

SECTION 8

ADCs FOR SIGNAL CONDITIONING

Walt Kester, James Bryant, Joe Buxton

The trend in ADCs and DACs is toward higher speeds and higher resolutions at reduced power levels. Modern data converters generally operate on $\pm 5V$ (dual supply) or $+5V$ (single supply). In fact, many new converters operate on a single $+3V$ supply. This trend has created a number of design and applications problems which were much less important in earlier data converters, where $\pm 15V$ supplies and $\pm 10V$ input ranges were the standard.

Lower supply voltages imply smaller input voltage ranges, and hence more susceptibility to noise from all potential sources: power supplies, references, digital signals, EMI/RFI, and probably most important, improper layout, grounding, and decoupling techniques. Single-supply ADCs often have an input range which is not referenced to ground. Finding compatible single-supply drive amplifiers and dealing with level shifting of the input signal in direct-coupled applications also becomes a challenge.

In spite of these issues, components are now available which allow extremely high resolutions at low supply voltages and low power. This section discusses the applications problems associated with such components and shows techniques for successfully designing them into systems.

The most popular precision signal conditioning ADCs are based on two fundamental architectures: *successive approximation* and *sigma-delta*. We have seen that the *tracking* ADC architecture is particularly suited for resolver-to-digital converters, but it is rarely used in other precision signal conditioning applications. The *flash* converter and the *subranging (or pipelined)* converter architectures are widely used where sampling frequencies extend into the megahertz and hundreds of megahertz region, but are overkill's in both speed and cost for low frequency precision signal conditioning applications.

Figure 8.1

Figure 8.2

SUCCESSIVE APPROXIMATION ADCs

The successive approximation ADC has been the mainstay of signal conditioning for many years. Recent design improvements have extended the sampling frequency of these ADCs into the megahertz region. The use of internal switched capacitor techniques along with auto calibration techniques extend the resolution of these ADCs to 16-bits on standard CMOS processes without the need for expensive thin-film laser trimming.

ADCs FOR SIGNAL CONDITIONING

The basic successive approximation ADC is shown in Figure 8.3. It performs conversions on command. On the assertion of the CONVERT START command, 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.

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.

Figure 8.3

An N-bit conversion takes N steps. It would seem on superficial examination that a 16-bit converter would have twice the conversion time of an 8-bit one, but this is not the case. In an 8-bit converter, the DAC must settle to 8-bit accuracy before the bit decision is made, whereas in a 16-bit converter, it must settle to 16-bit accuracy, which takes a lot longer. In practice, 8-bit successive approximation ADCs can convert in a few hundred nanoseconds, while 16-bit ones will generally take several microseconds.

Notice that the overall accuracy and linearity of the SAR ADC is determined primarily by the internal DAC. Until recently, most precision SAR ADCs used laser-trimmed thin-film DACs to achieve the desired accuracy and linearity. The thin-film resistor trimming process adds cost, and the thin-film resistor values may be affected when subjected to the mechanical stresses of packaging.

For these reasons, switched capacitor (or charge-redistribution) DACs have become popular in newer SAR ADCs. The advantage of the switched capacitor DAC is that the accuracy and linearity is primarily determined by photolithography, which in turn controls the capacitor plate area and the capacitance as well as matching. In addition, small capacitors can be placed in parallel with the main capacitors which can be switched in and out under control of autocalibration routines to achieve high accuracy and linearity without the need for thin-film laser trimming. Temperature tracking between the switched capacitors can be better than 1ppm/°C, thereby offering a high degree of temperature stability.

A simple 3-bit capacitor DAC is shown in Figure 8.4. 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.

Figure 8.4

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 (cap 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.

Because of their popularity, successive approximation ADCs are available in a wide variety of resolutions, sampling rates, input and output options, package styles, and costs. It would be impossible to attempt to list all types, but Figure 8.5 shows a number of recent Analog Devices' SAR ADCs which are representative. Note that many devices are complete data acquisition systems with input multiplexers which allow a single ADC core to process multiple analog channels.

Figure 8.5

While there are some variations, the fundamental timing of most SAR ADCs is similar and relatively straightforward (see Figure 8.6). The conversion process is initiated by asserting a CONVERT START signal. The \overline{CONVST} signal is a negative-going pulse whose positive-going edge actually initiates the conversion. The internal sample-and-hold (SHA) amplifier is placed in the hold mode on this edge, and the various bits are determined using the SAR algorithm. The negative-going edge of the \overline{CONVST} pulse causes the \overline{EOC} or BUSY line to go high. When the conversion is complete, the BUSY line goes low, indicating the completion of the conversion process. In most cases the trailing edge of the BUSY line can be used as an indication that the output data is valid and can be used to strobe the output data into an external register. However, because of the many variations in terminology and design, the individual data sheet should always be consulted when using with a specific ADC.

Figure 8.6

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

ADCs FOR SIGNAL CONDITIONING

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 allow single-shot conversion at any repetition rate from DC to the converter's maximum conversion rate.

In a SAR ADC, the output data for a particular cycle is valid at the end of the conversion interval. In other ADC architectures, such as sigma-delta or the two-stage subranging architecture shown in Figure 8.7, this is not the case. The subranging ADC shown in the figure is a two-stage *pipelined* or subranging 12-bit converter. The first conversion is done by the 6-bit ADC which drives a 6-bit DAC. The output of the 6-bit DAC represents a 6-bit approximation to the analog input. Note that SHA2 delays the analog signal while the 6-bit ADC makes its decision and the 6-bit DAC settles. The DAC approximation is then subtracted from the analog signal from SHA2, amplified, and digitized by a 7-bit ADC. The outputs of the two conversions are combined, and the extra bit used to correct errors made in the first conversion. The typical timing associated with this type of converter is shown in Figure 8.8. Note that the output data presented immediately after sample X actually corresponds to sample X-2, i.e., there is a two clock-cycle "pipeline" delay. The pipelined ADC architecture is generally associated with high speed ADCs, and in most cases the pipeline delay, or *latency*, is not a major system problem in most applications where this type of converter is used.

