loop filter configuration**
- **Select loop filter bandwidth and phase margin**
- **Run simulation**
- **Evaluate time and frequency domain results**
- **Optimize**

- **Works for integer-N or fractional-N but does not simulate fractional-N spurs. Phase noise prediction for fractional-N devices assumes the device is operating in the "lowest phase noise" mode.**

This figure shows the basic steps in creating a design with ADisimPLL.

ADIsimPLL accurately simulates locking behavior in the PLL, including the most significant non-linear effects. Unlike simple linear simulators based on Laplace transform solutions, ADIsimPLL includes the effects of phase detector cycle slipping, charge pump saturation, curvature in the VCO tuning law, and the sampling nature of the phase-frequency detector. As well as providing accurate simulation of frequency transients, giving detailed lock-time predictions for frequency and phase lock, ADIsimPLL also simulates the lock detect circuit. For the first time, designers can easily predict how the lock detect circuit will perform without having to resort to measurements.

The simulation engine in ADIsimPLL is fast, with all results typically updating "instantaneously," even transient simulations. As well as providing an interactive environment that enables the design to be easily optimized, it also encourages the designer to explore the wide range of design options and parameters available. Contrary to the traditional methods where to design, to build, and then to measure parameters takes days, ADIsimPLL enables the user to change the PLL circuit design and to observe instantly the performance changes. ADIsimPLL allows the designer to work at a higher level and to directly modify derived parameters such as the loop bandwidth; phase margin, pole locations, and the effects of the changes on performance are shown instantly (and without burning fingers with a soldering iron!).

Data Panel and Time Domain Results Panel

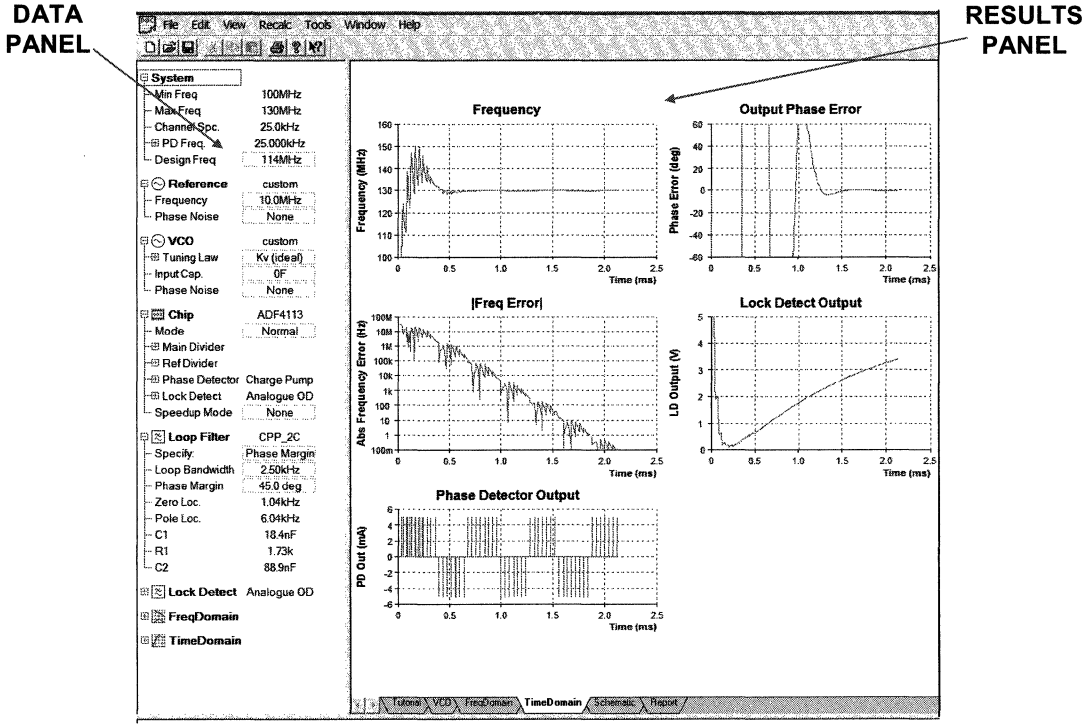

This figure shows the data panel on the left and the time domain results panel on the right. If values are changed in the data panel, the results of the changes are instantly computed and displayed in the results panel.

ADIsimPLL Version 3.0 includes many enhancements including:

1. The program includes a short-form selector guide for choosing the PLL chip, displaying short-form data for all chips, with built-in links to the product pages on the Analog Devices website.

2. Similar short-form selector guides are available for choosing the VCO, and these contain links to detailed device data on vendor's websites.

3. The chip-programming assistant enables rapid calculation of programming register values to set the chip to any specified frequency. This is also good for checking channels that cannot be reached due to prescaler restrictions.

4. The range of loop filters has been expanded to include a 4-pole passive filter and a non-inverting active filter. As with all loop filter designs in ADIsimPLL, these models accurately include the thermal noise from resistors, the op-amp voltage and current noise, as well as predicting reference spurs resulting from the op-amp bias current.

5. Phase jitter results can now be displayed in degrees, seconds or Error Vector Magnitude (EVM)

6. It is now possible to simulate the power-up frequency transient.

7. Support has been included for the new Analog Devices PLL chips with integrated VCOs.

8.The program works for both integer-N and fractional-N devices. However, Phase noise prediction for fractional-N devices assumes the device is operating in the "lowest phase noise" mode.

Frequency Domain Results

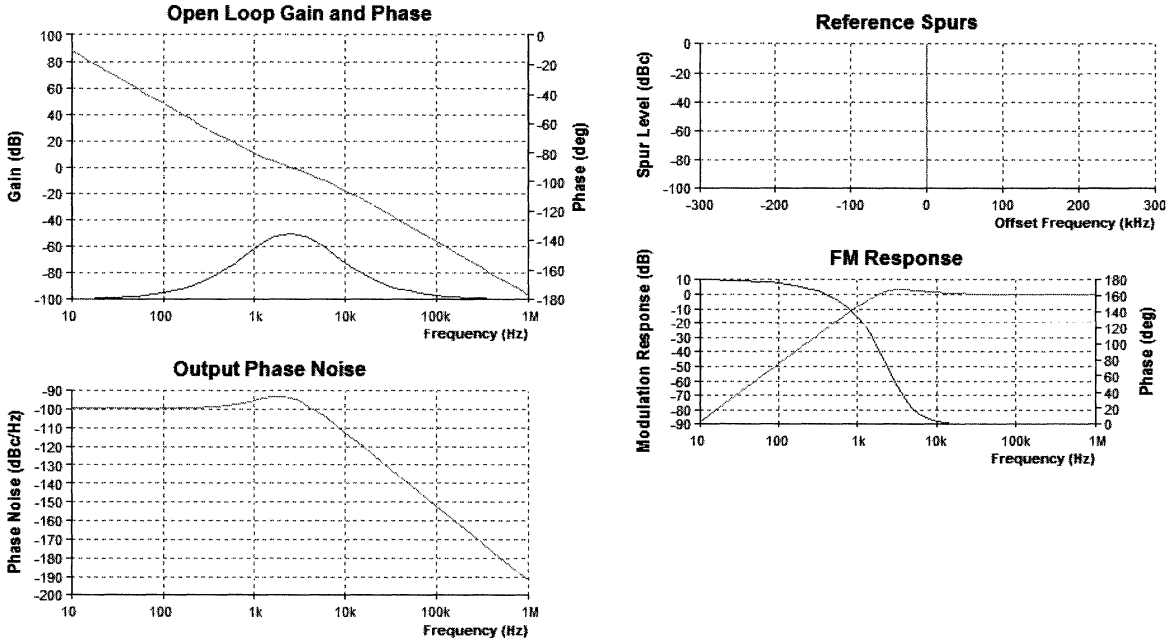

This figure shows a typical display of the frequency domain results in ADIsimPLL. Note that an ideal reference with no spurs was chosen for this particular program demonstration.

ADIsimPLL—Schematic

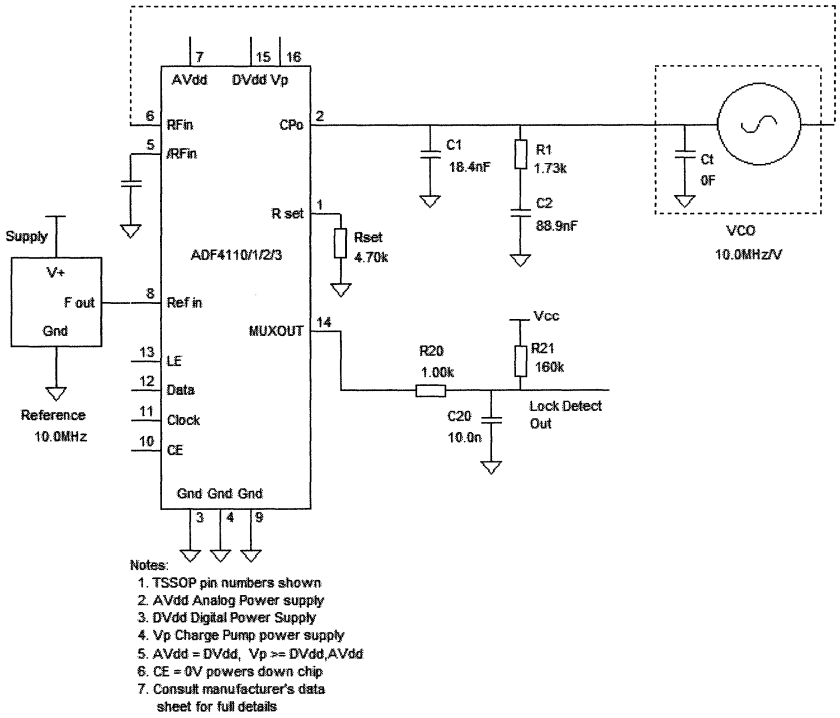

This figure shows the schematic output for the final design, including component values.

All the information required to build the actual PLL is available from the output of ADIsimPLL.

PLL References

1. Mike Curtin and Paul O'Brien, "Phase-Locked Loops for High-Frequency Receivers and Transmitters"
   Part 1, *Analog Dialogue*, 33-3, Analog Devices, 1999
   Part 2, *Analog Dialogue*, 33-5, Analog Devices, 1999
   Part 3, *Analog Dialogue*, 33-7, Analog Devices, 1999
2. Roland E. Best, *Phase Locked Loops, 5th Edition*, McGraw-Hill, 2003, ISBN: 0071412018.
3. Floyd M. Gardner, *Phaselock Techniques, 2nd Edition*, John Wiley, 1979, ISBN: 0471042943.
4. Dean Banerjee, *PLL Performance, Simulation and Design, 3rd Edition*, Dean Banerjee Publications, 2003, ISBN: 0970820712 .
5. Bar-Giora Goldberg, *Digital Frequency Synthesis Demystified*, Newnes, 1999, ISBN: 1878707477.
6. Brendan Daly, "Comparing Integer-N and Fractional-N Synthesizers," *Microwaves and RF*, September 2001, pp. 210-215.
7. Adrian Fox, "Ask The Applications Engineer-30 (Discussion of PLLs)," *Analog Dialogue*, 36-3, 2002.
8. Walt Kester, *Data Conversion Handbook*, Elsevier-Newnes, 2005, ISBN: 0750678410, Chapter 6, pp. 427-440.
9. Hank Zumbahlen, *Basic Linear Design*, Analog Devices, 2006.

This figure shows a list of key references on PLLs and PLL design. They should be very helpful for users inexperienced in PLLs. Reference 1 is a particularly good overview of PLL design. Reference 2 and 3 are classic textbooks on PLLs, and References 4-8 are more recent articles and books.

Clock Generation and Distribution

www.analog.com/clocks

Dimensions of Performance in Clock Generation and Distribution ICs

FUNCTIONS

- Clock Generation
- Clock Cleanup
- Clock Division
- Clock Multiplication
- Delay Adjustment
- Phase Adjustment
- Clock Distribution
- Fanout Buffers
- Logic Level Translation

SPECIFICATIONS

- Speed
- Slew Rate
- Edge Skew
- Rise/Fall Time
- Phase Noise
- Jitter
- Channel Isolation
- EMI/RFI
- Design Tools

System clock management has become much more complex in the last few years. In addition to the requirement for multiple synchronized clocks, the proliferation of data converters has added a new dimension of performance requirements.

Because of the effect of jitter on converter SNR, the converter sampling clock must have very low jitter regardless of how corrupt the other clocks in the system may be.

This figure shows the basic functions of clock generation and distribution and the important performance specifications. The need for these functions has resulted in an entirely new product family of clock and clock distribution ICs which did not exist a decade ago.

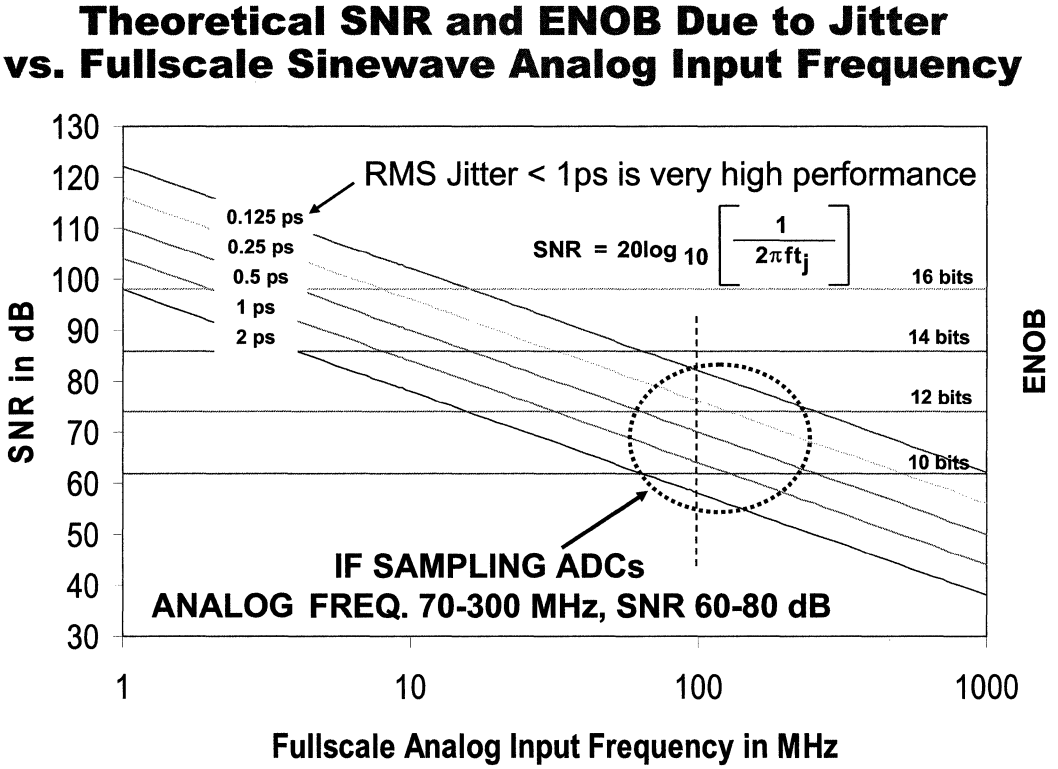

This figure shows the effects of sampling clock jitter on system SNR. It can be shown that jitter on the sampling clock degrades the overall SNR per the simple equation:

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

where f is the fullscale analog input frequency and t_j is the total clock jitter (the root-sum-square of the ADC aperture jitter and the external sampling clock jitter).

It should be noted that this equation assumes an ADC of infinite resolution, where the only error is the noise produced by the clock jitter.

This figure shows SNR on the left and ENOB on the right, related by the simple equation SNR = 6.02N + 1.76dB, where N is the resolution of the data converter.

Typical IF sampling ADCs have IF inputs between 70MHz and 300MHz. The SNR requirement is generally between 60dB and 80dB shown by the circled area in the figure. Note that in order to obtain 80dB SNR for a 100MHz input, the total jitter should be no more than 0.125ps which represents extremely high performance.

Another consideration is that clock jitter adds on an rms basis. The total jitter of a number of series-connected gates is equal the root-sum-square of the jitter associated with each gate:

$$\text{TOTAL JITTER} = \sqrt{J_1^2 + J_2^2 + J_3^2 + \ldots}$$

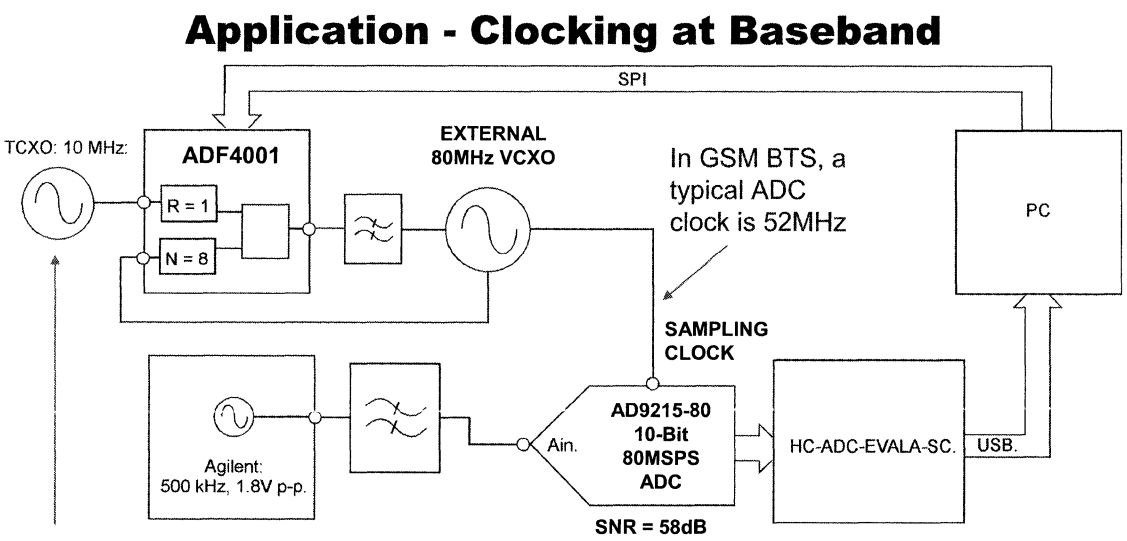

ADF4001 Converter Clock Test Setup

This figure shows a method of generating an ADC sampling clock using a PLL chip such as the ADF4001. The reference clock frequency is 10MHz, but in a GSM basestation it would be 13MHz. The ADF4001 is a clock generation PLL which operates up to 200MHz. It requires an external loop filter and VCO. The R-divider is set for 1, and the N-divider is set for 8 to produce the 80MHz output frequency. Note that in a GSM basestation, the ADC clock is typically 52MSPS. A VCXO rather than a VCO is used for lower phase noise.

The ADC in this application is the AD9215-80 10-bit, 80MSPS converter which has an SNR specification of 58dB. Since this is a baseband application, the maximum ADC input frequency is 40MHz. The SNR due to jitter should be at least 6dB better than the ADC SNR, or 64dB. From the equation on the previous page, SNR = $20\log_{10}[1/2\pi ft_j]$, the jitter corresponding to an SNR of 64dB and an input frequency of 40MHz is approximately 2.5ps.

The measured jitter of the PLL in the above circuit using the ADF4001 was 1.1ps rms, which is more than adequate for the 10-bit baseband application.

ADF4001 200MHz Clock Generator PLL

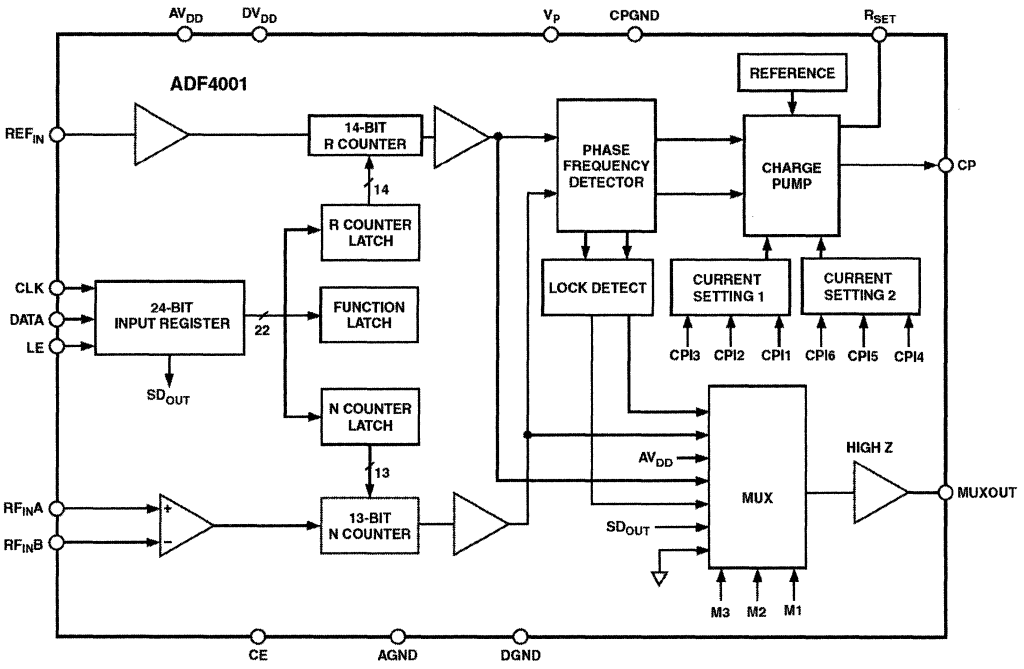

The ADF4001 clock generator can be used to implement clock sources for PLLs that require very low noise, stable reference signals as in the previous example. It consists of a low noise digital PFD (phase frequency detector), a precision charge pump, a programmable reference divider, and a programmable 13-bit N counter. In addition, the 14-bit reference counter (R counter) allows selectable REFIN frequencies at the PFD input. A complete PLL (phase-locked loop) can be implemented if the synthesizer is used with an external loop filter and VCO (voltage controlled oscillator) or VCXO (voltage controlled crystal oscillator). The N minimum value of 1 allows flexibility in clock generation.

The ADF4001 has 200MHz bandwidth, operates on a 2.7V to 5.5V power supply, and has a hardware and software power-down mode.

There is a separate charge pump supply, VP, which allows extended tuning voltage in 5V systems.

The device has programmable charge pump currents and a simple 3-wire serial control interface.

Typical operating current is 4.5mA, and the part comes in either a 16-lead TSSOP or a 20-lead LFCSP.

Reducing the Complexity of Clocking at Baseband by Using ADF4360-8 Integrated PLL

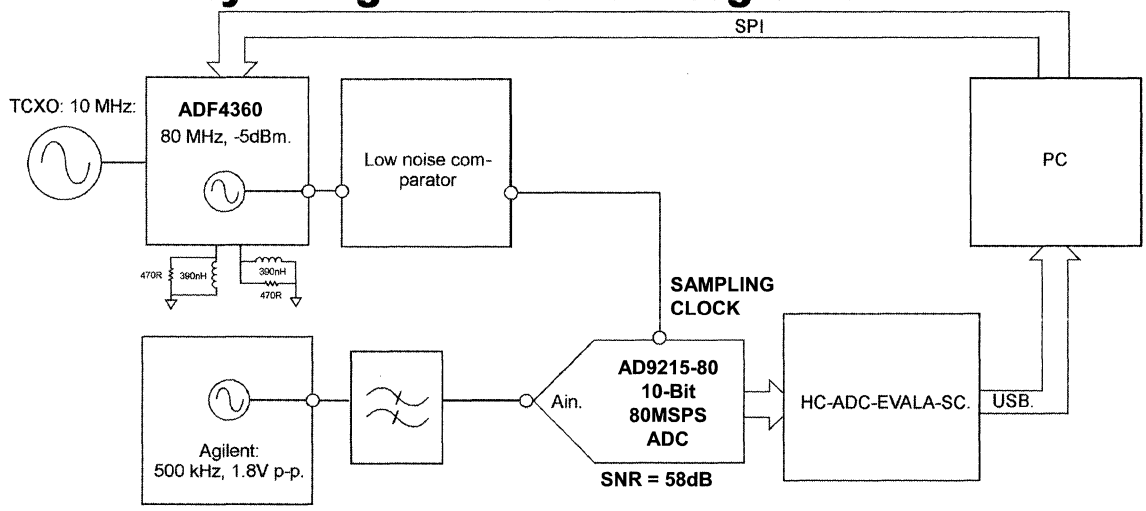

- **The ADF4360-8 integrates a VCO on-chip to further reduce cost and board space**
- **In this configuration, the ADF4360-based clock will produce 2.45 ps rms jitter. This is fine for most baseband applications.**

This baseband clock generator utilizes the ADF4360-8 to perform the PLL function. This is a simpler solution, because the ADF4360-series contains an integrated VCO. An external low noise comparator is used to boost the amplitude of the output of the ADF4360 to a level suitable for driving the ADC.

The measured jitter of this generator is 2.45ps which is approximately the system requirement for 64dB SNR (due to jitter only) with an input frequency of 40MHz. (See Figure 3.72).

ADF4360-8 Integrated Synthesizer and VCO Output: 65MHz to 400MHz

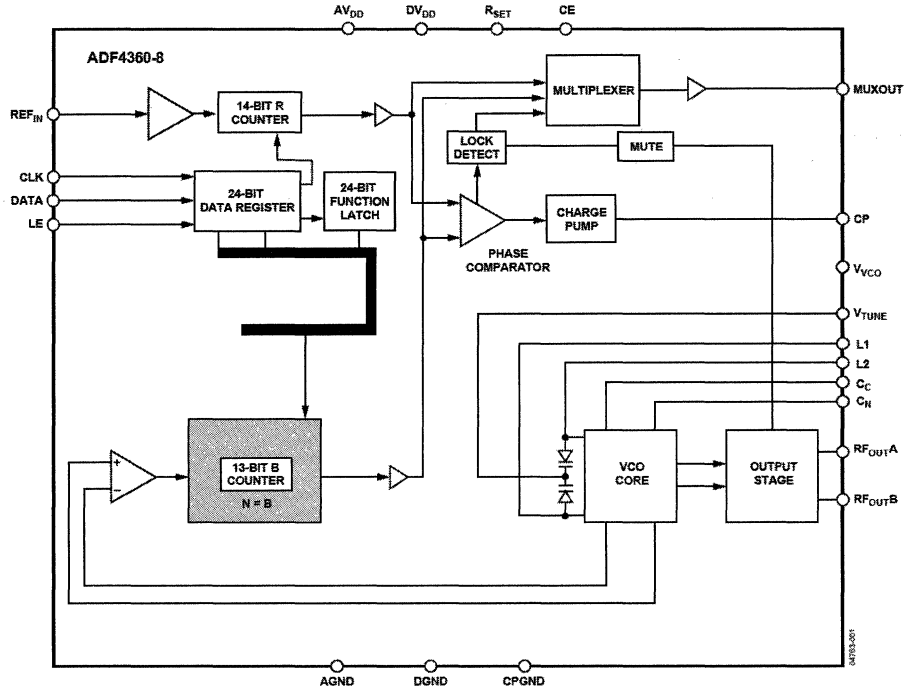

The ADF4360-8 is an integrated integer-N synthesizer and voltage controlled oscillator. The ADF4360-8 center frequency is set by external inductors. This allows a frequency range of between 65MHz and 400MHz.

Control of all the on-chip registers is through a simple 3-wire interface. The device operates with a power supply ranging from 3.0V to 3.6V and can be powered down when not in use.

ADF4001/ADF4360-8 Jitter and Phase Noise

- **ADF4001 with external VCXO, 1.1ps rms jitter**
- **ADF4360-8, 2.45ps rms jitter**
- **Both meet min jitter requirement for many baseband sampling applications**
- **Simulations shown using ADI's free PLL/VCO design tool, ADIsimPLL found at www.analog.com/ADIsimPLL**
- **The ADF4002 increases usable output bandwidth to 400MHz compared to 200MHz on the ADF4001.**

ADF4001

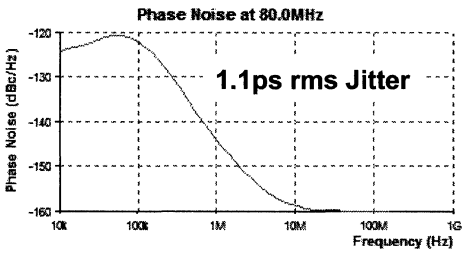

ADF4360-8

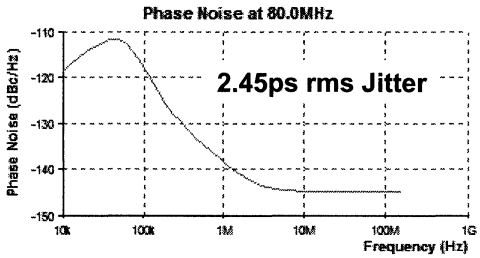

This figure compares the phase noise and jitter performance of the ADF4001 design (external VCXO) and the ADF4360-8 design (internal VCO). Both approaches meet the jitter requirement for many baseband sampling applications.

The ADF4002 is similar to the ADF4001, but increases the usable output bandwidth to 400MHz compared to 200MHz for the ADF4001.

AD9540 655 MHz Low Jitter Clock Generator

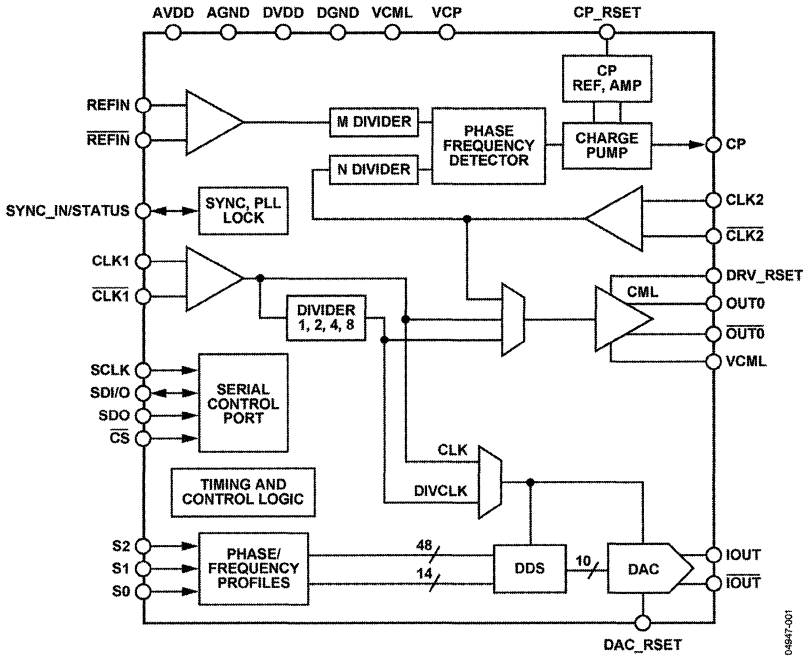

The AD9540 is Analog Devices' first dedicated clocking product specifically designed to support the extremely stringent clocking requirements of the highest performance data converters. The device features high performance PLL (phase-locked loop) circuitry, including a flexible 200MHz phase frequency detector and a digitally controlled charge pump current.

The device also provides a low jitter (720fs), 655MHz CML-mode), PECL-compliant output driver with programmable slew rates. External VCO rates up to 2.7GHz are supported. Extremely fine tuning resolution (steps less than 2.33μHz) is another feature supported by this device.

Information is loaded into the AD9540 via a serial I/O port that has a device write speed of 25Mbps. The AD9540 frequency divider block can also be programmed to support a spread spectrum mode of operation. The AD9540 is specified to operate over the extended automotive range of −40°C to +85°C.

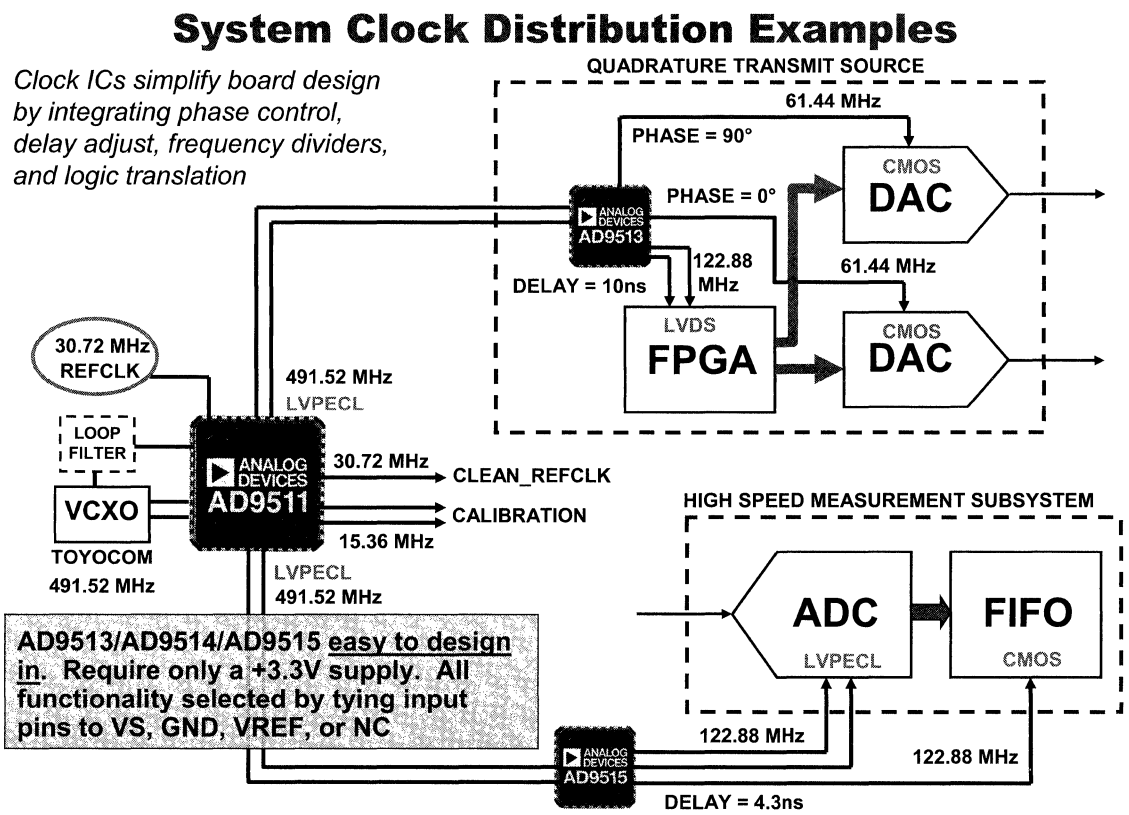

This figure illustrates a typical system clock generation and distribution application. The system reference clock is 30.72MHz and circled on the left. The PLL function of the AD9511 IC is used to generate a "clean" low jitter 491.52MHz clock from the reference clock. Note that the loop filter and the VCXO are external to the AD9511.

