SECTION 15

WIDE DYNAMIC RANGE ADC APPLICATIONS

- ACHIEVING WIDE DYNAMIC RANGE IN DIGITAL SPECTRAL ANALYSIS
- USE OF DITHER SIGNALS TO INCREASE ADC DYNAMIC RANGE
- DIRECT IF TO DIGITAL CONVERSION APPLICATIONS USING UNDERSAMPLING TECHNIQUES
- BANDPASS SIGMA-DELTA TECHNIQUES FOR DIRECT IF-TO-DIGITAL CONVERSION
- ACHIEVING WIDE DYNAMIC RANGE USING VARIABLE GAIN AMPLIFIERS
- WIDE DYNAMIC RANGE ULTRASOUND SYSTEMS

SECTION 15

WIDE DYNAMIC RANGE ADC APPLICATIONS

Walt Kester, Allen Hill, James Bryant

Many modern signal processing applications require ADCs with extremely wide dynamic range. The ADC is often the limiting factor in the performance of digital spectrum analyzers. In direct IF-to-digital receivers, the ADC must accurately digitize narrowband signals with a center frequency much greater than the Nyquist frequency of one-half the sampling rate. This is called *undersampling*, and often requires that a high-performance sample-and-hold be placed in front of the ADC to increase the dynamic range. The traditional sigma-delta architecture can be modified to yield a bandpass rather than a lowpass transfer function, thereby allowing it to be used to process IF signals well above 1MHz. Finally, variable gain, high speed, low distortion amplifiers may extend system dynamic range as in the case of ultrasound systems.

ACHIEVING WIDE DYNAMIC RANGE IN DIGITAL SPECTRAL ANALYSIS

Walt Kester, Allen Hill

Digital techniques are common in modern signal intelligence (SIGINT) and other high performance radios. The information is extracted from the signals using fast FFTs, digital filtering, and other powerful DSP techniques. One of the important and often limiting characteristics of such receivers is the inherent spectral purity and noise of the ADC. Digital spectral analysis is another application where the spectral purity and noise of the ADC may be the performance-limiting factor. In this section, we will examine the tradeoffs and illustrate some techniques which may be used to improve ADC performance.

High resolution, fast ADCs such as the AD9014 (see Figure 15.1) were designed to assist in spectral analysis by means of 14-bit digitization. The primary concern in spectral analysis applications is the spurious-free-dynamic-range (SFDR) of the ADC. The actual resolution (i.e., the number of bits) and the broadband rms noise level of the converter is often a secondary consideration. By using deep FFTs and averaging the results of a number of FFTs, the broadband random noise floor can be reduced, just as narrowing the bandwidth of an analog spectrum analyzer (usually at the expense of slower sweep rates) reduces the noise floor and allows small signals to be observed which were previously buried in the noise. For an M-point FFT, the level of the average noise component in each frequency bin ($\Delta f = f_S/M$) is $10\log_{10}(M/2)$ dB below the rms value of the random noise in the bandwidth dc to $f_S/2$.

AD9014 14-BIT, 10MSPS ADC BLOCK DIAGRAM

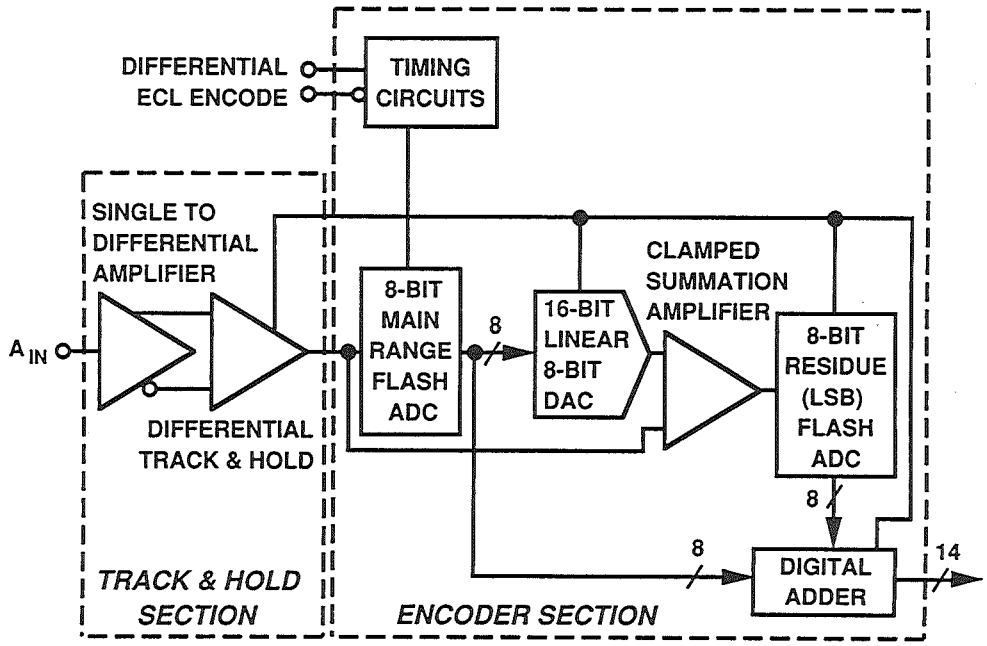

Figure 15.1

The SNR and SFDR of the AD9014 operating at 10, 12, and 14 bits is shown in Figure 15.2. Notice that the SFDR is relatively independent of the actual number of bits used, but the SNR drops to only 60dB at the 10-bit level. When selecting an ADC for spectral analysis applications, adding more bits does not always ensure better SFDR. On the other hand, there should be a sufficient number of bits present to ensure that samples from the ADC are not correlated to the input signal. There may still be correlation if the input frequency is an even sub-multiple of the sampling rate, which causes the quantization noise to be concentrated in the harmonics of the fundamental input signal frequency rather than being spread uniformly over the Nyquist bandwidth.

AD9014 SFDR AND SNR VERSUS FREQUENCY FOR 14, 12, AND 10 ACTIVE BITS

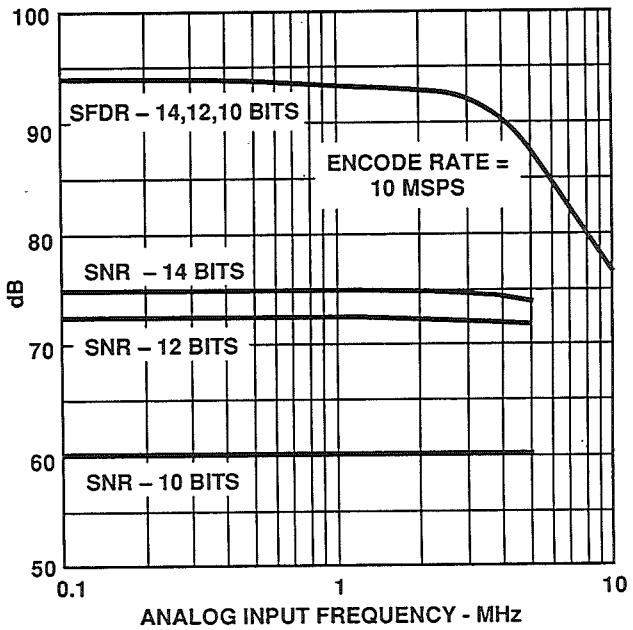

Figure 15.2

Achieving high SFDR with an ADC such as the AD9014 requires special attention to the drive amplifier. The circuit shown in Figure 15.3 is recommended in order to give ultra-low distortion levels. This circuit works well for analog input frequencies through 10MHz without degrading the AD9014's dynamic range. At 2.3MHz and 2V p-p output, all spurs are less than –100dBc. The signal path is through U3 and U4, which are set up in a series inverting configuration to cancel even-order harmonics that are generated as the loop gain diminishes with frequency. U1 and U2 supply the drive current for the outputs of U3 and U4, respectively. The closed-loop gain of U1 is 1.5 that of the closed-loop gain of U3. This gain, in conjunction with the 100Ω resistor ensures that U3 has to supply no dc output current. The value of R_p is chosen (depending on the load resistance) to ensure that U2 supplies all the output drive current of U4. The net effect is that the output stages of U3 and U4 are unloaded which minimizes their odd-harmonics. The overall gain of the driver circuit is +402Ω/R, and the input impedance is R/2.5. The output of the amplifier circuit is set up to drive either 2V p-p into 75Ω or 4V p-p into 150Ω by selecting R_p.

LOW DISTORTION DRIVE CIRCUIT (>100dBc) FOR AD9014 USING AD9617 OP AMPS

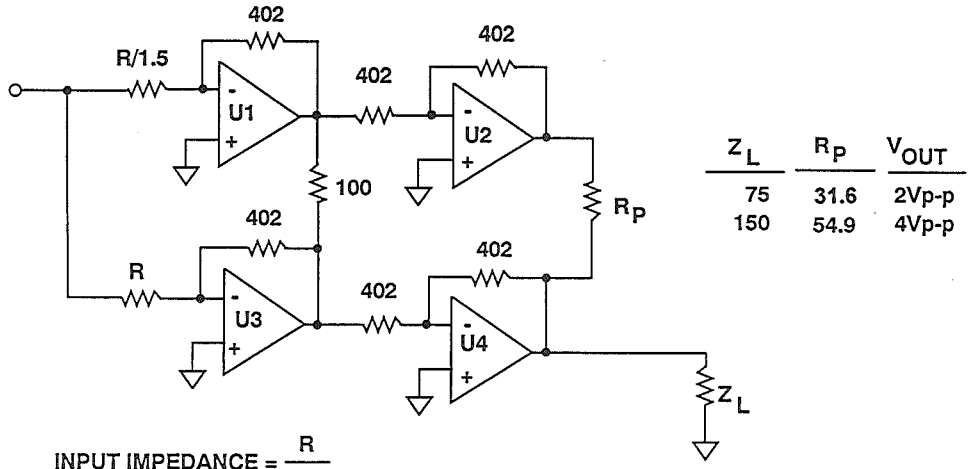

Z_L	R_P	V_{OUT}
75	31.6	2Vp-p
150	54.9	4Vp-p

INPUT IMPEDANCE = $\frac{R}{2.5}$

GAIN = $\frac{402}{R}$

RESISTOR VALUES IN Ω

Figure 15.3

AD9617 OUTPUT NOISE SPECTRAL DENSITY

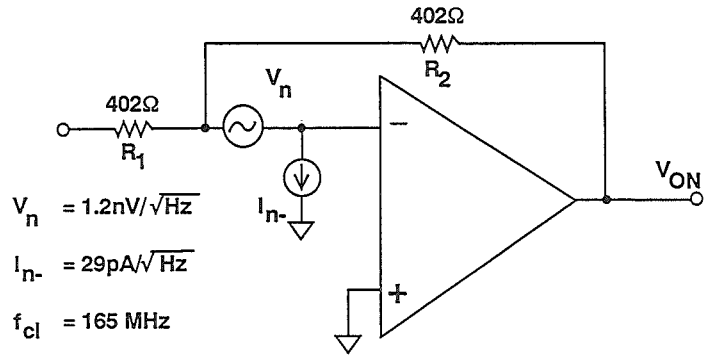

V_n = 1.2nV/$\sqrt{Hz}$

I_{n-} = 29pA/$\sqrt{Hz}$

f_{cl} = 165 MHz

$$V_{ON} = \sqrt{BW}\sqrt{V_n^2\left[1+\frac{R_2}{R_1}\right]^2 I_{n-}^2 R_2^2}$$

$$V_{ON} = \sqrt{BW}\sqrt{(1.2\times10^{-9})^2(2)^2 + (29\times10^{-12})^2(402)^2}$$

$$V_{ON} = \sqrt{BW}\sqrt{5.76\times10^{-18} + 136\times10^{-18}}$$

$$V_{ON} = \sqrt{BW}\ (11.9nV/\sqrt{Hz})$$

Figure 15.4

The output voltage noise spectral density calculations for a single AD9617 op amp is shown in Figure 15.4. Notice that the inverting input current noise dominates the output noise.

The approximate voltage noise spectral density for the four-amplifier driver circuit is twice that of a single AD9617, or 23.6nV/√Hz. Figure 15.5 shows the driver interfaced to the AD9014. If there is no filtering, the driver voltage noise must be integrated over the entire 60MHz input bandwidth of the AD9014. This yields an input noise of 230μV rms. However, an LSB at the input to the AD9014 is only 125μV. The 230μV rms input noise corresponds to a signal-to-noise ratio of 69.8dB for a 2V p-p input signal. The SNR specification of the AD9014 is 75dB, so the driver has become the limiting factor with respect to overall SNR performance.

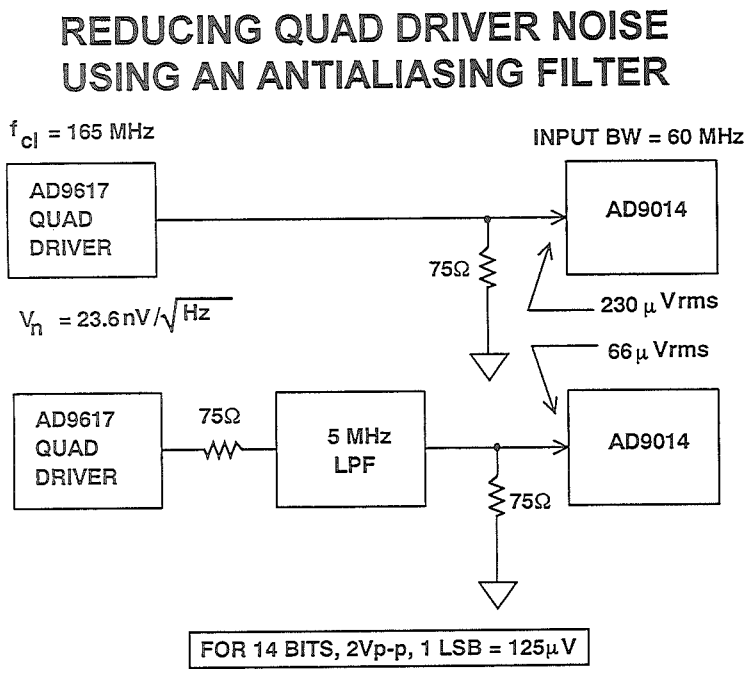

Figure 15.5

The noise contribution of the driver circuit can be significantly reduced by placing a 75Ω antialiasing filter between the driver and the AD9014 as shown in the lower half of Figure 15.5. the gain of the driver circuit must be increased by a factor of two to account for the attenuation of the filter series- and load-terminations. If the antialiasing filter is single-pole, and the cutoff frequency is 5MHz, the output noise of the AD9617 quad driver circuit is reduced to 66μV rms. This corresponds to a signal-to-noise ratio of 80.6dB, compared to a signal-to-noise ratio of 69.8dB with no filtering. Since 80.6dB > 75dB, the AD9014 is now the limiting noise contributor.

This example illustrates the significant noise reduction advantages of placing the antialiasing filter between the final driver stage and the ADC input rather than ahead of the ADC driver. This principle applies to most wide dynamic ADCs where the input bandwidth is significantly higher than the Nyquist frequency. If the extra bandwidth is not needed in the system, use the filter to reduce the high frequency noise output of the wideband drive amplifier.