Figure 8.7

Figure 8.8

Pipelined ADCs may have more than two clock-cycles latency depending on the particular architecture. For instance, the conversion could be done in three, or four, or perhaps even more pipelined stages causing additional latency in the output data.

Therefore, if the ADC is to be used in an event-triggered (or single-shot) mode where there must be a one-to-one time correspondence between each sample and the corresponding data, then the pipeline delay can be troublesome, and the SAR architecture is advantageous. Pipeline delay or latency can also be a problem in high speed servo-loop control systems or multiplexed applications. In addition, some pipelined converters have a *minimum* allowable conversion rate and must be kept running to prevent saturation of internal nodes.

Switched capacitor SAR ADCs generally have unbuffered input circuits similar to the circuit shown in Figure 8.9 for the AD7858/59 ADC. During the acquisition time, the analog input must charge the 20pF equivalent input capacitance to the correct value. If the input is a DC signal, then the source resistance, R_S , in series with the 125 Ω internal switch resistance creates a time constant. In order to settle to 12-bit accuracy, approximately 9 time constants must be allowed for settling, and this defines the minimum allowable acquisition time. (Settling to 14-bits requires about 10 time constants, and 16-bits requires about 11).

$$t_{ACQ} > 9 \times (R_S + 125)\Omega \times 20\text{pF}.$$

For example, if $R_S = 50\Omega$, the acquisition time per the above formula must be at least 310ns.

For AC applications, a low impedance source should be used to prevent distortion due to the non-linear ADC input circuit. In a single supply application, a fast settling rail-to-rail op amp such as the AD820 should be used. Fast settling allows the op amp to settle quickly from the transient currents induced on its input by the internal ADC switches. In Figure 8.9, the AD820 drives a lowpass filter consisting of the 50Ω series resistor and the 10nF capacitor (cutoff frequency approximately 320kHz). This filter removes high frequency components which could result in aliasing and increased noise.

Using a single supply op amp in this application requires special consideration of signal levels. The AD820 is connected in the inverting mode and has a signal gain of -1 . The noninverting input is biased at a common mode voltage of +1.3V with the $10.7k\Omega/10k\Omega$ divider, resulting in an output voltage of +2.6V for $V_{IN} = 0V$, and +0.1V for $V_{IN} = +2.5V$. This offset is provided because the AD820 output cannot go all the way to ground, but is limited to the V_{CESAT} of the output stage NPN transistor, which under these loading conditions is about 50mV. The input range of the ADC is also offset by +100mV by applying the +100mV offset from the $412\Omega/10k\Omega$ divider to the AIN $-$ input.

Figure 8.9

The AD789X-family of single supply SAR ADCs (as well as the AD974, AD976, and AD977) includes a thin film resistive attenuator and level shifter on the analog input to allow a variety of input range options, both bipolar and unipolar. A simplified diagram of the input circuit of the AD7890-10 12-bit, 8-channel ADC is shown in Figure 8.10. This arrangement allows the converter to digitize a $\pm 10V$ input while operating on a single +5V supply. The R1/R2/R3 thin film network provides the attenuation and level shifting to convert the $\pm 10V$ input to a 0V to +2.5V signal which is digitized by the internal ADC. This type of input requires no special drive circuitry because R1 isolates the input from the actual converter circuitry. Nevertheless, the source resistance, R_S , should be kept reasonably low to prevent gain errors caused by the $R_S/R1$ divider.

Figure 8.10

SAR ADCs WITH MULTIPLEXED INPUTS

Multiplexing is a fundamental part of many data acquisition systems, and a fundamental understanding of multiplexers is required to design a data acquisition system. Switches for data acquisition systems, especially when integrated into the IC, generally are CMOS-types shown in Figure 8.11. Utilizing the P-Channel and N-Channel MOSFET switches in parallel minimizes the change of on-resistance (R_{ON}) as a function of signal voltage. On-resistance can vary from less than 5Ω to several hundred ohms depending upon the device. Variation in on-resistance as a function of signal level (often called R_{ON} -modulation) can cause distortion if the multiplexer must drive a load, and therefore R_{ON} flatness is also an important specification.

ADCs FOR SIGNAL CONDITIONING

Figure 8.11

Because of non-zero R_{ON} and R_{ON} -modulation, multiplexer outputs should be isolated from the load with a suitable buffer amplifier. A separate buffer is not required if the multiplexer drives a high input impedance, such as a PGA, SHA or ADC - but beware! Some SHAs and ADCs draw high frequency pulse current at their sampling rate and cannot tolerate being driven by an unbuffered multiplexer.

The key multiplexer specifications are *switching time*, *on-resistance*, *on-resistance flatness*, and *off-channel isolation*, and *crosstalk*. Multiplexer switching time ranges from less than 20ns to over 1 μ s, R_{ON} from less than 5 Ω to several hundred ohms, and off-channel isolation from 50 to 90dB.

A number of CMOS switches can be connected to form a multiplexer as shown in Figure 8.12. The number of input channels typically ranges from 4 to 16, and some multiplexers have internal channel-address decoding logic and registers, while with others, these functions must be performed externally. Unused multiplexer inputs *must* be grounded or severe loss of system accuracy may result.

Switches and multiplexers may be optimized for various applications as shown in Figure 8.13.

Figure 8.12

Figure 8.13

An M-channel multiplexed data acquisition system is shown in Figure 8.14. The typical timing associated with the SAR ADC is also shown in the diagram. The conversion process is initiated on the positive-going edge of the \overline{CONVST} pulse. If maximum throughput is desired, the multiplexer is changed to the next channel at the same time. This allows nearly the entire sampling period ($1/f_s$) for the multiplexer to settle. Remember that it is possible to have a positive fullscale signal on one channel and a negative fullscale signal on the next, therefore the multiplexer output must settle from a fullscale output step change within the allocated time.