One of the 491.52MHz low voltage PECL (LVPECL—see footnote below) outputs is sent to the AD9515 clock distribution IC which contains two programmable divider channels (1 to 32) with a coarse delay adjustment in one channel. The AD9615 is used as the clock distribution chip for the high speed measurement system. The dividers are set to divide by 4 to generate the 122.88MHz sampling clock to the ADC and the data strobe to the FIFO. Note that the delay is set for 4.3ns in the FIFO strobe such that the ADC data is clocked into the FIFO at the proper time. The ADC is clocked with the LVPECL output, and the second output (with the delay) is set to generate a CMOS logic level for the FIFO. The additive jitter of the AD9515 is approximately 300fs.

A second 491.52MHz LVPECL output from the AD9511 is used as the master clock for the quadrature transmit source. It drives an AD9513 clock distribution IC which contains three divider channels (1 to 32), one of which includes a delay adjustment. In this application, the delay adjustment is used to properly time the 122.88MHz clock to the FPGA by introducing a delay of 10ns. The other two channels in the AD9513 are set to divide by 8 to generate the 61.44MHz clocks for the DACs. The coarse phase adjustment in the AD9513 is used to generate the 90° phase shift between the I/Q DAC clocks. The additive jitter of the AD9513 is approximately 300fs.

LVPECL is basically ECL logic operated between ground V_S = +3.3V. Under these conditions, a logic "1" is approximately V_S – 0.96V = +2.34V and a logic "0" is approximately V_S – 1.76V = +1.54V with 50Ω terminations to V_S – 2V = +1.3V.

AD9512 1.2GHz Clock Distribution IC

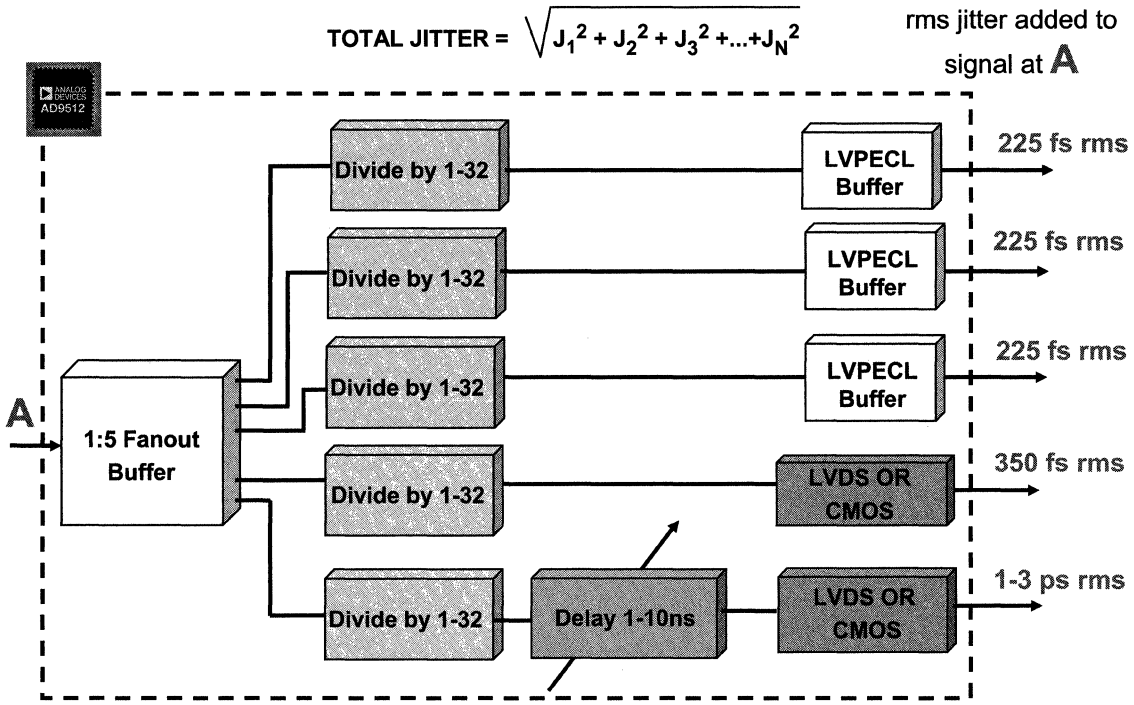

The AD9512 1.2GHz clock distribution IC provides a multi-output clock distribution in a design that emphasizes low jitter and low phase noise to maximize data converter performance. Other applications with demanding phase noise and jitter requirements can also benefit from this part. There are five independent clock outputs. Three outputs are LVPECL (1.2GHz), and two are selectable as either LVDS (800MHz) or CMOS (250MHz) levels.

Each output has a programmable divider that may be bypassed or set to divide by any integer up to 32. The phase of one clock output relative to another clock output may be varied by means of a divider phase select function that serves as a coarse timing adjustment. There is a small jitter penalty when using the variable phase output.

One of the LVDS/CMOS outputs features a programmable delay element with a range of up to 10 ns of delay. This fine tuning delay block has 5-bit resolution, giving 32 possible delays from which to choose. The AD9512 is ideally suited for data converter clocking applications where maximum converter performance is achieved by encode signals with subpicosecond jitter.

The AD9512 is available in a 48-lead LFCSP and can be operated from a single 3.3 V supply. The temperature range is −40°C to +85°C.

Clock Evaluation Kits

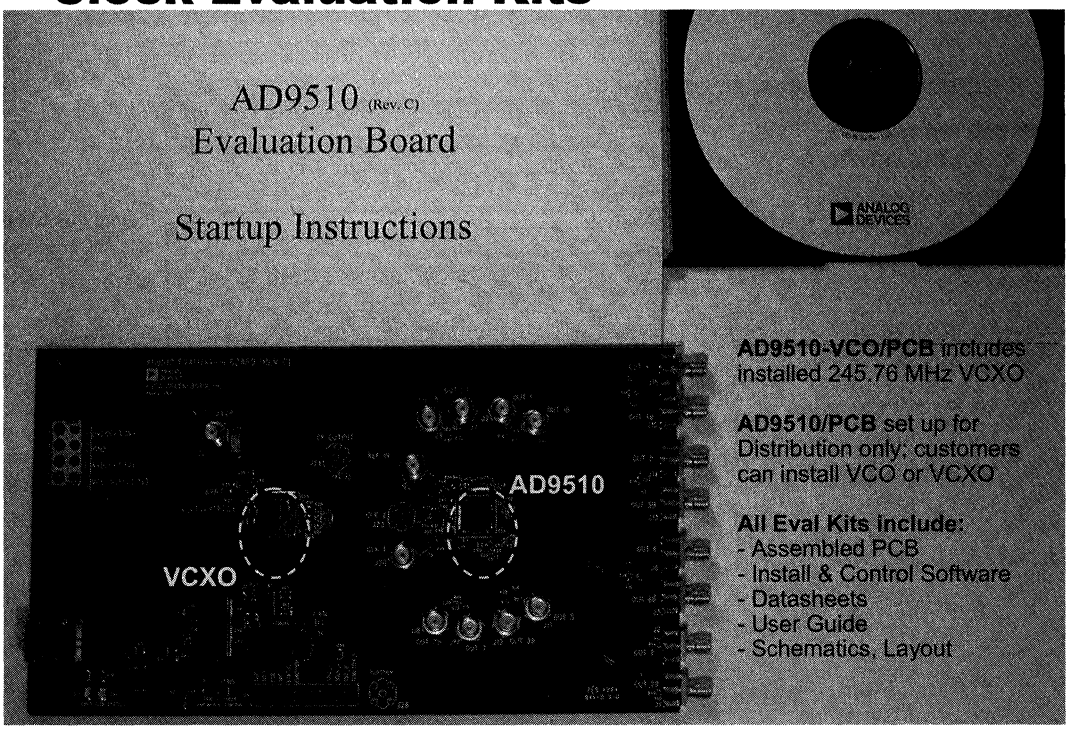

The Analog Devices' clock generator and distribution products have a complete set of evaluation boards. This shows the board for the AD9510 which includes an installed 245.76MHz VCXO. The board can also be set up for a user-installed VCO or VCXO.

Clock Evaluation Board User Interface Is Intuitive; Looks Like Block Diagram of Chip

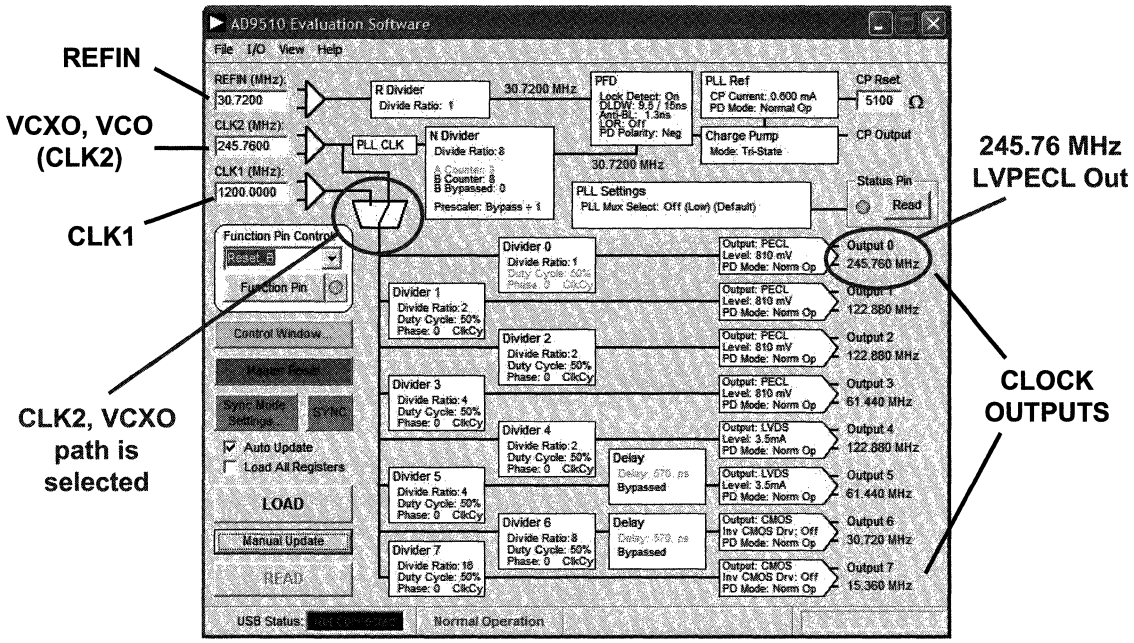

The clock evaluation board software has an intuitive user interface as shown here. The screen format follows the block diagram of the individual chip for ease of use.

ADIsimCLK Design and Simulation Software

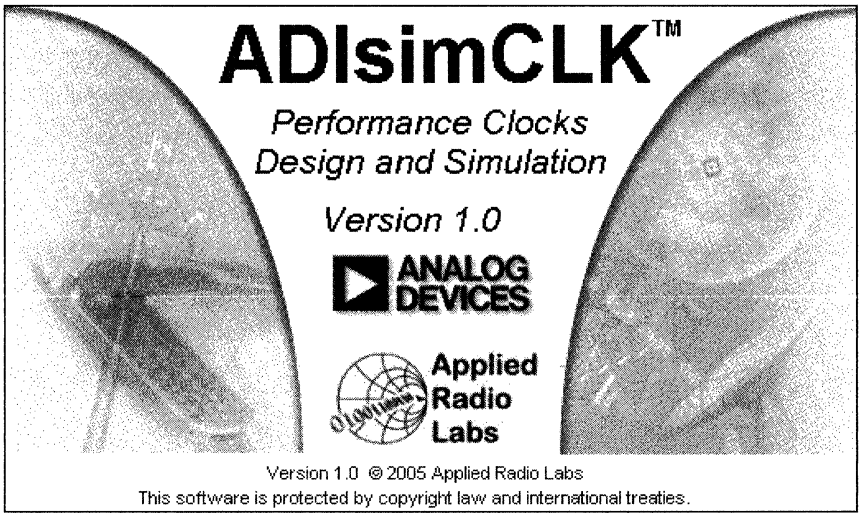

Free download from ADI website @ www.analog.com/clocks

The ADIsimCLK design and simulation software is designed to aid in the application of the clock product families.

Because many of the clock products contain PLLs, ADIsimCLK makes use of a large portion of ADIsimPLL for the PLL part of the design.

ADIsimCLK Design Block Diagram

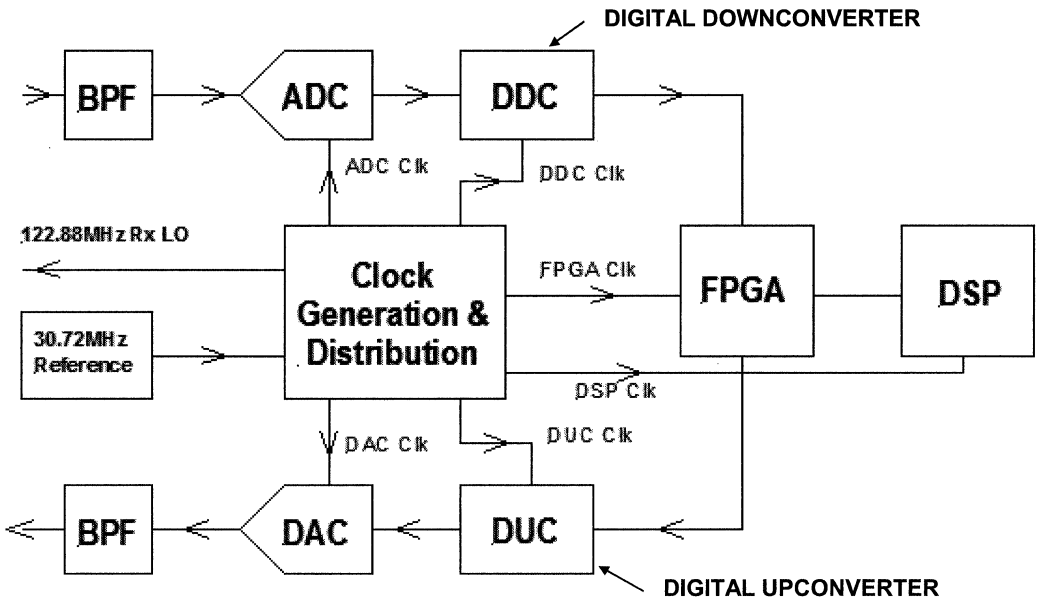

This shows a block diagram of the design interface for ADIsimCLK.

ADIsimCLK Is Schematic Driven

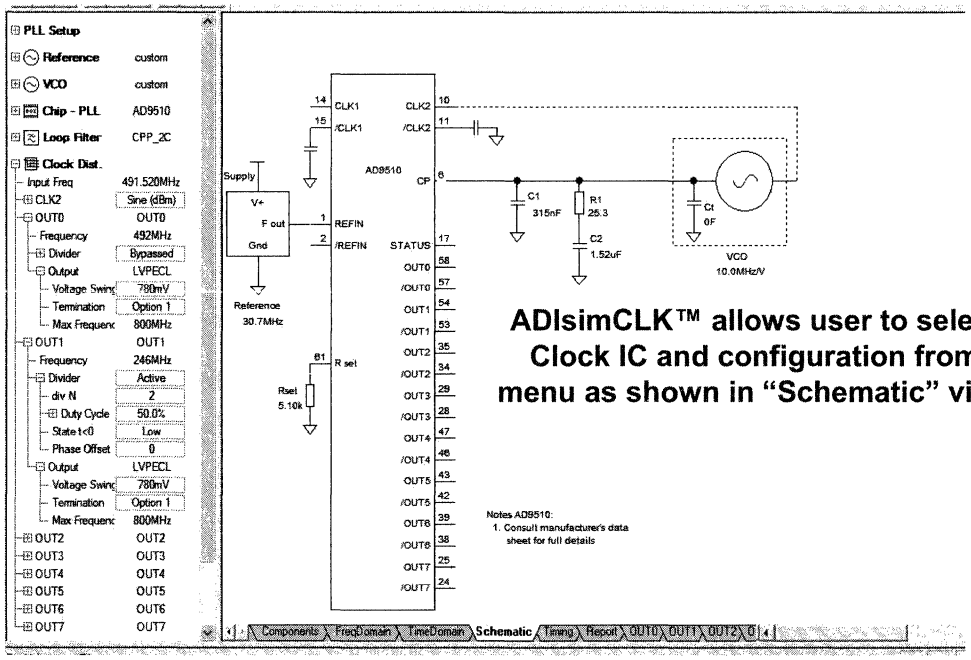

ADIsimCLK™ allows user to select Clock IC and configuration from menu as shown in "Schematic" view

The ADIsimCLK program is schematic driven and allows the user to select the clock IC and the configuration from the menu when in the "schematic" view.

Data in the table on the left can be modified and the results displayed instantaneously in the "results" table, similar to the ADIsimPLL program.

ADIsimCLK Configuration Options

(A) INTEGRATED PLL WITH CLOCK CLEANUP, FREQUENCY TRANSLATION AND CLOCK DISTRIBUTION

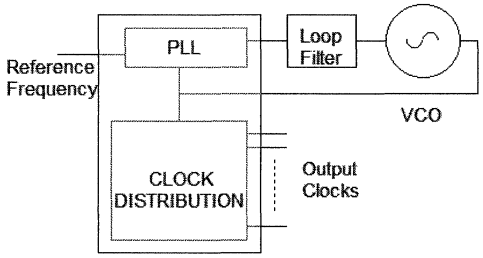

(B) ADDING AN EXTERNAL FILTER TO (A) FOR LOWER PHASE NOISE

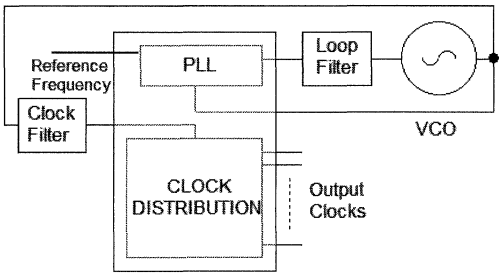

(C) CLOCK DISTRIBUTION CIRCUIT ONLY

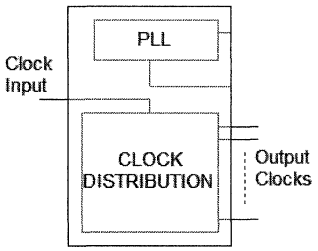

This shows the three basic configuration options currently available with ADIsimCLK.

The (A) configuration is for a clock generation product with an integrated PLL and clock distribution circuit.

The (B) configuration allows the addition of an external filter to the PLL output for lower phase noise and jitter.

The (C) configuration is for distribution products only.

ADIsimCLK Outputs

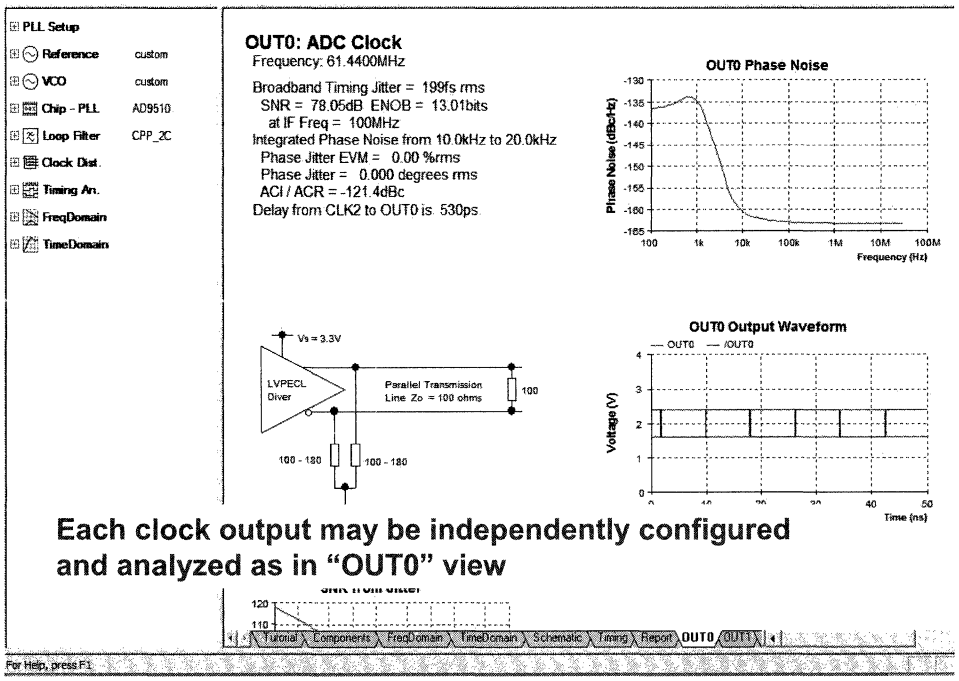

The various outputs of the clock product can be independently configured and analyzed as shown here.

Relative Timing of All Clocks Is Available in the "Timing" View

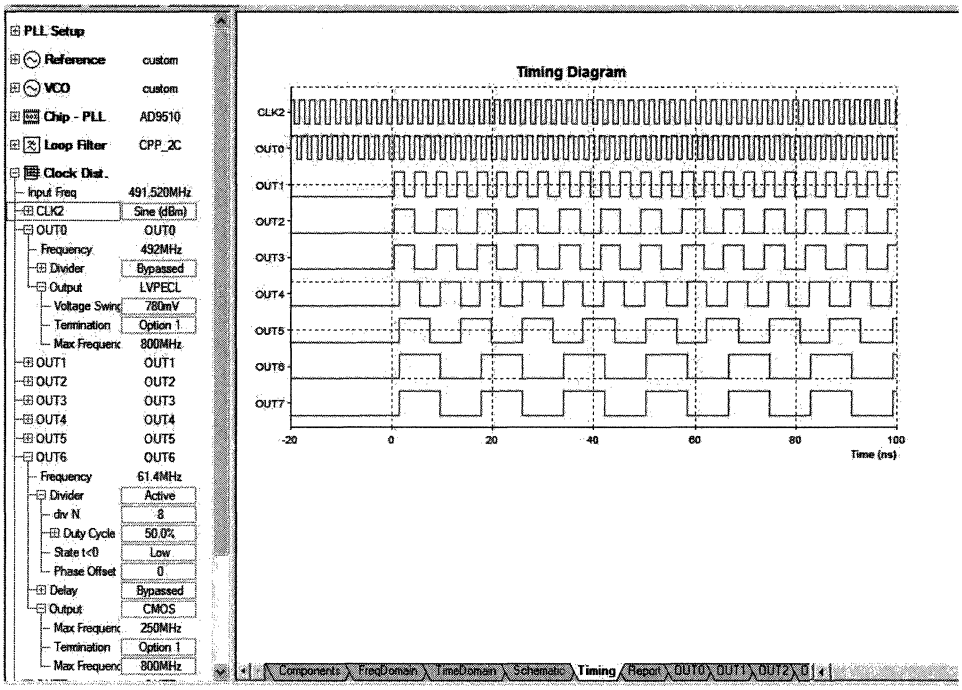

The relative timing of all the clocks associated with a particular clock IC is available in the "timing" view as shown here.

Generating Low Jitter Clocks Using DDS Systems

Reference:

David Brandon, "Direct Digital Synthesizers in Clocking Applications," Application Note AN-823, Analog Devices, 2006

Generating Clocks from a DDS

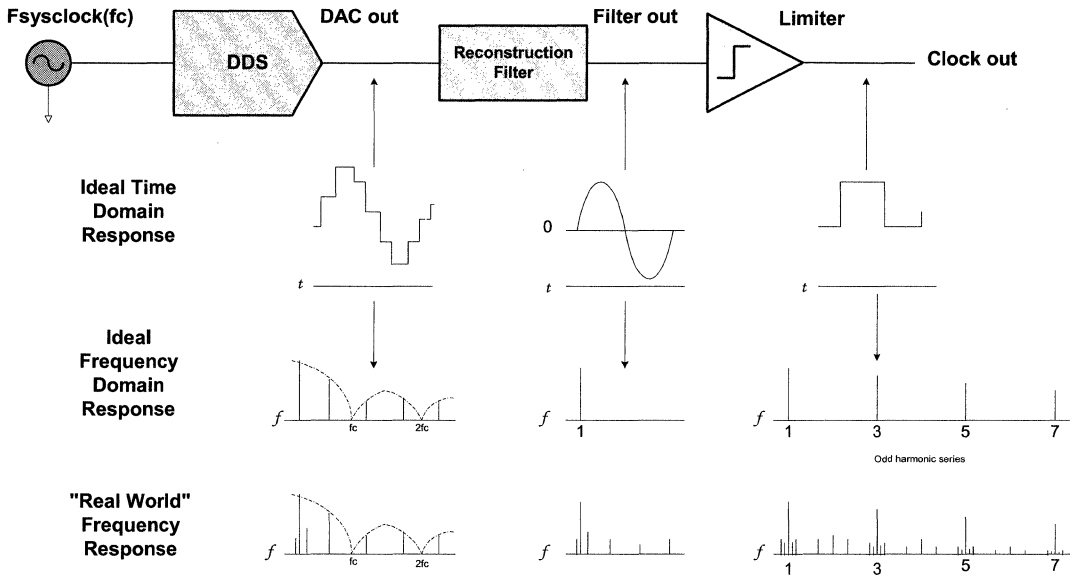

The DDS chip can synchronize to a user's reference. An on-chip clock multiplier can generate the fast clock needed to clock the NCO/DAC. A frequency tuning word may be written to set the output clock rate. External filtering removes unwanted images. A squaring function then converts sine wave to square wave.

The primary challenge in using DDS systems for clock generation is reducing the time jitter that is generated due to the discrete spurious components that are present on the DDS output signal. These spurious components are primarily due to the quantized nature of the DDS output signal and the phase noise produced because of phase truncation within the DDS.

This figure shows time and frequency domain representations at each point in the process. A real-world representation of the frequency domain is shown at the bottom of the figure. This response contains aliases of higher order harmonics of the output signal which fall close to the desired output signal, and filtering becomes difficult.

The overall quality of the final clock output of a DDS system is determined primarily by the filter and the limiter. AN-823 discusses the results of experiments conducted with a high performance DDS and clock distribution IC and shows the jitter associated with various filter configurations.

Test Set for Measuring Jitter of AD9958/AD9858 DDS Driving AD9515 Clock Distribution IC

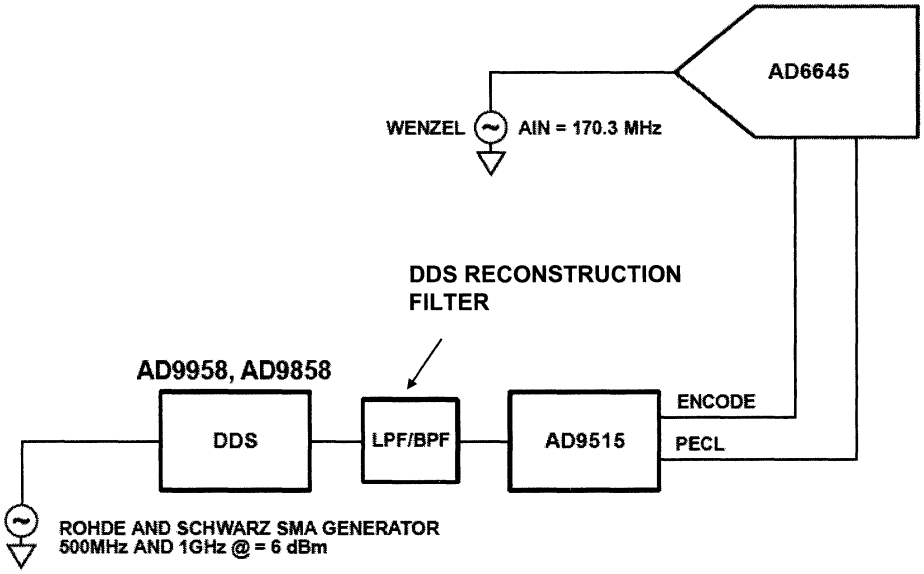

This is the basic test set used to measure the jitter performance of the AD9958/AD9858 DDS driving the AD9515 clock distribution IC. The jitter is calculated by FFT analysis of the AD6645 14-bit output data using the relationship between SNR and jitter,

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

The DDS sampling clock is derived from a Rohde and Schwarz SMA signal generator. Data was taken on two different DDSs, the AD9958 and the AD9858. The jitter measurement was made by using the clock derived by the DDS and the AD9515 as the sampling clock to the AD6645, a 14-bit, 100MSPS ADC.

The analog input to the ADC was derived from a Wenzel crystal oscillator (www.wenzel.com) known for its low jitter.

New DDS Chips Like the AD9958 Dual Channel and AD9959 Quad Have Very Low Phase Noise

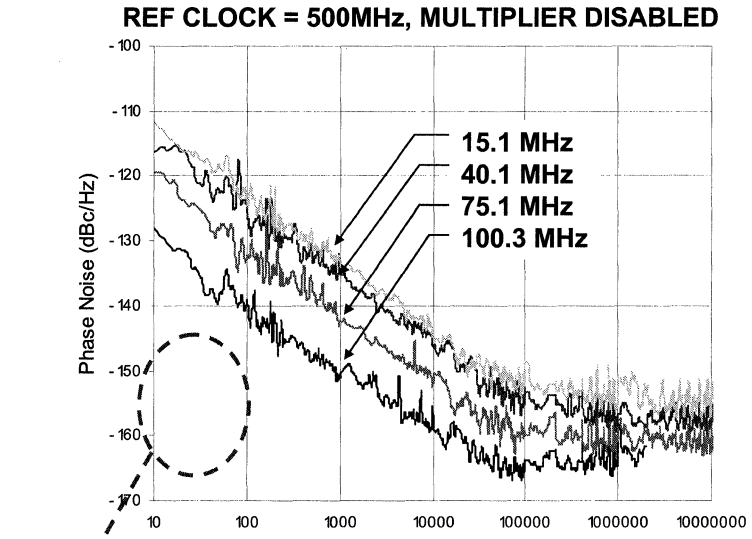

Power dissipation <200mW per channel

This shows the phase noise of the AD9958 DDS chip with a reference clock of 500MHz. Output frequencies of 15.1MHz, 40.1MHz, 75.1MHz, and 100.3MHz are shown.

Note that the broadband phase noise floor is below –150dBc/Hz in all cases.

Application Note AN-823 Details Performance of AD9958 Driving AD9515

AN-823 shows jitter < 1 ps rms possible with proper filtering

Table 1. Jitter Response AD9958 and AD9515 vs. Fout, Power, Frequency and Filter BW

Product	DDS Sample Rate (MHz)	DDS Output Frequency (MHz)	DDS Output Power (dBm)	DDS Reconstruction Filter (MHz)	AD9515 Divider Output Setting	AD9515 Output Frequency (MHz)	Jitter (rms) (ps)
AD9958/AD9515	500	38.88	–3.6	200 LPF	1	38.88	4.1
AD9958/AD9515	500	38.88	–3.6	200 LPF	2	19.44	4.1
AD9958/AD9515	500	38.88	–4.7	47 LPF	1	38.88	2.4
AD9958/AD9515	500	38.88	–4.7	47 LPF	2	19.44	2.4
AD9958/AD9515	500	38.88	–3.3	5% BPF	1	38.88	1.5
AD9958/AD9515	500	38.88	–3.3	5% BPF	2	19.44	1.5
AD9958/AD9515	500	77.76	–3.8	200 LPF	1	77.76	2.5
AD9958/AD9515	500	77.76	–3.8	200 LPF	2, 4	38.88, 19.44	2.5
AD9958/AD9515	500	77.76	–4.9	85 LPF	1	77.76	1.5
AD9958/AD9515	500	77.76	–4.9	85 LPF	2, 4	38.88, 19.44	1.5
AD9958/AD9515	500	77.76	–3.8	5% BPF	1	77.76	1.1
AD9958/AD9515	500	77.76	–3.8	5% BPF	2, 4	38.88, 19.44	1.1
AD9958/AD9515	500	155.52	–5.5	200 LPF	2	77.76	1.5
AD9958/AD9515	500	155.52	–5.5	200 LPF	4, 8	38.88, 19.44	1.5
AD9958/AD9515	500	155.52	–5.6	5% BPF	2	77.76	0.68
AD9958/AD9515	500	155.52	–5.6	5% BPF	4, 8	38.88, 19.44	0.68

This table shows the jitter of the final output clock for various DDS output frequencies, AD9515 divider settings, and reconstruction filter characteristics.

Note that the lowest jitter is obtained when the reconstruction filter is a bandpass filter having a bandwidth equal to approximately 5% of the DDS output frequency. This obviously limits the tuning range of the DDS, however.

Wider filter bandwidths allow wider tuning ranges, at the expense of increased jitter. For further details of the results of these experiments, consult the application note.

David Brandon, "Direct Digital Synthesizers in Clocking Applications," Application Note AN-823, Analog Devices, 2006.

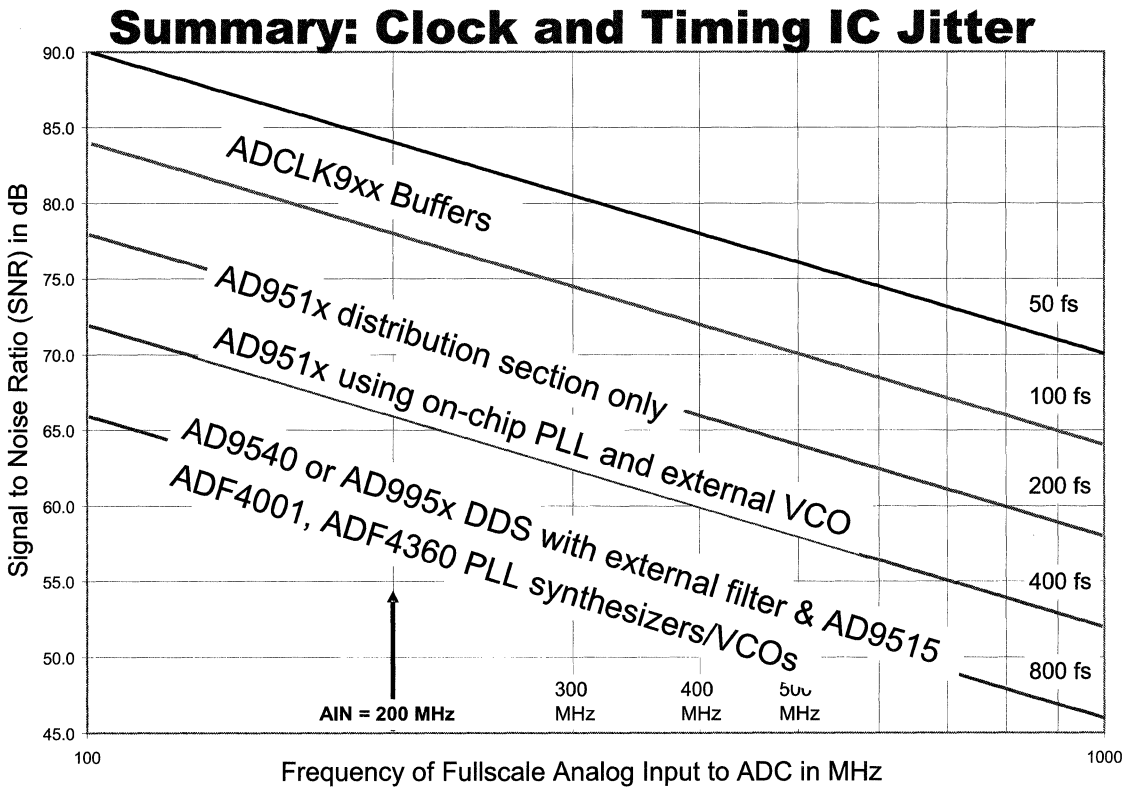

This figure summarizes the current jitter performance of the various IC solutions discussed in this section. The new ADCLK9xx family of clock buffers from Analog Devices are designed with additive jitter less than 100fs.