USE OF DITHER SIGNALS TO INCREASE ADC DYNAMIC RANGE

Walt Kester

In the development of classical ADC quantization noise theory, the assumption is usually made that the quantization error signal is uncorrelated with the ADC input signal. If this is true, then the quantization noise appears as random noise spread uniformly over the Nyquist bandwidth, dc to $f_S/2$, and it has an rms value equal to $q/\sqrt{12}$. If, however, the input signal is locked to an non-prime integer sub-multiple of f_S, the quantization noise will no longer appear as uniformly distributed random noise, but instead will appear as harmonics of the fundamental input sinewave. This is especially true if the input is an exact even submultiple of f_S. Figure 15.6 illustrates the point using FFT simulation for an ideal 12 bit ADC. The FFT record length was chosen to be 4096. The spectrum on the left shows the FFT output when the input signal is an exact even submultiple (1/32) of the sampling frequency (the frequency was chosen so that there were exactly 128 cycles per record). The SFDR is approximately 78dBc. The spectrum on the right shows the output when the input signal is such that there are exactly 127 cycles per record. The SFDR is now about 92dBc which is an improvement of 14dB. Signal-correlated quantization noise is highly undesirable in spectral analysis applications, where it becomes difficult to differentiate between real signals and system-induced spurious components, especially when searching the spectrum for the presence of low-level signals in the presence of large signals.

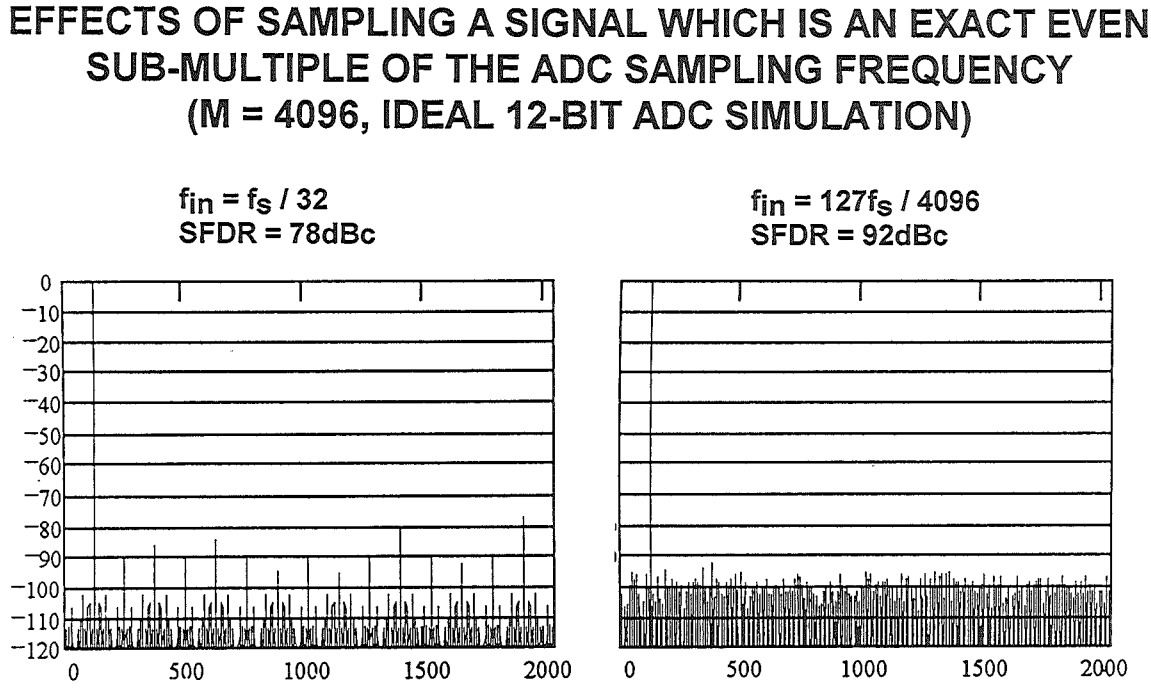

Figure 15.6

There are a number of ways to reduce this problem, but the easiest way is to add a small amount of broadband rms noise to the ADC input signal as shown in Figure 15.7. The rms value of this noise should be equal to about 1/2 LSB. The effect of this is to randomize the quantization noise and eliminate its possible signal-dependence. In many systems, there is usually enough random noise on the input signal and the sampling clock so that this happens automatically. This is especially likely when using high speed ADCs which have 12 or more bits of resolution and a relatively small input range of 2V p-p.

THE ADDITION OF GAUSSIAN WIDEBAND NOISE TO THE ADC INPUT RANDOMIZES QUANTIZATION NOISE AND REMOVES INPUT SIGNAL DEPENDENCE

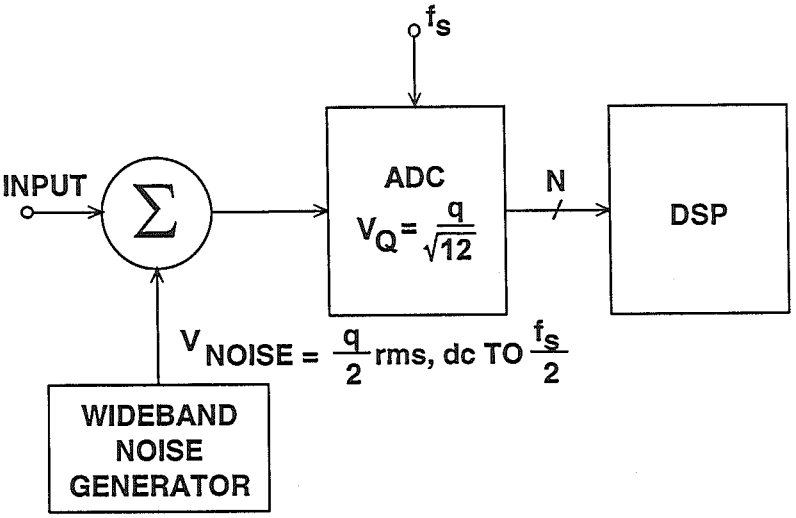

Figure 15.7

Another technique to achieve wider dynamic range in ADCs is the use of large-scale dithering as shown in Figure 15.8. This method serves to randomly distribute the ADC encoder static errors. A random number generator (RNG) provides a pseudo-random code which is loaded into a DAC and subtracted from the analog input signal. After the ADC processing is complete, the digital word is added to the resultant ADC digital word. The greater the randomization over the analog input range, the more linear the encoder becomes. However, this dynamic-range-expansion method also has its drawbacks. The encoder's SNR decreases as the linearization range increases, due to the reduced input levels. This is because the input signal level must now be chosen so that when the noise is added, the total signal never exceeds the ADC input range. The noise-generating DAC linearity also affects harmonics if the pseudo-random period is too small. In addition, this method requires the addition of several components with a resultant increase in system complexity.

USE OF LARGE-SCALE DIGITALLY-GENERATED DITHER SIGNAL TO ENHANCE ADC DYNAMIC RANGE

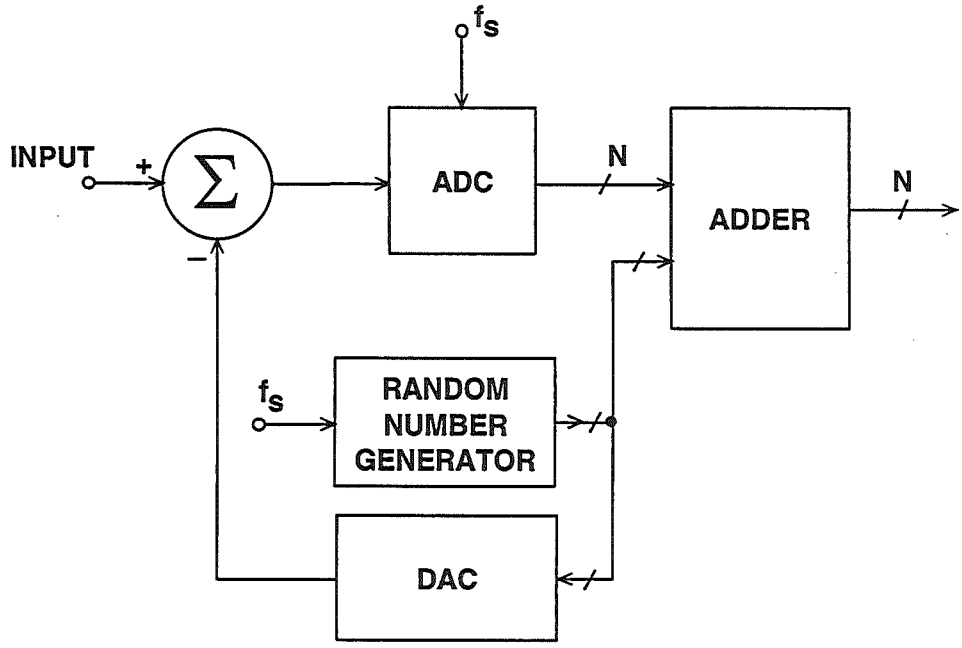

Figure 15.8

USING LOOK-UP TABLES TO CORRECT FOR ADC DYNAMIC ERRORS

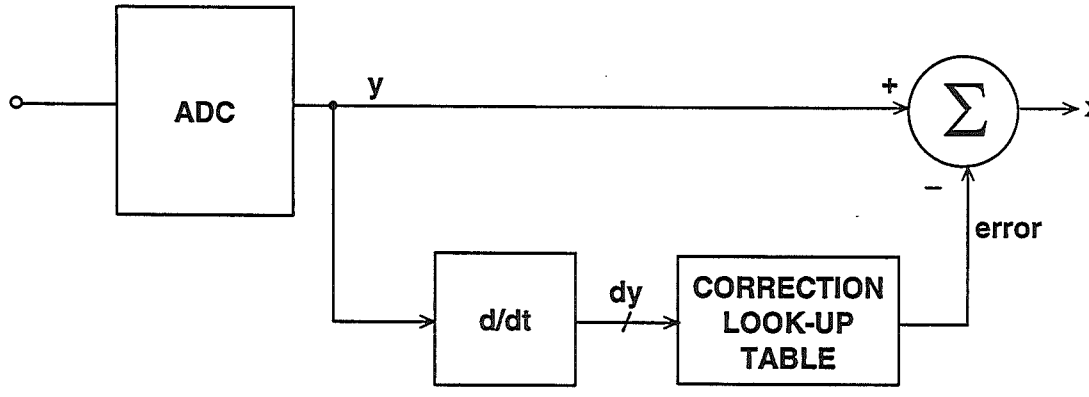

Figure 15.9

Some ADC users (see References 13, 14) have employed look-up tables (LUTs) to obtain higher performance from a given ADC. The technique consists of constructing a two-dimensional error correction table whereby the slewrate of the analog input signal serves as one axis of the table and the magnitude serves as the other axis (see Figure 15.9). Each address location holds the corrected digital output word. Once the LUT is constructed, a dedicated algorithm is applied to each digital word and to an address in the LUT to generate the corrected digital output word. This technique suffers several limitations. LUTs must be generated before data can be processed, which can be time consuming. Sources of nonlinearities within an ADC change as a function of temperature, requiring a new LUT when the temperature changes. Additional hardware, software, and circuit-board real estate to accommodate an LUT can be costly.

A final method for increasing dynamic range in ADCs for spectral analysis is the use of statistically-based continuous autocalibration techniques. The techniques involves a combination of a digitally-generated noise for error randomization and the use of continuously-updated LUTs to correct for internal encoder errors (see Reference 14). Obviously, this method involves a significant amount of additional digital circuitry and is impractical in a monolithic IC using current process technology.

DIRECT IF TO DIGITAL CONVERSION APPLICATIONS USING UNDERSAMPLING TECHNIQUES

Walt Kester, Allen Hill

Digital techniques have become widespread in radar receivers, broadband communications receivers, and mobile radio. A simplified block diagram of a traditional digital receiver using baseband sampling is shown in Figure 15.10. The mixer in the RF section of the receiver mixes the signal from the antenna with the RF frequency of the first local oscillator, LO1. The desired information is contained in relatively small bandwidth of frequencies Δf. In actual receivers, Δf may be as high as a few megahertz. The LO1 frequency is chosen such that the Δf band is centered about the IF frequency at the bandpass filter output. Popular IF frequencies are generally between 30 and 100MHz. The IF mixer then translates the Δf frequency band down to baseband where it is filtered and processed by a baseband ADC. Actual receivers generally have several stages of RF and IF processing, but the simple diagram serves to illustrate the concepts.

SIMPLIFIED DIGITAL RECEIVER USING BASEBAND SAMPLING

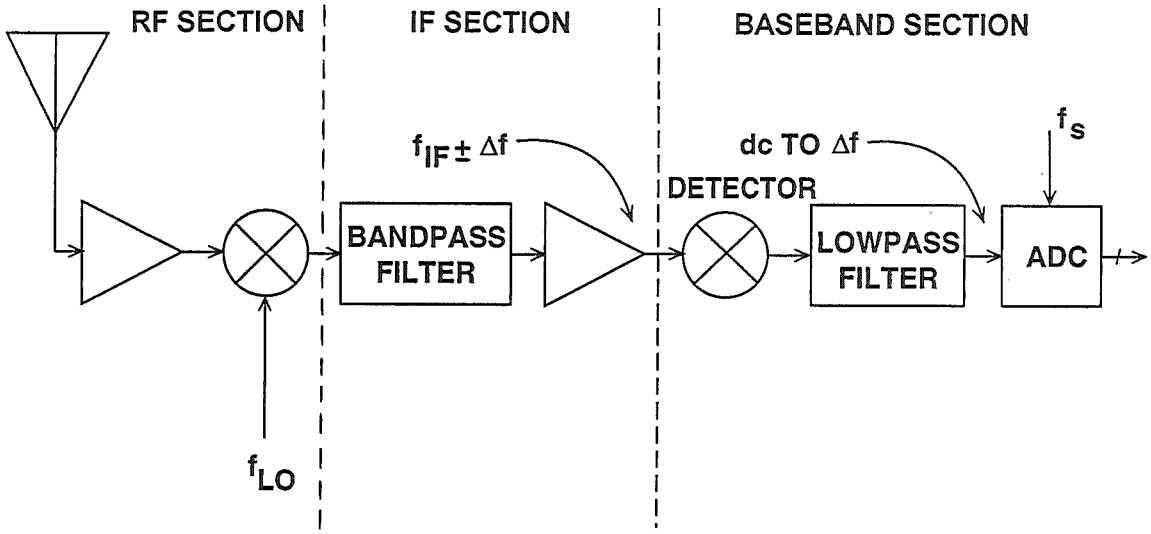

Figure 15.10

In a receiver which uses direct IF-to-digital techniques (IF sampling), the IF signal is applied directly to a wide bandwidth ADC as shown in Figure 15.11. The ADC sampling rate is chosen to be at least $2\Delta f$. The process of sampling the IF frequency at the proper rate causes one of the aliased components of Δf to appear in the dc to $f_S/2$ Nyquist bandwidth of the ADC output. DSP techniques can now be used to process the digital baseband signal. This approach may yield an improvement in overall signal-to-noise ratio by eliminating the detector stage. There is also more flexibility in the DSP because the ADC sampling rate can be shifted to tune the exact position of the Δf signal within the baseband. The obvious problem with this approach is that the ADC must now be able to accurately digitize signals which are well outside the dc to $f_S/2$ Nyquist bandwidth which most ADCs were designed to handle. Special techniques are available, however, which can extend the dynamic range of modern ADCs to include IF frequencies. Before examining the ADC problem, we will first look at the basic theory behind undersampling.