Figure 8.14

Also shown in Figure 8.14 are input filters on each channel. These filters serve as antialiasing filters to remove signals above one-half the effective per-channel sampling frequency. If the ADC is sampling at f_s , and the multiplexer is sequencing through all M channels, then the per-channel sampling rate is f_s/M . The input lowpass filters should have sufficient attenuation $f_s/2M$ to prevent dynamic range limitations due to aliasing.

It is not necessary, however, that each channel be sampled at the same rate, and the various input lowpass filters can be individually tailored for the actual sampling rate and signal bandwidth expected on each channel.

An optional lowpass filter is often placed between the multiplexer output and the SHA input, designated LPF_C in Figure 8.14. Care must be exercised in selecting its cutoff frequency because its time constant directly affects the multiplexer settling

time. If the filter is a single-pole, the number of time constants, n , required to settle to a desired accuracy is given in Figure 8.15.

Figure 8.15

If the time constant of LPF_C is τ , and its cutoff frequency f_c , then

$$f_c = \frac{1}{2\pi\tau}.$$

But the sampling frequency f_s is related to $n \cdot \tau$ by the equation:

$$f_s < \frac{1}{n \cdot \tau}.$$

Combining the two equations and solving for f_c in terms of n and f_s yields:

$$f_c > \frac{n \cdot f_s}{2\pi}.$$

As an example, assume that the ADC is a 12-bit one sampling at 100kSPS. From the table, $n = 8.32$, and therefore $f_c > 132\text{kSPS}$ per the above equation. While this filter will help prevent wideband noise from entering the SHA, it does not provide the same function as the antialiasing filters at the input of each channel, whose individual cutoff frequencies can be much lower.

For this reason, only a few integrated data acquisition ICs with on-board multiplexers give access to the multiplexer output and the SHA input. If access is offered and LPF_C is used, the settling time requirement must be observed in order to achieve the desired accuracy.

COMPLETE DATA ACQUISITION SYSTEMS ON A CHIP

VLSI mixed-signal processing allows the integration of large and complex data acquisition circuits on a single chip. Most signal conditioning circuits including multiplexers, PGAs, and SHAs, can now be manufactured on the same chip as the ADC. This high level of integration permits data acquisition systems (DASs) to be specified and tested as a single complex function.

Such functionality relieves the designer of most of the burden of testing and calculating error budgets. The DC and AC characteristics of a complete data acquisition system are specified as a complete function, which removes the necessity of calculating performance from a collection of individual worst case device specifications. A complete monolithic system should achieve a higher performance at much lower cost than would be possible with a system built up from discrete functions. Furthermore, system calibration is easier, and in fact many monolithic DASs are self calibrating, offering both internal and system calibration functions.

The AD7858 is an example of a highly integrated IC DAS (see Figure 8.16). The device operates on a single supply voltage of +3V to +5.5V and dissipates only

ADCs FOR SIGNAL CONDITIONING

15mW. The resolution is 12-bits, and the maximum sampling frequency is 200kSPS. The input multiplexer can be configured either as 8 single-ended inputs or 4 pseudo-differential inputs. The AD7858 requires an external 4MHz clock and initiates the conversion on the positive-going edge of the $\overline{\text{CONVST}}$ pulse which does not need to be synchronized to the high frequency clock. Conversion can also be initiated via software by setting a bit in the proper control register.

Figure 8.16

Figure 8.17

The AD7858 contains an on-chip 2.5V reference (which can be overridden with an external one), and the fullscale input voltage range is 0V to V_{REF} . The internal DAC is a switched capacitor type, and the ADC contains a self-calibration and system calibration option to ensure accurate operation over time and temperature. The input/output port is a serial one and is SPI, QSPI, 8051, and μP compatible.

The AD7858L is a lower power (5.5mW) version of the AD7858 which operates at a maximum sampling rate of 100kSPS.

SIGMA-DELTA ($\Sigma\Delta$) MEASUREMENT ADCs

James M. Bryant

Sigma-Delta Analog-Digital Converters ($\Sigma\Delta$ ADCs) have been known for nearly thirty years, but only recently has the technology (high-density digital VLSI) existed to manufacture them as inexpensive monolithic integrated circuits. They are now used in many applications where a low-cost, low-bandwidth, low-power, high-resolution ADC is required.

There have been innumerable descriptions of the architecture and theory of $\Sigma\Delta$ ADCs, but most commence with a maze of integrals and deteriorate from there. In the Applications Department at Analog Devices, we frequently encounter engineers who do not understand the theory of operation of $\Sigma\Delta$ ADCs and are convinced, from study of a typical published article, that it is too complex to comprehend easily.

There is nothing particularly difficult to understand about $\Sigma\Delta$ ADCs, as long as you avoid the detailed mathematics, and this section has been written in an attempt to clarify the subject. A $\Sigma\Delta$ ADC contains very simple analog electronics (a comparator, a switch, and one or more integrators and analog summing circuits), and quite complex digital computational circuitry. This circuitry consists of a digital signal processor (DSP) which acts as a filter (generally, but not invariably, a low pass filter). It is not necessary to know precisely how the filter works to appreciate what it does. To understand how a $\Sigma\Delta$ ADC works familiarity with the concepts of *over-sampling*, *quantization noise shaping*, *digital filtering*, and *decimation* is required.