High Speed System Applications

SECTION 4
PC BOARD LAYOUT AND DESIGN TOOLS

Grounding and Layout

One of the biggest problems in system design is how to handle grounding. There are several competing requirements that are dependent on the frequency and system complexity.

Unfortunately, there is no magic "cookbook" approach to grounding that will always guarantee success. What we will do here is present some of the effects that must be considered when designing the system.

The main thing is to look at how and where the dc and ac currents flow in a PCB.

More information on PC board design techniques can be found in the following two references:

Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN: 0916550273 Chapter 9. Also available as *Data Conversion Handbook*, Elsevier-Newnes, 2004, ISBN: 0750678410, Chapter 9.

Walt Jung, *Op Amp Applications*, Analog Devices, 2002, ISBN: 0-916550-26-5, Chapter 7. Also available as *Op Amp Applications Handbook*, Elsevier-Newnes, 2004, ISBN: 0-7506-7844-5, Chapter 7.

Kirchoff's Law Helps Analyze Voltage Drops Around a Complete Circuit

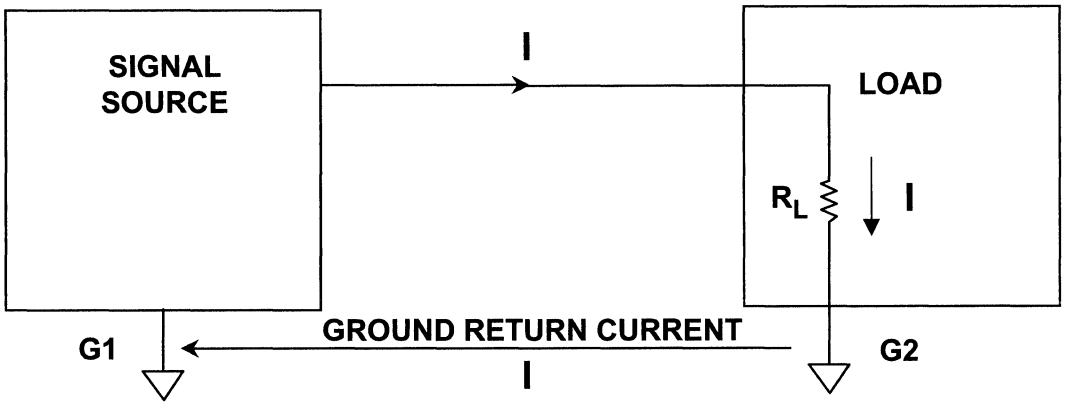

AT ANY POINT IN A CIRCUIT
THE ALGEBRAIC SUM OF THE CURRENTS IS ZERO
OR
WHAT GOES OUT MUST COME BACK
WHICH LEADS TO THE CONCLUSION THAT
ALL VOLTAGES ARE DIFFERENTIAL
(EVEN IF THEY'RE GROUNDED)

When we draw the ground symbol on a schematic, we assume that all ground points are at the same potential. This is rarely the case, unfortunately. Historically, ground was the reference level with which we measured various voltage levels in the circuit. However, ground has also become the power return not only for digital signals but for analog signals as well.

All signals that flow in a circuit must have a return path to complete the loop. Often we consider the forward path only, but there always must be a return to close the loop or current can not flow. This return path is often through the ground plane.

A More Realistic View of the Impedance Between Grounds

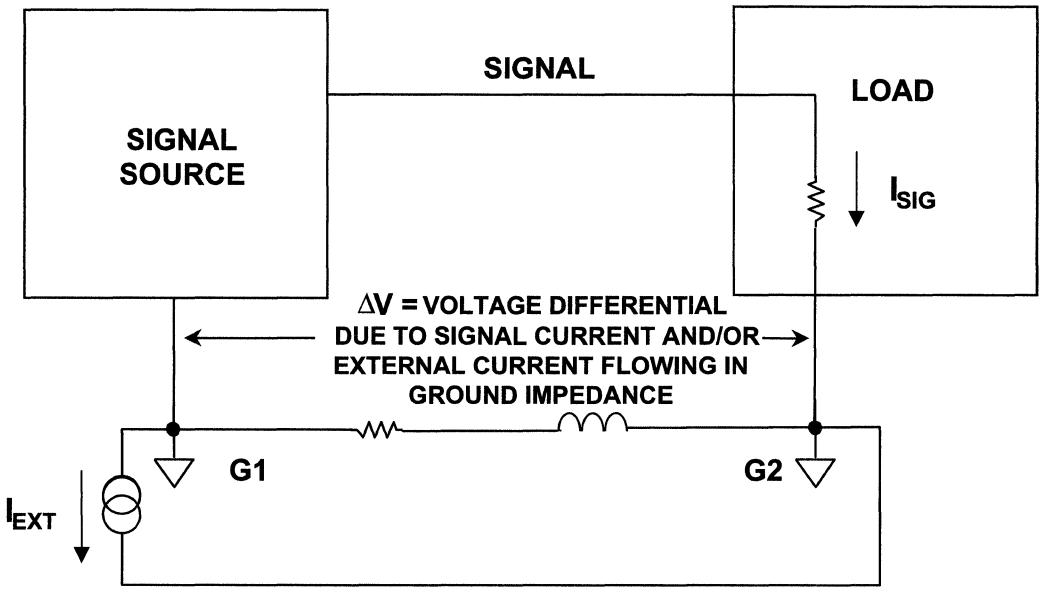

A ground is never zero impedance. There is always some resistance and inductance, even in a large area heavy ground plane. The magnitude of the impedance may be small, but it is not zero. And a current flowing through an impedance causes a voltage drop. This means that the two grounds in the diagram above will not be at the same potential.

It is important to consider the inductance of the ground as well as the resistance, especially as the frequency increases.

Digital Currents Flowing in Analog Return Path Create Error Voltages

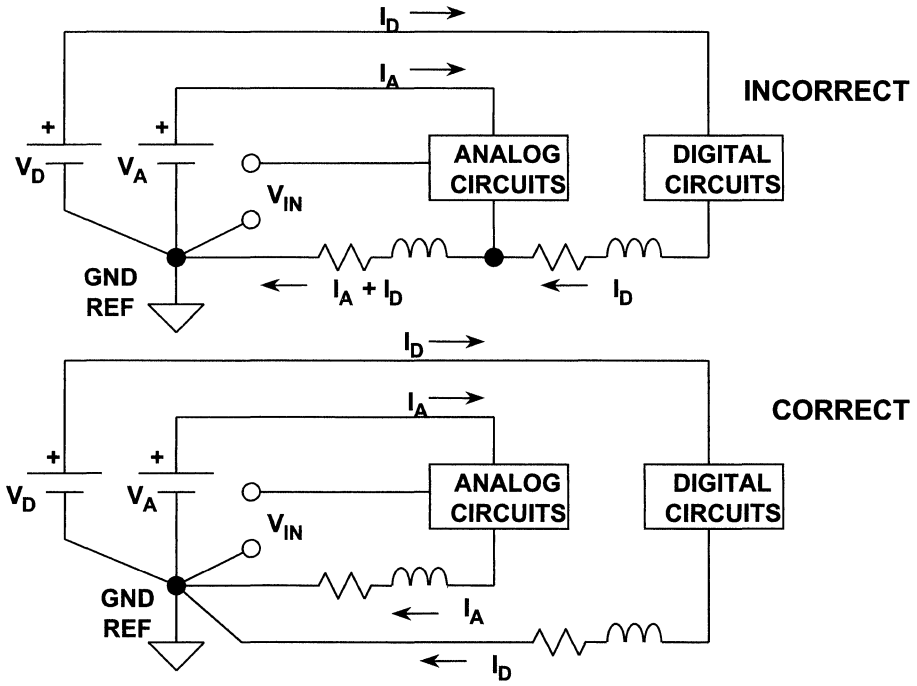

Because ground is the power return for all digital circuits, as well as many analog circuits, one of the most basic design philosophies is to separate digital ground from analog ground.

If the grounds are not separated, not only does the return from the analog circuitry flow through the analog ground impedance, but the digital ground current also flows through the analog ground, and the digital ground current is typically much greater than the analog ground current.

As the frequency of digital circuits increases, the noise generated on the ground increases dramatically. TTL and CMOS logic families are of the saturating types. This means that the logic transitions cause large transient currents on the power supply and ground. CMOS outputs basically connect the power to ground through a low impedance during the logic transitions.

And it's not just the basic clock rate that is a problem. Digital logic waveforms are basically rectangular waves, which implies many higher frequency harmonic components.

Calculation of Sheet Resistance and Linear Resistance

$$R = \frac{\rho Z}{XY}$$

ρ = RESISTIVITY

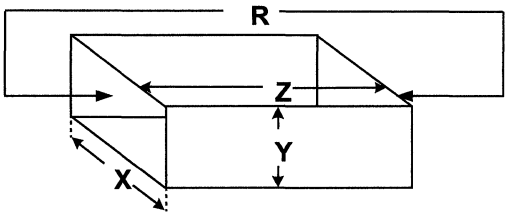

SHEET RESISTANCE CALCULATION FOR 1 OZ. COPPER CONDUCTOR:

$\rho = 1.724 \times 10^{-6}\ \Omega$ cm, Y = 0.0036cm

$$R = 0.48 \frac{Z}{X}\ m\Omega$$

$\frac{Z}{X}$ = NUMBER OF SQUARES

R = SHEET RESISTANCE OF 1 SQUARE (Z = X)
= 0.48mΩ/SQUARE

Resistance is the first ground component we will consider. All conductors have some resistance (at least when operating above 0°K). Using large area ground planes decreases the resistance, but cannot eliminate it. And from Ohm's law we know that a current flowing through a resistance will cause a voltage drop across the resistance.

The resistance of a trace (or a ground plane) can be calculated by taking the resistivity of the material, which will typically be given in a resistance per unit volume (squares) of the conductor material, and multiplying by the number of the squares.

In the above example, the sheet resistance of 1 oz. copper, which is a typical PC board material, is calculated as 0.48mΩ/square.

Even Small Common Ground Currents Can Degrade Precision Amplifier Accuracy

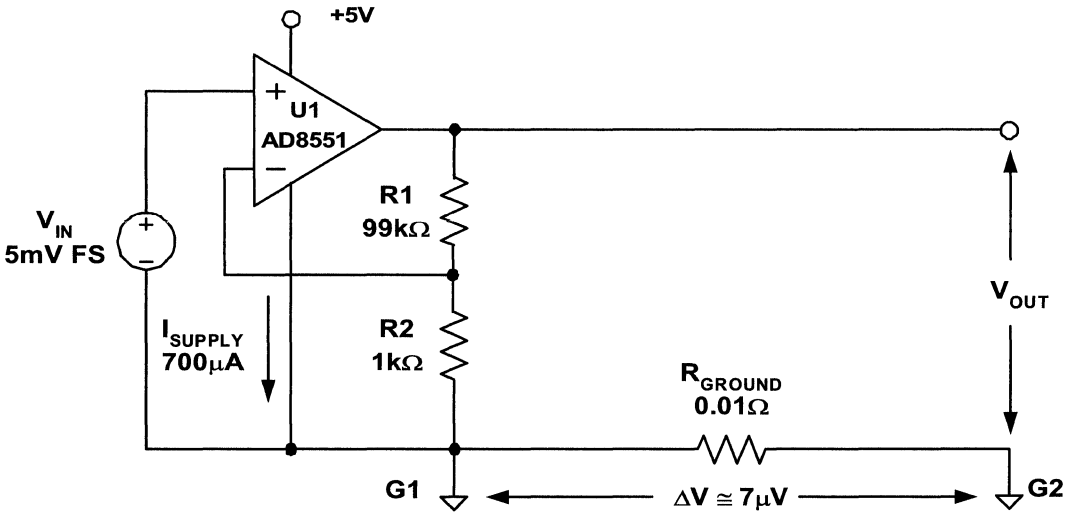

This is an example of how even a small amount of resistance can cause significant error.

Here the quiescent (supply) current of the AD8551 (700μA) flowing through the ground resistance (0.01Ω) causes an error at the point where the signal will be processed (Vout). Although 700μA is a relatively low current, and the ground resistance of 0.01Ω is also a relatively low value, this combination will cause a voltage drop, in this example, of 7μV, also a low value, but much greater than the offset voltage spec of 1μV of the AD8551.

The AD8551 is a chopper stabilized amplifier which, in addition to the voltage offset specification of 1μV, has a offset voltage drift spec of 0.005μV/°C.

Wire and Strip Inductance Calculations

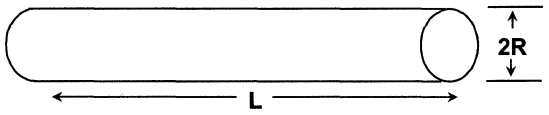

L, R in mm

$$\text{WIRE INDUCTANCE} = 0.0002L\left[\ln\left(\frac{2L}{R}\right) - 0.75\right]\ \mu H$$

EXAMPLE: 1cm of 0.5mm o.d. wire has an inductance of 7.26nH
(2R = 0.5mm, L = 1cm)

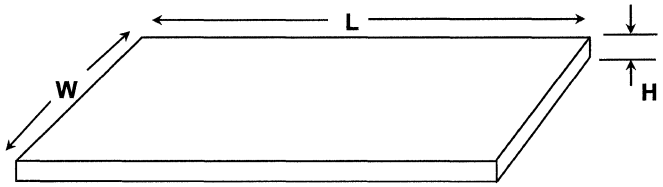

$$\text{STRIP INDUCTANCE} = 0.0002L\left[\ln\left(\frac{2L}{W+H}\right) + 0.2235\left(\frac{W+H}{L}\right) + 0.5\right]\ \mu H$$

EXAMPLE: 1cm of 0.25 mm PC track has an inductance of 9.59 nH
(H = 0.038mm, W = 0.25mm, L = 1cm)

In addition to the resistance (basically a dc spec), a trace (wire or ground plane) will have a frequency dependent impedance component (known as inductance). Inductive impedance increases linearly with frequency. This can become significant at higher frequencies.

The inductance of a trace (wire or ground plane) can be calculated from the information in this figure. The inductive impedance can be calculated from:

$$Z_L = j\omega L$$

where:

$$j = \sqrt{-1}$$

and

$$\omega = \text{radian frequency} = 2\pi \cdot \text{frequency in Hertz}$$

Basic Principles of Inductive Coupling

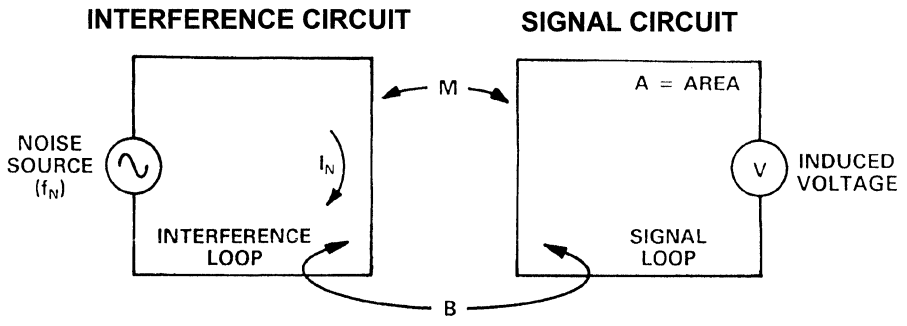

M = MUTUAL INDUCTANCE
B = MAGNETIC REFLUX DENSITY
A = AREA OF SIGNAL LOOP
$\omega_N = 2\pi f_N$ = FREQUENCY OF NOISE SOURCE

V = INDUCED VOLTAGE = $\omega_N M I_N = \omega AB$

The increase in impedance with frequency is only one of the issues with inductance. The other potential problem is the coupling of signal from one circuit to another via mutual inductance.

The amount of coupling will depend on the strength of the interference, the mutual inductance, the area enclosed by the signal loop (which is basically an antenna), and the frequency. Also, the mutual inductance will depend primarily on the physical proximity of the loops, as well as the permeability of the material.

Magnetic Field Lines and Inductive Loop (Right Hand Rule)

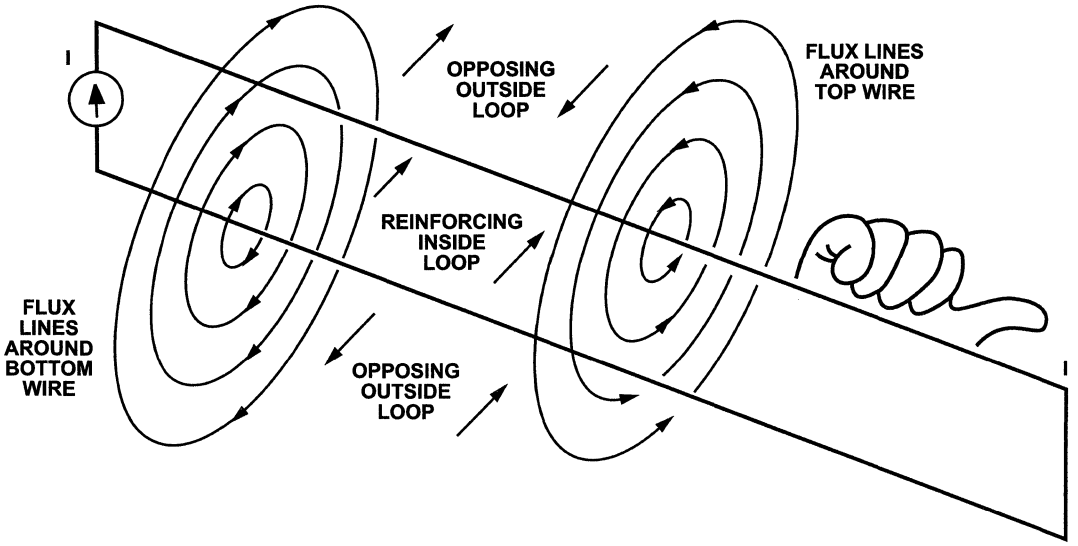

The right hand rule is useful in predicting the direction of the magnetic field lines produced by a current flowing in a conductor.

If you point the fingers of your right hand in the direction of the flux density, the induced signal will flow in the direction that your thumb is pointing.

Proper Signal Routing and Layout Can Reduce Inductive Coupling

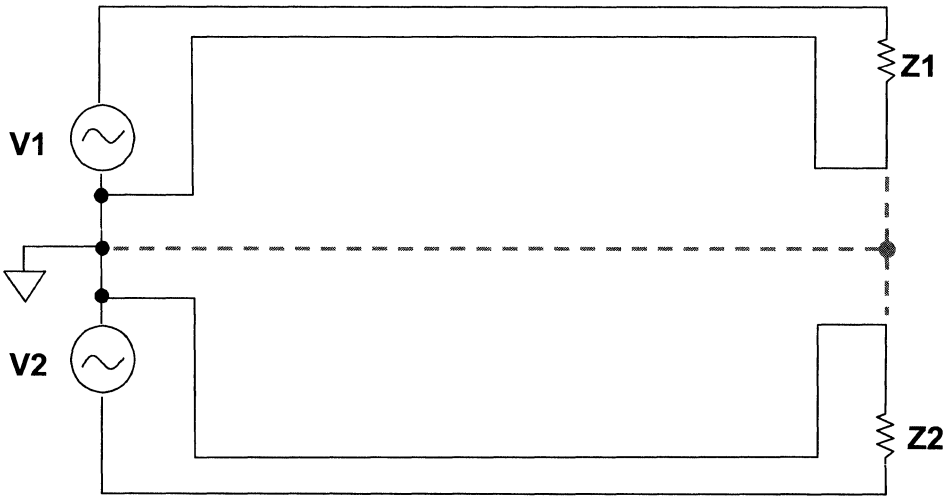

Here is an example of one technique to reduce inductive coupling.

Looking at the above circuit, at first it may seem logical to use a single trace as the return path for the two sources (shown by the dotted lines). However, this will cause the return currents for both signals to flow through the same impedance, which is not desirable. In addition, doing so will maximize the area of the interference loops and increase the mutual inductance by moving the loops close together. This will increase the mutual inductance and the coupling between the circuits.

Routing the traces in the manner shown in this figure minimizes the area enclosed by the loops and separates the return paths, thus separating the circuits and, in turn, minimizing the mutual inductance.

Capacitance of Two Parallel Plates

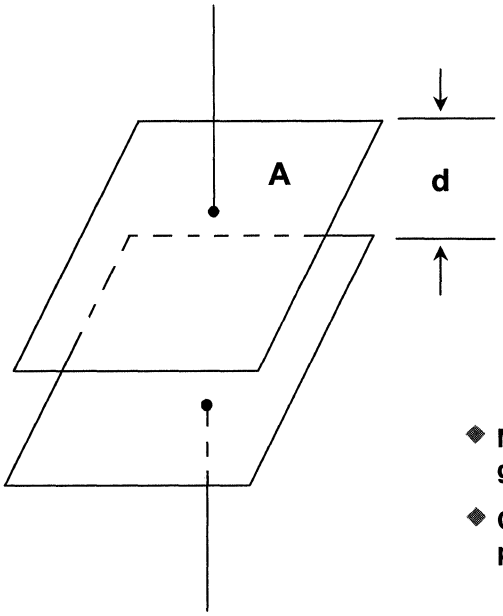

$$C = \frac{0.00885\ E_r\ A}{d}\ pF$$

A = plate area in mm^2

d = plate separation in mm

E_r = dielectric constant relative to air

- Most common PCB type uses 1.5mm glass-fiber epoxy material with E_r = 4.7
- Capacitance of PC track over ground plane is roughly 2.8 pF/cm^2

Another PCB parasitic to be considered is capacitance. While capacitance will not directly affect the impedance of the ground, it can be a major source of coupled interference in a system.

A capacitor consists of two conductors separated by an insulator. This can be as simple as a trace running over a ground, two traces running parallel to each other along a PC board, or two wires running next to each other in a cable bundle.

Capacitive impedance decreases linearly with frequency. This can become significant at higher frequencies.

The capacitance can be calculated from the information in this figure. It is dependant on the dimensions of conductors, the separation of the conductors, and the dielectric constant of the insulator. The capacitive impedance can be calculated from:

$$Z_C = 1/j\omega C$$

where:

$$j = \sqrt{-1}$$

and

$$\omega = \text{radian frequency} = 2\pi \cdot \text{frequency in Hertz}$$

Capacitive Coupling Equivalent Circuit Model

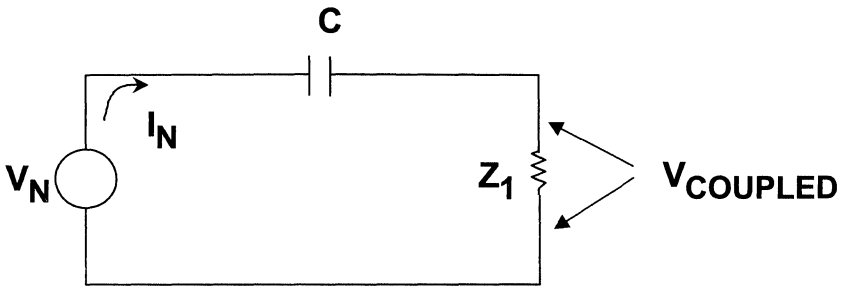

Z_1 = CIRCUIT IMPEDANCE
$Z_2 = 1/j\omega C$

$$V_{COUPLED} = V_N \left[\frac{Z_1}{Z_1 + Z_2} \right]$$

This illustrates the mechanism of capacitive coupling. The capacitor formed by the traces (C) forms a high pass filter with the impedance of the circuit into which the signal is being coupled (Z_1).

The corner frequency of this filter is:

$$f = 1/2\pi Z_1 C$$

If the impedance of the receiver circuit is low, the corner frequency is moved higher in frequency, This results in less coupling at lower frequencies.

This is one reason that most high frequency circuits tend to use low impedances in their design.

A High Speed Converter Directly Interfacing to a Digital Data Bus Showing Path of Injected Noise

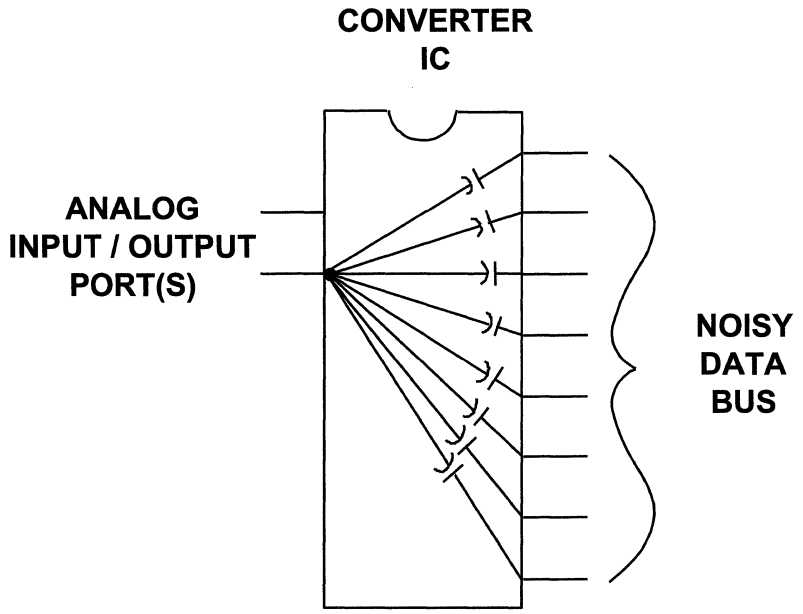

This figure shows how capacitive coupling can degrade the performance of data converters.

In this example, a converter is connected directly to the data bus. This data bus, in turn, is connected to the CPU (DSP, ASIC, FPGA, etc.), memory and other peripherals. The signals on the bus are generally high speed, with fast edges. This implies a great deal of high frequency energy.

Even if the data bus is not active, these signals still appear on the pins, and the bus presents a capacitive load to the converter data pins. The small capacitors in the diagram represent capacitance between the bond wires. This can typically be on the order of 0.2pF or so. As we have seen before, this is a relatively low value, but it becomes increasingly important as frequency increases.

This applies to DACs as well as ADCs, although it will typically be more of an issue with ADCs.

As a side note, this is not the only mechanism of noise injection into a data converter. The signals on the data bus will also get onto the die, even if there is an input latch (for data input to a DAC) or if the output drivers are tri-stated (for data output of an ADC). In either case, the high speed signals will still couple to the die through the stray capacitance.

Partial Solution to the Noisy Data Bus Issue

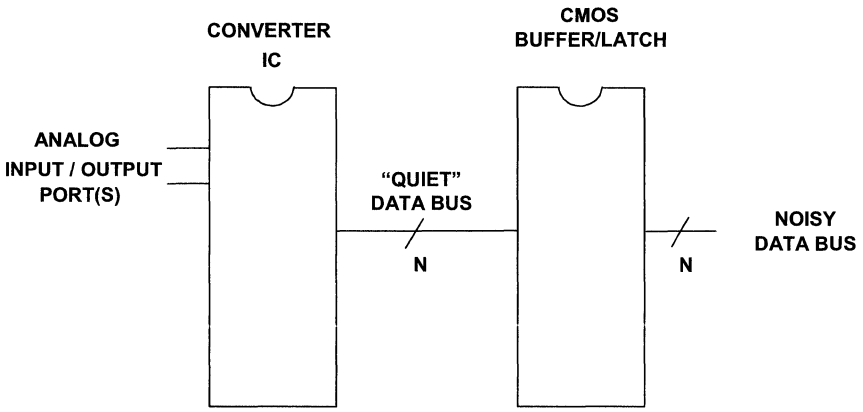

- THE OUTPUT BUFFER/LATCH ACTS AS A FARADAY SHIELD BETWEEN "N" LINES OF A FAST, NOISY DATA BUS AND A HIGH PERFORMANCE ADC.
- THIS MEASURE ADDS COST, BOARD AREA, POWER CONSUMPTION, RELIABILITY REDUCTION, DESIGN COMPLEXITY, AND MOST IMPORTANTLY, *IMPROVED PERFORMANCE!*

One part of the solution to the capacitively coupled noise issue is to isolate the converter from the data bus. This is accomplished by placing a Faraday shield, in the form of a buffer, between the converter and the data bus. This provides a degree of improvement due to the noise immunity provided by the buffer.

The bus between the buffer and the converter is lower noise since the load provided by these ICs requires less transient drive current (due to lower capacitance ~ 10pF).

The noise rejection of the "quiet" bus can be further enhanced by only making this bus active when actually writing to or reading from the converter. This is accomplished by replacing the buffer with a latch and providing some address decoding.

Further improvement could possible be achieved by connecting pull-up/pull-down resistors to the data lines.

Proper Grounding of Mixed-Signal ICs With Low Internal Digital Currents

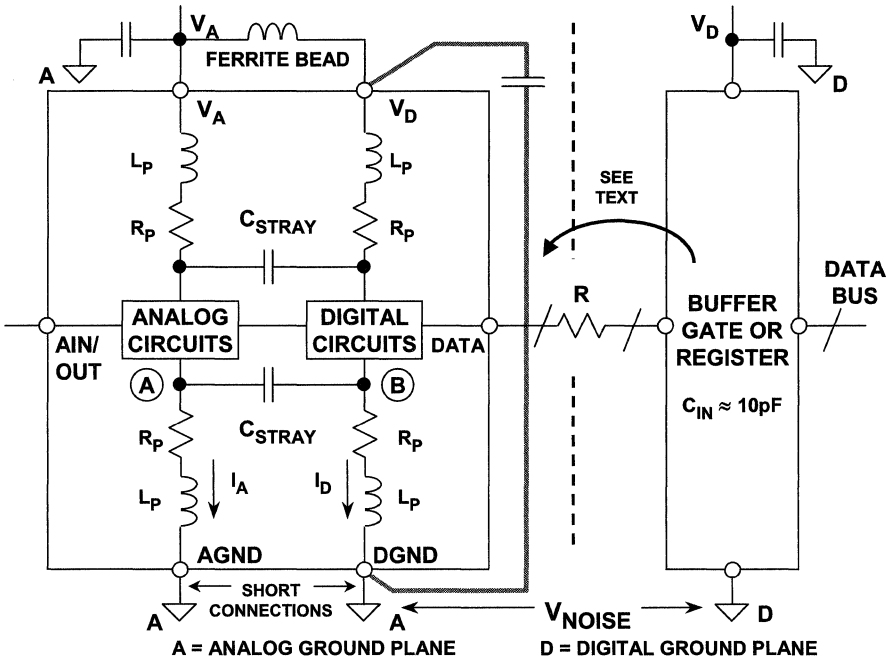

A more complete solution is shown here. In addition to the "quiet" bus solution of the previous slide, attention is paid to how the grounds are handled.

Historically converters have had two separate grounds. These were typically designated as "AGND" (for analog ground) and "DGND" (for digital ground). Conventional wisdom said to connect AGND to the system analog ground and DGND to the system digital ground. The "star" point (the point where the two grounds connect) was at the converter.

This worked if there is only one converter in the system. It also worked better with the relatively low speed logic of years ago. As systems have increased in both complexity and speed, this approach is no longer optimum.

As shown, both grounds of the converter are connected to the system analog ground. This is because it causes less problems for the relatively small amount of digital return current to be returned through the analog ground than it would to connect the converter to the much more noisy digital ground.

It can be seen that, in addition to the separate grounds, there are separate power supply pins. These should be connected to the analog supply as well. Note that even though the power is form the same source, there is a ferrite bead and a decoupling capacitor included in the digital supply pin. The intent here is to prevent the noise generated by the digital section of the converter from feeding back into the analog supply, rather than preventing noise from getting into the converter digital section.

Of course, there are occurrences where the digital power supply is different from the analog supply where this won't apply. The digital supply should still be developed from the analog source.

There is one concern here, however. The return path for the digital bus will be through the digital ground and then through the analog ground. This could result in a fairly large enclosed loop, resulting in a interference problem.

Grounding and Decoupling Points

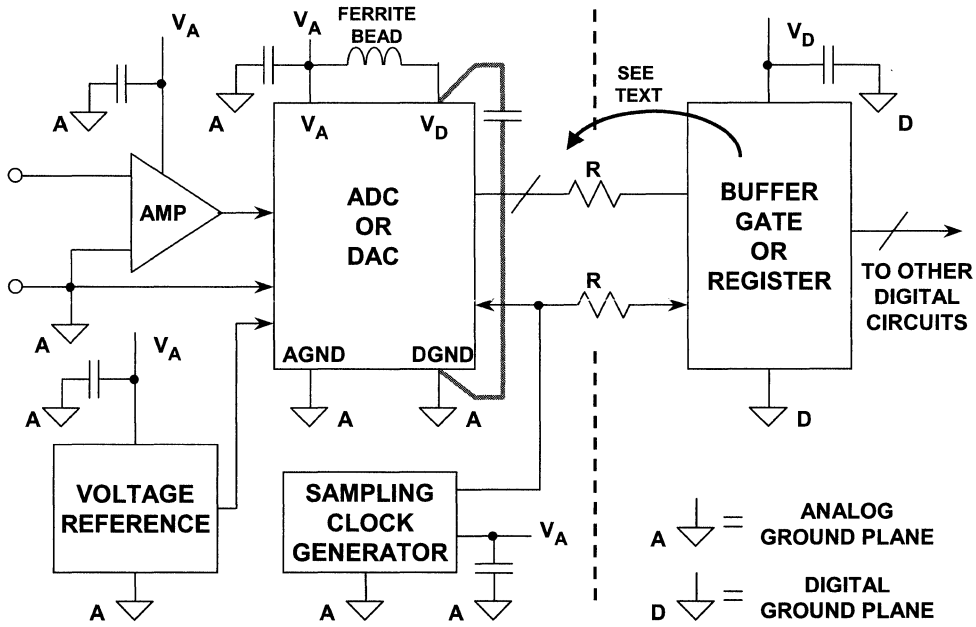

The clock, even though it is typically viewed as a digital signal, should be handled in the analog domain as well. The primary reason for this is to keep the jitter low. This is especially true as frequency goes up. As we have seen previously, high speed, high accuracy converters require a low jitter clock to realize their performance.