Figure 15.12 shows four cases where a signal having about a 1MHz bandwidth is located at different portions of the frequency spectrum. The minimum sampling rate required for no aliasing is also shown. In general, the sampling frequency must be at least twice the signal bandwidth, and the sampled signal must not cross an integer multiple of $f_S/2$.

SIMPLIFIED DIGITAL RECEIVER USING IF SAMPLING

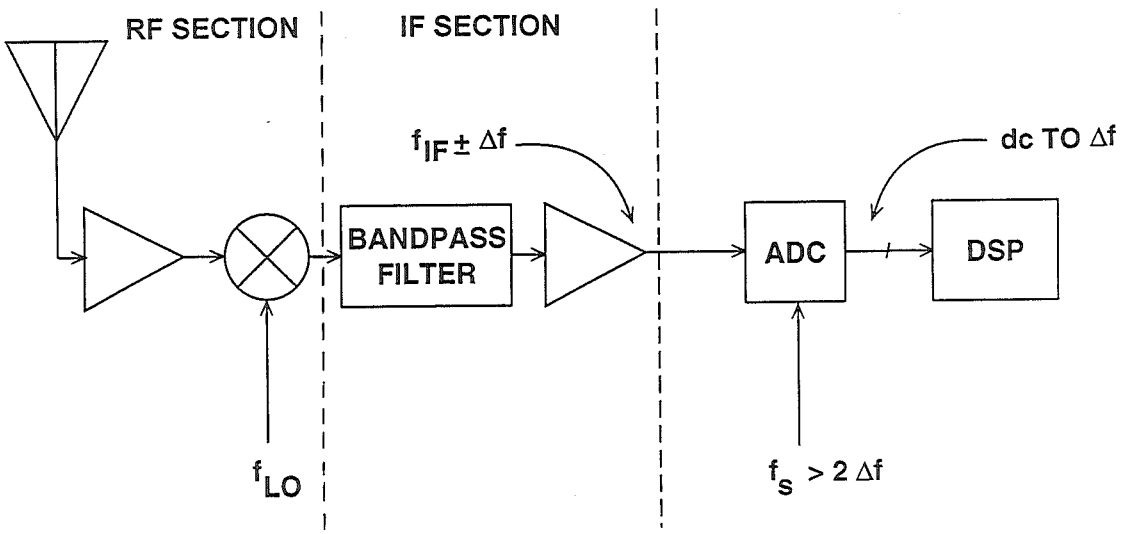

Figure 15.11

MINIMUM SAMPLING RATE REQUIRED FOR NO ALIASING OF A 1MHz BANDWIDTH SIGNAL

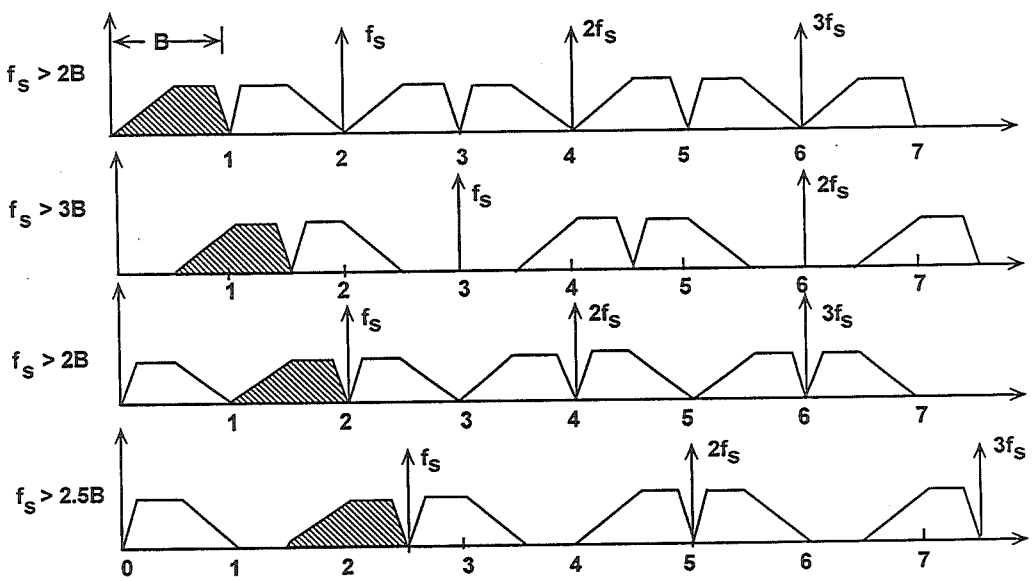

Figure 15.12

In the first case, the signal occupies a band from dc to 1MHz, and therefore must be sampled at greater than 2MSPS. The second case shows a 1MHz signal which occupies the band from 0.5 to 1.5MHz. Notice that this signal must be sampled at a minimum of 3MHz to avoid aliasing. In the third case, the signal occupies the band from 1 to 2MHz, and the minimum required sampling rate for no aliasing drops back to 2MHz. The last case shows a signal which occupies the band from 1.5 to 2.5MHz. This signal must be sampled at a minimum of 2.5MHz to avoid aliasing.

This analysis can be generalized as shown in Figure 15.13. The actual minimum required sampling rate is a function of the ratio of the highest frequency component to the total signal bandwidth.

Let us now examine a signal which occupies a band between 6 and 7MHz as shown in Figure 15.14. Assume the ADC sampling rate is 2MHz. Notice that the sampling process generates aliases of this signal around multiples of f_S. In the frequency spectrum, the alias component falling between 0 and 1MHz is an accurate representation of the original signal, assuming no ADC errors.

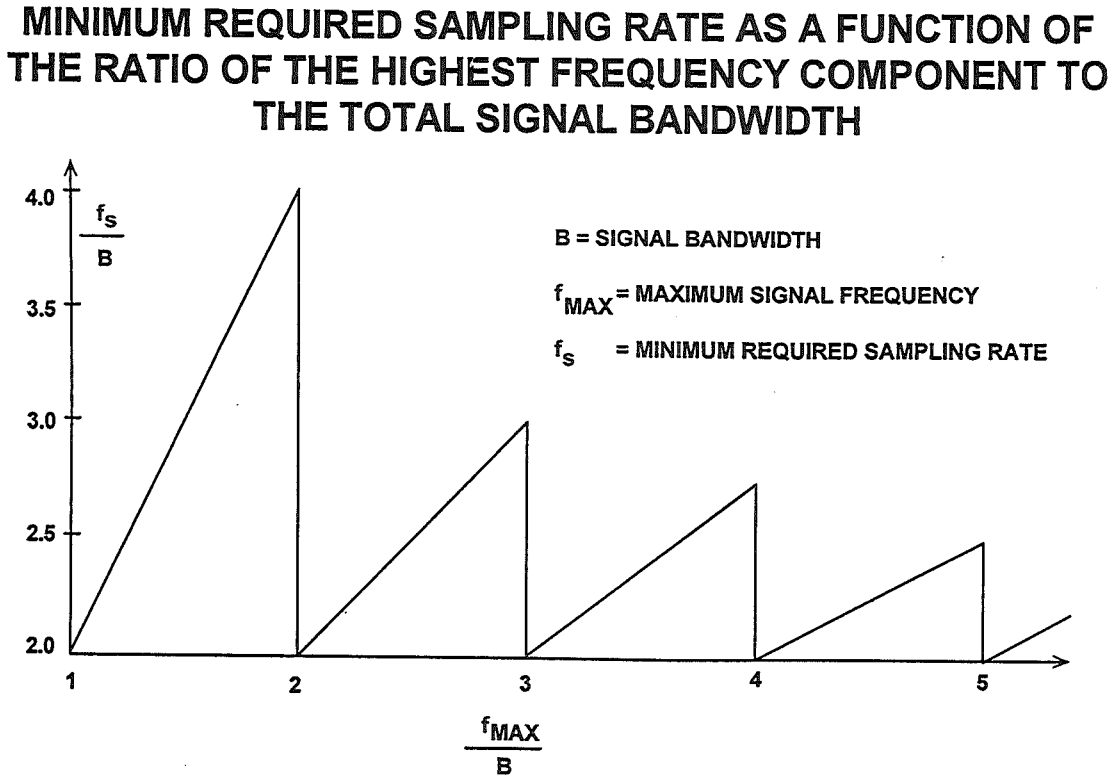

Figure 15.13

INTERMEDIATE FREQUENCY (IF) SIGNAL BETWEEN 6 AND 7 MHz IS ALIASED BETWEEN DC AND 1 MHz BY SAMPLING AT 2MSPS

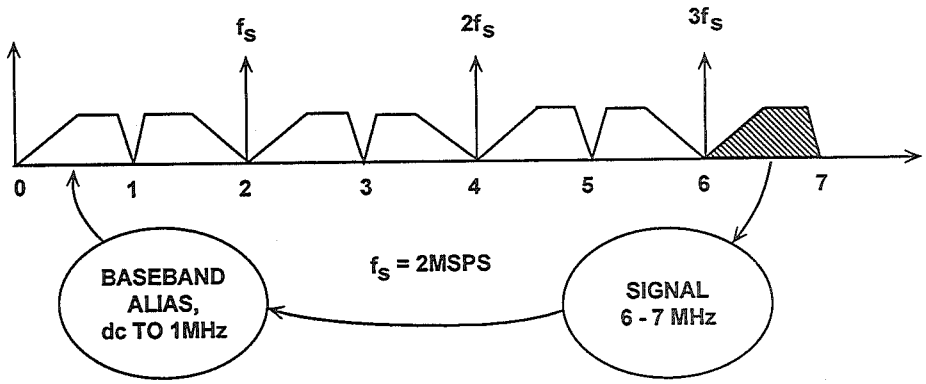

Figure 15.14

The above discussions show that although the concept of direct IF sampling is relatively straightforward, the implications for the ADC dynamic performance characteristics are significant.

Let us consider a typical example, where the IF frequency is 72.5MHz, and the desired signal occupies a bandwidth of 4MHz (B=4MHz), centered on the IF frequency (see Figure 15.15). We know from the previous discussion that the minimum sampling rate must be greater than 8MHz, probably on the order of 10MHz in order to prevent dynamic range limitations due to aliasing. If we place the sampling frequency at the lower band-edge of 70MHz (72.5–2.5), we will definitely recover the aliased component of the signal in the dc to 5MHz baseband. There is, however, no need to sample at this high rate, so we may choose any sampling frequency 10MHz or greater which is an integer sub-multiple of 70MHz, i.e, 70÷2 = 35.000MHz, 70÷3 = 23.333MHz, 70÷4 = 17.500MHz, 70÷5 = 14.000MHz, 70÷6 = 11.667MHz, or 70÷7 = 10.000MHz. We will therefore choose the lowest possible sampling rate of 10.000MHz (70÷7).

INTERMEDIATE FREQUENCY (IF) SIGNAL AT 72.5MHz (±2MHz) IS ALIASED BETWEEN DC AND 5MHz BY SAMPLING AT 10MSPS

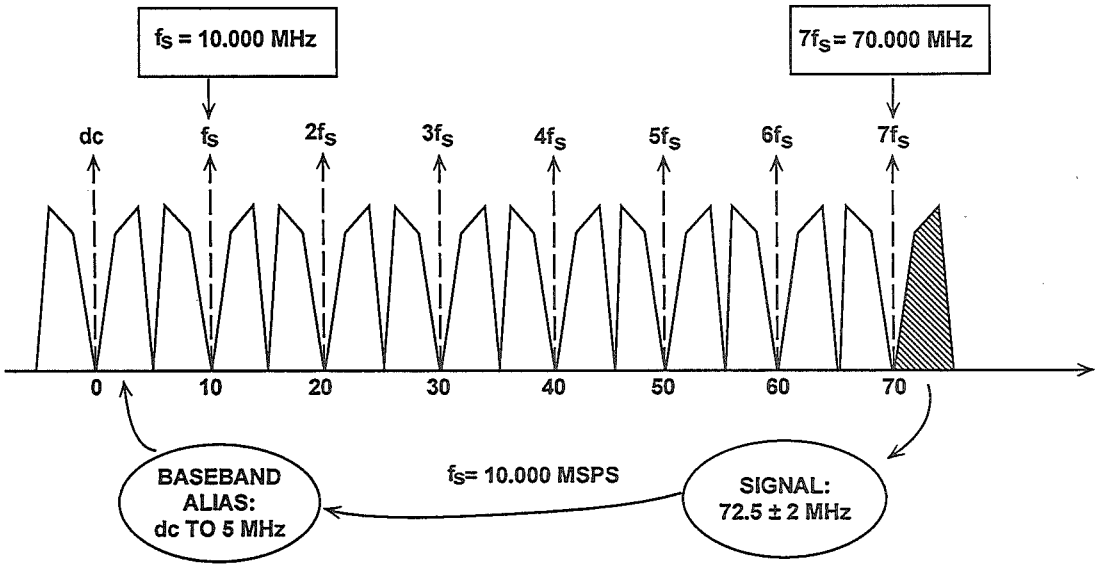

Figure 15.15

SPURIOUS FREE DYNAMIC RANGE OF THE AD9022 12-BIT, 20MSPS ADC FOR f_S = 10MSPS

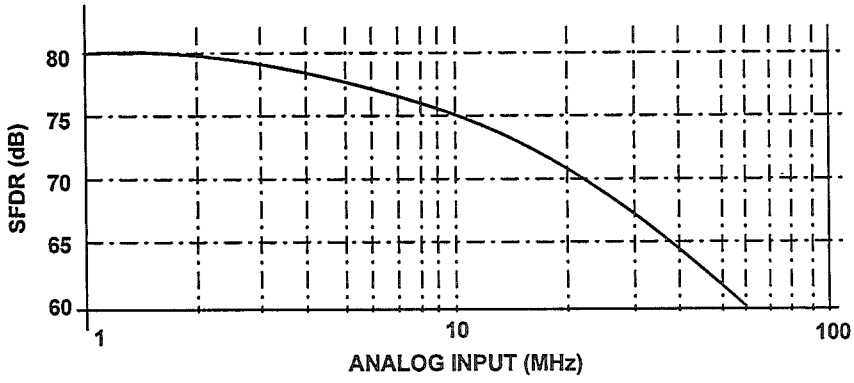

Figure 15.16

There is an advantage in choosing a sampling frequency which is an *odd* sub-multiple of the lower band-edge in that there is no frequency inversion in the baseband alias as would be the case for a sampling frequency equal to an *even* sub-multiple of the lower band-edge. (Frequency inversion can be easily dealt with in the DSP software should it occur, so the issue is not very important.)