Figure 8.18

Let us consider the technique of over-sampling with an analysis in the frequency domain. Where a DC conversion has a *quantization error* of up to $\frac{1}{2}$ LSB, a sampled

data system has *quantization noise*. A perfect classical N-bit sampling ADC has an RMS quantization noise of $q/\sqrt{12}$ uniformly distributed within the Nyquist band of DC to $f_s/2$ (where q is the value of an LSB and f_s is the sampling rate) as shown in Figure 8.19A. Therefore, its SNR with a full-scale sinewave input will be $(6.02N + 1.76)$ dB. If the ADC is less than perfect, and its noise is greater than its theoretical minimum quantization noise, then its *effective* resolution will be less than N-bits. Its actual resolution (often known as its Effective Number of Bits or ENOB) will be defined by

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

If we choose a much higher sampling rate, Kf_s (see Figure 8.19B), the quantization noise is distributed over a wider bandwidth DC to $Kf_s/2$. If we then apply a digital low pass filter (LPF) to the output, we remove much of the quantization noise, but do not affect the wanted signal - so the ENOB is improved. We have accomplished a high resolution A/D conversion with a low resolution ADC. The factor K is generally referred to as the *oversampling ratio*.

Figure 8.19

Since the bandwidth is reduced by the digital output filter, the output data rate may be lower than the original sampling rate (Kf_s) and still satisfy the Nyquist criterion. This may be achieved by passing every M th result to the output and discarding the remainder. The process is known as "decimation" by a factor of M . Despite the origins of the term (*decem* is Latin for ten), M can have any integer value, provided that the output data rate is more than twice the signal bandwidth. Decimation does not cause any loss of information (see Figure 8.19B).

If we simply use over-sampling to improve resolution, we must over-sample by a factor of 2^{2N} to obtain an N-bit increase in resolution. The $\Sigma\Delta$ converter does not need such a high over-sampling ratio because it not only limits the signal passband, but also shapes the quantization noise so that most of it falls outside this passband as shown in Figure 8.19C.

If we take a 1-bit ADC (generally known as a comparator), drive it with the output of an integrator, and feed the integrator with an input signal summed with the output of a 1-bit DAC fed from the ADC output, we have a first-order $\Sigma\Delta$ modulator as shown in Figure 8.20. Add a digital low pass filter (LPF) and decimator at the digital output, and we have a $\Sigma\Delta$ ADC: the $\Sigma\Delta$ modulator shapes the quantization noise so that it lies above the passband of the digital output filter, and the ENOB is therefore much larger than would otherwise be expected from the over-sampling ratio.

Figure 8.20

Intuitively, a $\Sigma\Delta$ ADC operates as follows. Assume a DC input at V_{IN} . The integrator is constantly ramping up or down at node A. The output of the comparator is fed back through a 1-bit DAC to the summing input at node B. The negative feedback loop from the comparator output through the 1-bit DAC back to the summing point will force the average DC voltage at node B to be equal to V_{IN} .

ADCs FOR SIGNAL CONDITIONING

This implies that the average DAC output voltage must equal to the input voltage V_{IN} . The average DAC output voltage is controlled by the *ones-density* in the 1-bit data stream from the comparator output. As the input signal increases towards $+V_{REF}$, the number of "ones" in the serial bit stream increases, and the number of "zeros" decreases. Similarly, as the signal goes negative towards $-V_{REF}$, the number of "ones" in the serial bit stream decreases, and the number of "zeros" increases. From a very simplistic standpoint, this analysis shows that the average value of the input voltage is contained in the serial bit stream out of the comparator. The digital filter and decimator process the serial bit stream and produce the final output data.

The concept of noise shaping is best explained in the frequency domain by considering the simple $\Sigma\Delta$ modulator model in Figure 8.21.

Figure 8.21

The integrator in the modulator is represented as an analog lowpass filter with a transfer function equal to $H(f) = 1/f$. This transfer function has an amplitude response which is inversely proportional to the input frequency. The 1-bit quantizer generates quantization noise, Q , which is injected into the output summing block. If we let the input signal be X , and the output Y , the signal coming out of the input summer must be $X - Y$. This is multiplied by the filter transfer function, $1/f$, and the result goes to one input to the output summer. By inspection, we can then write the expression for the output voltage Y as:

$$Y = \frac{1}{f}(X - Y) + Q.$$

This expression can easily be rearranged and solved for Y in terms of X , f , and Q :

$$Y = \frac{X}{f + 1} + \frac{Q \cdot f}{f + 1}.$$

Note that as the frequency f approaches zero, the output voltage Y approaches X with no noise component. At higher frequencies, the amplitude of the signal component decreases, and the noise component increases. At high frequency, the output consists primarily of quantization noise. In essence, the analog filter has a lowpass effect on the signal, and a highpass effect on the quantization noise. Thus the analog filter performs the noise shaping function in the $\Sigma\Delta$ modulator model.

For a given input frequency, higher order analog filters offer more attenuation. The same is true of $\Sigma\Delta$ modulators, provided certain precautions are taken.

By using more than one integration and summing stage in the $\Sigma\Delta$ modulator, we can achieve higher orders of quantization noise shaping and even better ENOB for a given over-sampling ratio as is shown in Figure 8.22 for both a first and second-order $\Sigma\Delta$ modulator. The block diagram for the second-order $\Sigma\Delta$ modulator is shown in Figure 8.23. Third, and higher, order $\Sigma\Delta$ ADCs were once thought to be potentially unstable at some values of input - recent analyses using *finite* rather than infinite gains in the comparator have shown that this is not necessarily so, but even if instability does start to occur, it is not important, since the DSP in the

digital filter and decimator can be made to recognize incipient instability and react to prevent it.

Figure 8.24 shows the relationship between the order of the $\Sigma\Delta$ modulator and the amount of over-sampling necessary to achieve a particular SNR. For instance, if the oversampling ratio is 64, an ideal second-order system is capable of providing an SNR of about 80dB. This implies approximately 13 effective number of bits (ENOB). Although the filtering done by the digital filter and decimator can be done to any degree of precision desirable, it would be pointless to carry more than 13 binary bits to the outside world. Additional bits would carry no useful signal information, and would be buried in the quantization noise unless post-filtering techniques were employed.