If required elsewhere in the system the clock can be distributed the same way as the converter digital input/output.

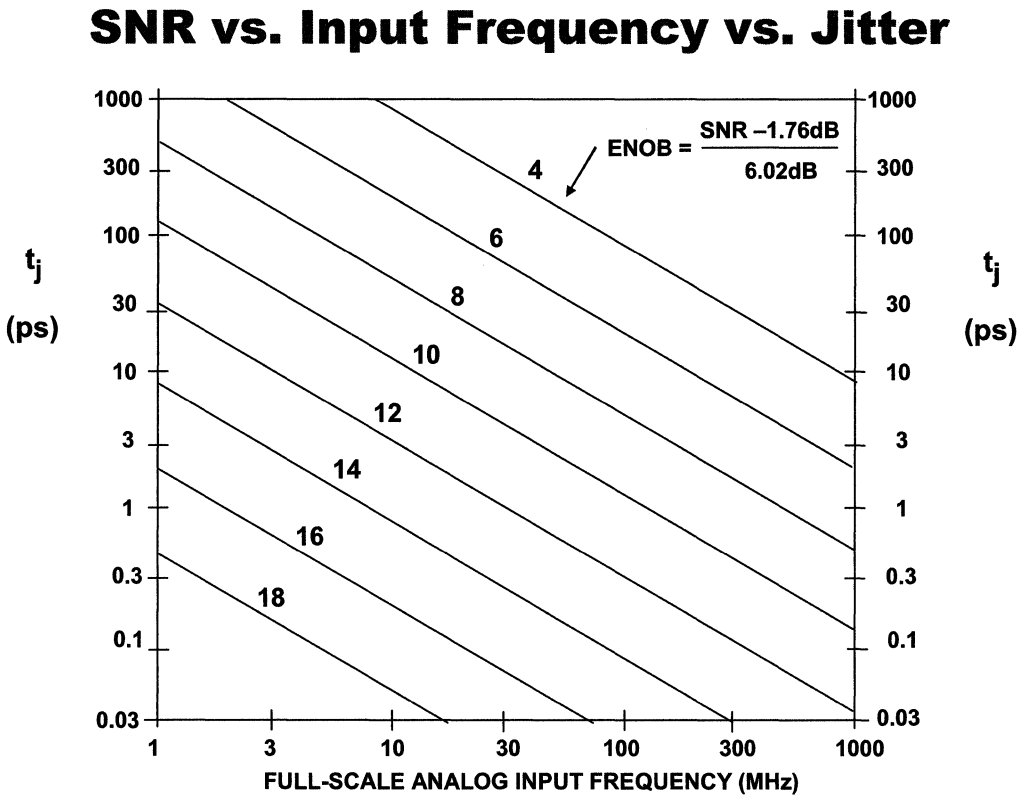

This figure is a reminder of the effect of sampling clock phase jitter on the SNR of an ADC.

This is the maximum signal to noise (SNR) performance possible with a given amount of phase jitter on the clock. Any other signal impairment will only reduce the SNR.

It is obvious that the issue of phase jitter becomes more important as the analog input frequency increases. This is because a given amount of jitter produces a larger error as the slew rate of the input signal increases. SNR is related to the analog input frequency and the clock jitter by the expression:

$$SNR = 20\log[1/2\pi ft_j].$$

AD9510 Clock Generation and Distribution

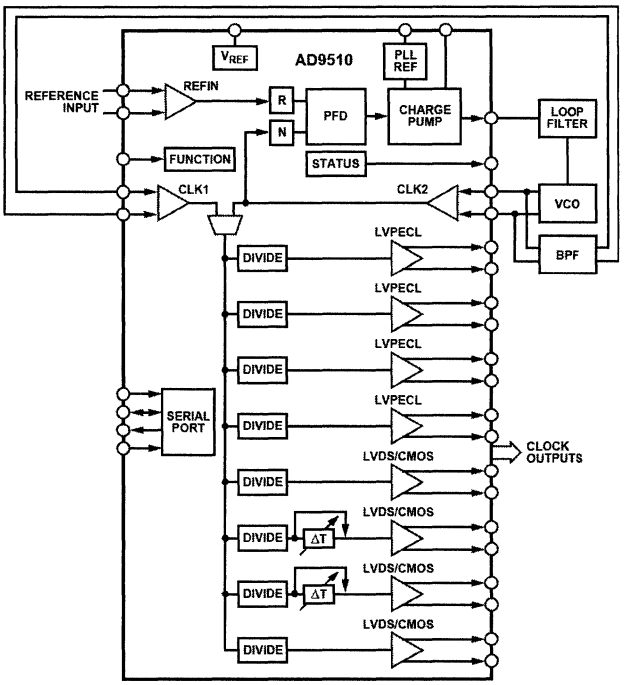

A good way of generating a low phase noise clock is to use a member of the AD851X clock generation and distribution circuits. These circuits are available as generation only, distribution only, and combination parts, as illustrated by the AD9510 in this figure.

The generation section is based on PLL and/or DDS technology .

The PLL section consists of a programmable reference divider (R); a low noise phase frequency detector (PFD); a precision charge pump (CP); and a programmable feedback divider (N).

By connecting an external VCXO or VCO to the CLK2/CLK2B pins, frequencies up to 1.6GHz may be synchronized to the input reference.

Each output has a programmable divider that may be bypassed or set to divide by any integer up to 32. The phase of one clock output relative to another clock output may be varied by means of a divider phase select function that serves as a coarse timing adjustment. Two of the LVDS/CMOS outputs feature programmable delay elements with full-scale ranges up to 10ns of delay. This fine tuning delay block has 5-bit resolution, giving 32 possible delays from which to choose for each full-scale setting.

An external band-pass filter may be used to try to improve the phase noise and spurious characteristics of the PLL output. This option is most appropriate to optimize cost by choosing a less expensive VCO combined with a moderately priced filter. Note that the BPF is shown outside of the VCO-to-N divider path, with the BP filter outputs routed to CLK1.

Grounding Mixed Signal ICs: Single PC Board (Typical Evaluation/Test Board)

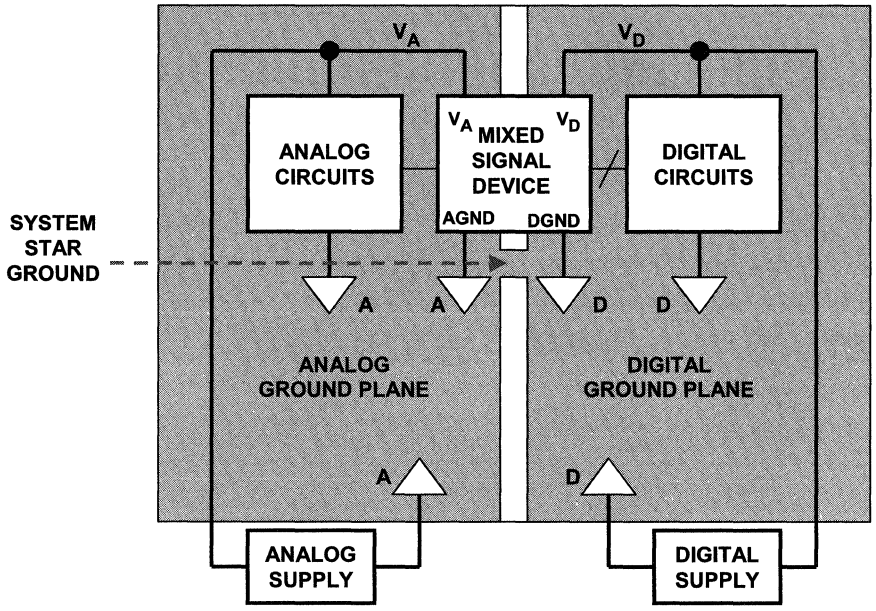

The technique shown in Figure 4.15 assumes a converter with a minimum of digital content. Low digital content implies low digital supply currents.

In some instances the converter will have significantly more digital content. An example of this would be a Sigma-Delta converter, which would include a digital filter as a subsection of the converter architecture. Other examples of high digital content components would be mixed signal parts, codecs, AFEs (analog front ends), and receive signal processors, In the case of higher digital content it may be desirable to keep the digital signals isolated from the analog ground. This is how the evaluation boards for these parts are typically designed. The star point for the system is at the converter.

The problem is that an evaluation board is a fairly small system. What is appropriate for this system may not be appropriate for a larger system containing many data converters. Connecting the ground planes together at each data converter generates multiple "star" points and can introduce ground loops.

Grounding Mixed Signal ICs with High Internal Digital Currents: Multiple PC Boards

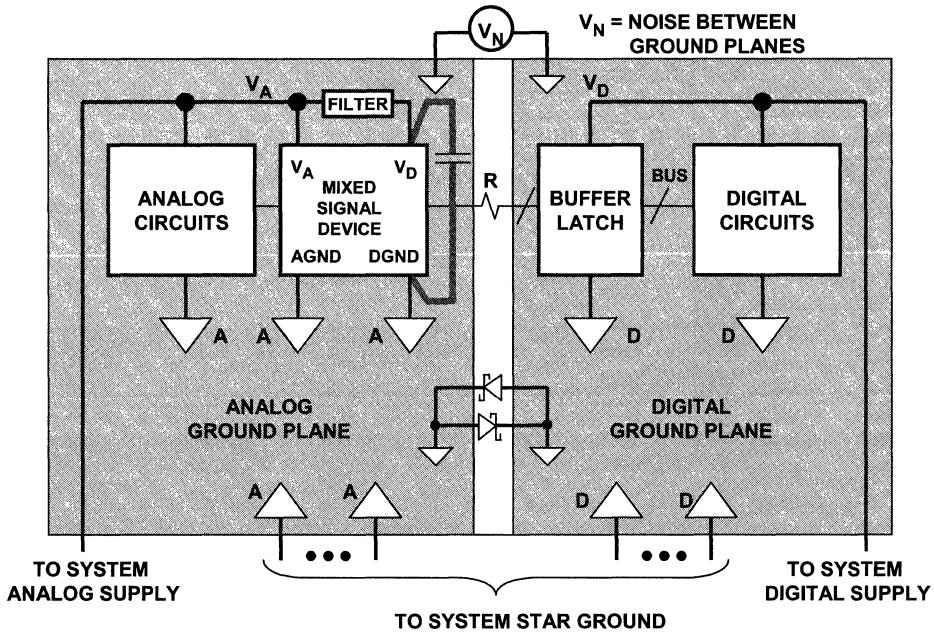

In some cases, even high digital content parts should be connected to the analog ground plane. This is especially true in systems where there are a number of converters, and the star ground point is usually located near the system power supplies.

The caveats noted in Figure 4.19 still apply. The return currents for the digital signals flow through the digital ground and the analog ground and can enclose a large loop.

The back-to-back Schottky diodes are there as protection. They will ensure that the grounds on the board will not drift too far away from each other, should there be a break in the connection between the board and the power supply. A typical specification for the maximum allowable voltage between the analog and digital grounds is ±300mV.

Balanced Sampling Clock or Data Distribution From Analog to Digital Ground Planes

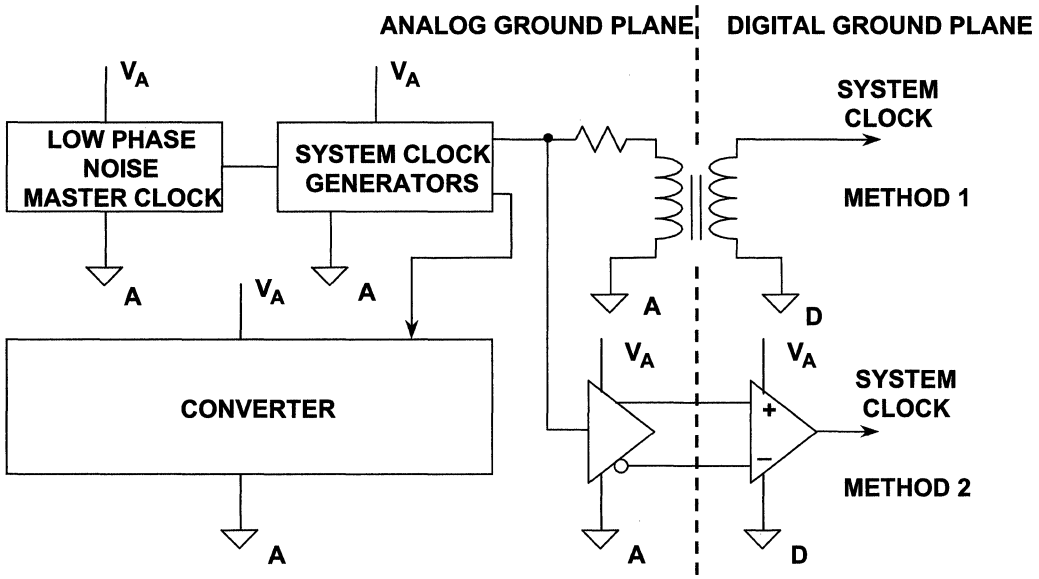

One possible solution to the problem of the return currents enclosing a large loop is to cross the boundary between the analog and digital ground plane with a balanced signal. In this case the return currents will form a very small loop.

This technique works for both clocks and digital signals.

LVDS Driver and Receiver

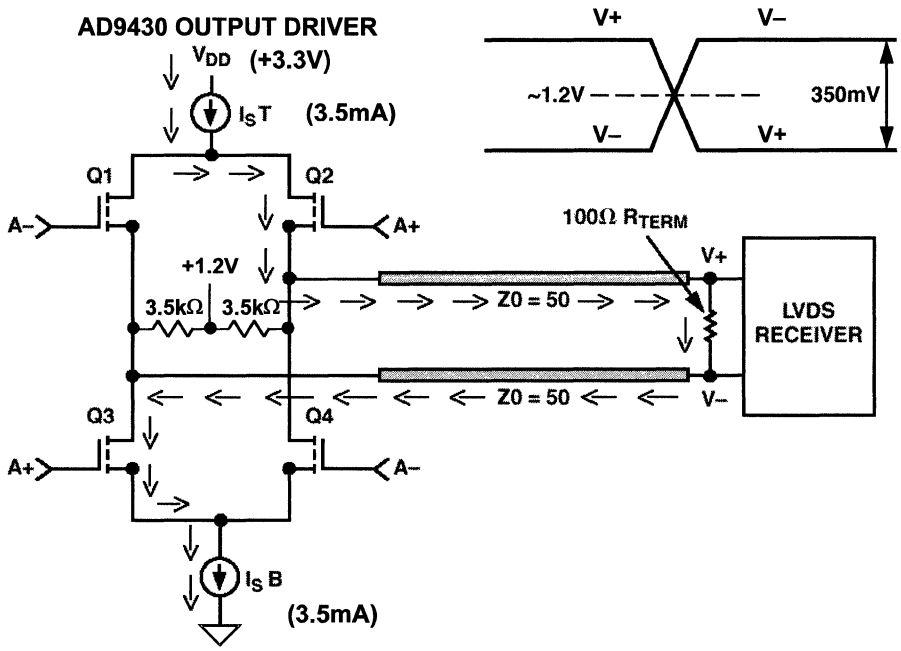

Another alternative for balanced signals might be the use of LVDS (Low Voltage Differential Signal) components.

LVDS is basically a descendant of emitter-coupled logic (ECL). It generates much less noise, since it is not saturating logic like TTL or CMOS. The current remains the constant, thereby minimizing transient switching noise.

*i*Coupler Block Diagram

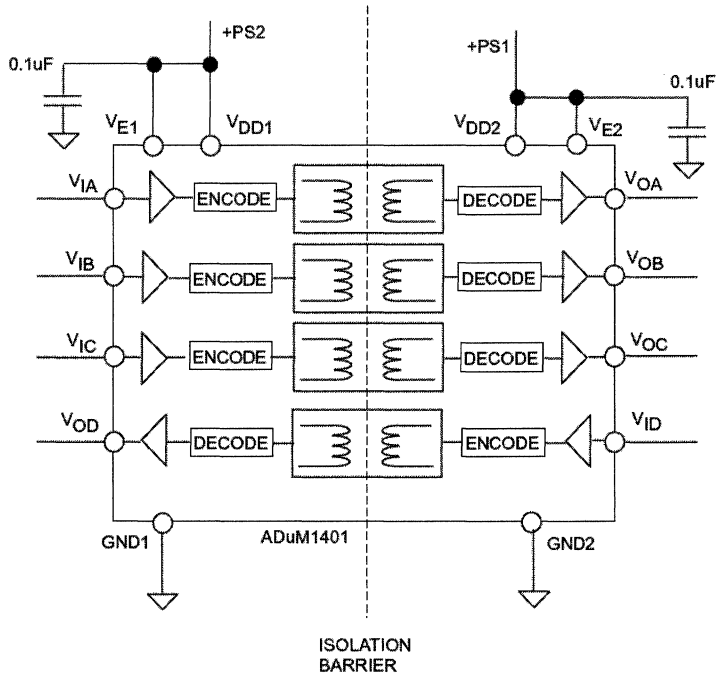

Another possibility is to use an isolation device to separate the grounds. This is demonstrated here by the *i*Coupler.

Designed primarily for galvanic isolation in industrial applications, the *i*Coupler isolator can be used to pass signals between ground planes at different potentials. No currents will circulate through the various grounds, since we are using transformers to isolate the grounds. The transfer rate through the *i*Coupler can be as high as 100MB/s.

Different models of the *i*Couplers have different configurations of forward and return paths.

*i*Coupler Operation

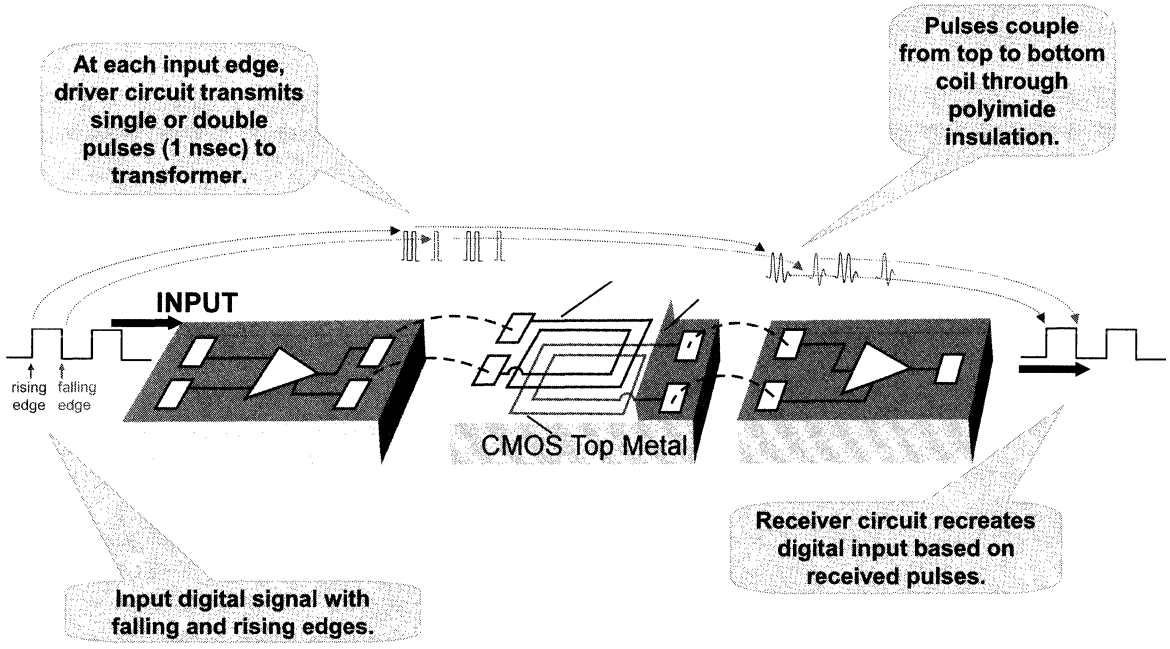

This figure illustrates *i*Coupler operation.

Transformer coupling is used for the isolation. There are separate transmit and receive circuits on either side. The planar transformers use CMOS metal layers, plus a gold layer that is placed on top of the passivation. A high breakdown polyimide layer underneath the gold layer insulates the top transformer coil from the bottom.

*i*Couplers have very good noise immunity since the area enclosed by the loop is very small, making it a very poor antenna, which is a good thing.

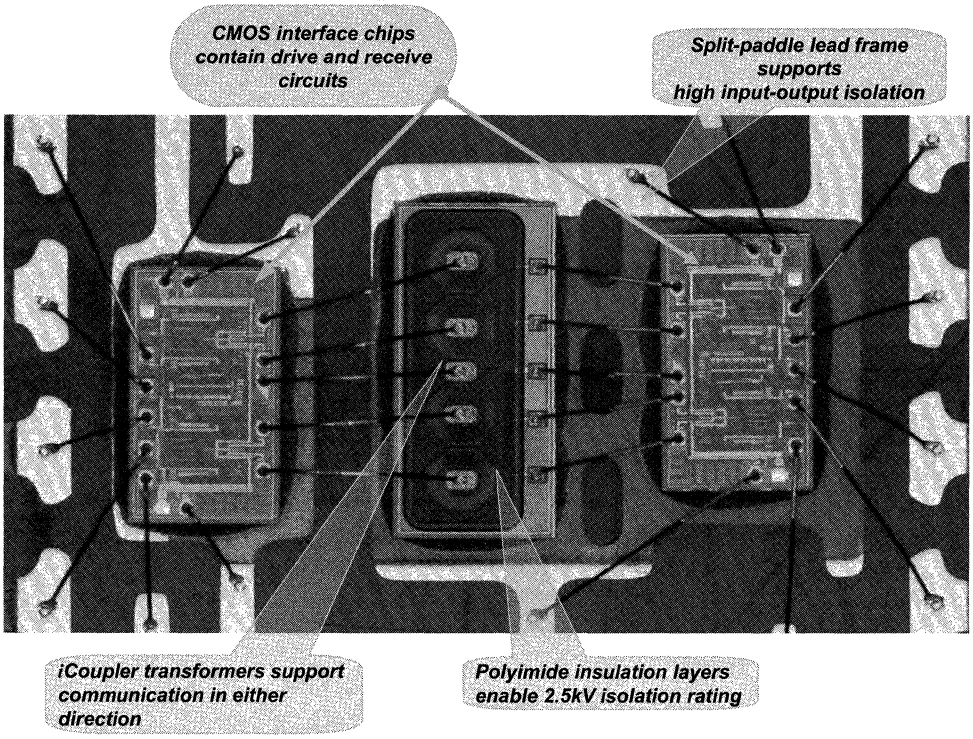

This is a die photo to show the three sections of the *i*Coupler inside the package. The internal spacing is needed where high isolation voltage is required. A high isolation voltage is not required for simply isolating grounds.

A Slit in the Ground Plane Can Reconfigure Current Flow for Better Accuracy

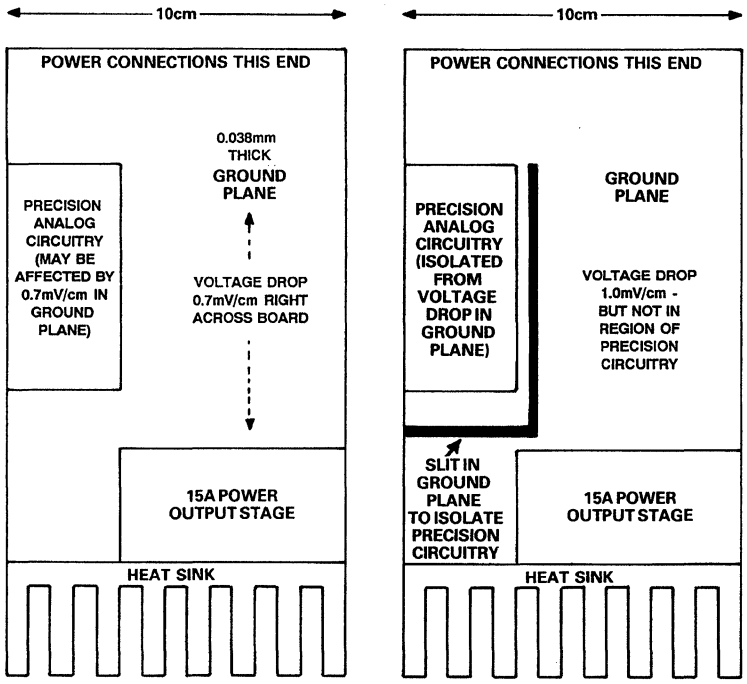

In most cases, a ground plane should be as free from breaks and crossovers as possible. An example of where this is not so is shown in this figure.

As originally designed this board had a solid ground plane. The location of the edge connector and the high current circuit was set by the physical constraints of the system. As originally configured, the relatively large current of the output section flowed through the sensitive precision circuitry. This caused an offset voltage error.

By adding the slit in the ground plane, these currents are forced to flow around the precision circuitry, thereby eliminating the offset error.

Effects of 10pF Stray Capacitance on the Inverting Input on Amplifier Pulse Response

NO CAPACITOR **WITH CAPACITOR**

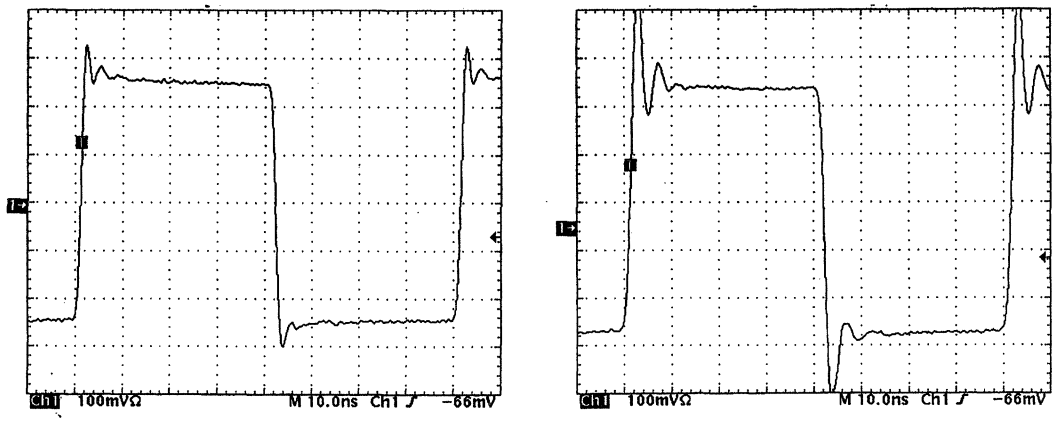

VERTICAL SCALE: 100mV/div
HORIZONTAL SCALE: 10ns/div

Another example of where less ground plane is better is shown here.

Stray capacitance on the inverting input of current feedback amplifiers causes peaking and instability.

These scope captures show the effect on pulse response of a small (10pF) stray capacitance on the inverting input of a typical current feedback amplifier. This capacitance can be caused by running the ground plane too close to the pin on the amplifier.

In most cases, the ground plane should be etched away directly underneath the inverting input pin so that parasitic capacitance is minimized.

Parasitics Plague Dynamic Response of PCB-Based Circuits

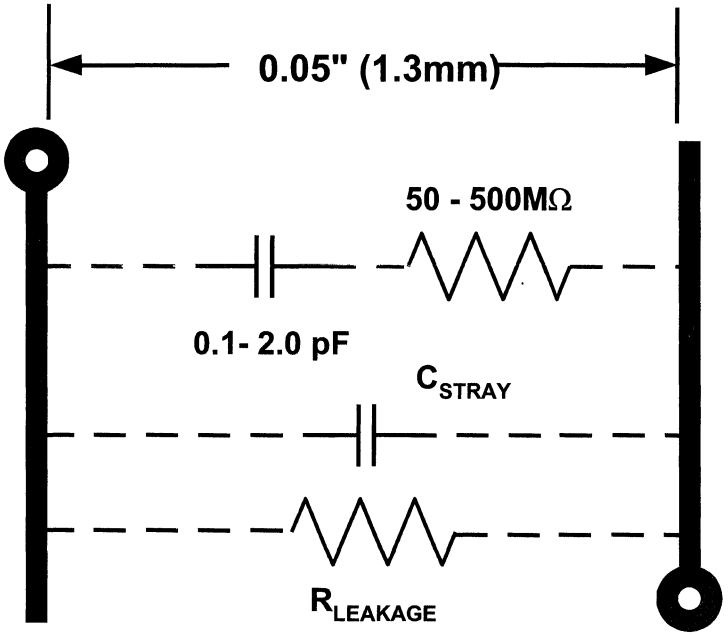

In addition to the parasitics created between a trace and ground, there can also be parasitics between two traces running parallel to each other.

This figure gives an idea of the magnitude of the parasitics between two traces 50 mils apart. Closer spaced traces will have more stray capacitance. In addition, the capacitance will vary with the PC board material and the solder mask used.

Dielectric absorption (DA) represents a more troublesome and still poorly understood circuit-board phenomenon. Like DA in discrete capacitors, DA in a printed-circuit board can be modeled by a series resistor and capacitor connecting two closely spaced nodes. Its effect is inverse with spacing and linear with length.

As shown in this figure, the RC model for this effective capacitance ranges from 0.1 to 2.0pF, with the resistance ranging from 50 to 500MΩ. Values of 0.5pF and 100MΩ are most common. Consequently, circuit-board DA interacts most strongly with high-impedance circuits.

Fortunately, there are solutions to DA. As in the case of capacitor DA, external components can be used to compensate for the effect. More importantly, surface guards that totally isolate sensitive nodes from parasitic coupling often eliminate the problem (note that these guards should be duplicated on both sides of the board, in cases of through-hole components). In addition, low-loss PCB dielectrics are also available.

Schematic and Layout of Current Source with U-shaped Trace on PC Board and Return Through Ground Plane

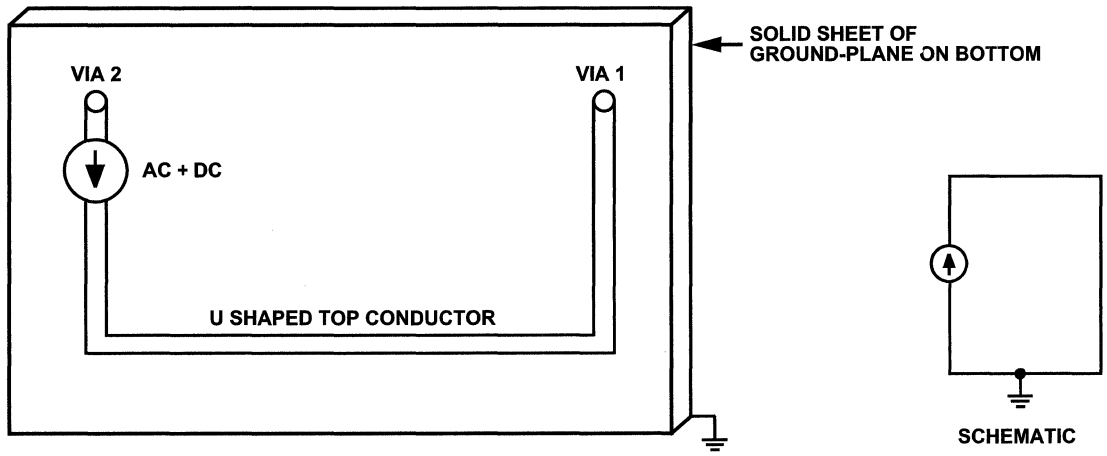

Ground planes are the most effective way to provide a low impedance current return path.

The following several figures investigate the return current flow in a circuit formed by a current source driving a current into a "U"-shaped trace on the top layer of a PC board with a ground plane for the return.

DC Current Flow for Figure 4.29

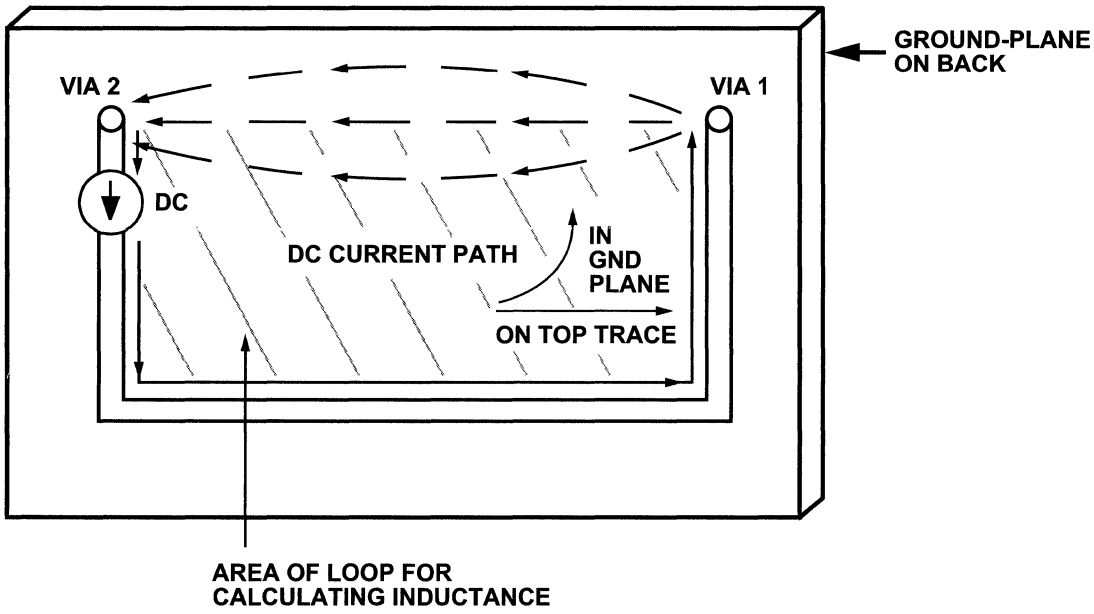

At dc and low frequencies, the current return path will be the path of least resistance, which will be the most direct path between the two vias. There will be some spreading of the return current path due to the finite resistance in the ground plane, but the return current basically takes the most direct path.

