Getting the Most out of Delta-Sigma Converters

by Russell Anderson
Texas Instruments Incorporated

Data converters have been making steady improvements in resolution and speed. I can remember sitting down in a meeting about 25 years ago at Tektronix and deciding as a group how data converters would progress in the future. I don't think we even considered the possibility that we would make a jump from 16 bits of resolution to 24 bits. But the architecture of the delta-sigma converter allows just such a significant improvement.

Delta-Sigma converters have the promise of delivering conversion results of 24 bits. As exciting as this sounds, there are many parameters of operation that need to be correctly selected to achieve the optimal results. The same data rate varies in performance as the decimation, modulation clock and PGA are adjusted. Understanding these trade-offs can go a long way in optimizing data conversion results. Additional areas of concern include input source impedance, filter response, anti-aliasing, and long term drift.

Introduction to Delta-Sigma Converters.

The benefit of a delta-sigma converter is that it moves most of the conversion process into the digital domain. This makes it easier to combine high-performance analog with digital processing. The analog components use a single comparator, integrator and 1-bit DAC. Since the 1-bit DAC has only two outputs, it is linear across the voltage range. This high level of linearity is one reason that delta-sigma converters achieve high levels of precision. The resulting absolute accuracy is mainly dependent on the accuracy of the voltage reference.

Delta-Sigma Modulator

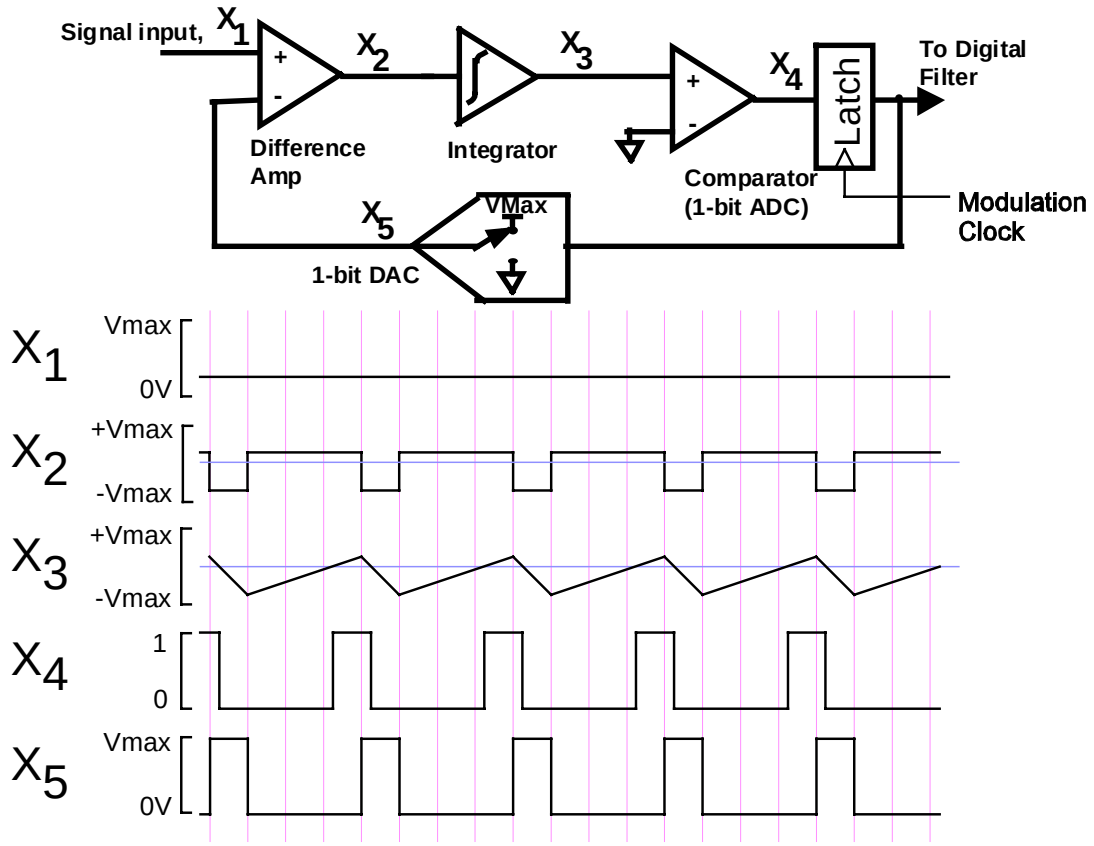

Fig. 1: Delta-Sigma Modulator

Let us review the waveforms found in a simple delta-sigma modulator (see Fig. 1.) The input signal, X1, is at 1/4 scale. The input signal minus the DAC output signal (X5) is a pulse train with one period low and three periods high (X2.) The latched comparator output (X4) is the serial bit stream that is fed to the digital filters with the ratio of ones to zeros directly correlating with the ratio of the input voltage to the full-scale input range.