Figure 8.22

Figure 8.23

Figure 8.24

The $\Sigma\Delta$ ADCs that we have described so far contain integrators, which are low pass filters, whose passband extends from DC. Thus, their quantization noise is pushed up in frequency. At present, most commercially available $\Sigma\Delta$ ADCs are of this type (although some which are intended for use in audio or telecommunications applications contain bandpass rather than lowpass digital filters to eliminate any system DC offsets). Sigma-delta ADCs are available with resolutions up to 24-bits for DC measurement applications (AD77XX-family), and with resolutions of 18-bits for high quality digital audio applications (AD1879).

But there is no particular reason why the filters of the $\Sigma\Delta$ modulator should be LPFs, 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 (BPFs), the quantization noise is moved up and down in frequency to leave a virtually noise-free region in the pass-band (see Reference 1). If the digital filter is then programmed to have its pass-band in this region, we have a $\Sigma\Delta$ ADC with a bandpass, rather than a lowpass characteristic. Although studies of this architecture are in their infancy, such ADCs would seem to be ideally suited for use in digital radio receivers, medical ultrasound, and a number of other applications.

A $\Sigma\Delta$ ADC works by over-sampling, where simple analog filters in the $\Sigma\Delta$ modulator shape the quantization noise so that the SNR *in the bandwidth of interest* is much greater than would otherwise be the case, and by using high performance digital filters and decimation to eliminate noise outside the required passband. Because the analog circuitry is so simple and undemanding, it may be built with the same digital VLSI process that is used to fabricate the DSP circuitry of the digital filter. Because the basic ADC is 1-bit (a comparator), the technique is inherently linear.

Although the detailed analysis of $\Sigma\Delta$ ADCs involves quite complex mathematics, their basic design can be understood without the necessity of any mathematics at all. For further discussion on $\Sigma\Delta$ ADCs, refer to References 2 and 3.

HIGH RESOLUTION, LOW-FREQUENCY SIGMA-DELTA MEASUREMENT ADCs

The AD7710, AD7711, AD7712, AD7713, and AD7714, AD7730, and AD7731 are members of a family of sigma-delta converters designed for high accuracy, low frequency measurements. They have no missing codes to 24-bits, and their effective resolutions extend to 22.5 bits depending upon the device, update rate, programmed filter bandwidth, PGA gain, post-filtering, etc. They all use similar sigma-delta cores, and their main differences are in their analog inputs, which are optimized for different transducers. Newer members of the family, such as the AD7714, AD7730/7730L, and the AD7731/7731L are designed and specified for single supply operation.

There are also similar 16-bit devices available (AD7705, AD7706, AD7715) which also operate on single supplies.

The AD1555/AD1556 is a 24-bit two-chip $\Sigma\Delta$ modulator/filter specifically designed for seismic data acquisition systems. This combination yields a dynamic range of 120dB. The AD1555 contains a PGA and a 4th-order $\Sigma\Delta$ modulator. The AD1555 outputs a serial 1-bit data stream to the AD1556 which contains the digital filter and decimator.

Because of the high resolution of these converters, the effects of noise must be fully understood and how it affects the ADC performance. This discussion also applies to ADCs of lower resolution, but is particularly important when dealing with 16-bit or greater $\Sigma\Delta$ ADCs.

Figure 8.25 shows the output code distribution, or histogram, for a typical high resolution ADC with a DC, or "grounded" input centered on a code. If there were no noise sources present, the ADC output would always yield the same code, regardless of how many samples were taken. Of course, if the DC input happened to be in a transition zone between two adjacent codes, then the distribution would be spread between these two codes, but no further. Various noise sources internal to the converter, however, cause a distribution of codes around a primary one as shown in the diagram.

This noise in the ADC is generated by unwanted signal coupling and by components such as resistors (Johnson noise) and active devices like switches (kT/C noise). In addition, there is residual quantization noise which is not removed by the digital filter. The total noise can be considered to be an input noise source which is summed with the input signal into an ideal noiseless ADC. It is sometimes called *input-referred noise*, or *effective input noise*. The distribution of the noise is primarily gaussian, and therefore an RMS noise value can be determined (i.e., the standard deviation of the distribution).

Figure 8.25

In order to characterize the input-referred noise, we introduce the concept of *Effective Resolution*, sometimes referred to as effective number of bits (ENOB). It should be noted, however, that ENOB is most often used to describe the dynamic

performance of higher speed ADCs with AC input signals, and is not often used with respect to precision low frequency $\Sigma\Delta$ ADCs.

Effective Resolution is defined by the following equation:

$$\text{Effective Resolution} = \log_2 \left[\frac{\text{Fullscale Range}}{\text{RMS Noise}} \right] \text{Bits} .$$

Noise-Free Code Resolution is defined by:

$$\text{Noise Free Code Resolution} = \log_2 \left[\frac{\text{Fullscale Range}}{\text{Peak to Peak Noise}} \right] \text{Bits} .$$

Peak-to-peak noise is approximately 6.6 times the RMS noise, so Noise-Free Code Resolution can be expressed as:

$$\begin{aligned} \text{Noise Free Code Resolution} &= \log_2 \left[\frac{\text{Fullscale Range}}{6.6 \times \text{RMS Noise}} \right] \text{Bits} . \\ &= \text{Effective Resolution} - 2.72 \text{ Bits} \end{aligned}$$

Noise Free Code Resolution is therefore the maximum number of ADC bits that can be used and still always get a single-code output distribution for a DC input placed on a code center, i.e., there is no *code flicker*. This does not say that the rest of the LSBs are unusable, it is only a way to define the noise amplitude and relate it to ADC resolution. It should also be noted that additional external post-filtering and averaging of the ADC output data can further reduce input referred noise and increase the effective resolution.