AC Current Path Without (left) and With (right) Resistance in the Ground Plane

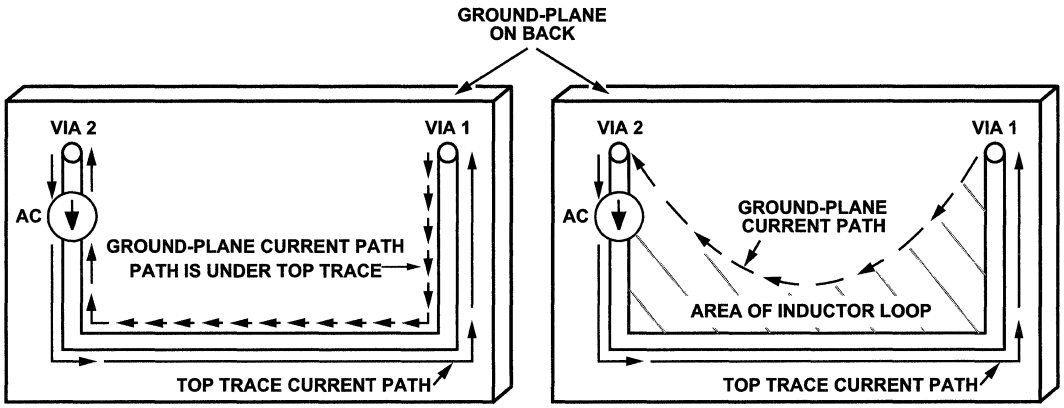

At higher frequencies, however, the return current follows the path of least *impedance* rather than the path of least resistance.

Inductance will be minimized when the enclosed area is reduced as much as possible. This occurs when the enclosed loop is at it's minimum, and that is when the return current flows directly under the forward path. Again, there will be some spreading of the path due to resistance, and there will also be frequency dependence. However, as frequency goes up, the return current will more closely follow the forward path.

The return current flow is nearly completely under the forward trace even at frequencies as low as 1 to 2MHz.

A Ground Plane Break Raises Circuit Inductance and Increases Vulnerability to External Fields

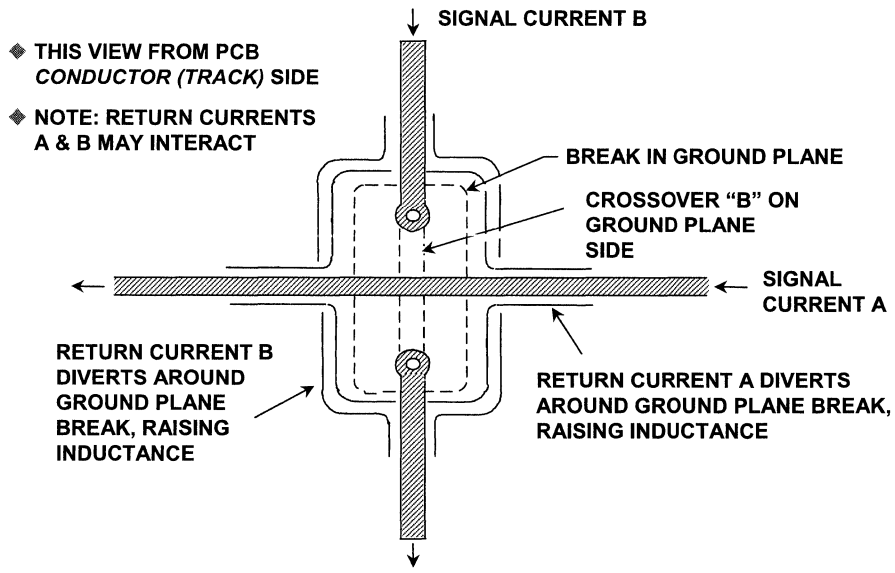

Wherever there is a break in the ground plane beneath a conductor, the ground plane return current must by necessity flow *around* the break. As a result, both the inductance and the vulnerability of the circuit to external fields are increased. This situation is diagrammed in this figure, where conductors A and B must cross one another.

Where such a break is made to allow a crossover of two perpendicular conductors, it would be far better if the second signal were carried across both the first and the ground plane by means of a piece of wire. The ground plane then acts as a shield between the two signal conductors, and the two ground return currents, flowing in opposite sides of the ground plane as a result of skin effects, do not interact.

With a multi-layer board, both the crossover and the continuous ground plane can be accommodated without the need for a wire link. Multi-layer PCBs are expensive and harder to trouble-shoot than more simple double-sided boards, but do offer even better shielding and signal routing. The principles involved remain unchanged but the range of layout options is increased.

The use of double-sided or multi-layer PCBs with at least one continuous ground plane is undoubtedly one of the most successful design approaches for high performance mixed signal circuitry. Often the impedance of such a ground plane is sufficiently low to permit the use of a single ground plane for both analog and digital parts of the system. However, whether or not this is possible does depend upon the resolution and bandwidth required, and the amount of digital noise present in the system.

Skin Depth in a PCB Conductor

- **HF Current flows only in thin surface layers**

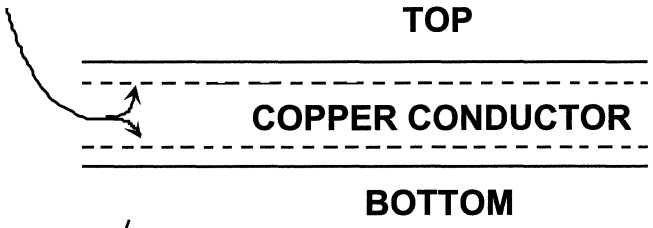

- **Skin Depth ≈ 6.61 / $\sqrt{f}$ cm, f in Hz**
- **Skin Resistance ≈ 2.6×10^{-7} $\sqrt{f}$ ohms per square, f in Hz**
- **Since skin currents flow in both sides of a PC track, the value of skin resistance in PCBs must take account of this**

At high frequencies, we must also consider *skin effect*, where inductive effects cause currents to flow only in the outer surface of conductors. Note that this is in contrast to the earlier discussions of this section on dc resistance of conductors.

The skin effect has the consequence of increasing the resistance of a conductor at high frequencies. Note also that this effect is separate from the increase in impedance due to the effects of the self-inductance of conductors as frequency is increased.

The equation for calculating the skin depth is given as:

$$\delta = \sqrt{\frac{2}{\mu \omega \sigma}} = \sqrt{\frac{\rho}{f \pi \mu}}$$

where:

μ = permeability ($4\pi \times 10^{-7}$ H/m), note: H = henries = $\Omega \cdot s$

δ = skin depth (m)

ρ = resistivity ($\Omega \cdot m$)

ω = radian frequency = $2\pi f$ (Hz)

σ = conductivity (mho/m), note: mho [℧] = siemen [S]

Obviously the skin depth calculation breaks down when the skin depth is greater than the conductor thickness (i.e., at lower frequencies).

Skin Effect with PCB Conductor and Ground Plane

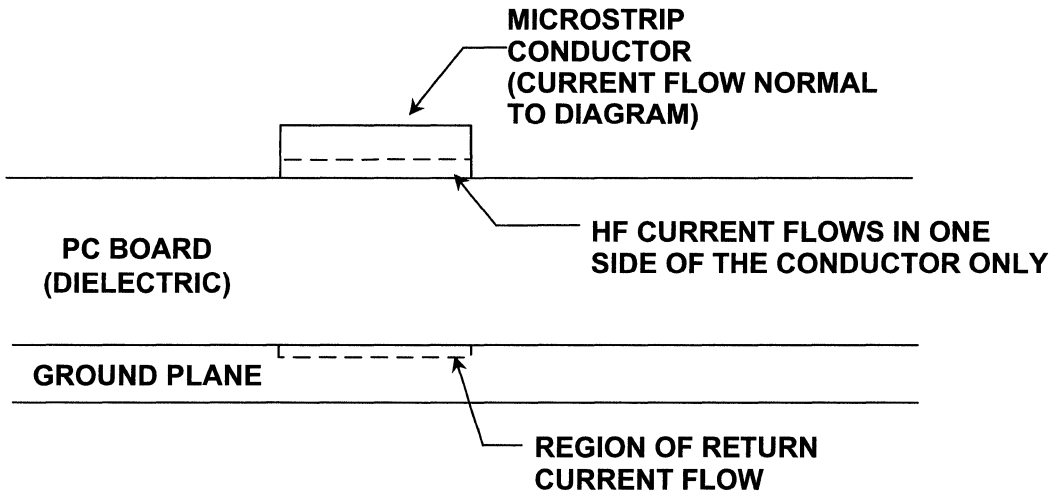

This figure shows the regions of current flow for higher frequencies, as it is reduced by the skin effect.

It is important to remember that the current will flow in both sides of the PCB trace. This is not necessarily the case with microstrips, however. These will be covered next.

In the case where current does flow in both edges of a trace, the resistivity is halved.

For copper, the equation for skin depth and skin resistance can be approximated by:

Skin Depth $\approx 6.61 / \sqrt{f}$ cm, where f is in Hz.

Skin Resistance $\approx 2.6 \times 10^{-7} \sqrt{f}$ ohms per square, where f is in Hz.

A Microstrip Transmission Line

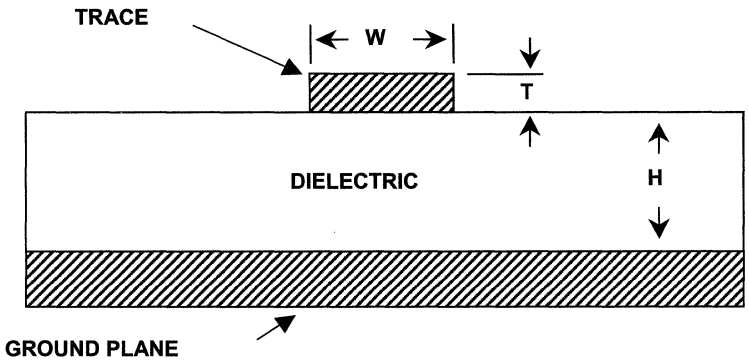

$$Z_o(\Omega) = \frac{87}{\sqrt{\varepsilon_r + 1.41}} \ln\left[\frac{5.98H}{(0.8W + T)}\right]$$

The characteristic impedance of a microstrip transmission line will depend on the width and thickness of the trace and the thickness and dielectric constant of the PCB material.

The characteristic impedance is unimportant at lower frequencies, but is important in maintaining proper termination of signal lines as frequency increases.

For a case of dielectric constant of 4.0 (FR-4), it turns out that when W/H is 2/1, the resulting impedance will be close to 50Ω.

Characteristic Capacitance and Propagation Delay in a Stripline

- In addition to the characteristic impedance (Z_0). the stripline also has a characteristic capacitance, which can be calculated in terms of pF/in:

$$C_o(pF/in) = \frac{0.67(\varepsilon_r + 1.41)}{\ln[5.98H/(0.8W + T)]}$$

- The propagation delay of the stripline is shown in slide 4.35

$$t_{pd}(ns/ft) = 1.017\sqrt{0.475\varepsilon_r + 0.67}$$

- or, in terms of ps:

$$t_{pd}(ps/in) = 85\sqrt{0.475\varepsilon_r + 0.67}$$

- Thus, for an example PCB dielectric constant of 4.0, it can be noted that a microstrip's delay constant is about 1.63 ns/ft, or 136 ps/in. These two additional approximations can be useful in designing the timing of signals across PCB trace runs.

In addition to the characteristic impedance we should be aware of the characteristic capacitance of the trace. This could become an issue in some applications.

Driving large capacitances requires more drive capability from the op amp. Large capacitive loads may also cause the op amp driver to become unstable.

A Symmetric Stripline Transmission Line

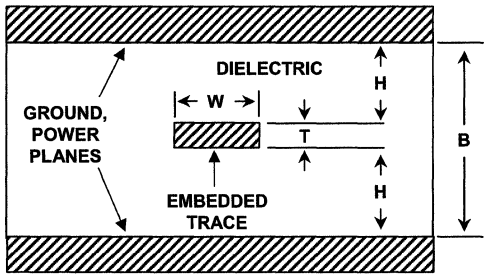

$$Z_o(\Omega) = \frac{60}{\sqrt{\varepsilon_r}} \ln\left[\frac{1.9(B)}{(0.8W + T)}\right]$$

A microstrip transmission line can also be embedded between two ground or power planes (remembering that a power plane is essentially a ground plane for ac signals).

In this case we must calculate the characteristic impedance including both planes.

Characteristic Capacitance and Propagation Delay in a Symmetrical Stripline

- The symmetric stripline also has a characteristic capacitance, which can be calculated in terms of pF/in:

$$C_o(pF/in) = \frac{1.41(\varepsilon_r)}{\ln[3.81H/(0.8W + T)]}$$

- The propagation delay of the symmetric stripline is shown in slide 4.37

$$t_{pd}(ns/ft) = 1.017\sqrt{\varepsilon_r}$$

- or, in terms of ps:

$$t_{pd}(ps/in) = 85\sqrt{\varepsilon_r}$$

- For a PCB dielectric constant of 4.0, it can be noted that the symmetric stripline's delay constant is almost exactly 2 ns/ft, or 170 ps/in.

As in the case of a stripline, in addition to the characteristic impedance of a symmetrical stripline we should be aware of the characteristic capacitance as well. This could become an issue in some applications. Driving larger capacitances requires more drive ability for the output of the driving amplifier. Again, stability may be an issue with the load capacitance affecting the driving op amp.

Microstrip PCB Layout for Two Pairs of LVDS Signals

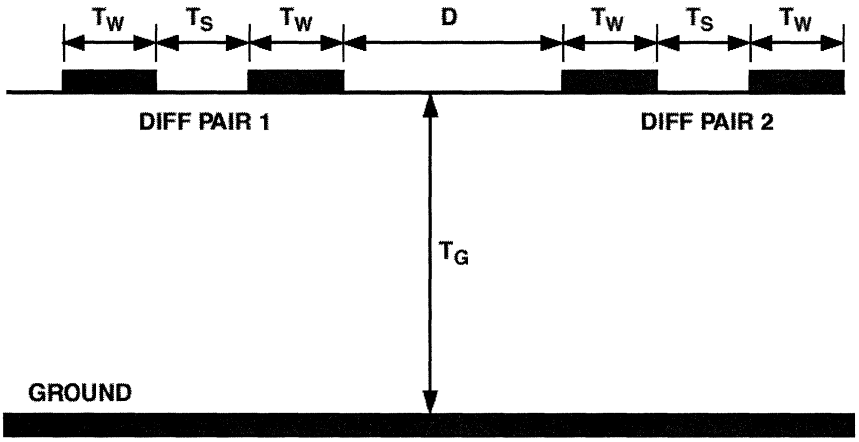

- Keep T_W, T_S, and D constant over the trace length
- Keep T_S ~ < $2T_W$
- Avoid use of vias if possible
- Keep D > $2T_S$
- Avoid 90° bends if possible
- Design T_W and T_G for ~ 50Ω

Some rules of thumb for laying out LVDS microstrip lines are given here.

LVDS outputs for high-performance ADCs should be treated differently than standard LVDS outputs used in digital logic. While standard LVDS can drive 1 to 10 meters in high-speed digital applications (dependent on data rate), it is not recommended to let a high-performance ADC drive that distance. It is recommended to keep the output trace lengths short (< 2 in.), minimizing the opportunity for any noise coupling onto the outputs from the adjacent circuitry, which may get back to the analog inputs. This also controls the power that the output drivers have to develop to drive the line, which keeps the internal noise of the converter down. The differential output traces should be routed close together, maximizing common-mode rejection, with the 100 Ω termination resistor close to the receiver. Users should pay attention to PCB trace lengths to minimize any delay skew.

The impedance can be determined by the information in Figure 4.35.

Decoupling

Power supply decoupling is important in any precision or high speed circuit. The power supply is part of the circuit and should be handled accordingly. The idea is to develop a low noise environment in which the circuit can operate. Improper decoupling can destroy the performance of an otherwise competent design.

Power Supply Rejection Ratio vs. Frequency for the AD8099

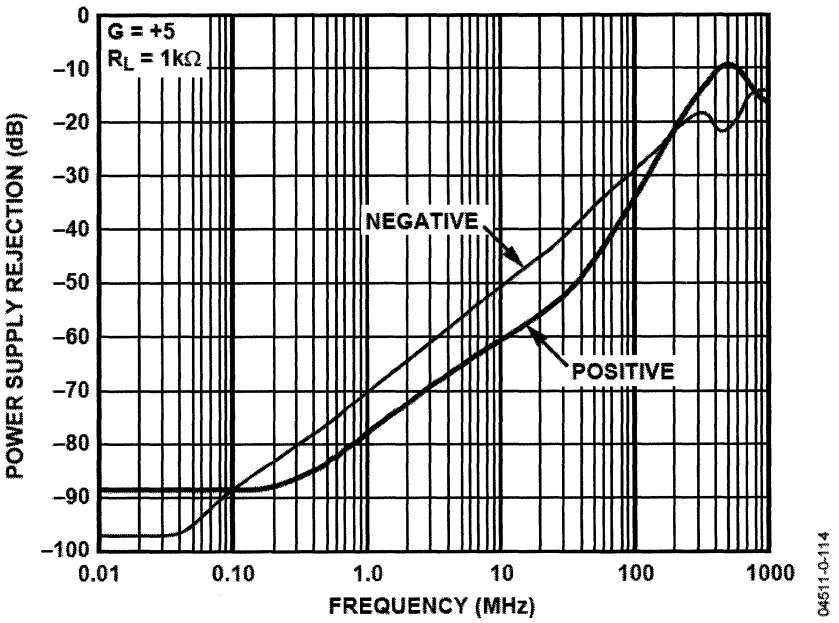

Why is decoupling necessary?

This graph shows how the power supply rejection ratio (PSRR) of an amplifier varies with frequency. The power supply pin is really in series with the output. Therefore, any high frequency energy on the power line will couple to the output directly. So it is necessary to keep this high frequency energy from entering the chip in the first place. This is done by using a small capacitor to short the high frequency signals away from the chip.

Power supply rejection of data converters is typically the same order of magnitude as shown in this figure.

Another aspect to decoupling is the lower frequency interference. Here we use larger electrolytic capacitors as shown in the next figure.

What Is Proper Decoupling?

- **A large electrolytic capacitor (typ. 10 µF – 100 µF) no more than 2 in. away from the chip.**
  - **The purpose of this capacitor is to be a reservoir of charge to supply the instantaneous charge requirements of the circuits locally so the charge need not come through the inductance of the power trace.**
- **A smaller cap (typ. 0.01 µF – 0.1 µF) as physically close to the power pins of the chip as is possible.**
  - **The purpose of this capacitor is to short the high frequency noise away from the chip.**
- **Optionally a small ferrite bead in series with the supply pin.**
  - **Localizes the noise in the system.**
  - **Keeps external high frequency noise from the IC.**
  - **Keeps internally generated noise from propagating to the rest of the system.**

As stated, the electrolytic type large value capacitors are used as local charge reservoirs. This means that the instantaneous current requirements do not have to be met by the power supply, which may be located an appreciable distance away with a considerable amount of inductance in the line.

The smaller value capacitors are used to short the high frequency interference away from the chip. Relevant parameters here are low equivalent series inductance (ESL) and equivalent series resistance (ESR). Quite often multilayer ceramics are excellent choices for these applications.

Ferrites (nonconductive ceramics manufactured from the oxides of nickel, zinc, manganese, or other compounds) are useful for decoupling in power supply filters. At low frequencies (<100kHz), ferrites are inductive; thus they are useful in low-pass LC filters. Above 100kHz, ferrites becomes resistive, an important characteristic in high-frequency filter designs. Ferrite impedance is a function of material, operating frequency range, dc bias current, number of turns, size, shape, and temperature.

The ferrite beads may not always be necessary, but they will add extra high frequency noise isolation and decoupling, which is often desirable. Possible caveats here would be to verify that the beads never saturate, especially when op amps are driving high output currents. When a ferrite saturates it becomes nonlinear and loses its filtering properties.

Note that some ferrites, even before full saturation occurs, can be nonlinear; so if a power stage is required to operate with a low distortion output, this should also be checked in a prototype.

A Nonideal Capacitor Equivalent Circuit Includes Parasitic Elements

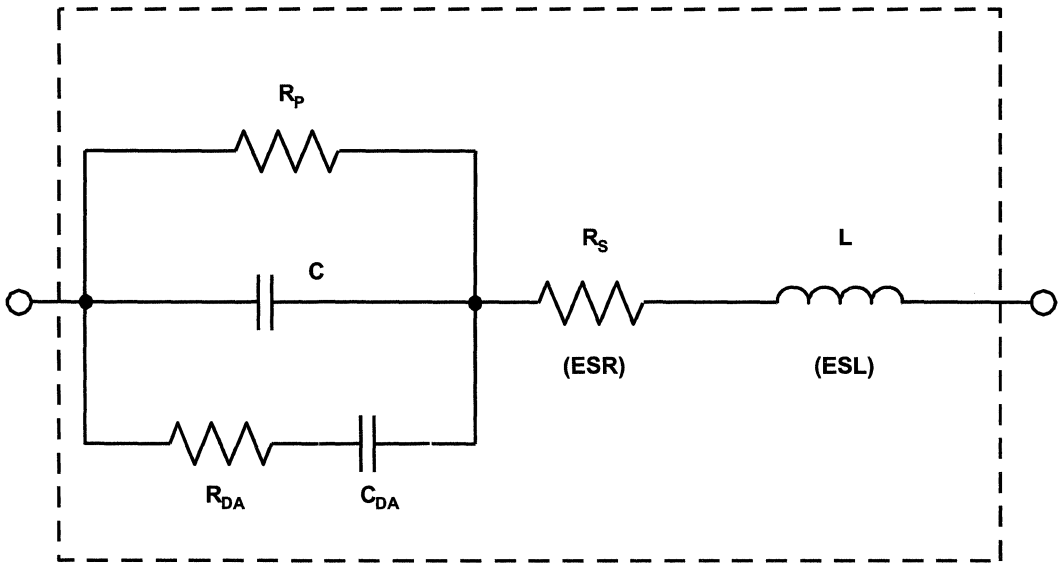

This is a workable model of a nonideal capacitor. The nominal capacitance, C, is shunted by a resistance RP, which represents insulation resistance or leakage. A second resistance, RS—equivalent series resistance, or ESR,—appears in series with the capacitor and represents the resistance of the capacitor leads and plates.

Note that capacitor phenomena aren't that easy to separate. The model is for convenience in explanation. Inductance, L—the equivalent series inductance, or ESL,—models the inductance of the leads and plates. Finally, resistance R_{DA} and capacitance C_{DA} together form a simplified model of a phenomenon known as dielectric absorption, or DA. It can ruin fast and slow circuit dynamic performance. In a real capacitor, R_{DA} and C_{DA} may actually consist of multiple parallel sets.

Impedance Z(Ω) vs. Frequency for 100μF Electrolytic Capacitors

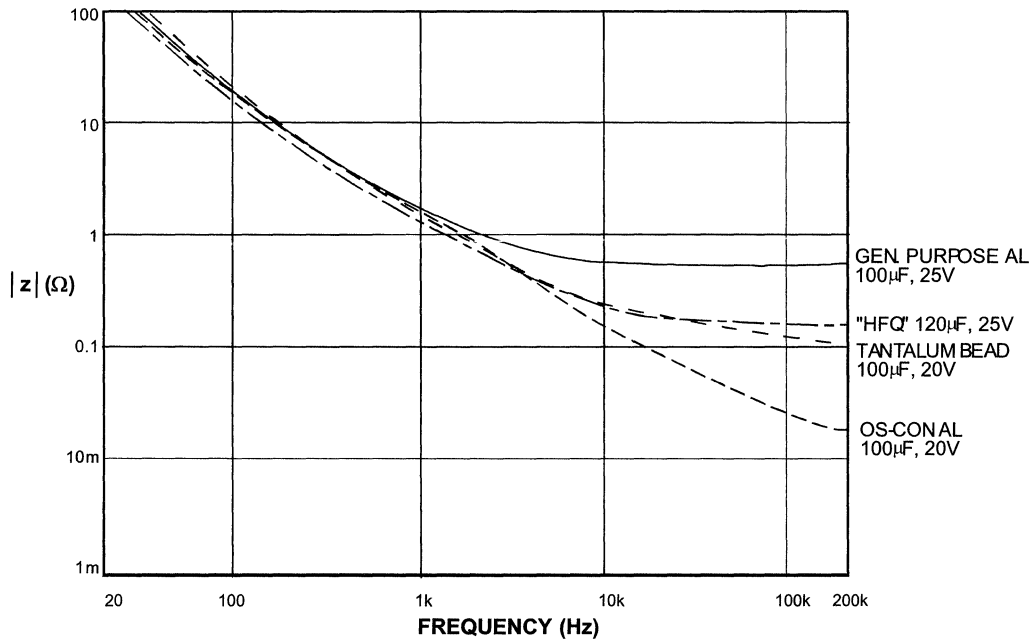

This figure illustrates the high frequency impedance characteristics of several electrolytic capacitor types, using nominal 100μF/20V samples. In these plots, the impedance, |Z|, vs. frequency over the 20Hz to 200kHz range is displayed using a high resolution 4-terminal setup.

Shown in this display are performance samples for a 100μF/25V general purpose aluminum unit, a 120μF/25V HFQ unit, a 100μF/20V tantalum bead type, and a 100μF/20V OS-CON unit (lowest curve). While the HFQ and tantalum samples are close in the specified 100kHz impedance, the general purpose unit is about four times worse. The OS-CON unit is nearly an order of magnitude lower in 100kHz impedance than the tantalum and switching electrolytic types.

100μF/20V Tantalum Capacitor Simplified Model Impedance (Ω) vs. Frequency (Hz)

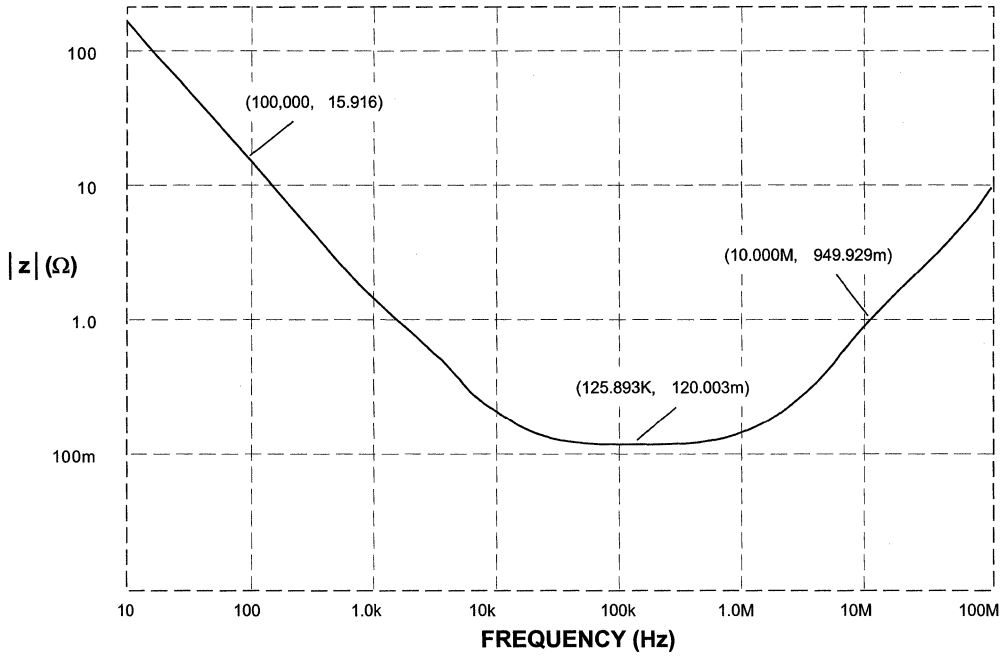

Here we have expanded the frequency range in the previous figure for the 100μF/20V tantalum electrolytic capacitor.

At low frequencies the net impedance is almost purely capacitive, as noted by the 100Hz impedance of 15.9Ω. At the bottom of this "bathtub" curve, the net impedance is determined by ESR, which is shown to be 0.12Ω at 125kHz. Above about 1MHz this capacitor becomes inductive, and impedance is dominated by the effect of ESL.

All electrolytics will display impedance curves which are similar in general shape. The exact values will be different, but the general shape stays the same. The minimum impedance will vary with the ESR, and the inductive region will vary with ESL (which in turn is strongly affected by package style).

Some Capacitor Dielectric Types for Decoupling Applications

Type	Typical DA	Advantages	Disadvantages
NPO Ceramics	< 0.1%	Small case size Inexpensive Many vendors Good stability Low inductance	DA generally low (may not be specified) Low maximum value
Monolithic Ceramic (High K)	>0.2%	Low inductance Wide range of values	Poor stability Poor DA High voltage coefficient
Multilayer Ceramic	< 0.1%	Very low inducatance Small case size	

Ceramic is often the capacitor material of choice above a few MHz, due to its compact size and low loss. But the characteristics of ceramic dielectrics varies widely. Some types are better than others for various applications, especially power supply decoupling. Ceramic dielectric capacitors are available in values up to several μF in the high-K dielectric formulations of X7R and Z5U, at voltage ratings up to 200V. NP0 (also called COG) types use a lower dielectric constant formulation, and have nominally zero TC, plus a low voltage coefficient (unlike the less stable high-K types). The NP0 types are limited in available values to 0.1μF or less, with 0.01μF representing a more practical upper limit.

Multilayer ceramic "chip caps" are increasingly popular for bypassing and filtering at 10MHz or more, because their very low inductance design allows near optimum RF bypassing. In smaller values, ceramic chip caps have an operating frequency range to 1GHz. For these and other capacitors for high frequency applications, a useful value can be ensured by selecting a value which has a self-resonant frequency above the highest frequency of interest.

In general, film type capacitors are not useful in power supply decoupling applications. This is due to their construction. They are generally wound, which increases their inductance.

Electrolytic Capacitor Types for Power Supply Applications

Type	Advantages	Disadvantages
Aluminum	Cost Wide variety of values and working voltages	High leakage
Switching type Aluminum	High frequency performance Broader range of values than Tant.	Slightly more expensive than standard Al.
Tantalum	Size Lower ESR Better high frequency performance	Availability Limited range of values Cost
OS-CON	Much lower ESR Much better high frequncy performance	Cost Availability

The electrolytic family provides an excellent, cost effective low-frequency filter component because of the wide range of values, a high capacitance-to-volume ratio, and a broad range of working voltages. It includes general-purpose aluminum electrolytic types, available in working voltages from below 10V up to about 500V, and in size from 1μF to several thousand μF (with proportional case sizes). All electrolytic capacitors are polarized, and thus cannot withstand more than a volt or so of reverse bias without damage. They have relatively high leakage currents (this can be tens of μA, but is strongly dependent upon specific family design, electrical size, and voltage rating versus applied voltage). However, this is not likely to be a major factor for basic filtering applications.

A subset of aluminum electrolytic capacitors is the switching type, which is designed and specified for handling high pulse currents at frequencies up to several hundred kHz with low losses. This type of capacitor competes directly with the tantalum type in high frequency filtering applications and has the advantage of a much broader range of available values.

Also included in the electrolytic family are tantalum types, which are generally limited to voltages of 100V or less, with capacitance of 500μF or less. In a given size, tantalums exhibit higher capacitance-to-volume ratios than do the general purpose electrolytics, and have both a higher frequency range and lower ESR. They are generally more expensive than standard electrolytics and must be carefully applied with respect to surge and ripple currents.

More recently, high performance aluminum electrolytic capacitors using an organic semiconductor electrolyte have appeared. These OS-CON families of capacitors feature appreciably lower ESR and higher frequency range than do the other electrolytic types, with an additional feature of minimal low-temperature ESR degradation.

High Frequency Supply Filter(s) Require Decoupling via Short Low-Inductance Path (Ground Plane)

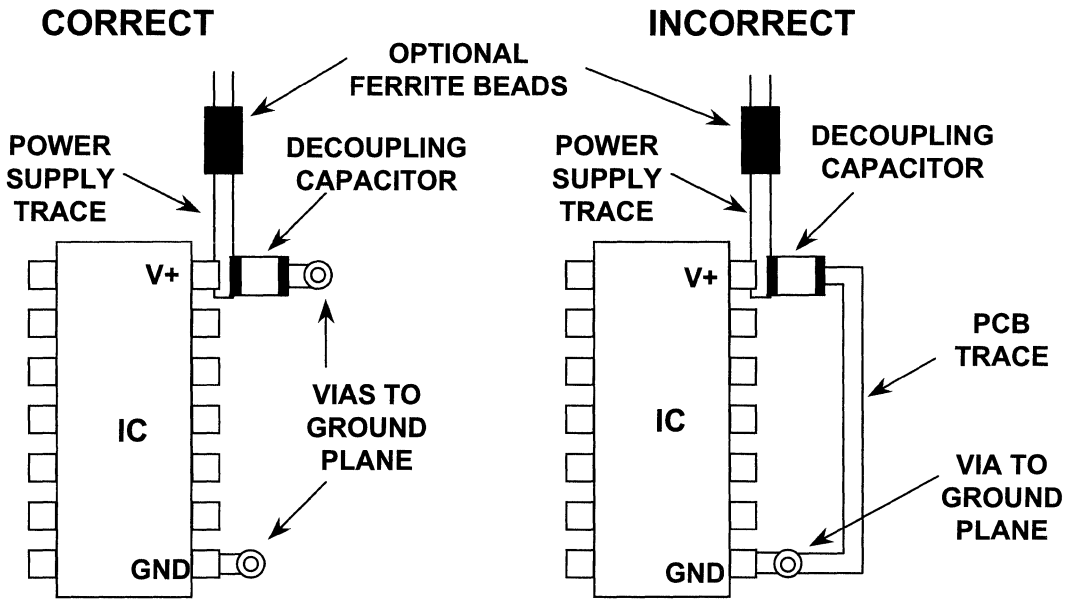

The decoupling capacitor must be as close to the chip as possible. If it is not, the inductance of the connecting trace will have a negative impact on the effectiveness of the decoupling.

In the left figure, the connection to both the power pin and the ground are a short as possible, so this would be the most effective configuration.

In the figure on the right, however, the extra inductance and resistance in the PCB trace will cause a decrease in the effectiveness of the decoupling scheme and may cause interference problems by increasing the enclosed loop.

Resonant Circuit Formed by Power Line Decoupling

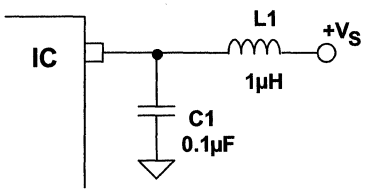

EQUIVALENT DECOUPLED POWER LINE CIRCUIT RESONATES AT:

$$f = \frac{1}{2\pi\sqrt{LC}}$$

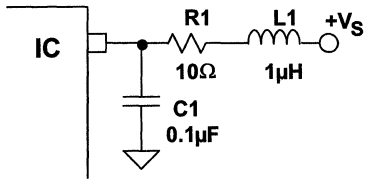

SMALL SERIES RESISTANCE CLOSE TO IC REDUCES Q

An inductor in series or parallel with a capacitor forms a resonant, or "tuned," circuit, whose key feature is that it shows marked change in impedance over a small range of frequency. Just how sharp the effect is depends on the relative Q (quality factor) of the tuned circuit. The Q of a resonant circuit is a measure of its reactance to its resistance.

$$Q = 2\,\pi\,f\,(L/R)$$

If stray inductance and capacitance in a circuit forms a tuned circuit, then that tuned circuit may be excited by signals in the circuit, and ring at its resonant frequency.