Each of the vertical lines align with where the comparator output is latched by the modulation clock. To analyze the operation, it is best to start with the output as if it were the driving signal and then close the loop later. The input voltage is ¼ Vmax. The DAC is controlled by the digital output and so it starts with an output of Vmax. The difference of Vmax and the input gives -¾ Vmax, which is applied to the integrator. As can be seen, this negative voltage causes the integrator to have a steep negative slope.

At the next clock, since X3 is negative, the output at X4 is zero. This is latched which causes the DAC to now output zero volts and the difference at X_2 is only a +¼ Vmax. As can be seen, this smaller positive slope takes several cycles before it crosses the comparator threshold. The positive integration continues to ramp positive until the next clock cycle, which then latches a one into the output and we are back where we started.

If we look at the frequency response of the delta-sigma modulator we have a response that is characterized by:

$$y = \frac{x}{f+1} + \frac{Qf}{f+1} \qquad x = input,\ Q = quantization\ noise$$

It can be noted that at low frequencies the output equals the input (x) while at high frequencies the output equals the quantization noise giving the noise spectrum of Fig. 2.

Delta-sigma converters use oversampling to spread the quantization noise across more frequency bins, and in conjunction with the delta-sigma modulator, shapes this noise so that most of it is not included in the signal measurement band. Shaping allows a low-pass digital filter to remove most of the noise and yield high precision voltage measurements.

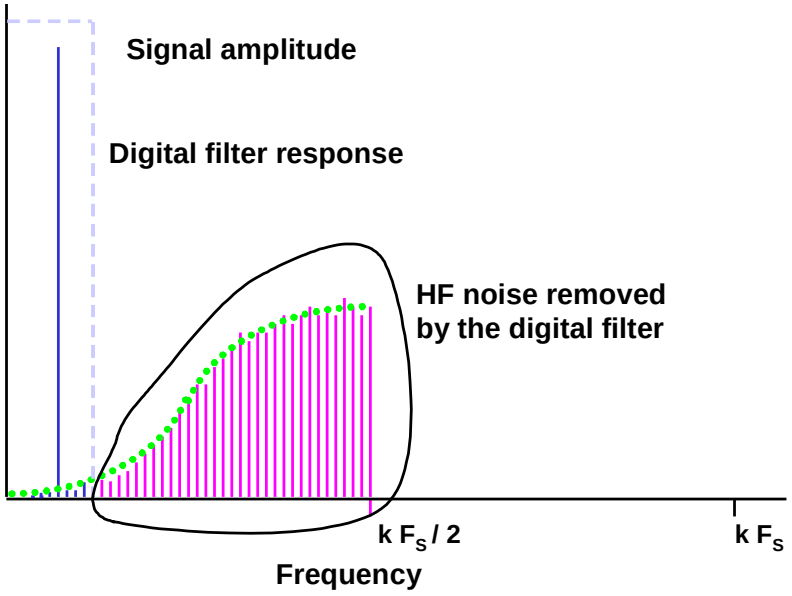

Fig. 2: Shaped Noise Spectrum

The output from the modulator goes into digital filters where the response is adjusted by the type of filter or the amount of decimation. The final output data rate is determined by:

Data rate = modulation clock ÷ decimation rate

ENOB

One figure of merit for an ADC is to express the noise as a ratio of the full-scale (FS) signal to rms noise, expressed as the Effective Number Of Bits (ENOB.) Using the standard deviation (σ) of the number of output codes we have, for a 24 bit converter:

$$2^{ENOB} = \frac{FullscaleV}{\sigma} = \frac{2^{24}}{\sigma}$$

Solving for ENOB:

$$\log_2\left(2^{ENOB}\right) = \log_2\left(\frac{2^{24}}{\sigma}\right) = \log_2\left(2^{24}\right) - \log_2(\sigma),$$

$$ENOB = 24 - \log_2(\sigma)$$

Or, if signal-to-noise ratio (SNR) is measured in dB, we can use:

$$ENOB = (SNR_{meas\ dB} - 1.76_{dB})/6.02_{dB}$$

A common type of filter used in delta-sigma converters are sinc filters. These have a deep attenuation notch at the output data rate and multiples of that rate meaning that a data rate of 60 Hz will effectively eliminate any 60 Hz signals from the measurement, and a10-Hz rate will eliminate both 50 and 60 Hz signals.

The ratio of the frequency of the input sample rate and the output data rate can be adjusted. This decimation ratio directly affects the Effective Number Of Bits (ENOB.) As the number of input samples per output result are increased, the ENOB improves and the resolution of the ADC effectively increases.