The next step is to select an ADC which has sufficient dynamic range with a 70 to 75MHz input to meet our system requirement. In most radar receivers, a SFDR of 60 to 80dB is desirable. Unfortunately, this requirement will not be met with standard Nyquist-sampling ADCs due to degradations which occur at high input frequencies. Figure 15.16 shows the SFDR of the AD9022 12 bit, 20MSPS ADC, which represents one of the best monolithic designs available. Note that at 1MHz the SFDR is 80dB, but at the IF frequency of 70MHz, the SFDR of the device is less than 60dB.

Meeting this performance requirement of 80dB SFDR requires the addition of an external wide-bandwidth, low distortion sample-and-hold, such as the AD9100 as shown in Figure 15.17. The external SHA serves to hold the signal constant during the ADC conversion cycle. The ADC sees a dc value during the hold-time of the external SHA. The process of optimizing the design for the best SFDR is not simple, and involves many tradeoffs.

THE ADDITION OF AN EXTERNAL WIDEBAND LOW DISTORTION SHA EXTENDS THE LOW FREQUENCY PERFORMANCE OF THE ADC TO HIGHER FREQUENCIES

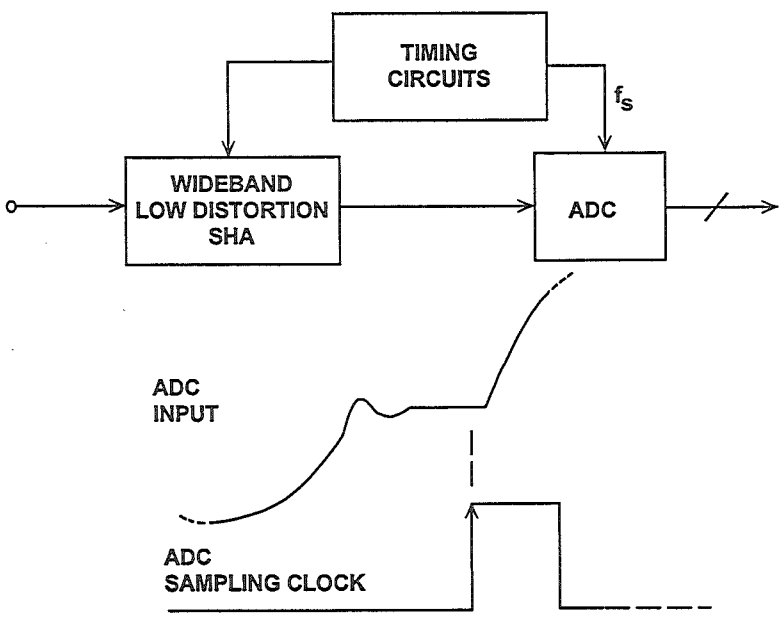

Figure 15.17

The first step in the process is to select an ADC which has sufficient SFDR at low frequencies. The low frequency distortion of an ADC is an indicator of the inherent non-linearity in the dc transfer function. The addition of an ideal external SHA will do nothing to improve on this basic performance. The best you can expect the SHA to do is to extend the low frequency performance of the ADC to higher frequencies. Because of its excellent low frequency distortion (–80dBc @ 1MHz), the AD9022 is good choice.

The next step is to select a low distortion SHA which will maintain sufficient dynamic performance at the IF frequency. Most SHAs are specified for distortion when operating in the track mode. What is of real interest, however, is the signal distortion in the hold mode when the SHA is operating dynamically. The AD9100 and AD9101 are ultra-fast SHAs and are specified in terms of hold-mode distortion. The measurement is done using a high performance low distortion ADC (such as the AD9014 14-bit, 10MSPS) to digitize the held value of the SHA output. An FFT is performed on the ADC output, and the distortion is measured digitally. For sampling rates greater than 10MSPS, the ADC is clocked at an integer sub-multiple of the SHA sampling frequency. This causes a frequency translation in the FFT output because of undersampling, but the distortion measurement still represents that of the SHA operating at the higher sampling rate.

The AD9100 is optimized for low distortion operation up to 30MSPS, while the AD9101 will provide low distortion performance up to a sampling rate of 125MSPS.
A block diagram of the AD9100 track-and-hold is shown in Figure 15.18. The switching bridge is integrated into the first stage closed-loop input amplifier. This innovation provides error (distortion) correction for both the switch and amplifier, while still achieving slewrates representative of a more traditional open-loop design. The device has a 250MHz input bandwidth and a 16ns acquisition time to 0.01%. Hold-mode distortion performance versus analog input frequency for a sampling rate of 30MSPS and a 2V p-p output is shown in Figure 15.19.

The hold-mode distortion shown in Figure 15.19 is clearly not good enough for the 70MHz IF application we are considering. However, the distortion performance of the AD9100 can be greatly improved by reducing the input signal level. This will decrease the signal-to-noise ratio, but will also reduce the distortion produced by the switching bridge nonlinearity.

The test configuration shown in Figure 15.20 was used to collect the AD9100 performance data for low amplitude input signals. The data shown in Figure 15.21 was taken at a 10MSPS sampling rate for three input amplitudes. For each amplitude, the gain of the AD9618 op amp was adjusted so that its output exactly filled the 2V p-p input range of the AD9014 ADC. Notice that the SFDR is optimum (70dBc) for a 200mV p-p input signal to the AD9100, and a corresponding post-amplifier gain of 10.

AD9100 MONOLITHIC 30MSPS TRACK-AND-HOLD

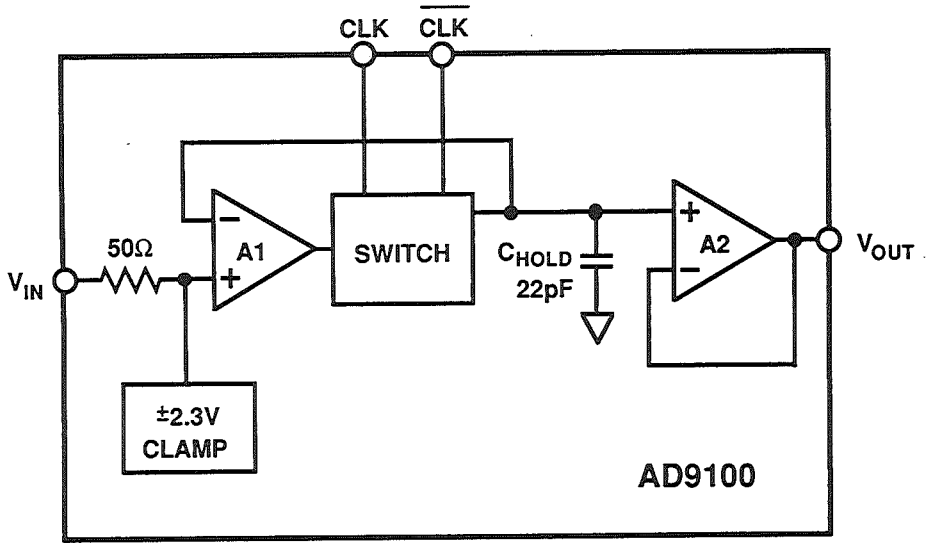

Figure 15.18

AD9100 HOLD-MODE SFDR VERSUS INPUT FREQUENCY FOR 30MSPS SAMPLING RATE, 2V P-P OUTPUT

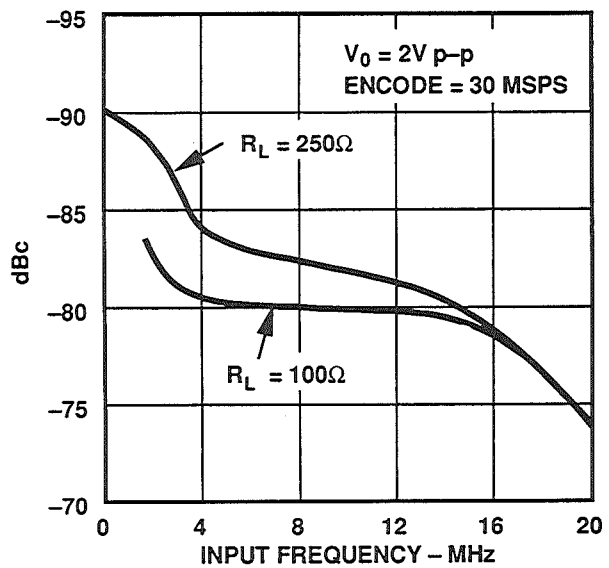

Figure 15.19

TEST CONFIGURATION AND TIMING FOR MEASURING AD9100 HOLD-MODE SFDR AT 10MSPS

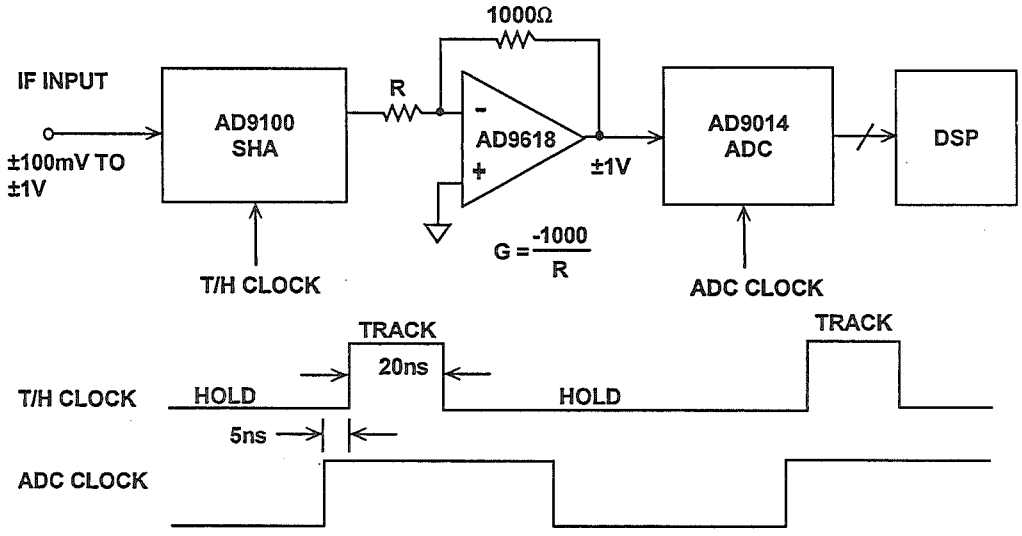

Figure 15.20

AD9100 HOLD-MODE SFDR MEASURED WITH AD9014 ADC SAMPLING AT 10MSPS

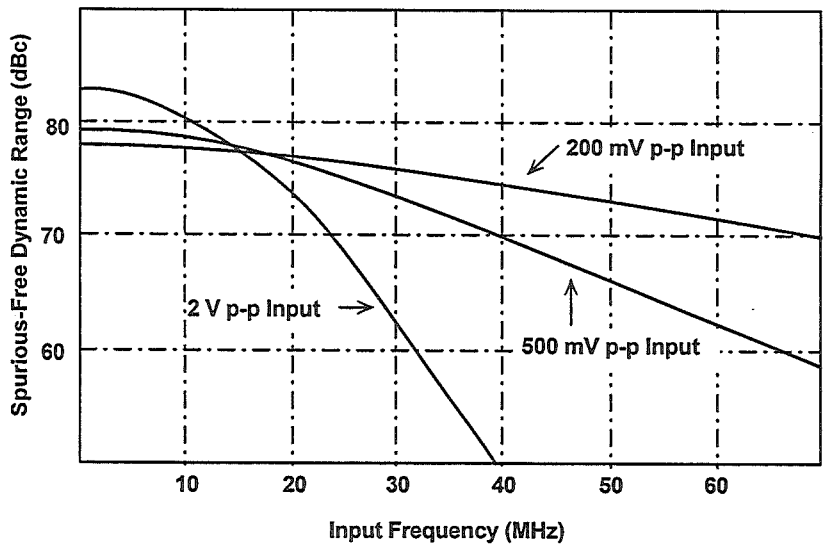

Figure 15.21

The performance of the AD9100 driving the AD9022 under identical conditions yielded the FFT spectrum shown in Figure 15.22. The 72MHz SFDR is greater than 70dBc, and the measured SNR is 62dB.

The tradeoff in optimizing the circuit for the best SFDR is the increase in overall noise resulting from the high-gain post-amplifier amplifying the hold-mode noise of the AD9100.

The other part of a successful wide-dynamic range design involves the optimization of the timing between the SHA and the ADC. This involves even more tradeoffs. The SHA acquisition time should be long enough to achieve the desired accuracy, but short enough to allow sufficient hold-time for the ADC front-end to settle and yield a low-distortion conversion. For the example above, the optimum performance was achieved using an acquisition time of 20ns and a track time of 80ns. As shown in Figure 15.20, the ADC is clocked close to the end of the SHA's hold time. Best performance in designs such as these is always achieved by optimizing the timing in the actual circuit.

Similar dynamic range improvements can be achieved with high speed flash converters at higher sampling rates using the AD9101 SHA whose block diagram is shown in Figure 15.23. The AD9101 is a track-and-hold with an internal post-amplifier. This configuration allows the front end sampler to operate at relatively low signal amplitudes, resulting in dramatic improvement in hold-mode distortion at high input frequencies and sampling rates up to 125MSPS. The AD9101 has an input bandwidth of 350MHz and an acquisition time of 7ns to 0.1% and 11ns to 0.01%.