While normal trace inductance and typical decoupling capacitances of 0.01µF to 0.1µF will resonate well above a few MHz, an example 0.1µF capacitor and 1µH of inductance resonates at 500kHz. Left unchecked, this could present a resonance problem, as shown in the left case. Should an undesired power line resonance be present, the effect may be minimized by lowering the Q of the inductance. This is most easily done by inserting a small resistance (~10Ω) in the power line close to the IC, as shown in the right case.

The resistance should be kept as low as possible to minimize the IR drop across the resistor. The resistor should be as large as needed, but no larger. An alternative to a resistor is a small ferrite bead which looks primarily resistive at the resonant frequency.

Effects of Decoupling on Performance of the AD8000 Op Amp

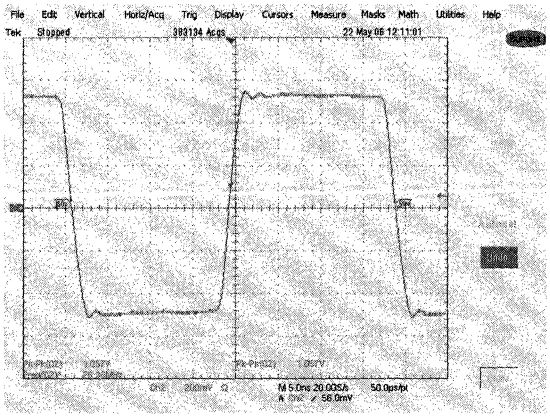

Proper decoupling

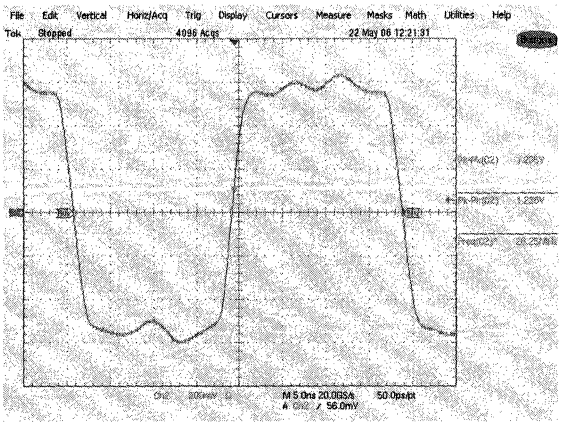

No decoupling

What is the effect of proper and improper decoupling?

Here is an example of what could happen to the response of an op amp with no decoupling.

Both of the oscilloscope graphs were taken on the same board. The difference is that on the right the decoupling caps were removed. Other than that everything remained the same. In this case the device was an AD8000, a 1.5GHz high speed current feedback op amp, but the effect will occur in most any high speed circuit.

AD8000 Power Supply Rejection Ratio (PSRR)

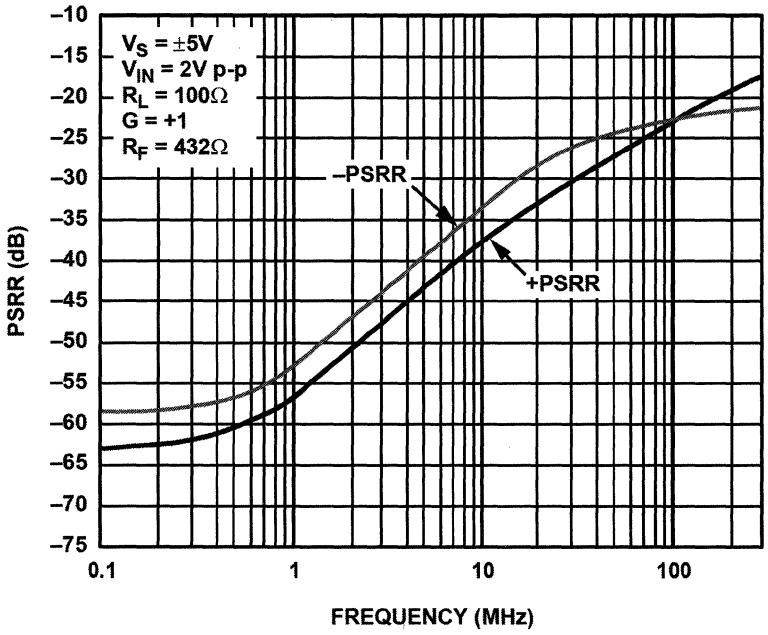

This figure shows the PSRR of the AD8000 as a function of frequency.

Note that the PSRR falls to a relatively low value in its working frequency range. This means that signals on the power line will propagate easily to the output.

As a side note, decoupling keeps signals generated internally in the op amp from propagating through the power lines to other circuits.

AD8000 Positive PSRR Test Circuit

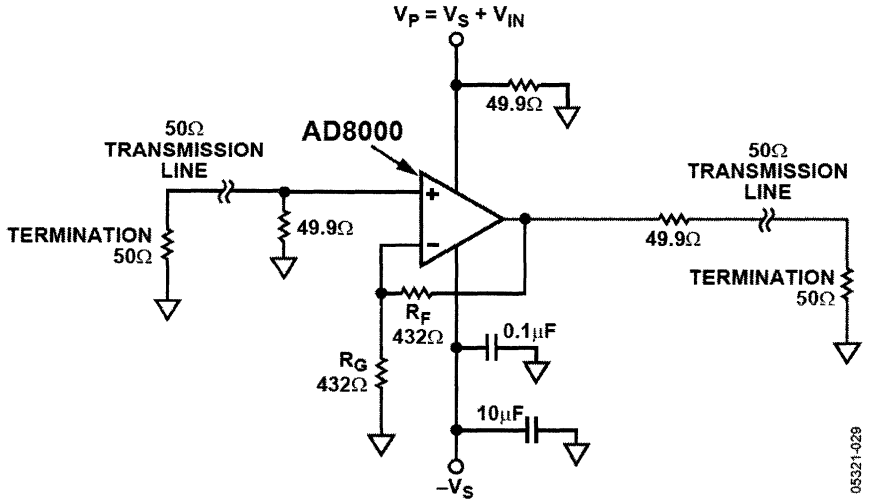

This figure shows a test circuit with which to measure PSRR. In this instance, we are measuring the positive supply rejection. Note that the negative rail is fully decoupled. Also note that we are operating the input and output into matched lines. The positive rail is driven with the signal generator, V_{IN}, which is terminated in its characteristic impedance (50Ω).

The negative rail test circuit is similar. The positive rail is decoupled, and the negative rail is driven with the signal generator.

SNR Plot for the AD9445 Evaluation Board with Proper Decoupling

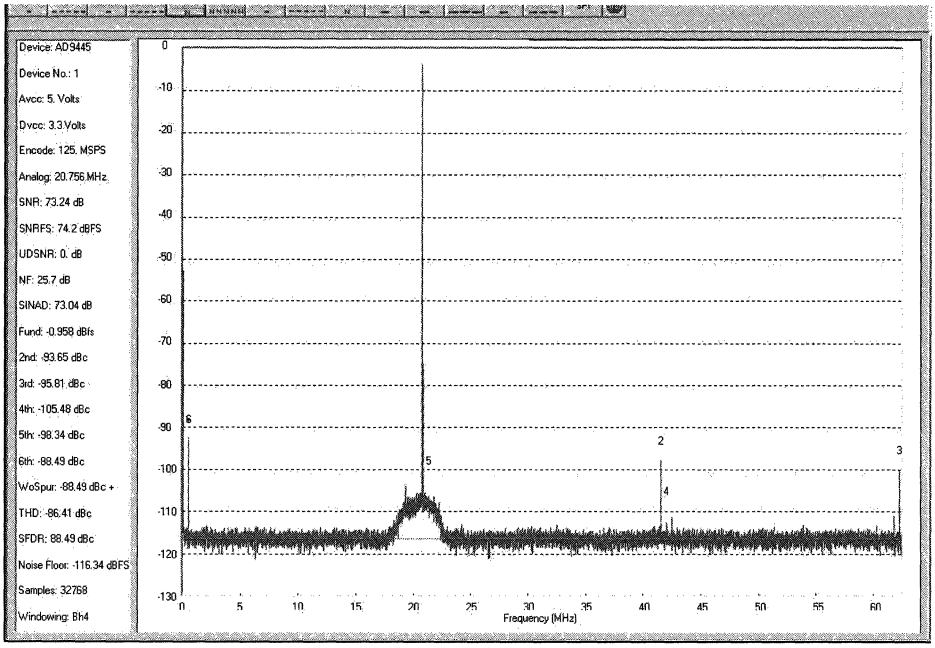

We will now examine the effect of proper and improper decoupling on a high performance data converter, the AD9445 14-bit, 105/125MSPS ADC.

While a converter will typically not have a PSRR specification, proper decoupling is still very important. Here is the FFT output of a properly designed circuit. In this case, we are using the evaluation board for the AD9445. Note the clean spectrum.

AD9445 Pinout Diagram

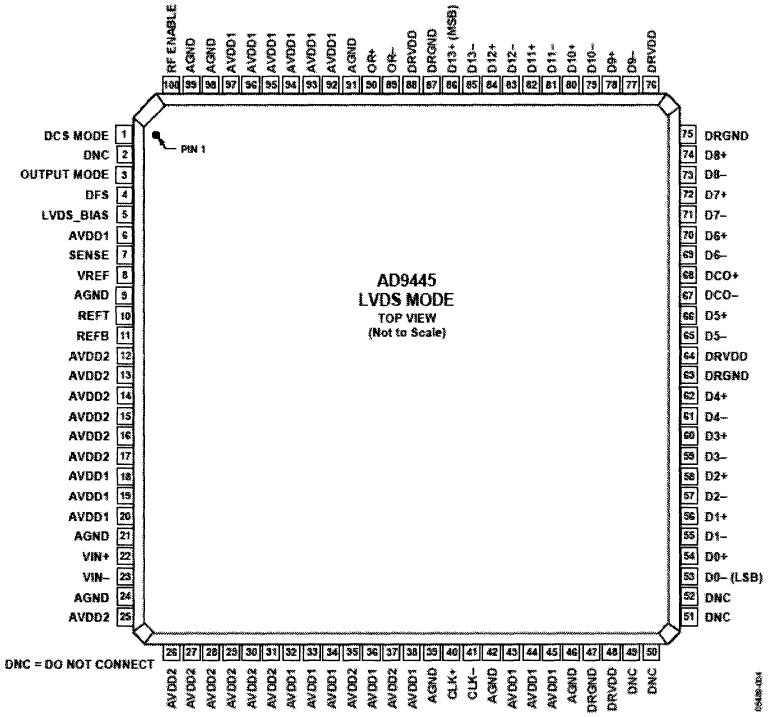

This is the pinout of the AD9445. Note that there are multiple power and ground pins. This is done to lower the impedance of the power supply (pins in parallel). There are 33 analog power pins. 18 pins are connected to AVDD1 (which is +3.3V ± 5 %) and 15 pins are connected to AVDD2 (which is +5V ± 5%). There are four DVDD (which is +5V ± 5 %) pins.

On the evaluation board used in this experiment, each pin has a decoupling cap. In addition, there are several 10μF electrolytic capacitors as well.

SNR Plot for an AD9445 Evaluation Board with Caps Removed from the Analog Supply

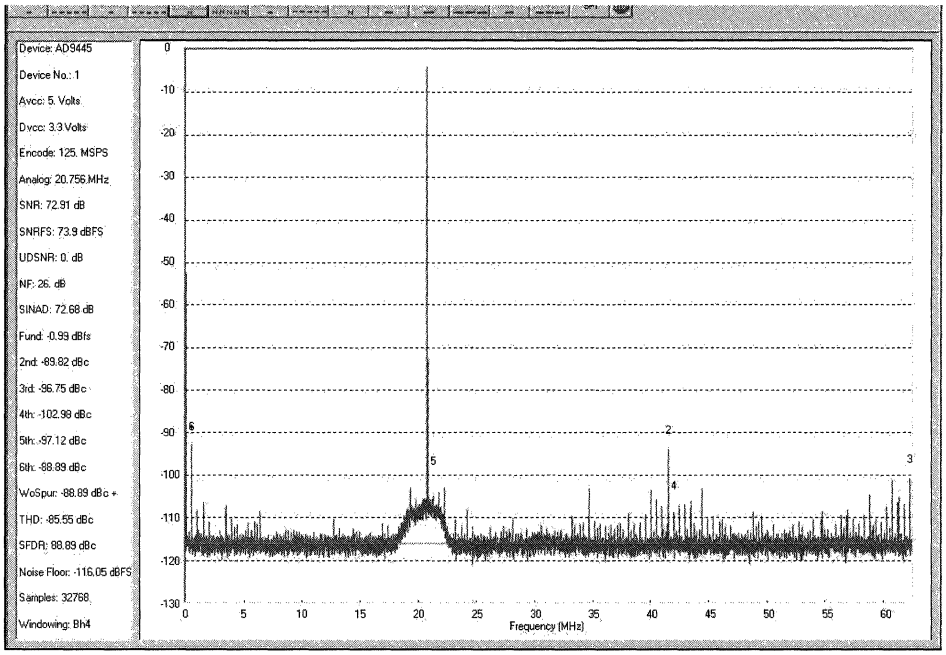

Here is the spectrum with the decoupling caps removed from the analog supply. Note the increase in high frequency spurious signals, as well as some intermodulation products (lower frequency components). The SNR of the signal has obviously degenerated.

The only difference between this figure and the last is removal of the decoupling capacitors. Again we used the AD9445 evaluation board.

SNR Plot for an AD9445 Evaluation Board with Caps Removed from the Digital Supply

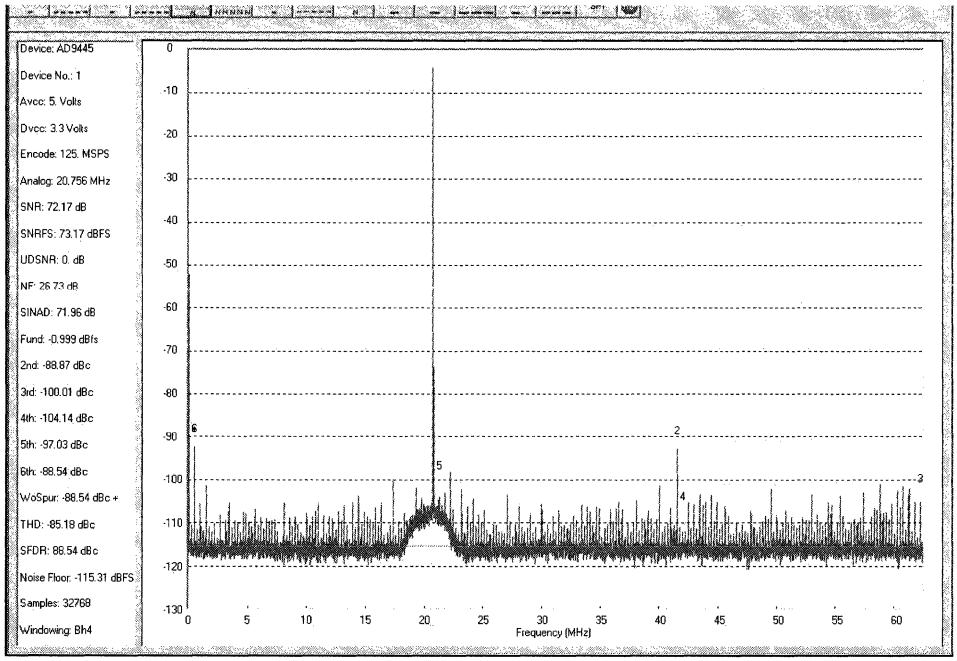

Here is the result of removing the decoupling caps from the digital supply. Again note the increase in spurs. Also note the frequency distribution of the spurs. Not only do these spurs occur at high frequencies, but across the spectrum.

This experiment was run with the LVDS version of the converter. We can assume that the CMOS version would be worse because LVDS is less noisy than saturating CMOS logic.

More information on PC board design techniques can be found in the following two references:

Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN: 0916550273 Chapter 9. Also available as *Data Conversion Handbook*, Elsevier-Newnes, 2005, ISBN: 0750678410, Chapter 9.

Walt Jung, *Op Amp Applications*, Analog Devices, 2002, ISBN: 0-916550-26-5, Chapter 7. Also available as *Op Amp Applications Handbook*, Elsevier-Newnes, 2004, ISBN: 0-7506-7844-5, Chapter 7.

Design Tools

www.analog.com/designcenter

Analog Devices' Online Design Center

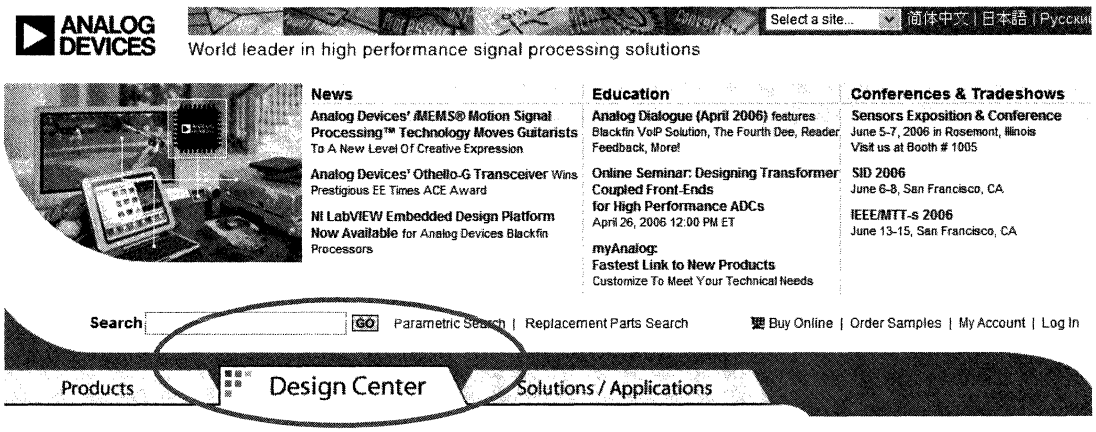

Design Tools ▶ **Learn from the Experts ▶** **Technical Library ▶**

Models ▶ **Evaluation Boards and Kits ▶**

Parametric Selection Tables ▶

Embedded Processing and DSP Development Tools ▶

By Product Category ▶

The Design Center tab of the Analog Device Home page is where the design tools reside.

Design Tools Overview

- **ADC Simulation: ADIsimADC**
- **DDS Simulation: ADIsimDDS**
- **PLL Simulation and Design: ADIsimPLL**
- **Clock Circuit Simulation: ADIsimCLK**
- **Amplifier Evaluation Tool: ADIsimOpAmp**
- **Design Assistants:**
  - **Analog Active Filter Design Assistant**
  - **ADC Aliasing Suppression**
  - **DAC Harmonic and Image Locator**
  - **Error Budget Calculation**
  - **Differential Amplifier Gain and Level Calculator**
  - **Analog Bridge Design Assistant**
  - **Analog Photodiode Preamp Design Assistant**
  - **Voltage Reference Wizard**
- **Register Configuration Assistants**
  - **Sigma-Delta ADCs, TxDACs, DDS, etc.**
- **Useful Calculators**
  - **SNR/THD/SINAD/ENOB, dBm/dBV/dBu, Die Temperature, etc.**

This is a partial list of the design tools available from Analog Devices.

More tools are being added all the time.

We have discussed the ADIsimADC, ADIsimDAC, ADIsimPLL and ADIsimCLK programs previously in this book.

ADIsimOpAmp

- **An On-Line tool to help with the selection, evaluation and troubleshooting of voltage feedback operational amplifiers (Op-Amps). Allows for two modes of Evaluation:**
  - **APET Mode Uses National Instruments LabVIEW® along with typical parametric data to mathematically model the general behavior of a selected amplifier. It allows a user to choose an amplifier, quickly configure a circuit, apply a signal, and evaluate its general performance.**
  - **SPICE Mode Uses the MultiSIM9® SPICE simulation engine allowing the user to perform additional testing in the SPICE environment.**
- **The tool is useful for quickly selecting and checking an amplifier's parametric performance such as Gain Bandwidth, Slew Rate, Input/Output Range, Differential Voltages, Gain Error, Load Current, Possible Stability Issues and dc Errors. The APET mode is limited to first-order approximations and additional evaluation should be performed using SPICE simulation and hardware testing.**

The ADIsimOpAmp is an on-line aid to help engineers find and evaluate an amplifier for their application. It can quickly search through hundreds of amplifiers to find the right one for the application. ADIsimOpAmp can be used in either the Amplifier Parametric Evaluation mode, which uses the National Instrument's LabVIEW environment, or the SPICE simulation mode.

The LabVIEW environment uses a mathematical model to evaluate the circuit. The spice mode uses a more traditional spice simulation engine.

How ADIsimOpAmp Is Used

- **Select the Circuit**
- **Enter the Circuit Component Values**
- **Select and Enter the Input Signal Parameters**
- **Select the Amplifier to be Evaluated**
  - **Parametric Search**
  - **Amplifier Wizard**
  - **Suggest Amplifier (Reverse Search-See Below)**
- **Analyze the Amplifier's Response**
  - **Run Model**
  - **View Results**
- **Suggest Amplifier uses the entered circuit requirements to perform parametric calculations on all amplifiers located in the data base. Once the calculations are complete a list of parts from best to least is generated. If an amplifier that meets all requirements can't be found, the search will suggest components that are close.**

Choose an Amplifier to be evaluated from the pull down, labelled "Select Amplifier," located at the top of the tool. Additional amplifiers can be added to the list by using one of the options below:

Parametric Search - If you know the required op-amp parameters the Parametric Search Engine can be used to find and suggest amplifiers for evaluation. To perform a search enter the values for the desired parameters in the input blocks provided and press search. The tool will display results that best meet the input criteria. Additional search parameters can be added by checking the desired box under "Add Searchable Parameters." If desired, amplifiers can be selected for further evaluation by selecting the "Add Part(s) to Amplifier Parametric Evaluation Tool," selecting the checkbox next to the desired part(s), then clicking the "Add to Tool" button at the bottom of the page.

Amplifier Wizard - If you're not sure of how to select an amplifier, let the Amplifier Wizard help suggest one. The Amplifier Selection Wizard will lead you through a few generic questions and based on the response information will search for and sugest amplifiers for further evaluation.

Reverse Search - Data derived from the configured circuit can be used to automatically load the parametric search engine. This allows the parametric search to look for an amplifier that will work in the circuit. If an amplifier that meets all requirements can't be found, the search will suggest components that are close.

Select the Circuit

Choose the a circuit from the pull down, labelled "Select Mode," located at the top of the tool. The selections are Inverting amp, Non-Inverting amp or Difference amp configurations.

Configure the Circuit, Set Signal Parameters, Run the Model, View Results

ADIsimOpAmp (Opening Screen)

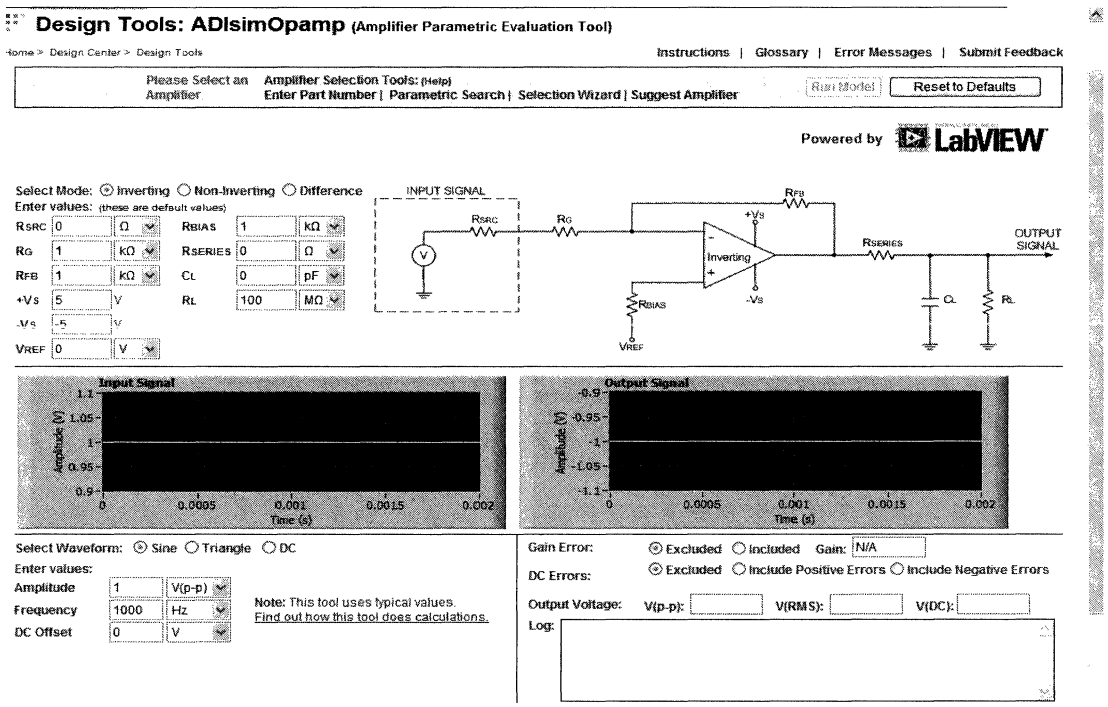

When first opened the ADIsimOpAmp will appear as shown above. This screen allows the user to configure a circuit, enter component values, and select the input signal type (square wave or sine wave), frequency, and amplitude. It also offers the option of either entering a desired part number or using the selection features to help find a part that will be suitable for the circuit.

ADIsimOpAmp: Selection Wizard

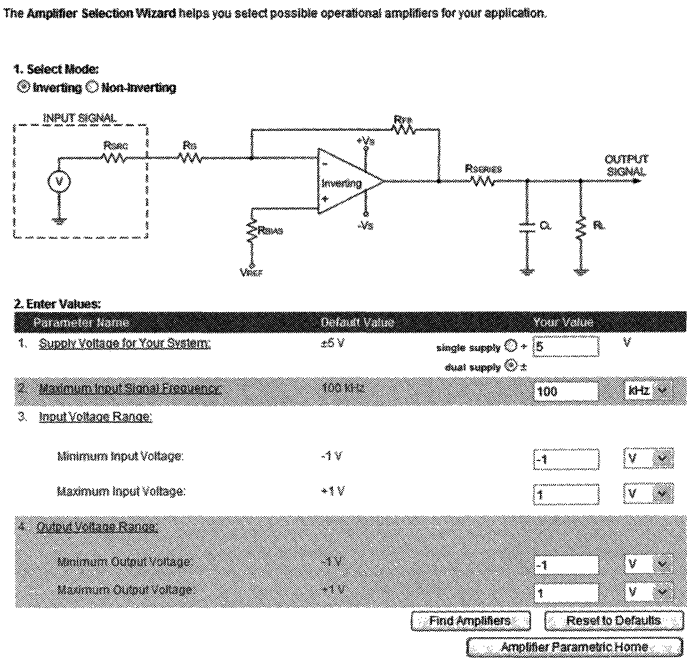

If assistance is needed with determining a circuit configuration or level shifting, the Selection Wizard offers the user additional help. By selecting the desired function and entering values in the blocks provided the Wizard can suggest a circuit and amplifiers for the application.

This shows the opening screen of the "Selection Wizard" portion of the Amplifier Evaluation tool.

First select the mode of operation: inverting or non-inverting.

Then select the power supply values, the maximum input frequency, minimum and maximum input voltage range, and minimum and maximum output voltage range.

Next, click on the button "Find Amplifiers."

ADIsimOpAmp: Selection List

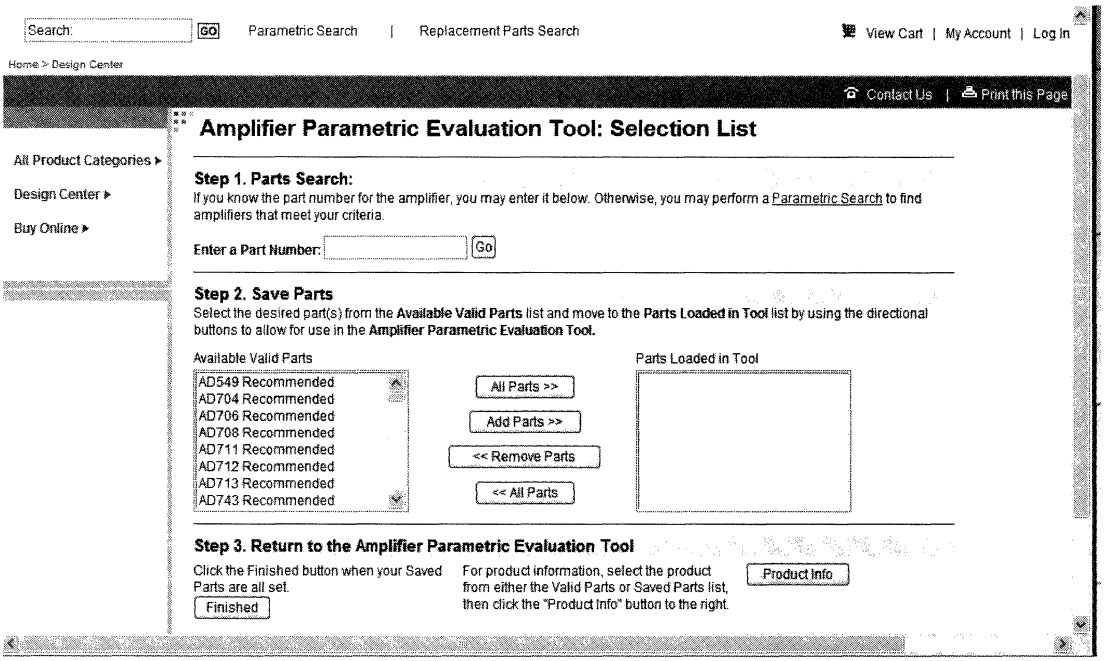

Shown above are the results of a Suggest Amplifier (reverse search) List. Note that the part numbers listed are "Recommended." Recommended means that the device will work in that application as defined earlier. Desired parts can be selected from the list and loaded into the tool for evaluation.

ADIsimOpAmp Screen – Normal Results

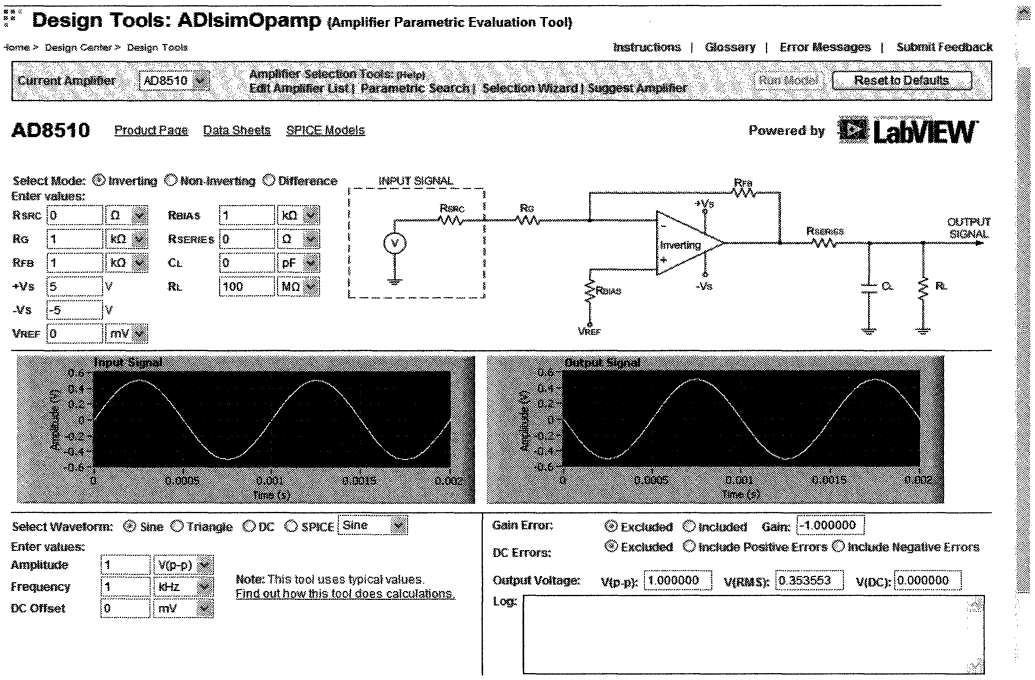

This figure shows the result of a circuit simulation using the AD8510, in the APET mode. No failures were detected. At this point the user could then perform further evaluation of the AD8510 in SPICE mode by clicking the SPICE radio button and then selecting a Sine or Square wave for the input signal.

ADIsimOpAmp Screen - Warning

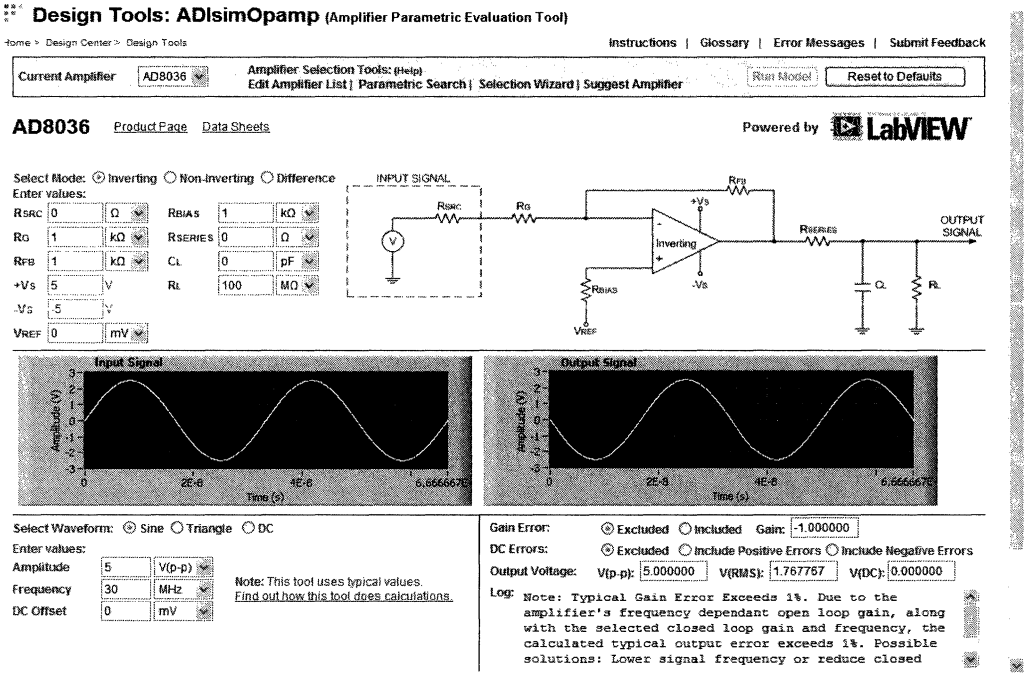

In this slide the amplifier gain error exceeds a set limit and so the tool returns a warning. This warning is in the log. Note that the log box is framed in yellow to emphasize the warning. It is then left to the designer to determine if this warning is a problem that needs correcting. Note that several corrective actions are suggested.