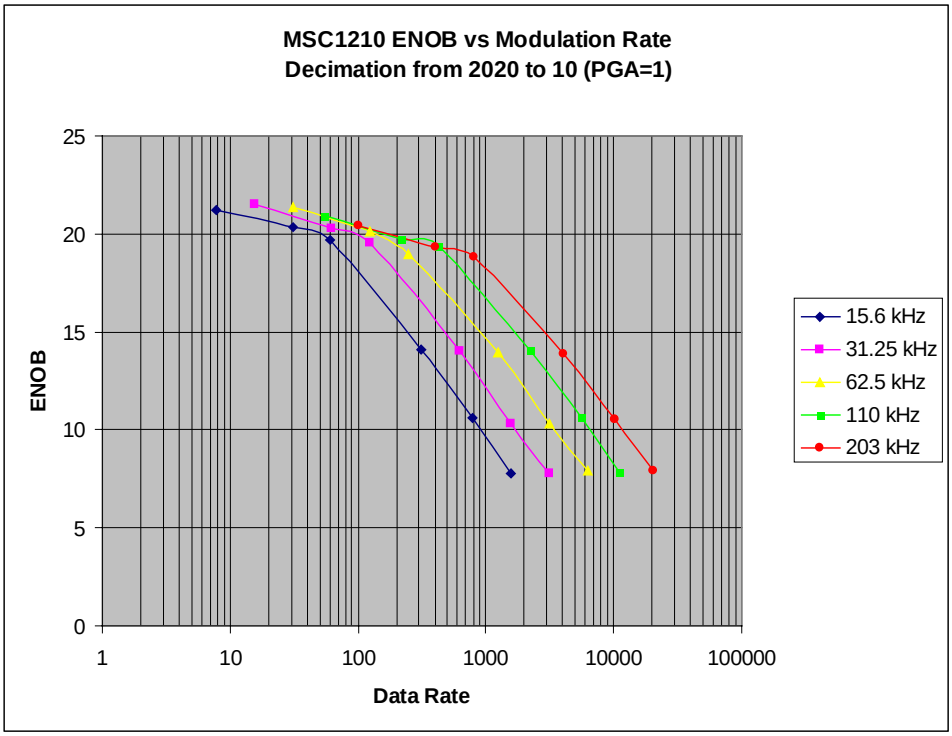

Fig. 3: ENOB Vs. Modulation Rates In MSC1210

Some delta-sigma converters have a fixed data output rate which can only be adjusted in a narrow range, while in others there is some flexibility allowing adjustment of the decimation ratio with some adjustment of the modulator clock rate. When combined with a 8051 microprocessor (in TI's MSC1210) there is even more flexibility in how these parameters interact, and we can easily adjust and evaluate the performance of a converter at various modulation clocks and decimation values. Each line (see Fig. 3) is a different clock rate while the points on the lines are decimation ratios of 2020, 500, 255, 50, 20 and 10. Note that measurement of ENOB is essentially determined by the decimation ratio, and a specified level of performance can be changed by adjusting the modulation

clock. As might be expected, at the highest modulation clock rates there is some reduction in the ENOB performance for the highest decimation ratios.

So the question can be asked, if the performance doesn't change much for different clock rates, why not just use the highest rates and get faster data conversion results? One reason is that power in CMOS circuits goes up quickly as clock rate increases.

If power isn't an issue then a number of samples at a faster output rate can be averaged together to further improve the level of performance. This is easier in the MSC1210 with a 32-bit accumulator that can be set to average 256 samples without processor interaction.

Input Impedance & Chopper Stabilization

The analog input of a delta-sigma converter can be viewed as a switch and a capacitor. The switching frequency has the effect of simulating a resistor continuously connected to the internal capacitor with its impedance and, hence, the input impedance of the converter directly related to the switching frequency. For the MSC1210 the input impedance is:

$$INPEDANCE = \left(\frac{1MHz}{64 \bullet SAMPLE\_CLOCK} \right) \bullet \left(\frac{5M\Omega}{PGA} \right)$$

With a sample rate of 15.625 kHz and PGA of 1, the input impedance is 5 MΩ. Higher sample rates and PGA values will reduce that value and to eliminate this effect many delta-sigma converters provide a buffer on-chip. Even with the buffer enabled there is still some sampling of the input signal which is a chopping to provide high dc accuracy.

Programmable Gain Amplifiers (PGAs)

Many delta-sigma converters provide on-chip PGAs but they all do not provide the same, or expected, benefits. Some of the higher gains are effectively just a shifting of the digital data or a multiplication by 2 with is essentially no benefit. These can be observed by carefully examining the data sheet. If increasing the PGA by a factor of 2 similarly decreases the ENOB there has been no real net gain and it just means that the noise covers more levels in the output.

Sometimes gain can be increased by using a smaller reference voltage as it determines the FS signal range. Decreasing the reference by 50% gains the input signal by two, but this type of gain improvement runs into noise limitations for a low reference voltage.