Figure 8.26

The AD7730 is one of the newest members of the AD77XX family and is shown in Figure 8.27. This ADC was specifically designed to interface directly to bridge outputs in weigh scale applications. The device accepts low-level signals directly from a bridge and outputs a serial digital word. There are two buffered differential inputs which are multiplexed, buffered, and drive a PGA. The PGA can be programmed for four differential unipolar analog input ranges: 0V to +10mV, 0V to +20mV, 0V to +40mV, and 0V to +80mV and four differential bipolar input ranges: $\pm 10\text{mV}$, $\pm 20\text{mV}$, $\pm 40\text{mV}$, and $\pm 80\text{mV}$. The maximum peak-to-peak, or noise-free resolution achievable is 1 in 230,000 counts, or approximately 18-bits. It should be noted that the noise-free resolution is a function of input voltage range, filter cutoff, and output word rate. Noise is greater using the smaller input ranges where the PGA gain must be increased. Higher output word rates and associated higher filter cutoff frequencies will also increase the noise.

The analog inputs are buffered on-chip allowing relatively high source impedances. Both analog channels are differential, with a common mode voltage range that comes within 1.2V of AGND and 0.95V of AVDD. The reference input is also differential, and the common mode range is from AGND to AVDD.

ADCs FOR SIGNAL CONDITIONING

The 6-bit DAC is controlled by on-chip registers and can remove TARE (pan weight) values of up to $\pm 80\text{mV}$ from the analog input signal range. The resolution of the TARE function is 1.25mV with a $+2.5\text{V}$ reference and 2.5mV with a $+5\text{V}$ reference.

The output of the PGA is applied to the $\Sigma\Delta$ modulator and programmable digital filter. The serial interface can be configured for three-wire operation and is compatible with microcontrollers and digital signal processors. The AD7730 contains self-calibration and system-calibration options and has an offset drift of less than $5\text{nV}/^\circ\text{C}$ and a gain drift of less than $2\text{ppm}/^\circ\text{C}$. This low offset drift is obtained using a *chop* mode which operates similarly to a chopper-stabilized amplifier.

The oversampling frequency of the AD7730 is 4.9152MHz , and the output data rate can be set from 50Hz to 1200Hz . The clock source can be provided via an external clock or by connecting a crystal oscillator across the MCLK IN and MCLK OUT pins.

The AD7730 can accept input signals from a DC-excited bridge. It can also handle input signals from an AC-excited bridge by using the AC excitation clock signals (ACX and $\overline{\text{ACX}}$). These are non-overlapping clock signals used to synchronize the external switches which drive the bridge. The ACX clocks are demodulated on the AD7730 input.

The AD7730 contains two 100nA constant current generators, one source current from AVDD to AIN(+) and one sink current from AIN(-) to AGND. The currents are switched to the selected analog input pair under the control of a bit in the Mode Register. These currents can be used in checking that a sensor is still operational before attempting to take measurements on that channel. If the currents are turned on and a fullscale reading is obtained, then the sensor has gone open circuit. If the measurement is 0V , the sensor has gone short circuit. In normal operation, the burnout currents are turned off by setting the proper bit in the Mode Register to 0.

Figure 8.27

Figure 8.28

The AD7730 contains an internal programmable digital filter. The filter consists of two sections: a first stage filter, and a second stage filter. The first stage is a sinc^3 lowpass filter. The cutoff frequency and output rate of this first stage filter is programmable. The second stage filter has three modes of operation. In its normal mode, it is a 22-tap FIR filter that processes the output of the first stage filter. When a step change is detected on the analog input, the second stage filter enters a second mode (FASTStep™) where it performs a variable number of averages for some time after the step change, and then the second stage filter switches back to the FIR filter mode. The third option for the second stage filter (SKIP mode) is that it is completely bypassed so the only filtering provided on the AD7730 is the first stage. Both the FASTStep mode and SKIP mode can be enabled or disabled via bits in the control register.

Figure 8.29 shows the full frequency response of the AD7730 when the second stage filter is set for normal FIR operation. This response is with the chop mode enabled

and an output word rate of 200Hz and a clock frequency of 4.9152MHz. The response is shown from DC to 100Hz. The rejection at $50\text{Hz} \pm 1\text{Hz}$ and $60\text{Hz} \pm 1\text{Hz}$ is better than 88dB.

Figure 8.30 shows the step response of the AD7730 with and without the FASTStep mode enabled. The vertical axis shows the code value and indicates the settling of the output to the input step change. The horizontal axis shows the number of output words required for that settling to occur. The positive input step change occurs at the 5th output. In the normal mode (FASTStep disabled), the output has not reached its final value until the 23rd output word. In FASTStep mode with chopping enabled, the output has settled to the final value by the 7th output word. Between the 7th and the 23rd output, the FASTStep mode produces a settled result, but with additional noise compared to the specified noise level for normal operating conditions. It starts at a noise level comparable to the SKIP mode, and as the averaging increases ends up at the specified noise level. The complete settling time required for the part to return to the specified noise level is the same for FASTStep mode and normal mode.

Figure 8.29

Figure 8.30

The FASTStep mode gives a much earlier indication of where the output channel is going and its new value. This feature is very useful in weigh scale applications to give a much earlier indication of the weight, or in an application scanning multiple channels where the user does not have to wait the full settling time to see if a channel has changed.

Note, however, that the FASTStep mode is not particularly suitable for multiplexed applications because of the excess noise associated with the settling time. For multiplexed applications, the full 23-cycle output word interval should be allowed for settling to a new channel. This points out the fundamental issue of using $\Sigma\Delta$ ADCs in multiplexed applications. There is no reason why they won't work, provided the internal digital filter is allowed to settle fully after switching channels.

The calibration modes of the AD7730 are given in Figure 8.31. A calibration cycle may be initiated at any time by writing to the appropriate bits of the Mode Register. Calibration removes offset and gain errors from the device.