ADIsimOpAmp Screen – Warning (Cont.)

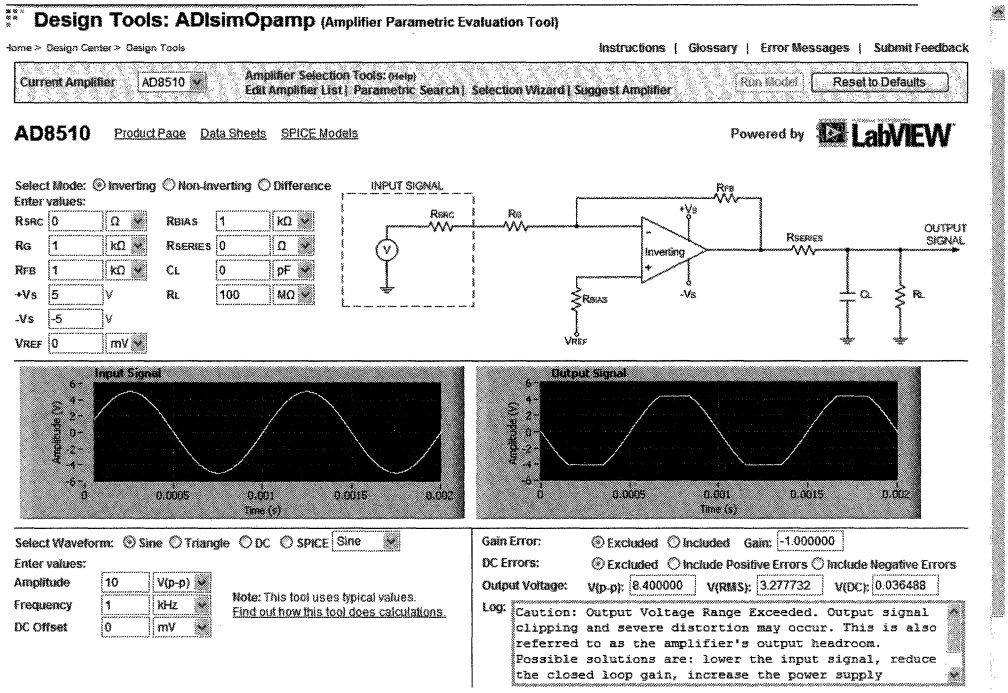

This slide demonstrates another level of warning. Note that the log box in this case is framed in orange. The issue here is that the amplifier is clipping. Again, note that several corrective actions are suggested.

ADIsimOpAmp Screen – Error

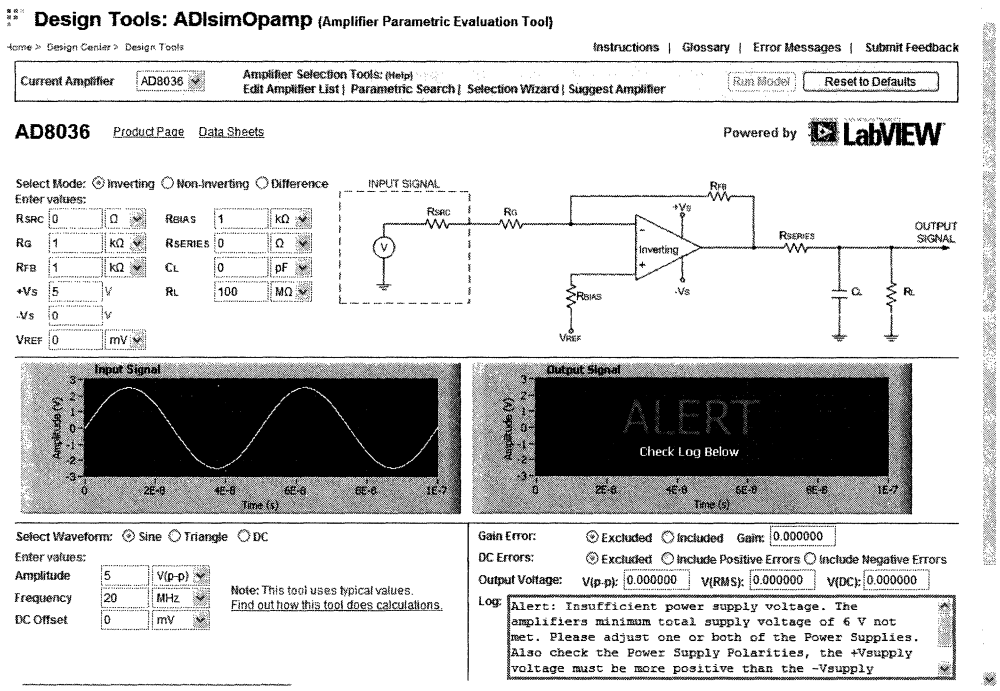

A more serious error is detected here. The minimum power supply requirements have not been met. Note that the log box is framed in red this time. Once more, corrective actions are suggested.

ADIsimOpAmp Screen – Spice Mode

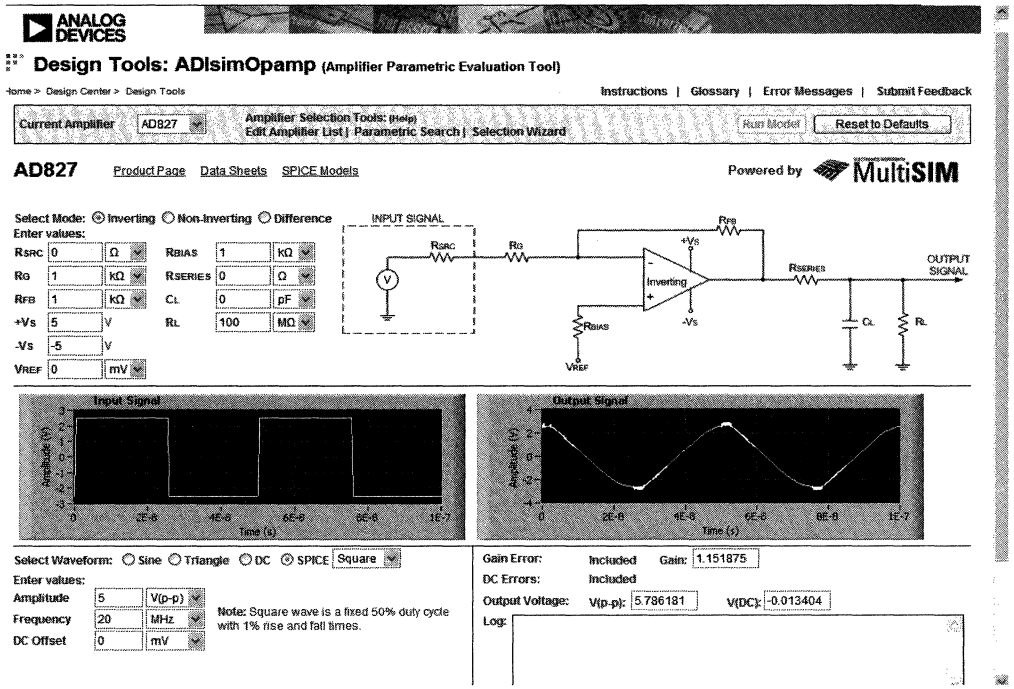

In this case the spice mode has been selected. In the spice mode, some of the error checking is not used, since this is a function of the LabVIEW environment. In the spice mode the only error that will be reported is if the circuit fails to converge. The output is whatever it would be. In this case the amplifier is slew limited. The amplifier is a 50MHz part and the input is a 20MHz square wave. Being a square wave, the harmonic content is largely outside the range of the op amp. Also the amplitude of the input is in excess of the slew rate of the amplifier.

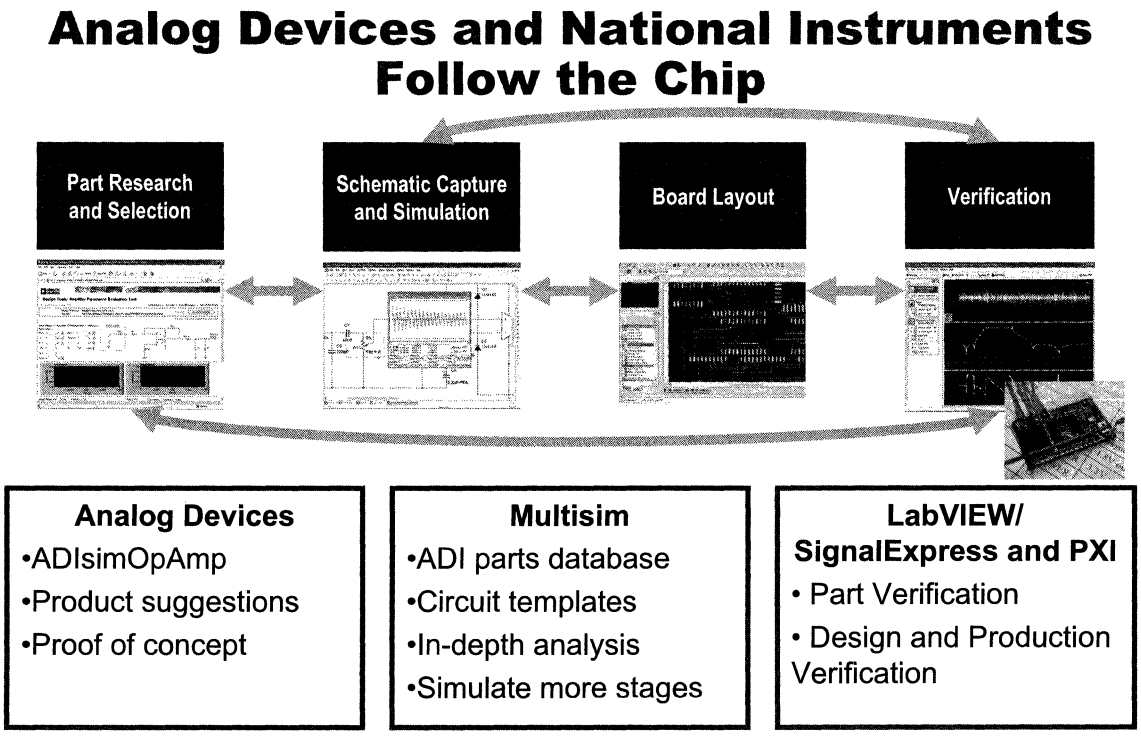

Any design has four major stages

Part Research and Selection

Schematic Capture and sometime simulation

Board Layout and

Verification

Traditionally many of these stages were done in isolation with little automation or connection between the tools used for each stage – now ADI and National Instruments are working to integrate these steps.

Tools like the ADIsimOpamp help you find and select parts far better than just data sheets.

Multisim supports the ADIsimOpamp and gives you an easy transition to working in Multisim by providing circuit templates that complement the ADI tool.

You can also find a complete database of all available ADI op amp models for Multisim. Information on how to obtain these models for the Freeware edition of Multisim are included in the CD package.

Multisim lets you evaluate your components in more detail as well as simulate your complete designs.

You can also use LabVIEW and SignalExpress in conjunction with test equipment to verify component or design operation. Once you set up a project, you can save and reuse it; this helps to automate your evaluation process.

Multisim Product Overview

- **Integrated schematic capture and simulation**
- **Interactive simulation**
- **Advanced analyses**
- **Useful for SPICE and VHDL* simulation and co-simulation**
- **Tight integration with real-world measurements**
- **Analog Devices circuit templates**

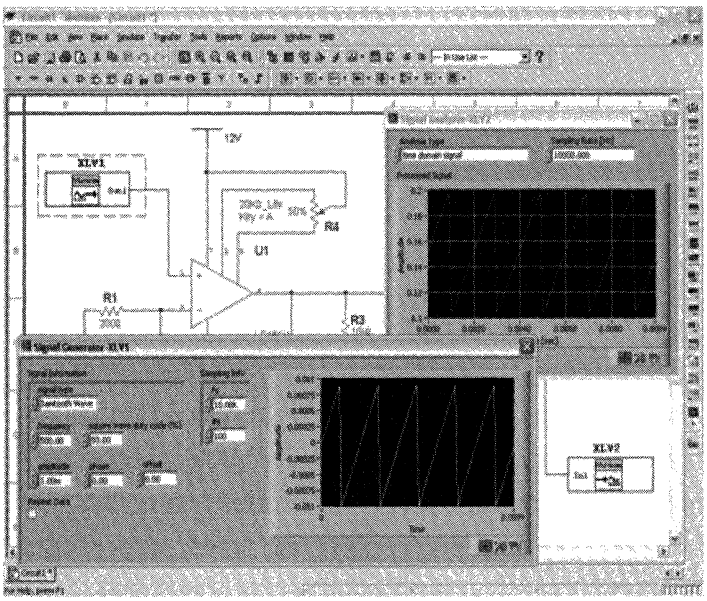

Multisim provides you with an integrated schematic capture and simulation environment. You can simulate right from the schematic environment. You can also immediately run simulations without going through complicated set-up procedures.

You can make use of Virtual Instruments, like oscilloscopes and function generators, to simulate your designs. These instruments connect to your schematic like real instruments connect to test circuits. You can even create your own custom sources and instruments with NI LabVIEW and use them in Multisim.

Virtual instruments, together with interactive components like potentiometers, enable an interactive simulation environment that closely matches the lab experience. Available analyses range from ac and dc sweeps to noise analysis and worst case analysis. If you have data acquisition devices or want to import scope data, you can make use of a wide range a data import functions. You can also export your results to NI SignalExpress and compare your simulations to measurements from your prototypes.

To help you get started, you can use some of the circuit templates that complement the ADIsimOpAmp tool.

You will receive a freeware edition of Electronics Workbench software. This edition limits your design size to 50 components and 2-layer boards. The simulation capabilities will run for 45 days after you run Multisim for the first time.

Active Filter Design Using Filter Wizard

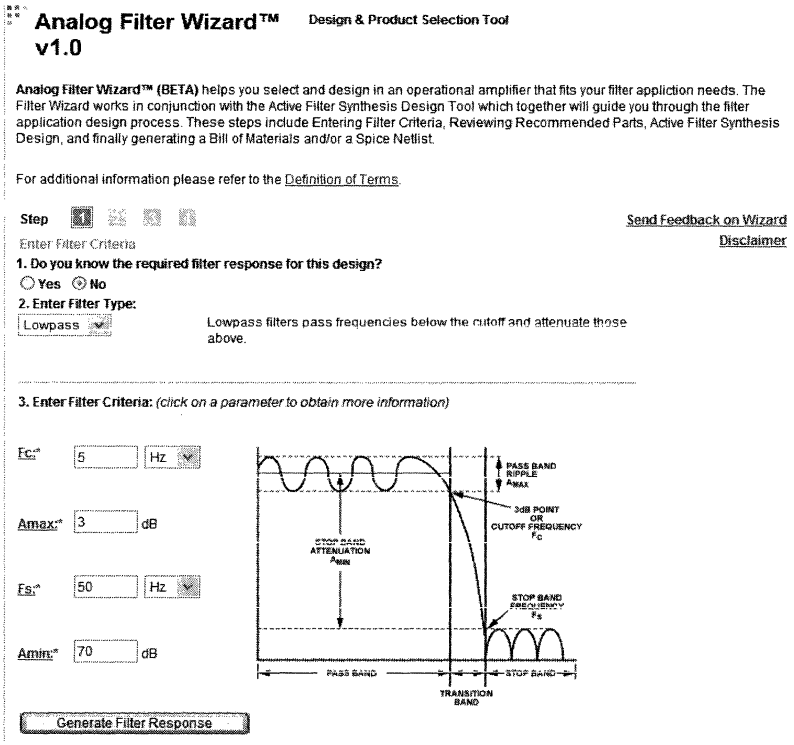

The Filter Design Wizard is designed to assist in the design and part selection for active filters.

There are two modes of operation for the Filter Wizard. The first is the "expert" mode. In this mode the designer knows the type of filter to be designed. That is he knows that he needs, for instance, a fifth-order 0.5dB Chebyshev.

Otherwise, you must describe the response of the filter in terms of the passband ripple, corner frequency, stopband frequency, and minimum attenuation. Then you will be asked to choose between the various filter types.

Selecting the Filter Response, Topology, Power Supply, CM Voltage, and Gain

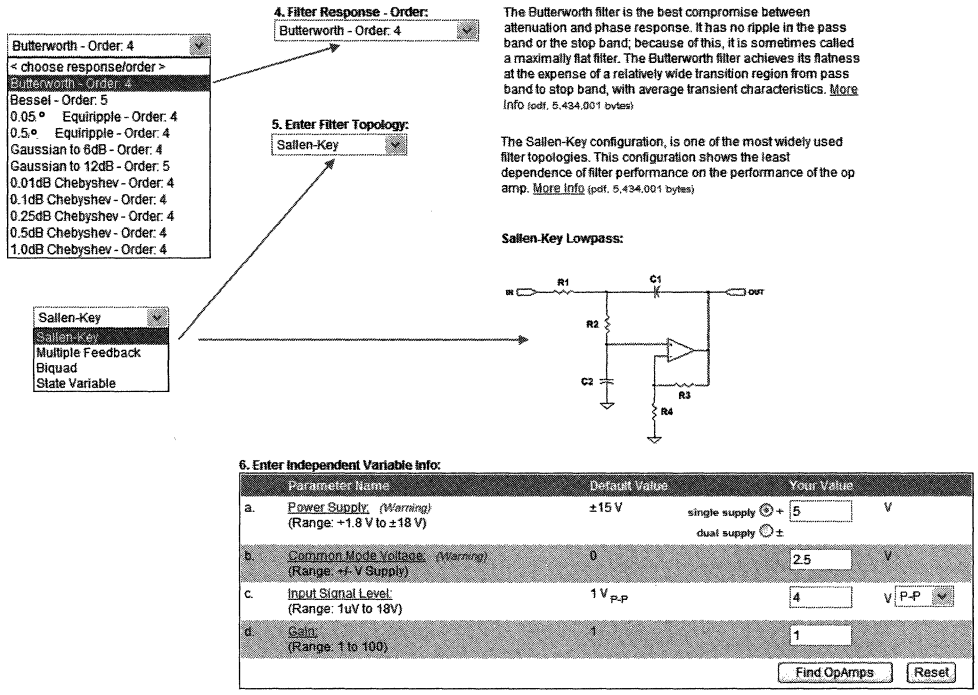

The wizard will then come back with several possible filters that will satisfy the requirements in a pull down menu. Beside the filter response pull down is a discussion that gives a brief description of the tradeoffs of each of the possible choices. There is also a link to the filter section which will give a much more detailed description of each filter type.

Designing a filter is a two step process. The first is to decide what it is you want to build. This means to determine the response characteristics and order of the filter, which is what we have just done. The next step is to determine how to build it. What this means is to determine the circuit topology. Again, several choices are available in a pull down box. Next to the pull down, again, is a brief description of the circuit, with a link to a more detailed description. A schematic is also included.

The Wizard then suggests op amps that best fit the application performance values you entered. "Best-fit" is determined by the results of a parametric search of Analog Devices' product database. The order of parametric preference is: input bias current, voltage noise density, current noise density, input offset voltage, open-loop gain, and power supply voltage.

Each product number is a link to the product page. The product page provides links to the product description, data sheet, package/price, samples, and purchasing information.

Once the selection of circuit topology has been made, the designer enters information on a few more variables, such as supply voltages and signal levels, and the wizard then returns a selection of op amps that will work in the circuit.

Since the open loop response of an op amp is, in itself, a low pass filter, the op amps are chosen so that the open loop gain is large enough that the op amp response will not materially affect the response of the filter.

Filter Wizard Stage Design (top)

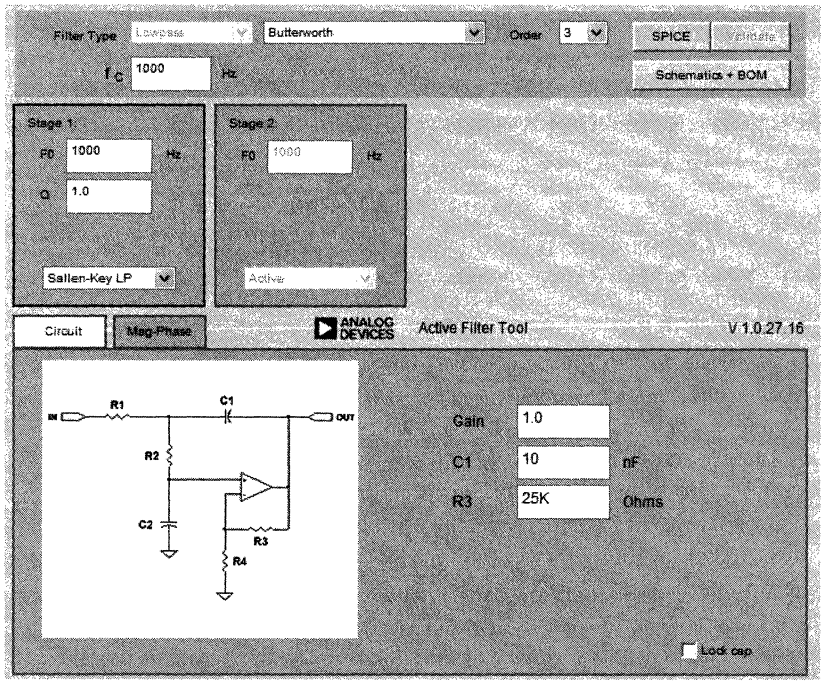

Once the op amp is chosen from the list provided, the designer enters the third page. Here, the component values for the filter are determined.

Multipole filters are made up of cascaded first-order and second-order sections. While it is possible to design a third-order section with only one active element, the circuit value sensitivities increase, so we have limited the choices here to first- and second-order sections.

The pertinent parameters (Fo and Q) for each of the filter section are loaded into the wizard. Since all of the component values are ratiometric, one value has to be chosen to set the rest of the variables. We typically choose a capacitor first, since there are fewer choices of values for capacitors than there are for resistors.

Filter Wizard Stage Design (bottom)

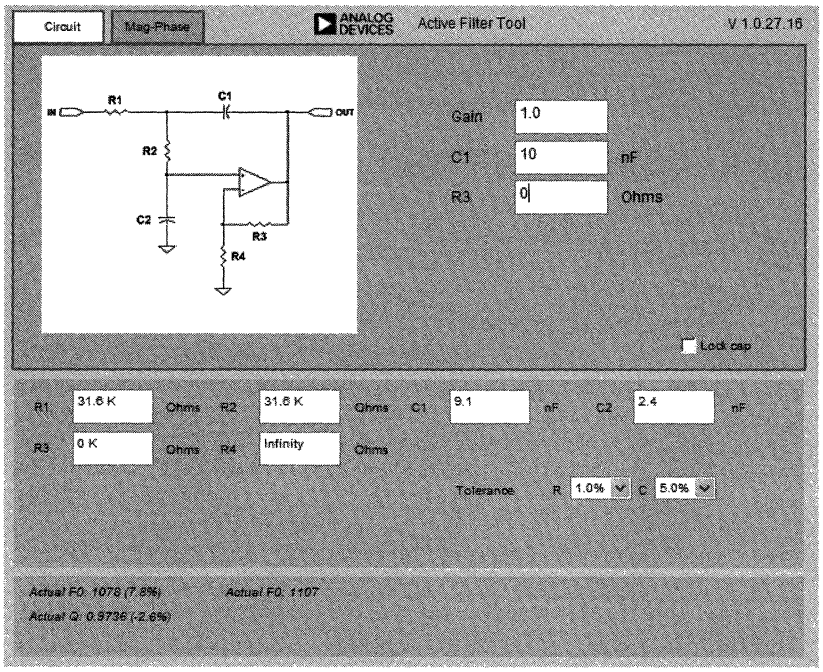

The values returned for the component values are exact values. There is the option to choose standard values (0.5% to 2% for resistors and 2% to 5% for capacitors). The designer has the ability to overwrite these values as well. Since changing the resistor and capacitor values will cause a small change in the Fo and Q of the filter, the wizard will recalculate these values and return the error introduced by going to standard values.

Once this process is completed for each section of the filter, the design is finished.

Filter Design with Component List

Stage 1: AD820 Sallen-Key LP

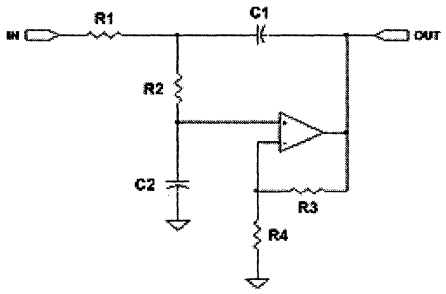

AD820
R1 3.445e4
C1 1.000e-6
R2 3.445e4
C2 8.535e-7
R3 2.500e4
R4 Inf.

Stage 2: AD820 Sallen-Key LP

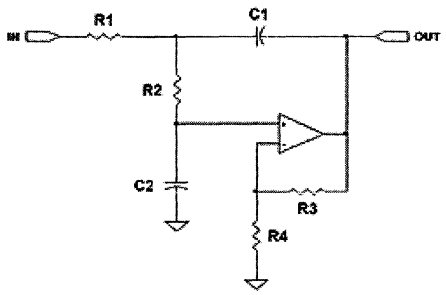

AD820
R1 8.321e4
C1 1.000e-6
R2 8.321e4
C2 1.463e-7
R3 2.500e4
R4 Inf.

The possible outputs of the wizard include magnitude and phase plots, a schematic page with component values list, and a Spice deck. With the Spice deck, further characterization of the filter is possible. This includes Monte-Carlo analysis of the component values and possibly including the op amp response with the filter response.

ADC Anti-Aliasing Suppression Assistant

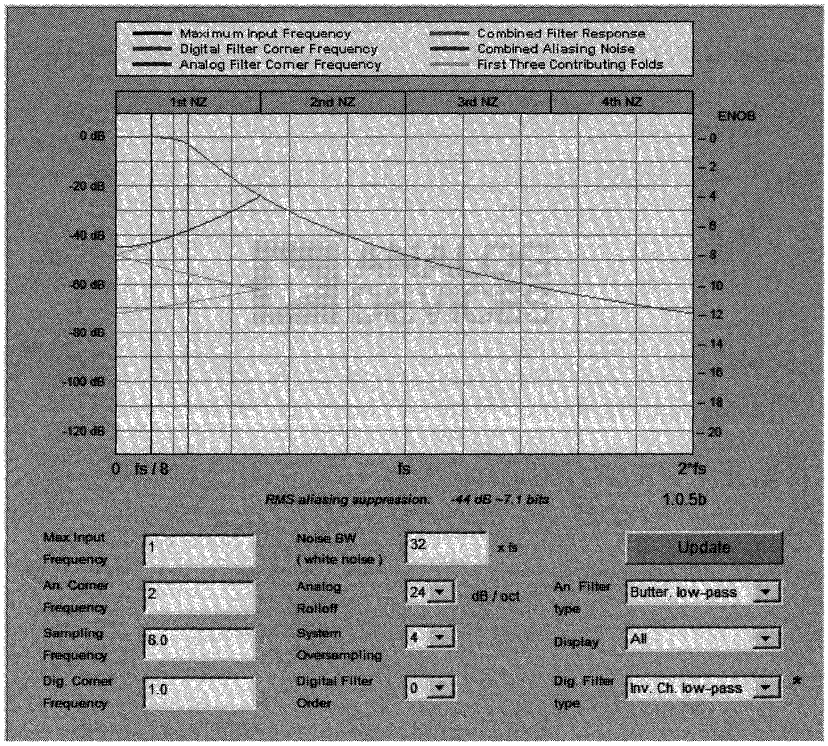

A design assistant is available to help in designing the anti-aliasing filter of an ADC system. This applet illustrates aliasing and its suppression through filtering and oversampling in a classic (non-sigma-delta) A/D converter. An ideal ADC is assumed – distortion free, unlimited bandwidth, etc. – in order to focus solely on aliasing effects. The input signal is also assumed to be noise free, but the most practical use of this applet is to find a combination of filtering and oversampling that pushes the aliased terms below the noise floor of the input, or the overall system .

The finite rolloff of practical analog filters means there are always some undersampled high frequency components that fold into the passband, or "Nyquist Zone" (NZ), and appear in the sampled signal as noise. For the simple case of an input composed of bandlimited white noise, this applet gives an estimate of how much out-of-band signal will be folded into the baseband (1st NZ). In many real-world situations, the out-of-band signal is of lesser amplitude and so this estimate is too conservative; in other cases it might not be conservative enough – it depends on the out-of-band signal level. The bandwidth limit for the input white noise signal is a multiple of the sampling frequency. By default, that multiple is 32×; the maximum is 256×.

Aliasing suppression in oversampled systems is achieved through a combination of analog and digital filtering, but digital filtering cannot replace a high-quality analog filter, because it cannot remove aliasing noise after it's been folded into the passband by the sampling process. Instead, oversampling must be used to put enough octaves and attenuation between the Nyquist frequency and the highest passband frequency of interest. Digital filtering can then be used on the sampled signal to eliminate frequencies between passband and Nyquist. For simplicity, the additional aliasing noise that results from downsampling the sampled signal is not shown.

Harmonic Image DAC Design Tool

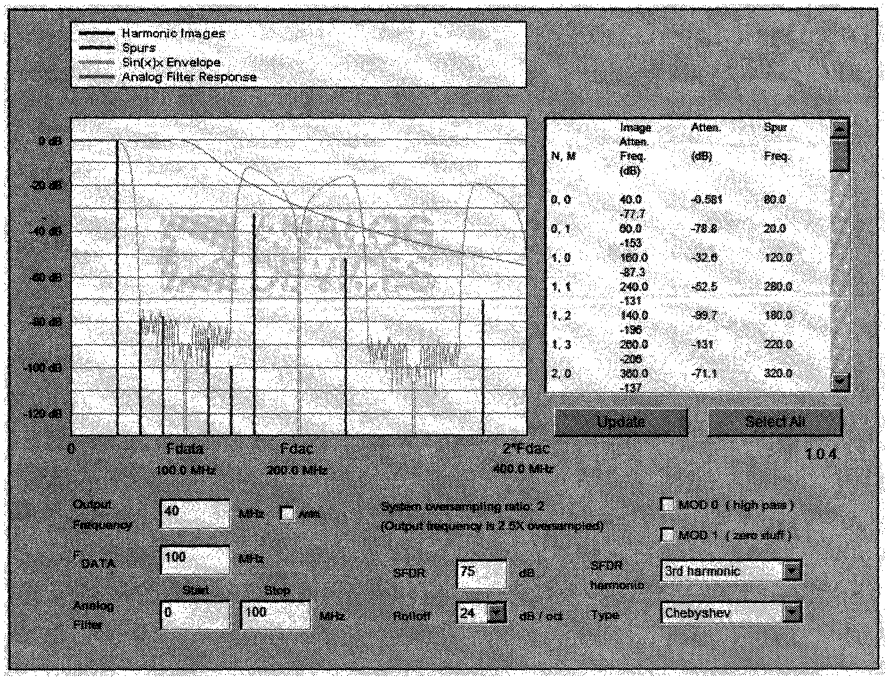

The DAC image applet shows the harmonic images and spurs for single frequency output from an DAC, an AD9772 in this example. The model of the AD9772 is simplified and idealized – only SFDR is modeled, and it is assumed frequency-independent. The response characteristics of the internal digital filter have been approximated. See the data sheet for actual performance data.

For an ordinary DAC, images are located at N·FDAC ± FOUT. The AD9772 contains an integral interpolator which doubles the input data rate creating an image of the output frequency mirrored about FDATA/2. An interpolation filter suppresses the upper image in low-pass mode (MOD0=0), or suppresses the fundamental when in high-pass mode (MOD0=1). Both the filtered and unfiltered images then create further images and spurs at the DAC data rate, according to the N·FDAC ± FOUT rule.

The AD9772 also has a "zero-stuffing" mode (MOD1=1) which allows the data stream to be doubled a second time by inserting zeroes between each sample. Zero-stuffing doubles again the number of images per FDAC harmonic, but FDAC is twice what it would be without zero-stuffing and the mathematics work so that the location of the images is the same as without zero-stuffing – only the amplitudes change. These new images are not filtered internally, so the upper image can be used for direct IF synthesis. MOD0 and MOD1 are often used together for this purpose.

Spurious 2nd or 3rd harmonics of each image are assumed to result from D/A nonlinearities and so are folded within the first Nyquist Zone (NZ) of FDAC. These spurs then have their own harmonic images that roll off as sin(x)/x (where x = π·FSPUR / FDAC). The magnitude response of the AD9772, combining its internal interpolation with the sin(x)/x envelope, is shown.

To show external selection/suppression of desired/undesired images and spurs, the applet can apply a simulated post-DAC analog filter.

In Amp Gain and Level Calculator—AD623

(A) Differential Input Voltage = 0V

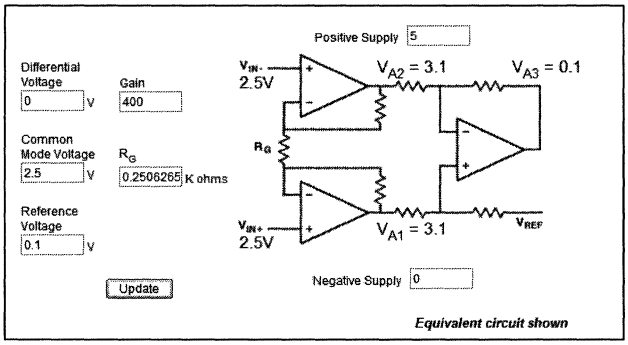

All Nodes Within Allowable Range

(B) Differential Input Voltage = +10mV

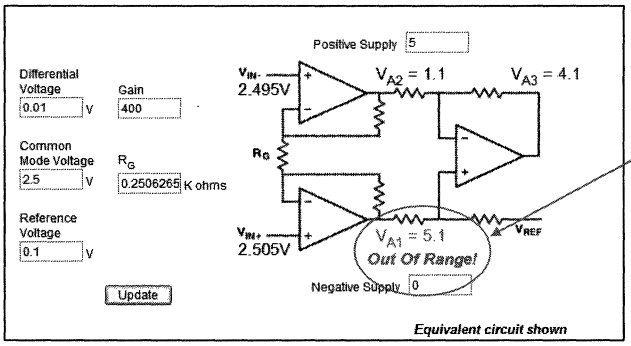

Out of Range Error at Node V_{A1}

The Instrumentation Amplifier Gain Calculator calculates the gain of an in-amp circuit given the gain-setting resistor or conversely gives the value of the resistor need for a particular gain. It also checks to ensure all node, even those internal, are kept in their operational range.