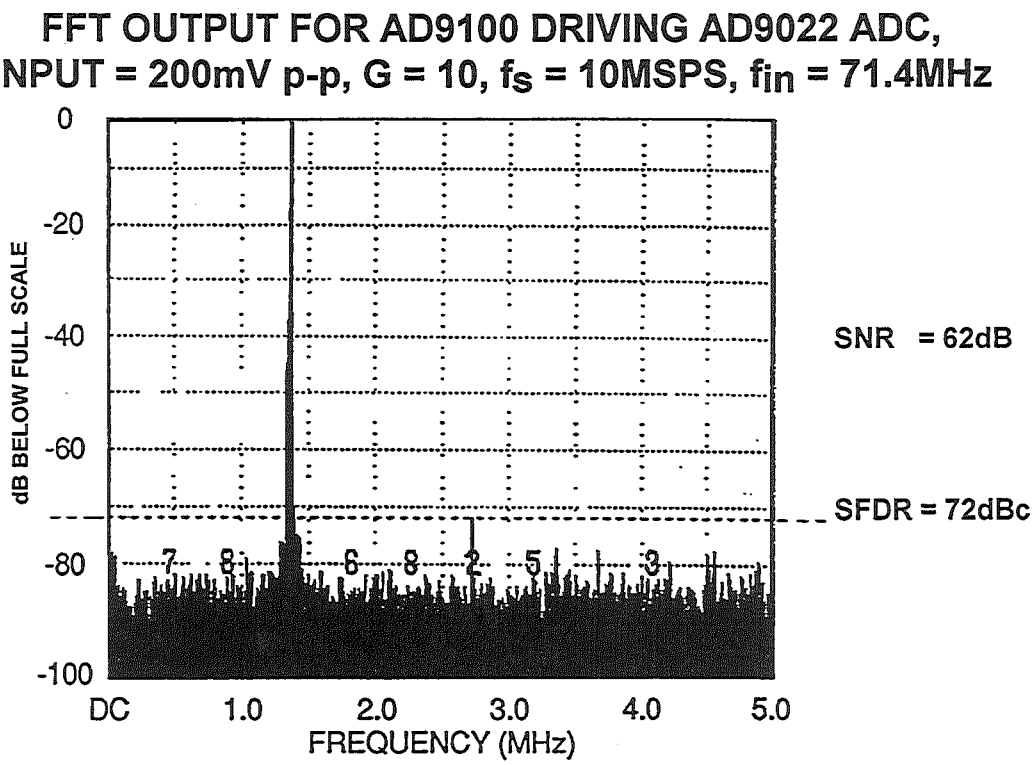

Figure 15.22

AD9101 125MSPS SAMPLING AMPLIFIER

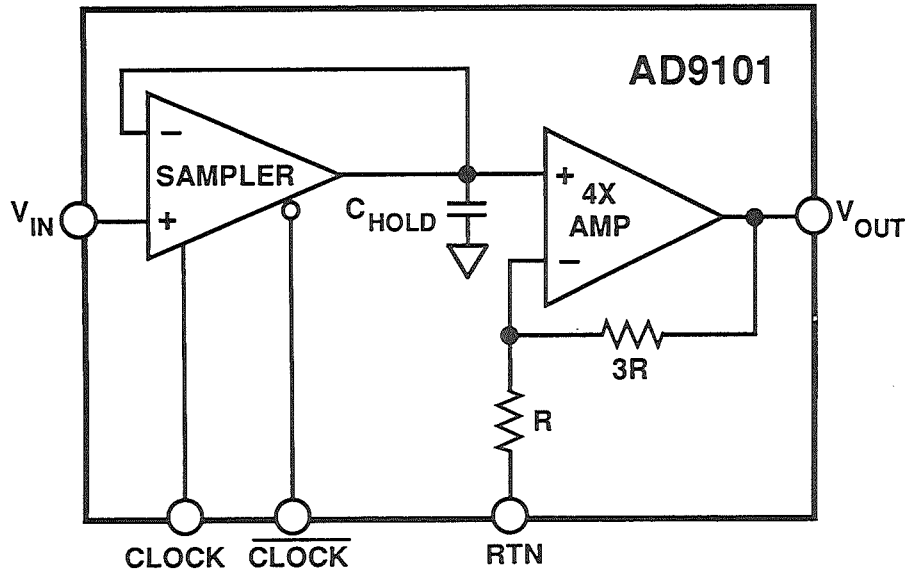

Figure 15.23

A block diagram and a timing diagram is shown for the AD9101 driving the AD9002 8 bit flash converter at 125MSPS (see Figure 15.24). The corresponding dynamic range with and without the AD9101 is shown in Figure 15.25.

AD9101 SHA DRIVING AD9002 8-BIT, 125MSPS FLASH CONVERTER FOR IMPROVED DYNAMIC RANGE

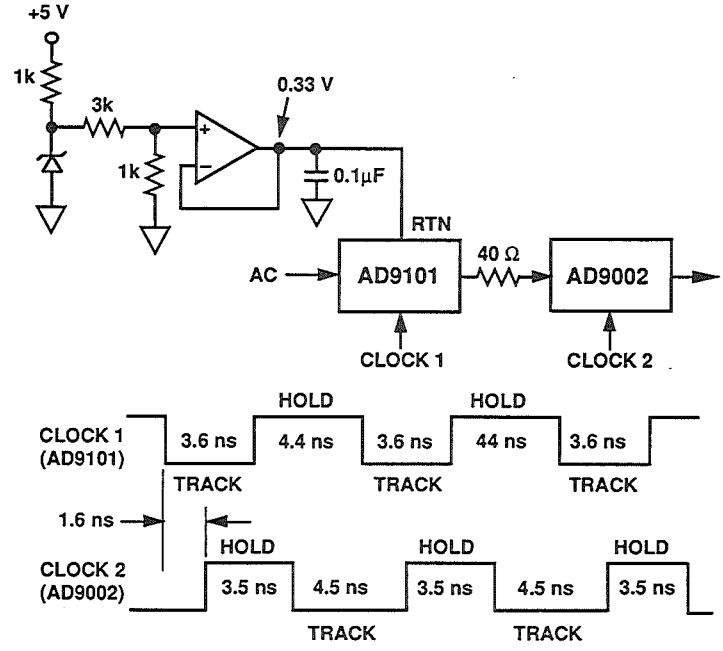

Figure 15.24

AD9002 DYNAMIC PERFORMANCE WITH AND WITHOUT AD9101

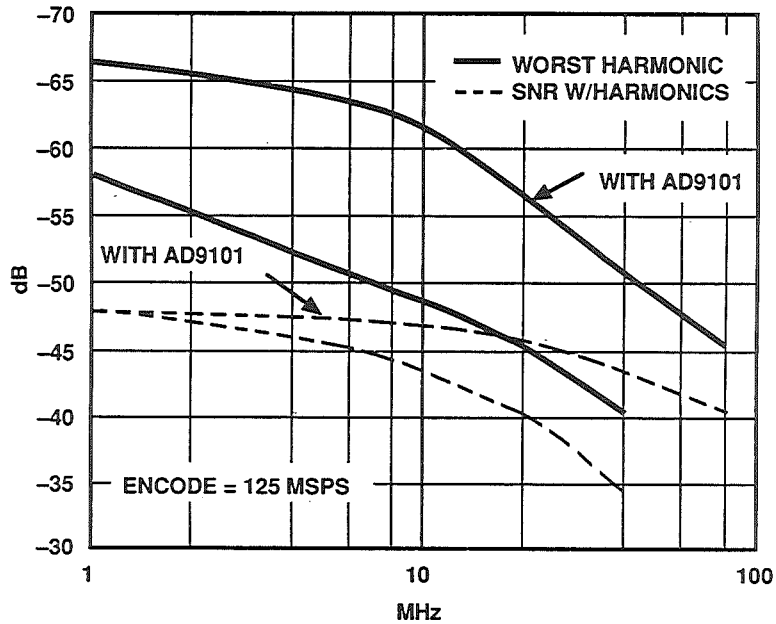

Figure 15.25

The increase in performance is also dramatic for the AD9060 10 bit flash converter. The block diagram and timing is shown in Figure 15.26, and the corresponding performance comparison in Figure 15.27.

AD9101 SHA DRIVING AD9060 10-BIT FLASH CONVERTER AT 60MSPS

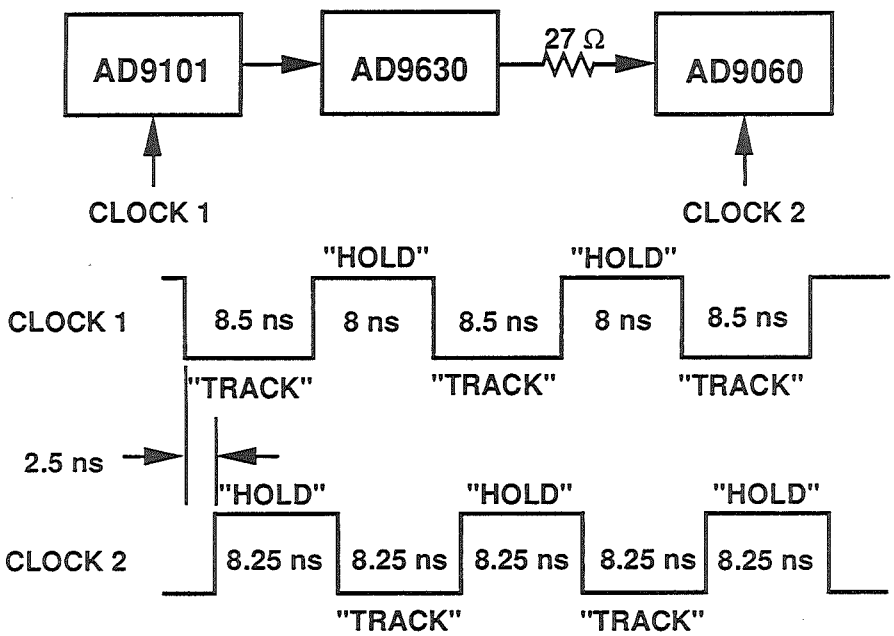

Figure 15.26

15

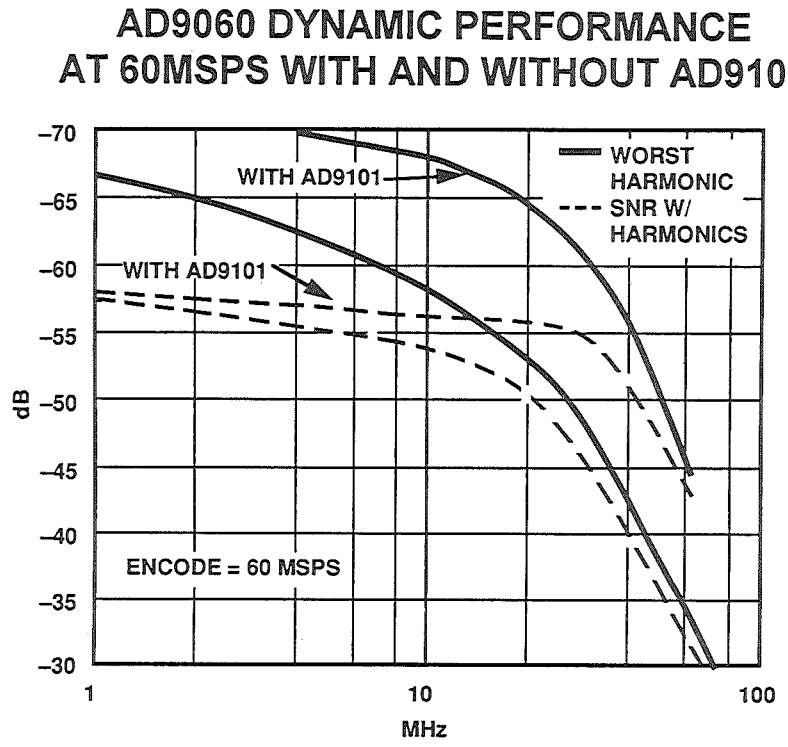

Figure 15.27

BANDPASS SIGMA-DELTA TECHNIQUES FOR DIRECT IF-TO-DIGITAL CONVERSION

James Bryant

The sigma-delta ADCs that were discussed in the previous section contain integrators, which are lowpass filters. They therefore have a passband extending from DC, and the quantization noise is pushed up in frequency. At present all commercially available sigma-delta ADCs are of this type (although some which are intended for use in audio or telecommunications applications contain bandpass digital filters to eliminate any DC response).

There is no particular reason why the filters of the sigma-delta modulator should be lowpass filters, except that traditionally ADCs have been thought of as being baseband devices, and that integrators are somewhat easier to construct than bandpass filters. If we replace the integrators in a sigma-delta ADC with bandpass filters as shown in Figure 15.28, the quantization noise is moved up and down in frequency to leave a virtually noise-free region in the passband (see Figure 15.29. If the digital filter is then programmed to have its passband in the region, we have a sigma-delta ADC with a bandpass, rather than a lowpass characteristic.

REPLACING INTEGRATORS WITH BANDPASS FILTERS GIVES A BANDPASS SIGMA-DELTA ADC

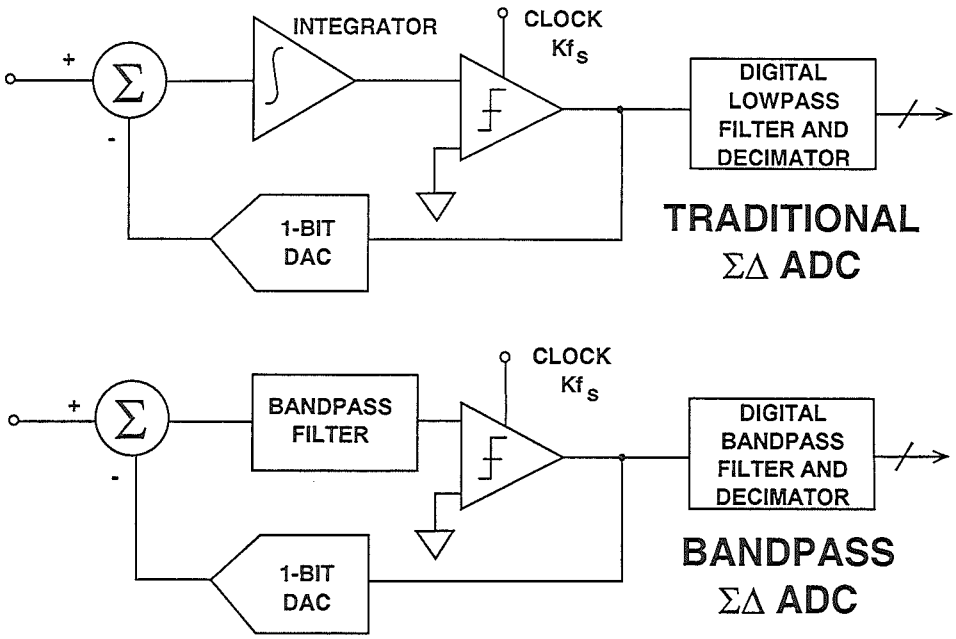

Figure 15.28

NOISE SHAPING FUNCTIONS FOR TRADITIONAL AND BANDPASS SIGMA-DELTA ADCs

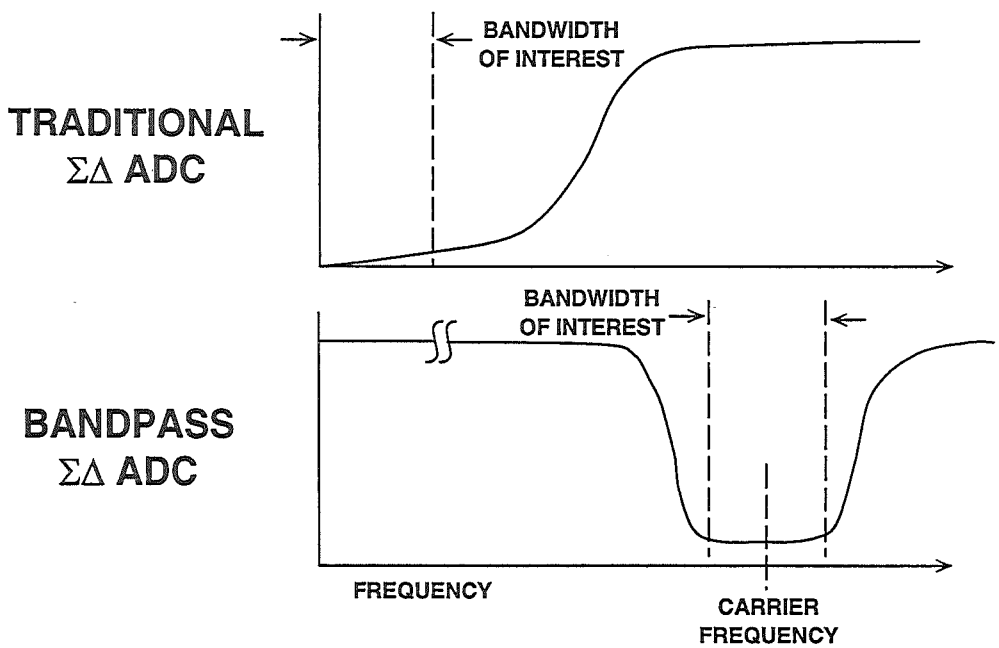

Figure 15.29

The theory is straightforward, but the development of a sigma-delta ADC is expensive, and there is no universal agreement on ideal characteristics for such a bandpass sigma-delta ADC, so developing such a converter from scratch to verify the theory would be unlikely to yield a commercial product. Researchers at Analog Devices and the University of Toronto (See References 16, 17, and 18) have therefore modified a commercial baseband (audio) sigma-delta ADC chip by rewiring its integrators (see Figure 15.30) as switched capacitor bandpass filters and reprogramming its digital filter and decimator. This has provided a fast, and comparatively inexpensive, proof of the concept, but at the expense of relative low Effective Bits (11-bits), the result of less than ideal bandpass filters. Nevertheless the results are extremely encouraging and open the way to the design of purpose-built bandpass sigma-delta ADC chips for specific ASIC applications, especially, but not exclusively, radio receivers.