Figure 8.31

The AD7730 gives the user access to the on-chip calibration registers allowing an external microprocessor to read the device's calibration coefficients and also to write its own calibration coefficients to the part from prestored values in external E²PROM. This gives the microprocessor much greater control over the AD7730's calibration procedure. It also means that the user can verify that the device has performed its calibration correctly by comparing the coefficients after calibration with prestored values in E²PROM. Since the calibration coefficients are derived by performing a conversion on the input voltage provided, the accuracy of the calibration can only be as good as the noise level the part provides in the normal mode. To optimize calibration accuracy, it is recommended to calibrate the part at

ADCs FOR SIGNAL CONDITIONING

its lowest output rate where the noise level is lowest. The coefficients generated at any output rate will be valid for all selected output update rates. This scheme of calibrating at the lowest output data rate does mean that the duration of the calibration interval is longer.

The AD7730 requires an external voltage reference, however, the power supply may be used as the reference in the ratiometric bridge application shown in Figure 8.32. In this configuration, the bridge output voltage is directly proportional to the bridge drive voltage which is also used to establish the reference voltages to the AD7730. Variations in the supply voltage will not affect the accuracy. The SENSE outputs of the bridge are used for the AD7730 reference voltages in order to eliminate errors caused by voltage drops in the lead resistances.

Figure 8.32

The AD7730 has a high impedance input buffer which isolates the analog inputs from switching transients generated in the PGA and the sigma-delta modulator. Therefore, no special precautions are required in driving the analog inputs. Other members of the AD77XX family, however, either do not have the input buffer, or if one is included on-chip, it can be switched either in or out under program control. Bypassing the buffer offers a slight improvement in noise performance. The equivalent input circuit of the AD77XX family without an input buffer is shown in Figure 8.33. The input switch alternates between the 10pF sampling capacitor and ground. The 7k Ω internal resistance, R_{INT} , is the on-resistance of the input multiplexer. The switching frequency is dependent on the frequency of the input clock and also the PGA gain. If the converter is working to an accuracy of 20-bits, the 10pF internal capacitor, C_{INT} , must charge to 20-bit accuracy during the time the switch connects the capacitor to the input. This interval is one-half the period of the switching signal (it has a 50% duty cycle). The input RC time constant due to the 7k Ω resistor and the 10pF sampling capacitor is 70ns. If the charge is to achieve 20-bit accuracy, the capacitor must charge for at least 14 time constants, or 980ns. Any external resistance in series with the input will increase this time constant. There are tables on the data sheets for the various AD77XX ADCs which give the maximum allowable values of R_{EXT} in order maintain a given level of accuracy. These tables should be consulted if the external source resistance is more than a few k Ω .

Figure 8.33

Simultaneous sampling of multiple channels is relatively common in data acquisition systems. If sigma-delta ADCs are used as shown in Figure 8.34, their outputs must be synchronized. Although the inputs are sampled at the same instant at a rate Kf_s , the decimated output word rate, f_s , is generally derived internally in each ADC by dividing the input sampling frequency by K . The output data must therefore be synchronized by the same clock at the f_s frequency. The SYNC input of the AD77XX family can be used for this purpose.

Figure 8.34

Products such as the AD7716 include multiple sigma-delta ADCs in a single IC, and provide the synchronization automatically. The AD7716 is a quad sigma-delta ADC

with up to 22-bit resolution and an input oversampling rate of 570kSPS. A functional diagram of the AD7716 is shown in Figure 8.35, and key specifications in Figure 8.36. The cutoff frequency of the digital filters (which may be changed during operation, but only at the cost of a loss of valid data for a short time while the filters clear) is programmed by the data written to the control register. The output word rate depends on the cutoff frequency chosen. The AD7716 contains an auto-zeroing system to minimize input offset drift.

Figure 8.35**Figure 8.36**

APPLICATIONS OF SIGMA-DELTA ADCs IN POWER METERS

While electromechanical energy meters have been popular for over 50 years, a solid-state energy meter delivers far more accuracy and flexibility. Just as important, a well designed solid-state meter will have a longer useful life. The AD7750 Product-to-Frequency Converter is the first of a family of ICs designed to implement this type of meter.

We must first consider the fundamentals of power measurement (see Figure 8.37). Instantaneous AC voltage is given by the expression $v(t) = V \times \cos(\omega t)$, and the current (assuming it is in phase with the voltage) by $i(t) = I \times \cos(\omega t)$. The *instantaneous power* is the product of $v(t)$ and $i(t)$:

$$p(t) = V \times I \times \cos^2(\omega t)$$

Using the trigonometric identity, $2\cos^2(\omega t) = 1 + \cos(2\omega t)$,

$$p(t) = \frac{V \times I}{2} [1 + \cos(2\omega t)] = \text{Instantaneous Power.}$$

The *instantaneous real power* is simply the average value of $p(t)$. It can be shown that computing the instantaneous real power in this manner gives accurate results even if the current is not in phase with the voltage (i.e., the power factor is not unity. By definition, the power factor is equal to $\cos\theta$, where θ is the phase angle between the voltage and the current). It also gives the correct real power if the waveforms are non-sinusoidal.

Figure 8.37

The AD7750 implements these calculations, and a block diagram is shown in Figure 8.38. There are two inputs to the device. The differential voltage between $V1+$ and $V1-$ is a voltage corresponding to the instantaneous current. It is usually derived from a small transformer placed in series with the line. The AD7750 is designed with a switched capacitor architecture that allows a bipolar analog input with a single +5V supply. The input voltage passes through a PGA which can be set for a gain of 1 or 16. The gain of 16 option allows for low values of shunt impedances in

ADCs FOR SIGNAL CONDITIONING

the current monitoring circuit. The output of the PGA drives a 2nd order 16-bit sigma-delta modulator which samples the signal at a 900kHz rate. The serial bit stream from the modulator is passed through a digital highpass filter to remove any DC component. The highpass filter has a phase lead of 2.58° at 50Hz. In order to equalize the phase difference between the two channels, a fixed delay of $143\mu\text{s}$ is then introduced in the signal path. Because the time delay is fixed, external phase compensation will be required if the line frequency differs from 50Hz. There are several ways to accomplish this, and they are described in detail in the AD7750 data sheet.