To use the tool, simply enter data in the fields provided. If an input is out of range an alert will appear (in red, along with an "error" message).

Click "Update" or tab to another field to see node voltages annotate the schematic on the right. Internal node voltages are for the equivalent 3-op amp (or 2-op amp) circuit, assuming a 0.5V level shift. This may not reflect the exact internal implementation, but is instead a simplified schematic.

If the combination of inputs causes internal or external limits to be exceeded, the problem node value will be highlighted in red and an "Out of Range!" message will appear. When this message is present, all node values should be considered invalid. These input conditions could include such things as overranging the output or, especially in single-supply applications, putting an internal node out of range. The fact that the tool will show the conditions of internal nodes is useful, since you cannot probe those nodes directly and may be mislead into believing the circuit is working when, in fact, it is not.

Only one of "Gain" or "RG" can be specified - the other is computed automatically.

In Amp Gain and Level Calculator—AD627

(A) Differential Input Voltage = 0V

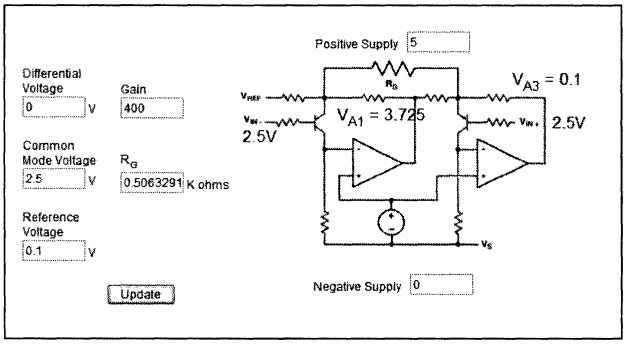

All Nodes Within Allowable Range

(B) Differential Input Voltage = +10mV

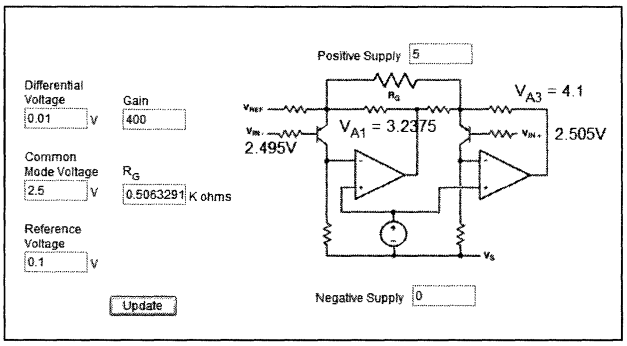

All Nodes Within Allowable Range

The same basic tool exists for the two amplifier configuration instrumentation amps. Special care is needed here to make sure that internal nodes are within their operational range. The tool will detect these conditions and flag them.

Differential Amplifier Interactive Tool

Interactive Design Tools: Differential Amplifiers :
DiffAmp Common-Mode Range / Gain / Noise Calculator
An online tool to select values for R_G & R_F and to illustrate the maximum differential and common mode voltages allowable.

AD8138 single supply

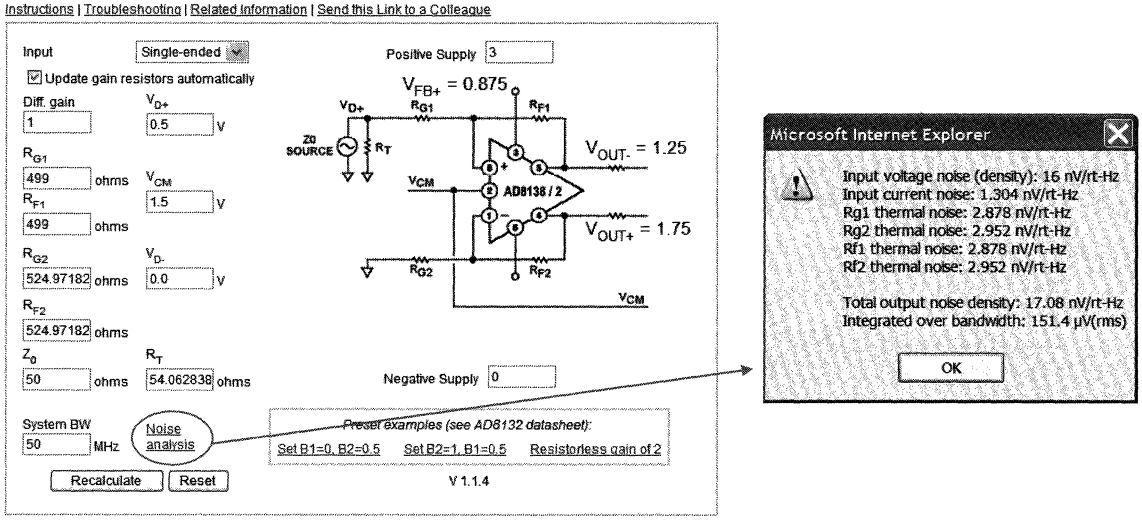

The Differential Amplifier Calculator has two basic modes: manual and automatic. The default is automatic in which it is assumed you want to calculate resistor values (incl. termination) to match a source impedance. In the automatic mode, RF is computed from RG and Gain: changing either of the latter will affect the former. Changing ZO affects all the gain resistors (which should always be a minimum of 10 × ZO), as well as RT. RT is the value for the termination resistor taking into account the impedance of the differential amplifier's gain network and is recalculated when either of RF or RG are changed (changing RT, however, affects nothing). Single-ended termination resistances are calculated assuming VCM = 0. For other fixed voltages, the input impedance is nonlinear.

By unchecking "Update resistor values automatically," the calculator is placed in the manual mode where each resistor can be set independently. Gain and the node voltages are the only thing computed when "Recalculate" is clicked or a new field is entered. For open circuits or "infinite" resistors use a large number instead, such as 1e99. Do not leave any fields blank.

The schematic shows a matched Thevenin source driving a terminated line, however, input voltages are actually set independently for each value of RG. Consequently ZO and RT don't affect the calculation of node voltages. For unterminated calculations set ZO to zero.

Note: For clarity, this calculator shows ideal behavior and does not show the effects of input offset currents and voltages.

Active Feedback Amplifier Common Mode Range/ Gain Calculator

Interactive Design Tools: Differential Amplifiers :
AD8129/30 Common-Mode Range / Gain Calculator

An online tool to select values for R_G, R_F and to illustrate the maximum differential and common mode voltages allowable.

AD8129 +/- 12V

Instructions | Troubleshooting | Related Information | Send this Link to a Colleague

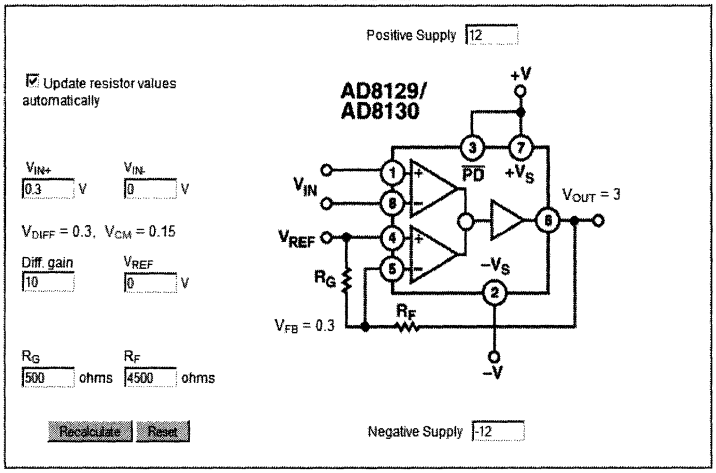

There is a novel amplifier topology called active feedback.

These amplifiers have two sets of fully differential inputs, VX1(1) - VX2(2) and VY1(4)- VY2(5). Internally, the outputs of the two GM stages are summed and drive a buffer output stage.

In this device the overall feedback loop forces the internal currents IX and IY to be equal. These currents are the outputs of the input pin summing amps. This condition forces the differential voltages VX1 - VX2 and VY1 - VY2 to be equal and opposite in polarity. Feedback is taken from the output back to one input differential pair, while the other pair is driven directly by an input differential input signal.

An important point of this architecture is that high CM rejection is provided by the two differential input pairs, so CMR is not dependent on resistor bridges and their associated matching problems. The inherently wideband balanced circuit and the quasi-floating operation of the driven input provide the high CMR, which is typically 100dB at dc.

One way to view this topology is as a standard op amp in a non-inverting mode with a pair of differential inputs in place of the op amps standard inverting and non-inverting inputs. The general expression for the stage's gain "G" is like a non-inverting op amp, or:

$$G = \frac{V_{OUT}}{V_{IN}} = 1 + \frac{R2}{R1}$$

As should be noted, this expression is identical to the gain of a non-inverting op amp stage, with R2 and R1 in analogous positions.

This tool calculates the common mode and differential mode gain and makes sure all internal nodes are within allowable ranges.

Op Amp Error Budget Calculator -1

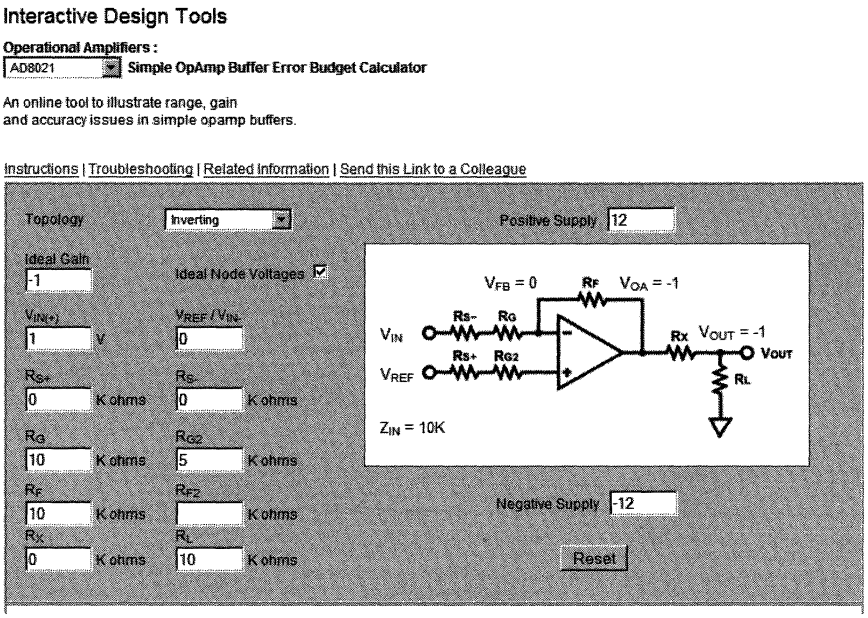

There are a number of sources of potential error in an op amp design. The Error Budget Calculator has two parts: an annotated schematic at top and a table of contributing error sources at bottom. Op amp parametric data is automatically entered in the appropriate fields of the bottom table, and default values for the application parameters have been assigned to the application-specific fields at the top. All input data can be manually overridden, however, output fields (surrounded by light gray) cannot be changed. This is a great time savings from having to enter all the data by hand.

After entering data in a field, hit tab or click "Update" to compute derived values and see node voltages updated on the schematic. If the inputs are out of range an alert will appear. If the combination of inputs causes internal or external output limits to be exceeded, the problem node value will be highlighted in red and an "Out of Range!" message will appear. When this message is present, all node values should be considered invalid. Do not leave fields blank: if you see "NaN" (Not a Number), this means that insufficient data was entered to compute a value.

"Gain" and "RF" are computed automatically from one another, based on the value of "RG." The calculation is ideal and doesn't reflect RS, RX and RL, for example.

Equations listed in the "Calculation" column are approximate and reflect the worst case between the three amplifier configuration choices. Modifications to the equation for particular configuration types are indicated in ( ). For example (1/2 : noninv) means an additional factor of 1/2 should be used to compute this quantity for the noninverting configuration.

Op Amp Error Budget Calculator -2

Application Parameters

Operating Temp., T_A	85 °C		Update		
Supply Variability (ripple+load reg.)	1 %				
Error Source	**Specification**	**Approx. Calculation**	**Absolute Error**	**Drift/Gain Error**	**Resolution Error**
Resistor Tolerance	0.1 %		2000 ppm		
Resistor Drift, TC_R	25 ppm / °C	~ (1/2 : noninv) $TC_R \times T_{DIFF}$		125 ppm	
Temp. difference, T_{DIFF}	5 °C				
Nom. Open Loop Gain, A_{OL}	25 V/mV		80 ppm		
Min. Open Loop Gain	16 V/mV				45 ppm
Input Offset Voltage, V_{OSI}	1 mV	$V_{OSI} / (V_{IN}-V_{REF})$	2000 ppm		
Input Offset Voltage Drift, V_{OSI_TC}	0.2 µV / °C	(2 : inv.) $V_{OSI\_TC} \times (T_A-25) / (V_{IN}-V_{REF})$		24 ppm	

Specs shown are worst-case for the selected part, if available, otherwise typical values are used. If no spec is available, "N/S" will appear in that field and an ideal spec (usually zero) will be used for the calculation. Please note that it is highly unlikely that all worst-case specs would ever be present at the same time in the same part. The designer should always refer to the appropriate data sheet and substitute numbers most appropriate to the application. All calculations are approximations, with errors displayed and summed in absolute PPM, even though in some scenarios the actual values would be negative.

The error calculated is separated into two parts, as discussed in the error source section on op amps. It is separated into "resolution error," which are errors which cannot be adjusted out of the system, and drift/gain errors, which can be adjusted out with the proper circuitry. Please refer to the op amp section on error sources for more information.

Op Amp Error Budget Calculator -3

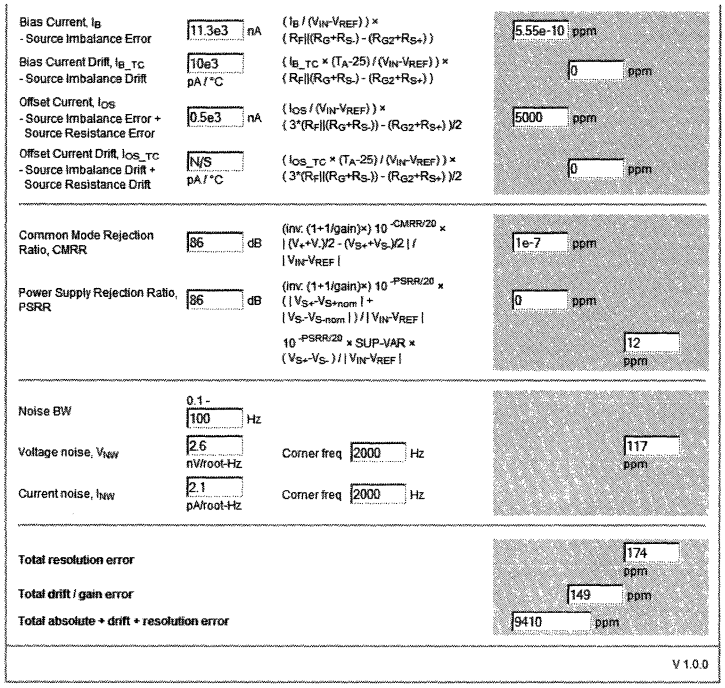

This shows the remaining parameters calculated by the Op Amp Error Budget Calculator.

Instrumentation Amp Error Budget Calculator -1

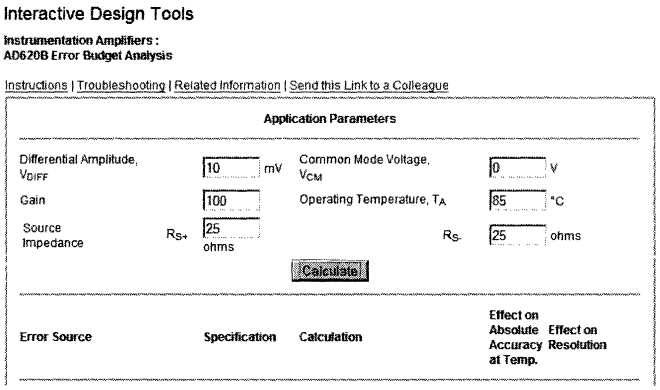
Interactive Design Tools

Instrumentation Amplifiers :
AD620B Error Budget Analysis

Instructions | Troubleshooting | Related Information | Send this Link to a Colleague

Application Parameters

Differential Amplitude, V_{DIFF}	10 mV		Common Mode Voltage, V_{CM}	0 V	
Gain	100		Operating Temperature, T_A	85 °C	
Source Impedance	R_{S+} 25 ohms			R_{S-} 25 ohms	

Calculate

Error Source	Specification	Calculation	Effect on Absolute Accuracy at Temp.	Effect on Resolution

A similar error budget calculator also exists for instrumentation amplifiers. For this tool the pertinent data is entered into the top field of the calculator. The tool automatically enters the specification data for the particular instrumentation amplifier selected and then calculates the error. Again, the error is separated into "resolution error," which is the non-reducible part and the drift/gain error.

Instrumentation Amp Error Budget Calculator -2

Gain Error	0.5	%		5000 ppm
Gain Drift, G_{TC}	-50 ppm / °C		$G_{TC} * (T_A-25)$	3000 ppm
Gain Nonlinearity	0.0095	%		95 ppm
Input Offset Voltage, V_{OSI}	85	µV	V_{OSI} / V_{DIFF}	8500 ppm
Input Offset Voltage Drift, V_{OSI_TC}	0.6 µV / °C		$(V_{OSI\_TC} / V_{DIFF}) * (T_A-25)$	3600 ppm
Output Offset Voltage, V_{OSO}	1	mV	$V_{OSO} / (GAIN * V_{DIFF})$	1000 ppm
Output Offset Voltage Drift, V_{OSO_TC}	7 µV / °C		$(V_{OSO\_TC} / (GAIN * V_{DIFF})) * (T_A-25)$	420 ppm
Bias Current, I_B - Source Imbalance Error	1.5	nA	$I_B * (R_{S+} - R_{S-}) / V_{DIFF}$	0 ppm
Bias Current Drift, I_{B_TC} - Source Imbalance Drift	3.0 pA / °C		$I_{B\_TC} * (R_{S+} - R_{S-}) * (T_A-25) / V_{DIFF}$	0 ppm
Offset Current, I_{OS} - Source Resistance + Imbalance Error	0.75	nA	$I_{OS} * MAX(R_{S+}, R_{S-}) / V_{DIFF}$	0 ppm
Offset Current Drift, I_{OS_TC} - Source Resistance + Imbalance Drift	1.5 pA / °C		$I_{OS\_TC} * MAX(R_{S+}, R_{S-}) * (T_A-25) / V_{DIFF}$	0 ppm
Common Mode Rejection, CMRR	80	dB	$V_{CM} / (10^{CMRR/20} * V_{DIFF})$	0 ppm
Noise, RTI (0.1 Hz - 10 Hz)	6 µV p-p			600 ppm
TOTALS				21520 ppm
				695 ppm

V 1.0.6

This shows the various errors calculated by the Instrumentation Amp Error Budget Calculator.

Register Configuration Assistants

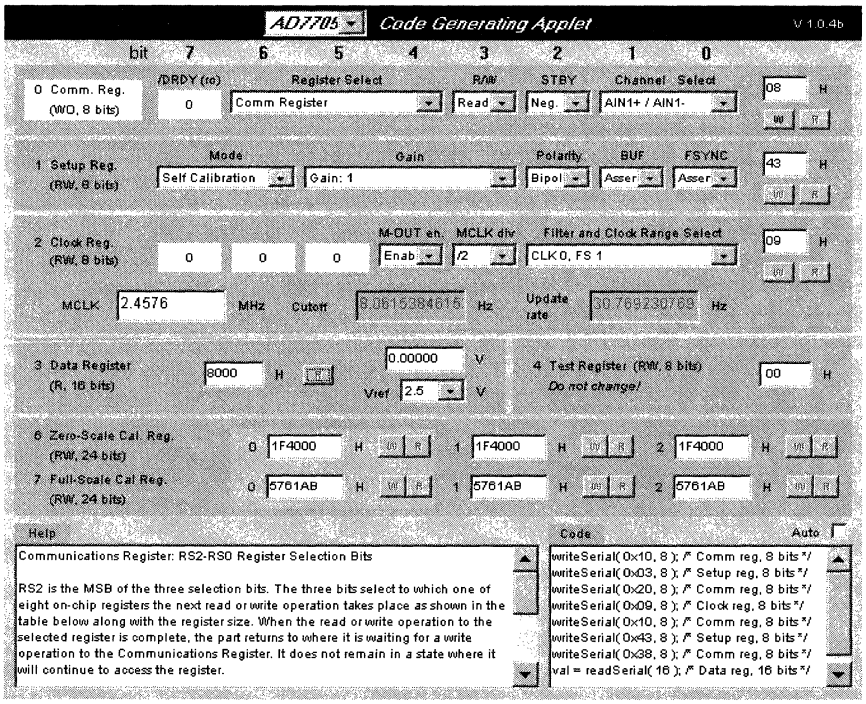

Another type of "on line" design assistant are the configuration assistants. Many modern converters are actually small systems, usually with significant digital content. This digital content is used to set operating conditions (which channel of a multi-channel input ADC for instance), While it is possible to take the register descriptions in the data sheets and use them to set the correct bits in the correct registers, it is much easier to use a configuration assistant.

An example is the configuration assistant for the AD7730 Σ-Δ ADC. Each dark gray rectangle shows a single register's content in two different ways. Most of the space is taken up with pull-down menus and a few hex input fields that display bit fields within the register, MSB to LSB. 8 bits are displayed per row, with the bits aligned in numbered columns. At the right is the combined hex code corresponding to values selected for the individual bit fields.

Changing either the hex code or the bit fields updates the other. Registers that can't be broken into separate fields are shown for completeness, even though not much can be done with them.

By default, reading and writing register pseudocode is, like the real AD7730, a two-step process: the Comm register must first be set up to select a target for the next read or write. After a register has been configured to your liking, select it in the Comm register and click "W." The selected register is now highlighted in yellow (and the channel in pink). To complete the cycle, click on the enabled "W" or "R." This will also clear the Register Select bits which must be manually reset for each access. Please note that the multiple read settings are nonfunctional – an ordinary read will be performed instead.

Both operations append pseudo-code instructions to a ticker of instructions at the bottom right of the applet. "writeSerial(val, length)" is an abstract subroutine that takes the quantity "val" and sends "length" bits of it serially to the AD7730, MSB first. The code list can be copied and pasted into another application (on most platforms).

Checking "Auto" above the Code output window puts code generation into a second mode, enabling all registers and automatically prepending the appropriate Comm register write when "W" or "R" for any register is clicked. Please note that in this mode the register select bits in the Comm register are ignored.

Direct Digital Synthesis (DDS) Device Configuration Assistant -1

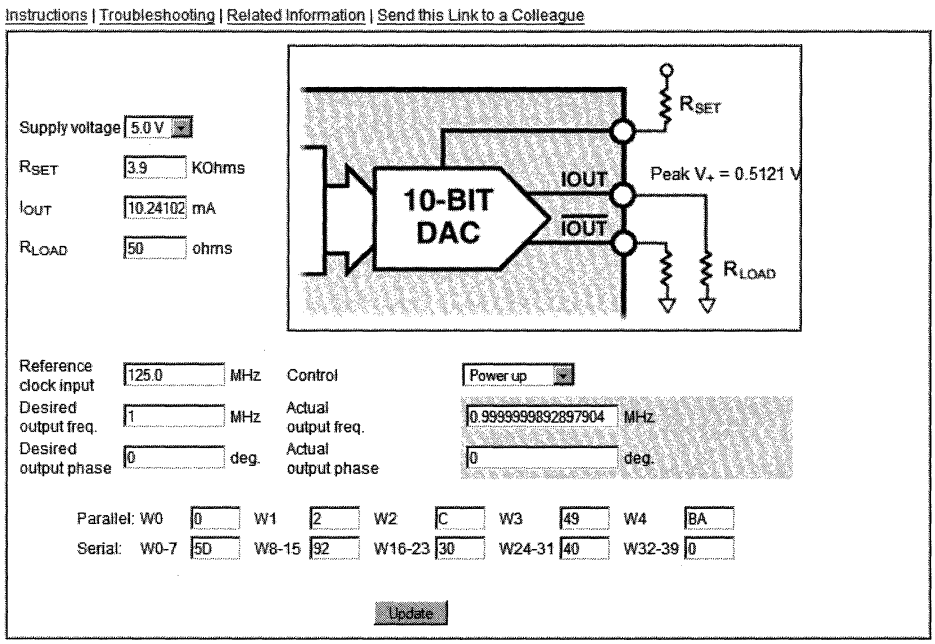

Another example of a register configuration assistant is the one for an AD9850 Direct Digital Synthesis (DDS) system.

This calculator has several distinct functions. First, it's a tool for selecting a value for RSET, which sets the output current, and checking that the output level remains within limits for a given load.

The AD9850 has complementary current output structures which limits the current and voltage that can be supplied and still meet other data sheet specifications. The output current level, IOUT, is set by a single external resistor, RSET, and the two are related by an equation. Changing one of these fields in the calculator updates the other automatically. If too high a current is selected, an error is noted. The IOUT current develops a voltage into the selected RLOAD, shown on the schematic, and is checked against the AD9850's compliance voltage.

Second, it's an assistant for selecting a 32-bit tuning word, given a reference clock and desired output frequency. Third, it shows the tuning word and other configuring bits encoded as a sequence of hex codes for use in programming the AD9850 via its parallel or serial interface.

A tuning word is selected by simply entering the desired REFCLK and output frequency. REFCLK has a maximum frequency that depends on supply voltage, selected at the top of the screen. Because the tuning word is limited to 32 bits there is typically a small deviation between the desired and actual output frequencies, which is shown in a field at right. The actual output frequency is what is encoded as the tuning word and this comprises the last four bytes of the parallel hex codes and the first four of the serial codes. Tuning words greater than 7FFFFFFF H exceed the Nyquist frequency and will cause error messages to appear.

The AD9850 has five bits of programmable phase which is selected in a manner similar to the desired output frequency. The closest available phase setting appears in the field at the right and in the corresponding hex code with the power down bit. The hex code fields are bidirectional, and a known set of hex codes can be entered to retrieve the programmed frequency and phase.

Direct Digital Synthesis (DDS) Device Configuration Assistant -2

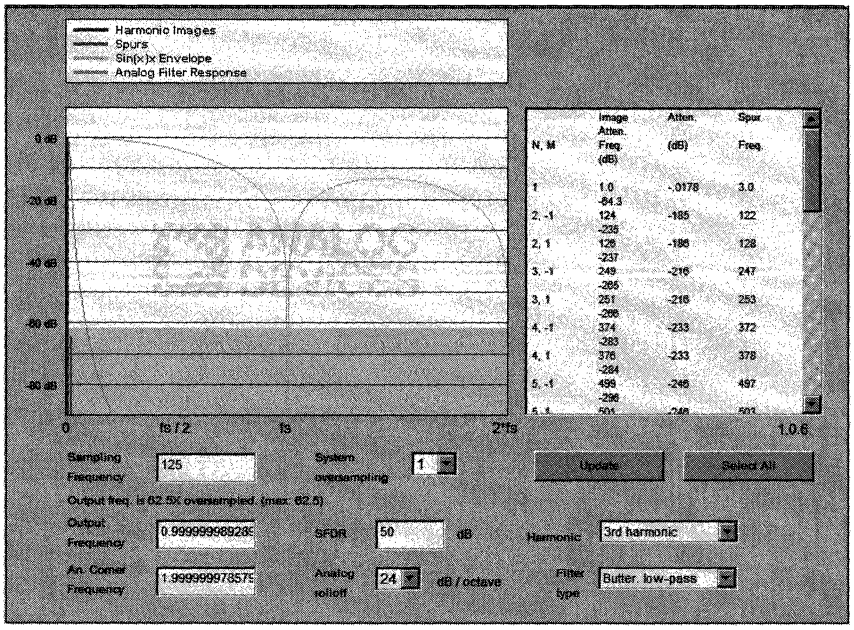

Lastly, output harmonics are shown for the selected reference clock and output frequency after an external reconstruction filter has been applied.

Suppression of images and spurs (waveform reconstruction) can be simulated by selecting corner frequency, filter order, and type of a simple analog filter (last line of images calculator; magnitude rolloff). A region corresponding to 10 bits of quantization noise is shaded at the bottom of the graph for reference. This calculator is based on a highly simplified model of the AD9850: check data sheet for parameters appropriate to your application.

Notes:

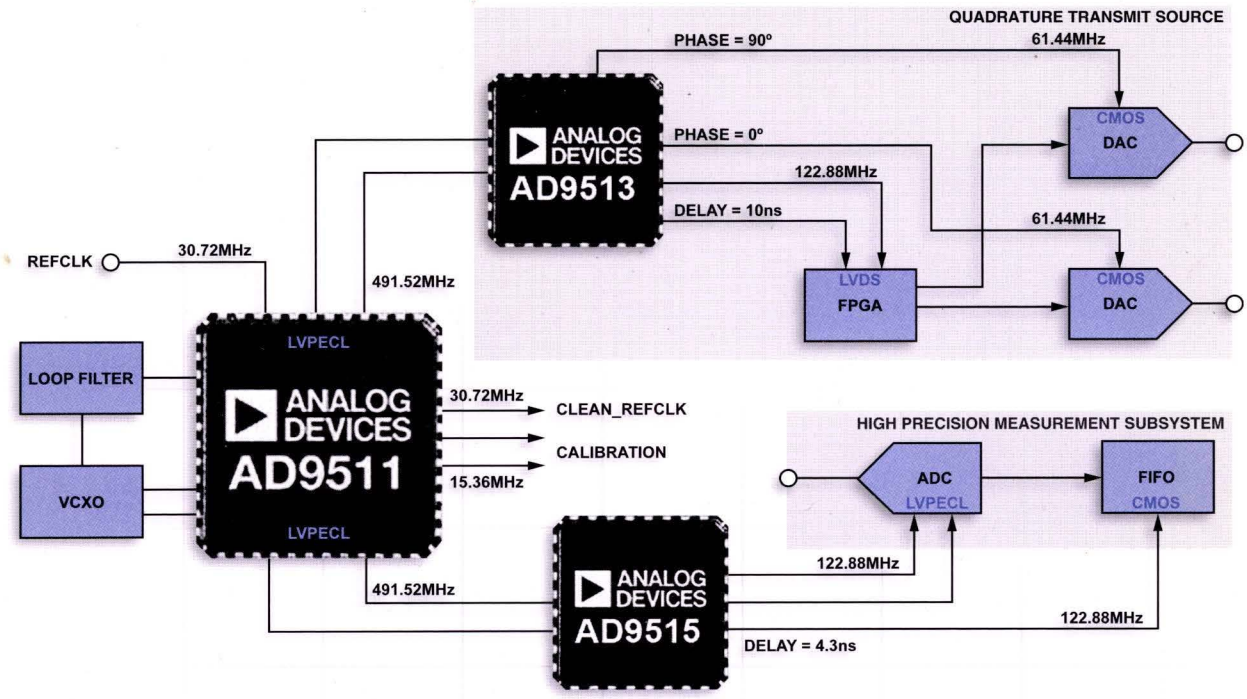

Analog Devices Design Tools

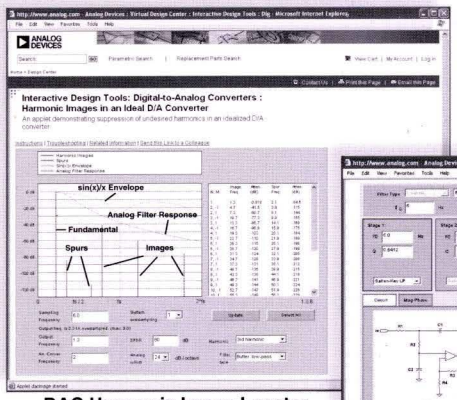

DAC Harmonic Image Locator

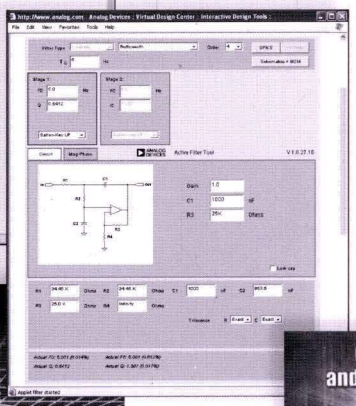

Filter Wizard

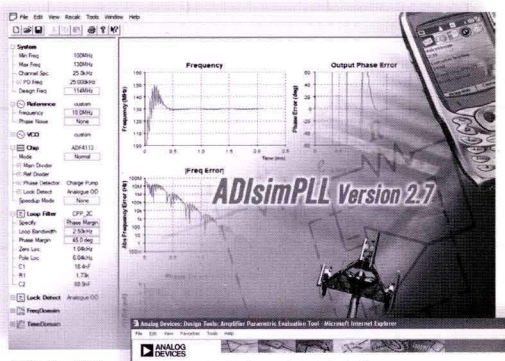

ADIsimPLL

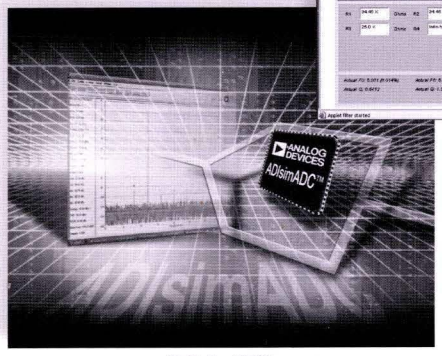

ADIsimADC

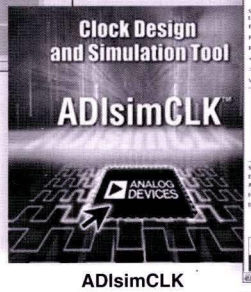

ADIsimCLK

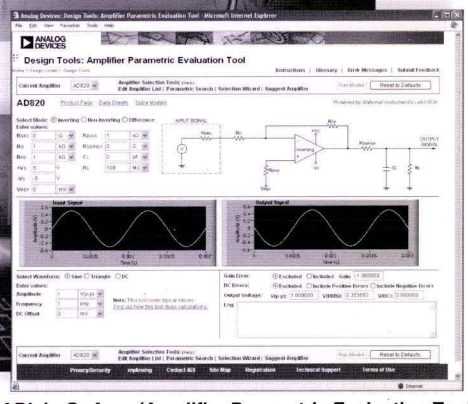

ADIsimOpAmp (Amplifier Parametric Evaluation Tool)

ANALOG DEVICES