The modulator configuration is shown in Figure 15.30, and the overall performance characteristics of the experimental ADC are shown in Figure 15.31. The device was designed to digitize the popular radio IF frequency of 455kHz.

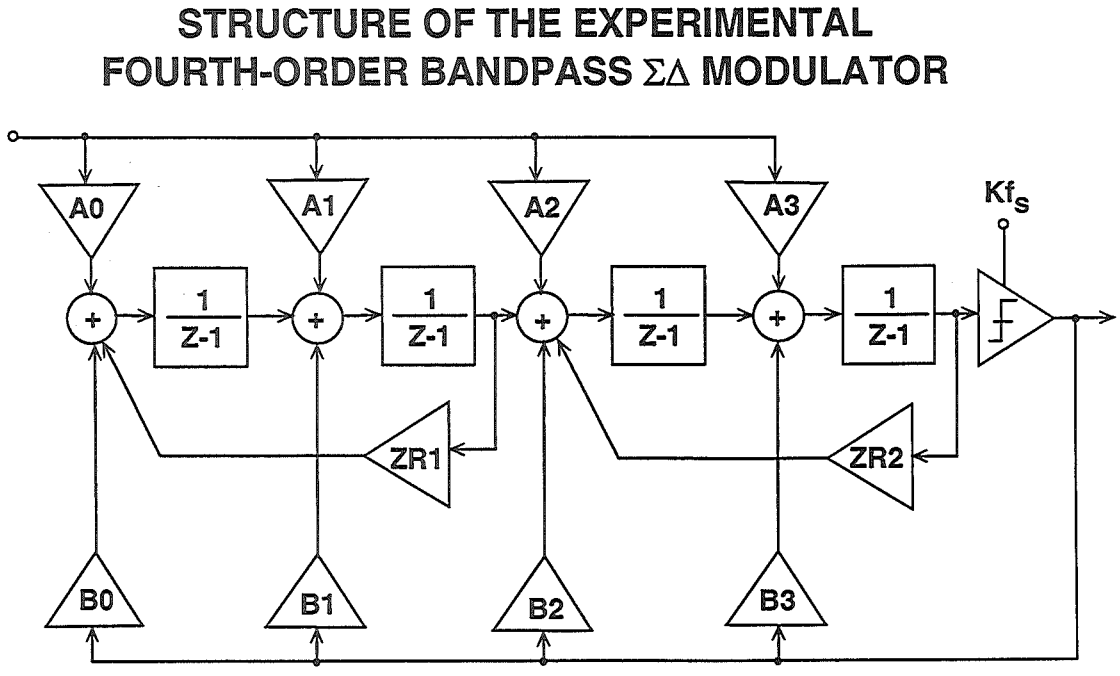

Figure 15.30

SUMMARY OF RESULTS FOR EXPERIMENTAL BANDPASS SIGMA-DELTA ADC

- Center Frequency: 455kHz
- Bandwidth: 10kHz
- Sampling Rate: 1.852MSPS
- Oversampling Ratio: 91
- SNR in Specified Band: 65dB
- Supply: ±5V, Power: 750mW
- Process: 3µm CMOS, Active Area: 1.8 x 3.4mm

Figure 15.31

In the future it may be possible to have such bandpass sigma-delta ADCs with user-programmable digital filter coefficients, so that the passband of a receiver could be modified during operation in response to the characteristics of the signal (and the interference!) being received. Such a function is very attractive, but difficult to implement, since it would involve loading, and storing, several hundreds or even thousands of 16-22 bit filter coefficients, and would considerably increase the size, and cost, of the converter.

A feature which could be added comparatively easily to a sigma-delta ADC is a more complex digital filter with separate reference (I) and quadrature (Q) outputs. Such a feature would be most valuable in many types of radio receivers.

ACHIEVING WIDE DYNAMIC RANGE USING VARIABLE GAIN AMPLIFIERS

Walt Kester, Barrie Gilbert, Bob Clarke

Another method to extend the dynamic range in a system is to precede the ADC with a wide bandwidth, variable gain amplifier (VGA). There are many ways to make VGAs, and some relatively good ones can be built using analog multipliers such as the AD834, 500MHz 4-quadrant multiplier shown in Figure 15.32. A diagram of a 90MHz voltage controlled amplifier (VCA) is shown in Figure 15.33, and the corresponding frequency response in Figure 15.34. The signal is applied to the Y input of the AD834, and the control voltage to the X input. The AD811 transimpedance amplifier level-shifts the differential outputs of the AD834 to a single-ended signal and acts as a current-to-voltage converter.

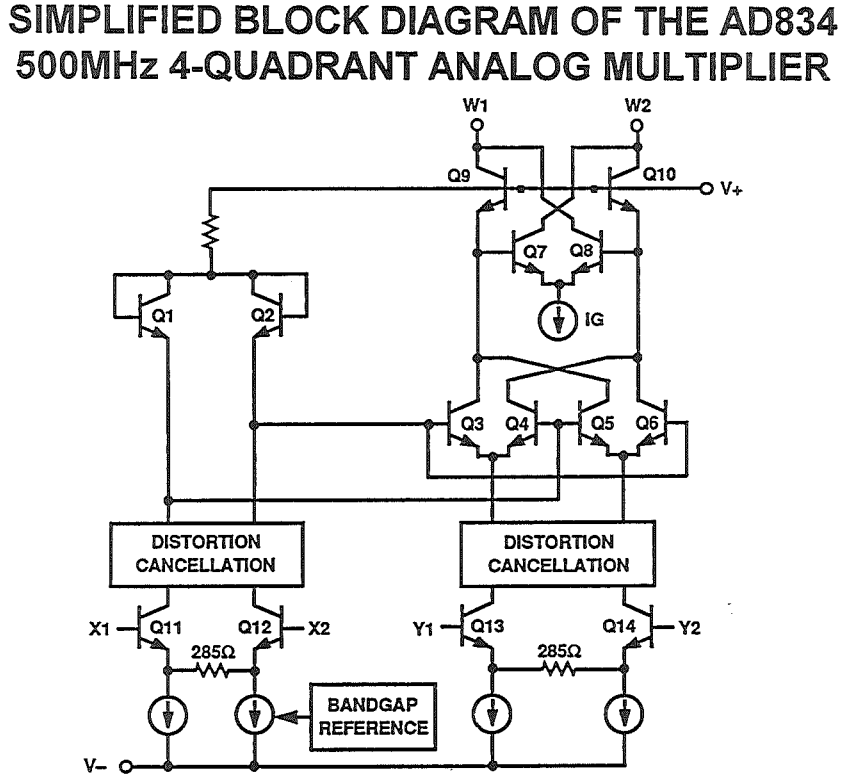

Figure 15.32

A 90MHz VCA USING THE AD834 AND THE AD811

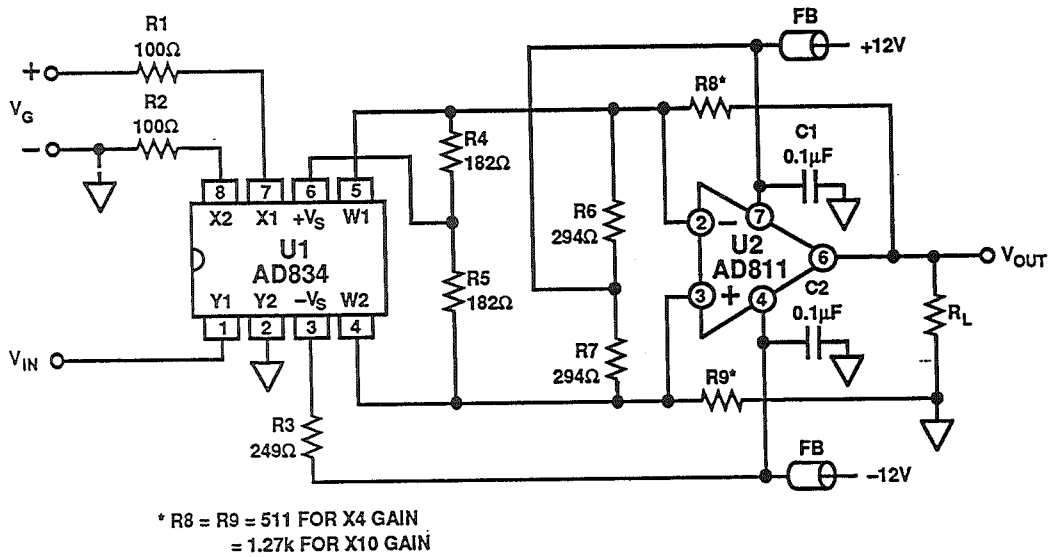

Figure 15.33

FREQUENCY RESPONSE OF THE 90MHz VCA

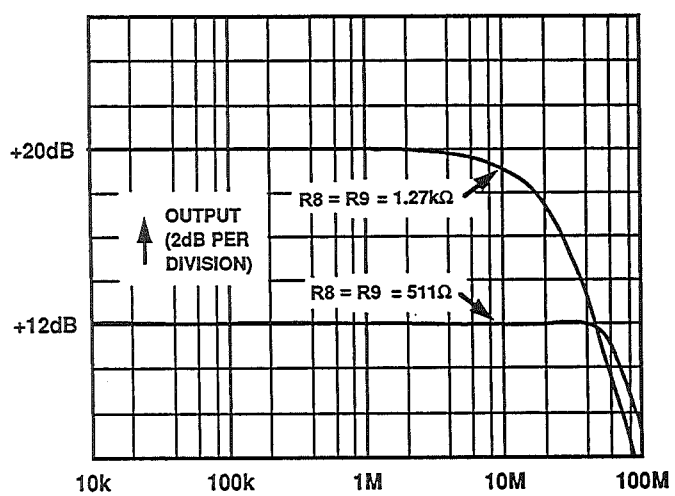

Figure 15.34

Most VCAs made with analog multipliers have gain which is *linear in volts* with respect to the control voltage, moreover they tend to be noisy. There is a demand, however, for a VCA which combines a wide gain range with constant bandwidth and phase, low noise with large signal-handling capabilities, and low distortion with low power consumption, while providing accurate, stable, *linear-in-dB* gain. The AD600 and AD602 achieve these demanding and conflicting objectives with a unique and elegant solution - the X-AMP™ (for *exponential amplifier*). The concept is simple: a fixed-gain amplifier follows a passive, broadband attenuator equipped with special means to alter its attenuation under the control of a voltage (see Figure 15.35). The amplifier is optimized for low input noise, and negative feedback is used to accurately define its moderately high gain (about 30 to 40dB) and minimize distortion. Since this amplifier's gain is fixed, so also are its ac and transient response characteristics, including distortion and group delay; since its gain is high, its input is never driven beyond a few millivolts. Therefore, it is always operating within its small signal response range.

SINGLE CHANNEL OF THE DUAL 30MHz AD600/AD602 X-AMP

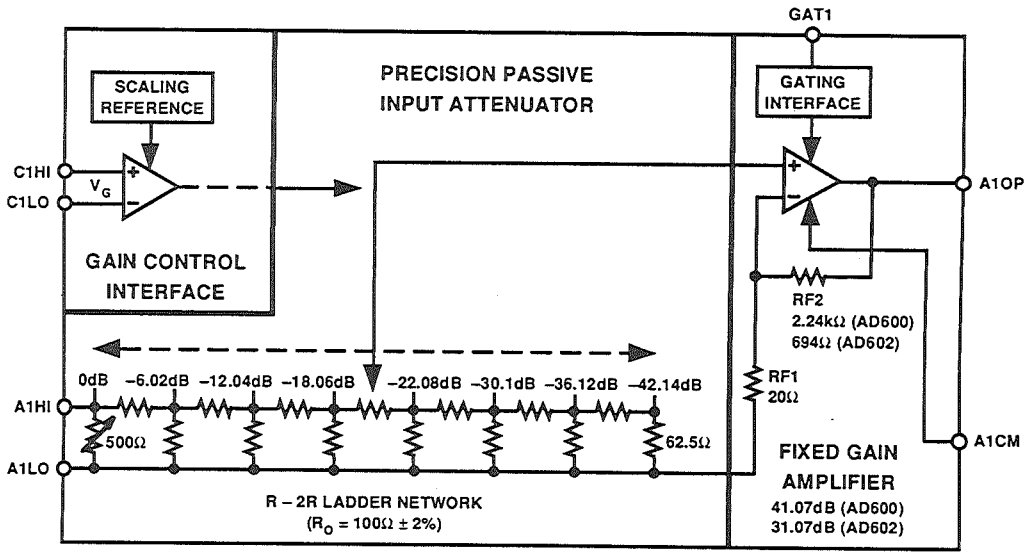

Figure 15.35

The attenuator is a 7-section (8-tap) R-2R ladder network. The voltage ratio between all adjacent taps is exactly 2, or 6.02dB. This provides the basis for the precise linear-in-dB behavior. The overall attenuation is 42.14dB. As will be shown, the amplifier's input can be connected to any one of these taps, or even *interpolated* between them, with only a small deviation error of about ±0.2dB. The overall gain can be varied all the way from the fixed (maximum) gain to a value 42.14dB less. For example, in the AD600, the fixed gain is 41.07dB (a voltage gain of 113); using this choice, the full gain range is –1.07dB to +41.07dB. The gain is related to the control voltage by the relationship $G_{dB} = 32V_G + 20$ where V_G is in volts. For the AD602, the fixed gain is 31.07dB (a voltage gain of 35.8), and the gain is given by $G_{dB} = 32V_G + 10$.

The gain at $V_G = 0$ is laser trimmed to an absolute accuracy of ±0.2dB. The gain scaling is determined by an on-chip bandgap reference (shared by both channels), laser trimmed for high accuracy and low temperature coefficient. Figure 15.36 shows the gain versus the differential control voltage for both the AD600 and the AD602.