Figure 8.38

The differential voltage applied between $V2+$ and $V2-$ represents the voltage waveform (scaled to the AD7750 input range). It is passed through a gain of 2 amplifier and a second sigma-delta modulator. The voltage and current outputs are then multiplied digitally yielding the *instantaneous power*. The *instantaneous real power* is then obtained by passing the instantaneous power through a digital lowpass filter. The low frequency outputs F1 and F2 are generated by accumulating this real power information. The F1 and F2 outputs provide two alternating low-going pulses. This low frequency inherently means a long accumulation time between output pulses. The output frequency is therefore proportional to the *average real power*. This average real power information can in turn be accumulated (e.g., an electromechanical pulse counter or full stepping two phase stepper-motor) to generate real energy information. The pulse width is set at 275ms. The frequency of these pulses is 0Hz to about 14Hz.

Because of its high output frequency and hence shorter integration time, the FOUT output is proportional to the *instantaneous* real power. This is useful for system calibration purposes which would take place under steady load conditions.

The error in the real power measurement is less than 0.2% over a dynamic range of 500:1 and less than 0.4% over a dynamic range of 1000:1.

A single-phase power meter application is shown in Figure 8.39. The ground for the entire circuit is referenced to the neutral line. The +5V power for the circuit is derived from an AC to DC supply which is powered from the phase (hot) line. This can be simple half-wave diode rectifier followed by a filter capacitor. The F1 and F2 outputs drive the kW-Hr counter which displays the energy usage.

The REVP output (reverse polarity) drives an LED and goes high when negative power is detected (i.e., when the voltage and current signals are 180° out of phase). This condition would generally indicate a potential mis-wiring condition.

The AD7751 Energy Metering IC operates in a similar fashion to the AD7750 but has enhanced performance features. It has on-chip fault detection circuits which monitor the current in both the phase (hot) and neutral line. A fault is indicated when these currents differ by more than 12.5%, and billing is continued using the larger of the two currents.

Figure 8.39

REFERENCES

1. S. A. Jantzi, M. Snelgrove & P. F. Ferguson Jr., *A 4th-Order Bandpass Sigma-Delta Modulator*, **IEEE Journal of Solid State Circuits**, Vol. 38, No. 3, March 1993, pp.282-291.
2. **System Applications Guide**, Analog Devices, Inc., 1993, Section 14.
3. **Mixed Signal Design Seminar**, Analog Devices, Inc., 1991, Section 6.
4. **AD77XX-Series Data Sheets**, Analog Devices, <http://www.analog.com>.
5. **Linear Design Seminar**, Analog Devices, Inc., 1995, Section 8.
6. J. Dattorro, A. Charpentier, D. Andreas, *The Implementation of a One-Stage Multirate 64:1 FIR Decimator for use in One-Bit Sigma-Delta A/D Applications*, **AES 7th International Conference**, May 1989.
7. W.L. Lee and C.G. Sodini, *A Topology for Higher-Order Interpolative Coders*, **ISCAS PROC. 1987**.
8. P.F. Ferguson, Jr., A. Ganesan and R. W. Adams, *One Bit Higher Order Sigma-Delta A/D Converters*, **ISCAS PROC. 1990**, Vol. 2, pp. 890-893.
9. R. Koch, B. Heise, F. Eckbauer, E. Engelhardt, J. Fisher, and F. Parzefall, *A 12-bit Sigma-Delta Analog-to-Digital Converter with a 15MHz Clock Rate*, **IEEE Journal of Solid-State Circuits**, Vol. SC-21, No. 6, December 1986.
10. Wai Laing Lee, *A Novel Higher Order Interpolative Modulator Topology for High Resolution Oversampling A/D Converters*, **MIT Masters Thesis**, June 1987.
11. D. R. Welland, B. P. Del Signore and E. J. Swanson, *A Stereo 16-Bit Delta-Sigma A/D Converter for Digital Audio*, **J. Audio Engineering Society**, Vol. 37, No. 6, June 1989, pp. 476-485.
12. R. W. Adams, *Design and Implementation of an Audio 18-Bit Analog-to-Digital Converter Using Oversampling Techniques*, **J. Audio Engineering Society**, Vol. 34, March 1986, pp. 153-166.
13. B. Boser and Bruce Wooley, *The Design of Sigma-Delta Modulation Analog-to-Digital Converters*, **IEEE Journal of Solid-State Circuits**, Vol. 23, No. 6, December 1988, pp. 1298-1308.
14. Y. Matsuya, et. al., *A 16-Bit Oversampling A/D Conversion Technology Using Triple-Integration Noise Shaping*, **IEEE Journal of Solid-State Circuits**, Vol. SC-22, No. 6, December 1987, pp. 921-929.

15. Y. Matsuya, et. al., *A 17-Bit Oversampling D/A Conversion Technology Using Multistage Noise Shaping*, **IEEE Journal of Solid-State Circuits**, Vol. 24, No. 4, August 1989, pp. 969-975.
16. P. Ferguson, Jr., A. Ganesan, R. Adams, et. al., *An 18-Bit 20-kHz Dual Sigma-Delta A/D Converter*, **ISSCC Digest of Technical Papers**, February 1991.
17. Steven Harris, *The Effects of Sampling Clock Jitter on Nyquist Sampling Analog-to-Digital Converters and on Oversampling Delta Sigma ADCs*, **Audio Engineering Society Reprint 2844 (F-4)**, October, 1989.
18. Max W. Hauser, *Principles of Oversampling A/D Conversion*, **Journal Audio Engineering Society**, Vol. 39, No. 1/2, January/February 1991, pp. 3-26.
19. *Designing a Watt-Hour Energy Meter Based on the AD7750*, AN-545, Analog Devices, Inc., <http://www.analog.com>.
20. Daniel H. Sheingold, **Analog-Digital Conversion Handbook**, Third Edition, Prentice-Hall, 1986.