Settling Time

Settling time is another issue that might be important for data throughput in a multiple-channel system. To achieve high performance the delta-sigma converter will usually use a FIR filter such as a $sinc^3$ filter, with the advantage that the signal delay through the filter is fixed; it can also be easily adjusted to achieve variable decimation levels using additional stages of delayed samples. With more stages the output data rate is slower and a $sinc^3$ filter will take three conversion cycles to fully settle to the expected precision.

Settling time has the effect that the first few samples after the channel is switched will have settling error which is solved in the MSC1210 by having three types of filters and an auto-mode that selects the best filter after channel changing: For the first two samples after the multiplexer is changed the fast settling filter is used, then a $sinc^2$ filter and finally the $sinc^3$ filter for all future samples; so all conversion results are fully settled.

For a multiplexed data system one way to get around the problem of settling time is to run the data rate much faster and average the outputs. Assume, for example, a desire to have a 60 Hz data output rate to get the benefits of the 60 Hz filtering notch. This could be achieved with a 240 Hz sampling rate with four samples averaged to give the resultant 60 Hz data rate with the advantage that now the filter settling time has been reduced from up to 4 samples (non-synchronized channel switch) at 60 Hz (66.6 ms) to 4 samples at 240 Hz (16.6 ms.) Settling is now one sample period of the 60 Hz data rate and retains the benefit of the 60-Hz filter notch. In the MSC1210 the 32-bit accumulator is setup to average 4 samples and the first result after switching the channels is discarded (assuming that the channel switch was synchronized to the 60-Hz output rate.)

Anti-Aliasing

There are two main types of filter responses used in data acquisition systems, flat pass-band and sinc. The flat pass-band filters have low attenuation up to the cut-off frequency, then a large stop-band attenuation until you get to the Nyquist rate. This makes it much easier to design an anti-aliasing filter because this Nyquist frequency is usually 64 times higher than the cut-off frequency. A simple R-C filter might be all that is required.

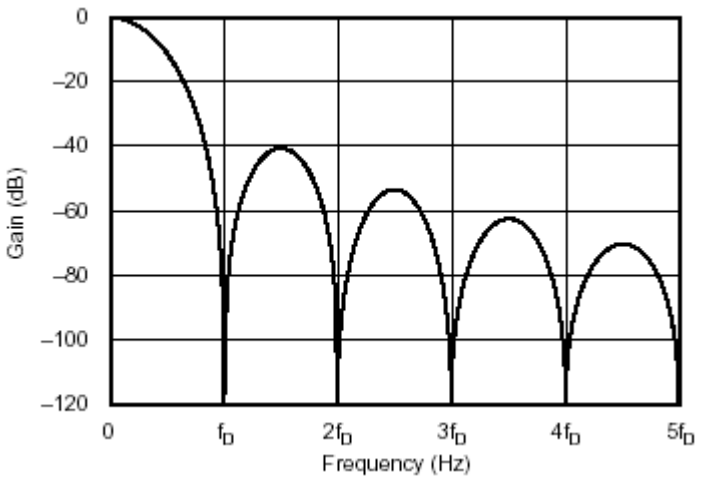

Fig. 4: Lobes Of A Sinc Filter

The other type of filter, the sinc filter, doesn't provide that same high attenuation from the data rate to the Nyquist rate (see Fig. 3) with several lobes that follow the sample rate. If you are trying to achieve 100-dB stop-band attenuation the filter must be designed to filter out those frequency components where the sinc filter attenuation is only down 40 dB. However, in designing the anti-aliasing filter it is important to remember that the

higher-frequency signals are not full amplitude. If the expected alias signal components are already at a maximum level of -20 dB, then to achieve 100 dB of attenuation for the sinc filter (Fig. 4 again) the anti-aliasing filter will only need to be down 40 dB. That is because the sinc filter provides 40 dB, the signal is assumed to be a maximum of -20 dB, which means that the anti-aliasing filter only needs to add an additional 40 dB of attenuation: But that can still be a significant requirement if you want the passband to include frequencies close to the data rate.

Drift

There are several noise sources for very low frequencies, one of them known as 1/f noise. The input chopping is effective at removing most of this noise but there are several other factors that can introduce low-frequency drift in high-performance systems. Care must be observed in how components are soldered on the board to avoid mechanical stresses and thermal gradients, thermocouple junctions, package orientations, etc. all can have an effect on signal quality that shows up as a drift. Techniques such as Allen Variance can be used to observe some of these effects and analyze success in removing them from the system.

About The Author

Russell Anderson is a Senior Applications Engineer with Texas Instruments' Data Acquisition Products group in Tucson, and specializes in microsystem converters (MSC1210.) He earned a BS at Brigham Young University in Provo, Utah. Anderson can be reached at anderson_russ@ti.com

as published in...
analog ZONE