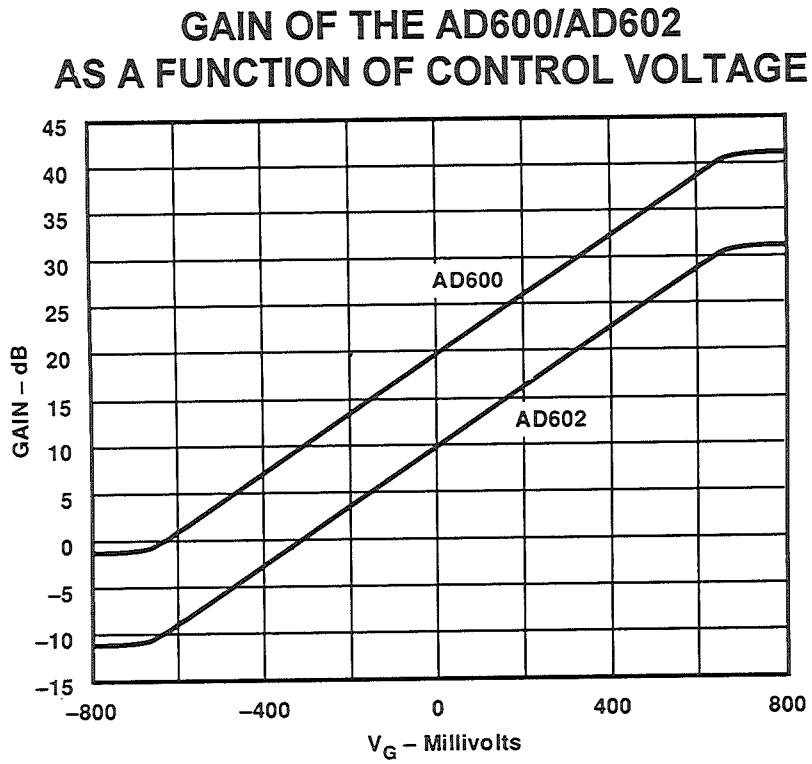

Figure 15.36

In order to understand the operation of the AD600/AD602, consider the simplified diagram shown in Figure 15.37. Notice that each of the eight taps is connected to an input of one of eight bipolar differential pairs, used as current-controlled transconductance (g_m) stages; the other input of all these gm stages is connected to the amplifier's gain-determining feedback network, R_{F1}/R_{F2}. When the emitter bias current, I_E, is directed to one of the 8 transistor pairs (by means not shown here), it becomes the input stage for the complete amplifier.

CONTINUOUS INTERPOLATION BETWEEN TAPS IN THE X-AMP IS PERFORMED WITH CURRENT-CONTROLLED gm STAGES

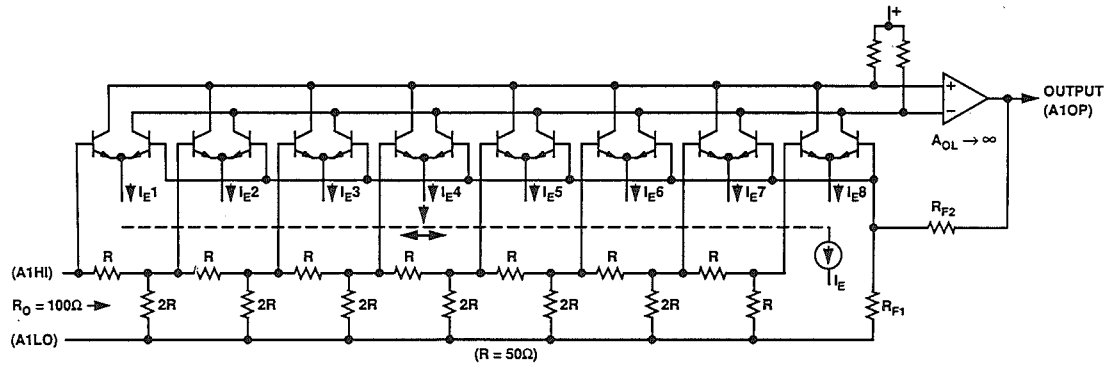

Figure 15.37

KEY FEATURES OF THE AD600/AD602 X-AMPS

- Precise Decibel-Scaled Gain Control
- Accurate Absolute Gain Calibration
- Low Input-Referred Noise (1.4nV/√Hz)
- Constant Bandwidth (dc to 35MHz)
- Low Distortion: –60dBc THD at ±1V Output
- Stable Group Delay (±2ns Over Gain Range)
- Response Time: Less than 1µs for 40dB Gain Change
- Low Power (125mW per channel maximum)
- Differential Control Inputs

Figure 15.38

When I_E is connected to the pair on the left-hand side, the signal input is connected directly to the amplifier, giving the maximum gain. The distortion is very low, even at high frequencies, due to the careful open-loop design, aided by the negative feedback. If I_E were now to be abruptly switched to the second pair, the overall gain would drop by exactly 6.02dB, and the distortion would remain low, because only one g_m stage remains active.

In reality, the bias current is *gradually* transferred from the first pair to the second. When I_E is equally divided between two gm stages, both are active, and the situation arises where we have an op amp with two input stages fighting for control of the loop, one getting the full signal, and the other getting a signal exactly half as large.

Analysis shows that the effective gain is reduced, not by 3dB, as one might first expect, but rather by 20log1.5, or 3.52dB. This error, when divided equally over the whole range, would amount to a gain ripple of ±0.25dB; however, the interpolation circuit actually generates a Gaussian distribution of bias currents, and a significant fraction of I_E always flows in adjacent stages. This smoothes the gain function and actually lowers the ripple (see Reference 19). As I_E moves further to the right, the overall gain progressively drops.

The total input-referred noise of the X-AMP™ is 1.4nV/√Hz; only slightly more than the thermal noise of a 100Ω resistor which is 1.29nV/√Hz at 25°C. The input-referred noise is constant regardless of the attenuator setting, therefore the output noise is always constant and independent of gain. For the AD600, the amplifier gain is 113 and the output noise spectral density is therefore 1.4nV/√Hz×113 or 158nV/√Hz. Referred to its maximum output of 2V rms, the signal-to-noise ratio would be 82dB in a 1MHz bandwidth. The corresponding signal-to-noise ratio of the AD602 is 10dB greater, or 92dB.

WIDE DYNAMIC RANGE ULTRASOUND SYSTEMS

Walt Kester

A block diagram of a typical ultrasound system is shown in Figure 15.39. A burst of ultrasound energy (1 to 13MHz) is generated in a piezoelectric transducer which physically contacts the outer body surface. The velocity of propagation of the ultrasound waves in most soft-body tissues (air and bones are the exception) is about 1500m/sec. Echoes are produced at interfaces between various types of soft-body structures. The round-trip time of each echo is used to determine its distance from the transducer.

B-SCAN ULTRASOUND SYSTEM BLOCK DIAGRAM

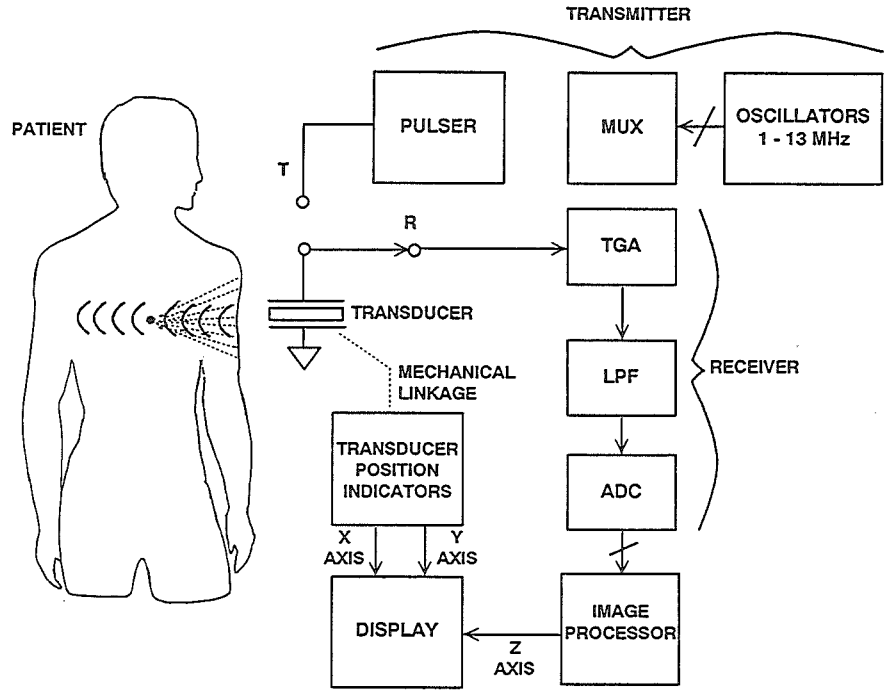

Figure 15.39

Soft-body tissues attenuate the burst of ultrasound energy by approximately 1dB/cm/MHz. For the thicker parts of the body, as in abdominal imaging, frequencies of 1 to 2MHz are common. For imaging of shorter path lengths, as in studies of the eye or other superficial structures, frequencies as high as 20MHz can be used. Because of soft-body tissue attenuation, the receiving transducer will see a dynamic range of 100dB when scanning from 1 to 10cm at 10MHz, independent of the tissue variations that need to be observed. Add the 50dB dynamic range typical for variations in tissue, and the transducer must have close to 150dB dynamic range. For this reason, the transducer output is usually applied to a Time Gain Amplifier (TGA) whose gain in dB is directly proportional to the amount of time elapsed from the transmission of the burst (see Figure 15.40). The AD600 X-AMP™ previously discussed is exactly such a device.

TIME GAIN AMPLIFIER

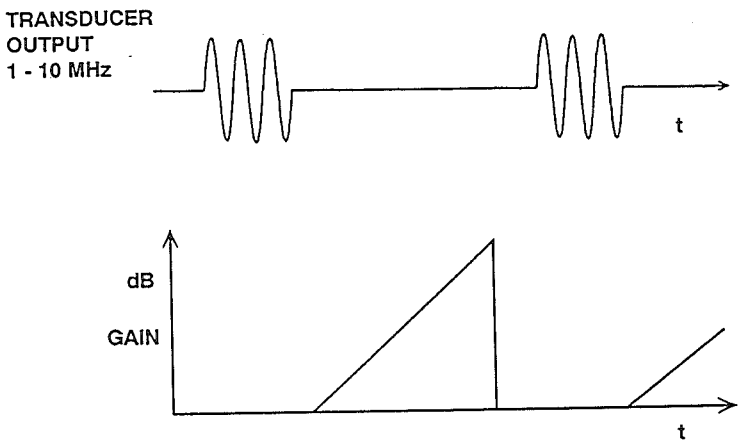

Figure 15.40

For reflections that are near the surface, there will be little attenuation. For deep signal returns, gain is applied to compensate for the path attenuation. The TGA thus compensates for normal signal attenuation associated with delay/distance. The receiver (ADC) therefore only sees the intensity variations associated with the different tissue types. In scans where the propagation path is primarily soft tissue structures of comparable attenuation, such as the abdomen, a fixed gain versus time function is usually adequate. In other cases involving blood pools or fixed regions, it is often desirable to vary the gain versus time function. Many commercial systems make this option available. In some cases it is even desirable for the operator to calibrate the TGA on a per-patient basis in order to achieve the best diagnostic image.

In phased array ultrasound systems, the angular information is precisely determined by phasing the delays from a number of transducers (transmitted and received) to electronically select the angle to be processed. The first generation of phased array elements used analog beam forming techniques as shown in Figure 15.41. Delays at the transmitter and receiver are adjusted using variable delay filters. The next generation of phased arrays will be digital. Low cost, low power, high performance ADCs and DSPs make it practical to digitize the rf directly and digitally control the delay requirements as shown in Figure 15.42. This technique is often referred to as digital beamforming and is also used in some modern radar and sonar systems.

ANALOG BEAMFORMING

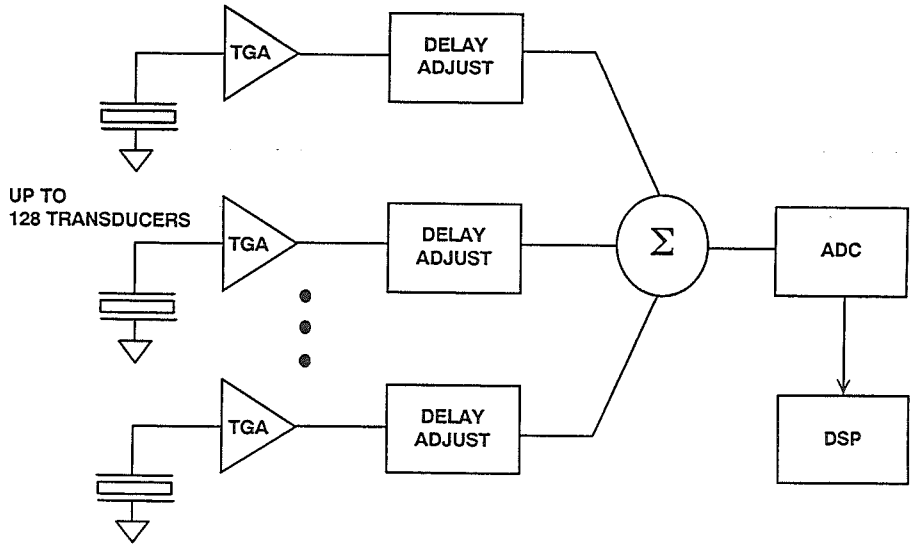

Figure 15.41

DIGITAL BEAMFORMING

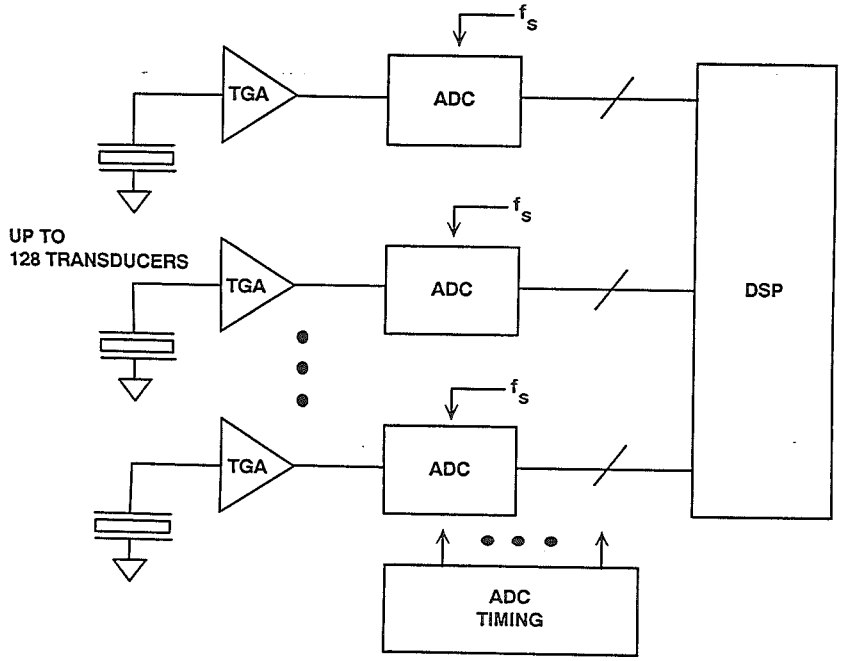

Figure 15.42

Figure 15.43 shows the complete signal path for a single channel of a phased array ultrasound system. The peak signal from the transducer is about 1V p-p into a 100Ω load. The signal then passes through a gain-of-two low noise pre-amplifier which drives the 100Ω input of the AD600. The appropriate gain function is generated by cascading each half of the AD600 and controlling the gain with the calibrated ramp voltage. The gain of the two stages can be changed from 0dB to 80dB. The interstage bandpass filter blocks the dc offset from the first gain stage and limits the out-of-band noise. The output of the second gain stage drives another bandpass filter and the AD9040 ADC input. The AD9040 output goes to the memory and the DSP circuits for further processing. This system is capable of maintaining a SFDR of greater than 55dB at ultrasound frequencies up to 15MHz.

Extreme care in layout, decoupling, grounding, and signal routing is essential in order to prevent oscillation in these circuits. A gain of 80dB (10,000) at a bandwidth of 1MHz corresponds to an effective *gain-bandwidth* product of 10GHZ!

Figure 15.44 shows the calculations to determine the proper gain function for the TGA assuming a 5MHz burst into soft body tissue. The corresponding control voltage ramp for the AD600 is shown in Figure 15.45.

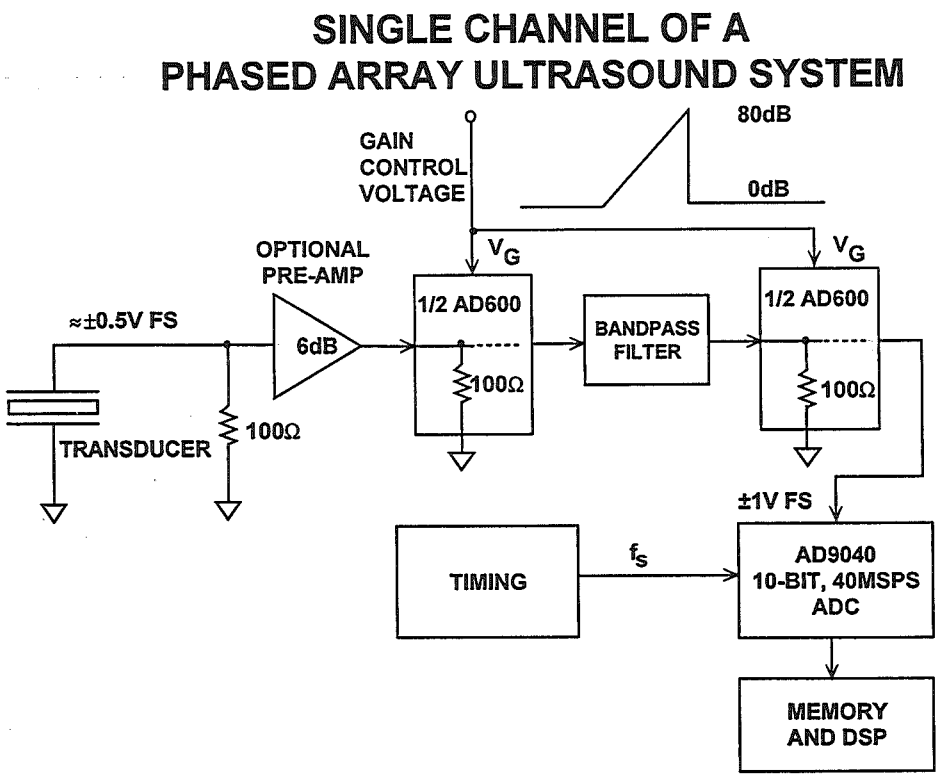

Figure 15.43

CALCULATION OF TGA COMPENSATION RAMP SLOPE FOR 5MHz ULTRASOUND BURST INTO SOFT BODY TISSUE

- Acoustic Velocity = 1500m/s (Soft Body Tissue)
- Propagation Delay = 6.7µs/cm (One Way)
  = 13.4µs/cm (Round Trip)
- Attenuation = 1dB/cm/MHz (One Way)
  = 5dB/cm @ 5MHz (One Way)
  = 10dB/cm @ 5MHz (Round Trip)
- TGA Ramp Slope = (10dB/cm)÷(13.4µs/cm) = 0.746dB/µs

Figure 15.44

TYPICAL CONTROL VOLTAGE FOR ULTRASOUND TIME GAIN AMPLIFIER, f = 5MHz, INTO SOFT BODY TISSUE

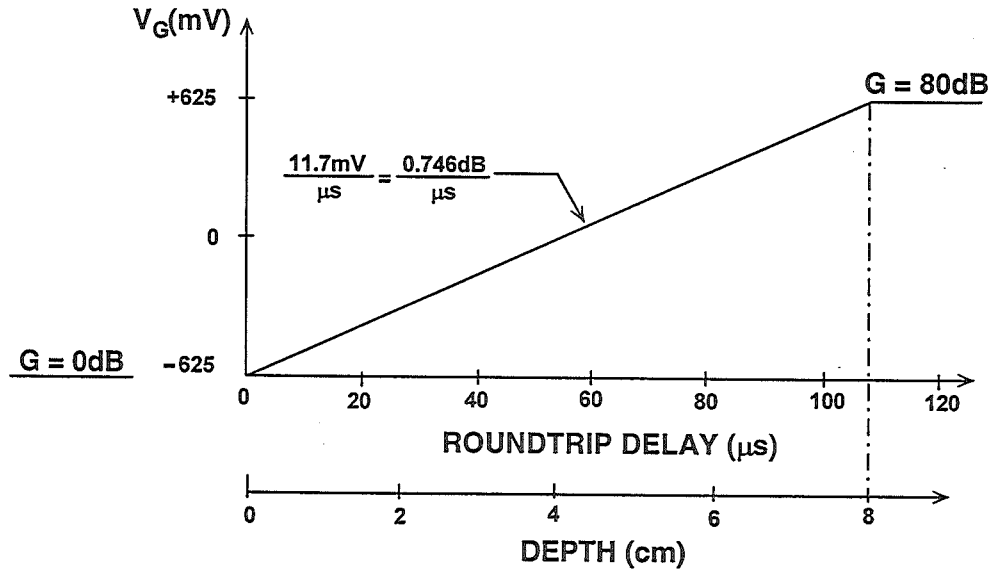

Figure 15.45

In some ultrasound applications, the user may wish to use a high input impedance preamplifier to avoid the signal attenuation that would result from loading the transducer by the 100Ω input resistance of the X-AMP. High gain cannot be tolerated, because the peak transducer signal is typically ±0.5V, while the peak input capability of the AD600/AD602 is only slightly more than ±1V. A gain of two is a suitable choice. In order to maintain the 1.4nV/√Hz noise performance of the AD600/AD602, the preamplifier's input noise should be less than 1.2nV/√Hz.

An inexpensive circuit, using complementary transistor types chosen for the low $r_{bb'}$, is shown in Figure 15.46. The gain is determined by the ratio of the net collector load resistance to the net emitter resistance, that is, it is an open-loop amplifier. The gain will be ×2 (6dB) only into a 100Ω load, assumed to be provided by the input to the X-AMP.; R2 and R7 are in shunt with this load, and their value is important in defining the gain. For small-signal inputs, both transistors contribute an equal transconductance, which is rendered less sensitive to signal level by the emitter resistors R4 and R5, which also play a dominant role in setting the gain.

This is a Class AB amplifier. As V_{IN} increases in a positive direction, Q1 conducts more heavily and its r_e becomes lower while that of Q2 increases. Conversely, more negative values of V_{IN} result in the r_e of Q2 decreasing, while that of Q1 increases. The design is chosen such that the net emitter resistance is essentially independent of the instantaneous value of V_{IN}, resulting in moderately low distortion. Low values of resistance and moderately high bias currents are important in achieving the low noise, wide bandwidth, and low distortion of this preamplifier. Heavy decoupling prevents noise on the power supply lines from being conveyed to the input of the X-AMP.

A LOW-NOISE ULTRASOUND PREAMPLIFIER FOR THE AD600 AND AD602

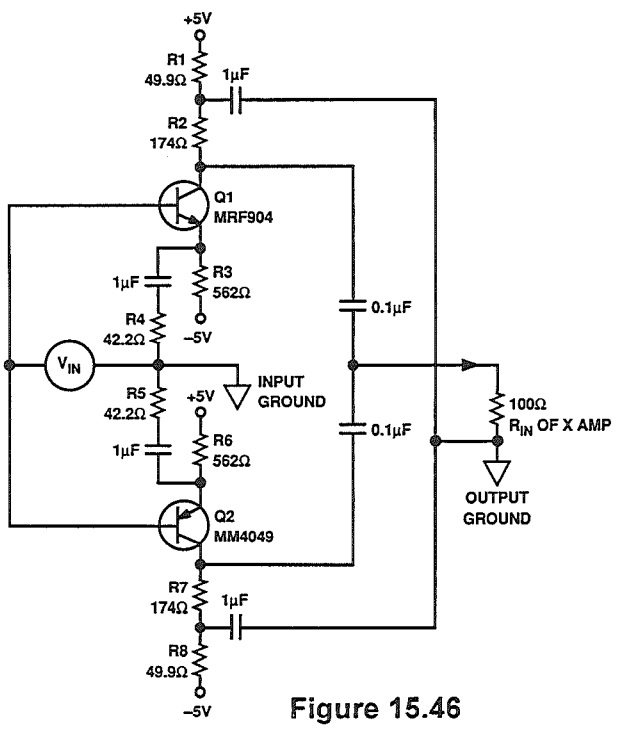

Figure 15.46

ULTRASOUND PREAMPLIFIER PERFORMANCE

Measurement	Value
Gain (f = 30MHz)	6dB
Bandwidth (–3dB)	250MHz
Input Signal for 1dB Compression	1V p-p
Second Harmonic Distortion, V_{IN} = 200mV p-p	–51dBc
Third Harmonic Distortion, V_{IN} = 200mV p-p	–57dBc
Second Harmonic Distortion, V_{IN} = 500mV p-p	–47dBc
Third Harmonic Distortion, V_{IN} = 500mV p-p	–45dBc
System Input NSD (Including X-AMP)	1.03nV/√Hz
Input Resistance	1.4kΩ
Input Capacitance	15pF
Input Bias Current	±150µA
Power Supply Voltage	±5V
Quiescent Current	15mA

Figure 15.47

For ultrasound frequencies below 1MHz, the AD797 op amp makes a good choice for the ultrasound pre-amplifier.

This example clearly demonstrates the power of the combination of analog and digital signal processing to solve a system problem. In order to provide the same dynamic range without the TGA (i.e. the ADC interfaces directly to the transducer via a fixed-gain preamp), the ADC would have to have a dynamic range of approximately 100dB. This implies a 16 bit ADC which would have to operate at a sampling frequency of at least 30MSPS — a requirement which is clearly beyond the present state of the art in ADC technology!

REFERENCES

1. Richard Groshong and Stephen Ruscak, *Undersampling Techniques Simplify Digital Radio*, **Electronic Design**, May 23, 1991, pp. 67-78.

2. Richard Groshong and Stephen Ruscak, *Exploit Digital Advantages In An SSB Receiver*, **Electronic Design**, June 13, 1991, pp. 89-96.

3. Richard G. Lyons, *How Fast Must You Sample?*, **Test and Measurement World**, November, 1988, pp. 47-57.

4. Richard C. Webb, *IF Signal Sampling Improves Receiver Detection Accuracy*, **Microwaves and RF**, March, 1989, pp. 99-103.

5. Jeff Kirsten and Tarlton Fleming, *Undersampling Reduces Data-Acquisition Costs For Select Applications*, EDN, June 21, 1990, pp 217-228.

6. Hans Steyskal and John F. Rose, *Digital Beamforming For Radar Systems*, **Microwave Journal**, January, 1989, pp. 121-136.

7. Tom Gratzek and Frank Murden, *Optimize ADCs For Enhanced Signal Processing*, **Microwaves and RF**, Vol. 30, No.3, March, 1991, pp. 129-136.

8. Howard Hilton, *10-MHz ADC With 110dB Linearity*, High Speed ADC Conference, Las Vegas, April 21-22, 1992.

9. Howard Hilton, *10MSample/Second ADC With Filter And Memory*, **High Speed ADC Conference**, Las Vegas, April 21-22, 1992.

10. Dan Asta, *Recent Dynamic Range Characterization of Analog-to-Digital Converters for Spectral Analysis Applications*, **Project Report AST-14**, MIT Lincoln Laboratory, Lexington, MA, July 9, 1991.

11. Fred H. Irons and T.A. Rebold, *Characterization of High-Frequency Analog-to-Digital Converters for Spectral Analysis Applications*, **Project Report AST-2**, MIT Lincoln Laboratory, November 28, 1986.

12. F. H. Irons, *Dynamic Characterization and Compensation of Analog-to-Digital Converters*, **IEEE International Symposium on Circuits and Systems**, May 1986, San Jose, CA. Catalog No. CH2255-8/86/0000-1273.

13. T.A. Rebold and F. H. Irons, *A Phase-Plane Approach to the Compensation of High-Speed Analog-to-Digital Converters*, **IEEE International Symposium on Circuits and Systems**, Philadelphia, PA, May, 1987.

14. Dan Asta and Fred H. Irons, *Dynamic Error Compensation of Analog-to-Digital Converters*, **The Lincoln Laboratory Journal**, Vol. 2, No. 2, 1989, pp. 161-182.

15. James M. Bryant, *Bandpass Sigma-Delta ADCs for Direct IF Conversion*, **DSP- The Enabling Technology for Communications**, RAI Congrescentrum, Amsterdam, Netherlands, 9-10 March 1993.

16. S.A. Jantzi, M. Snelgrove, P.F. Ferguson, Jr., *A 4th-Order Bandpass Sigma-Delta Modulator*, **Proceedings of the IEEE 1992 Custom Integrated Circuits Conference**, pp. 16.5.1-4.

17. S.A. Jantzi, R. Schreier, and M. Snelgrove, *A Bandpass Sigma-Delta Convertor for a Digital AM Receiver*, **Proceedings of the IEE International Conference on Analogue-to-Digital and Digital-to-Analogue Conversion**, Swansea, UK., September, 1991, pp. 75-80.

18. S.A. Jantzi, W. Martin Snelgrove, and Paul F. Ferguson, Jr., *A Fourth-Order Bandpass Sigma-Delta Modulator*, **IEEE Journal of Solid-State Circuits**, Vol. 38, No. 3, March 1993, pp. 282-291.

19. Barrie Gilbert, *A Low Noise Wideband Variable-Gain Amplifier Using an Interpolated Ladder Attenuator*, **IEEE ISSCC Technical Digest**, 1991, pages 280, 281, 330.

20. Barrie Gilbert, *A Monolithic Microsystem for Analog Synthesis of Trigonometric Functions and their Inverses*, **IEEE Journal of Solid State Circuits**, Vol. SC-17, No. 6, December 1982, pp. 1179-1191.

21. Albert Macovski, **Medical Imaging Systems**, Prentice-Hall, Inc., New Jersey, 1983.

22. **1992 Amplifier Applications Guide**, Analog Devices, Norwood MA, 1992.