Section 3

Applying Oversampling Data Converters

A Precision Temperature Control Circuit

Delta-Sigma Review

- Overview of Delta-Sigma
  - Delta-Sigma concept
  - 1/f noise
  - Quantization noise
  - Noise shaping and filtering
  - ENOB
- Delta-Sigma Products
- Applications
  - Temperature Measurement System
  - Ratiometric Measurements

The delta-sigma converter is used, because it shapes the noise in such a manner that most of the noise is moved from the passband to the stopband. A 1-bit sampling system has high quantization noise relative to the input signal, thus noise shaping filters are used to suppress the quantization noise. The quantization noise is spread over the frequency spectrum resulting in a spectrally filtered digital bit stream. This bit stream is processed by digital filters that take advantage of the shaped noise spectrum.

These converters are usually used in low frequency applications where accuracy and cost are important, but there are delta-sigma converters that have data output rates in excess of 1 MHz.

Delta-Sigma Overview

- What is a delta-sigma ADC?
  - A 1-bit converter that uses oversampling (can be multi-bit)
  - "Delta" = comparison with 1-bit DAC
  - "Sigma" = integration of the Delta measurement
- What is the advantage of delta-sigma?
  - Essentially digital parts which result in low cost
  - High resolution
- What are the disadvantages?
  - Limited frequency response
  - Most effective with continuous inputs
  - Latency

A delta-sigma ADC is a one-bit converter that uses oversampling implemented with a 1-bit ADC, a 1-bit DAC and an integrator to resolve the analog voltage into a digital equivalent. The basic conversion element of any converter must be more accurate than the required measurement. A 1-bit ADC is as linear as we can make because there are no differential non-linearities. The delta-sigma converter lends itself to low power supply or non-analog processes because the 1-bit converter can be made in any process at low supply voltages. Furthermore, a 1-bit ADC outputs a stream of ones and zeros which requires lots of data shifting, digital filtering, and logic, thus the 1-bit converter can be made efficiently from a digital process.

The benefit of delta-sigma is that it moves most of the process into the digital domain. 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.

The Delta-Sigma Modulator

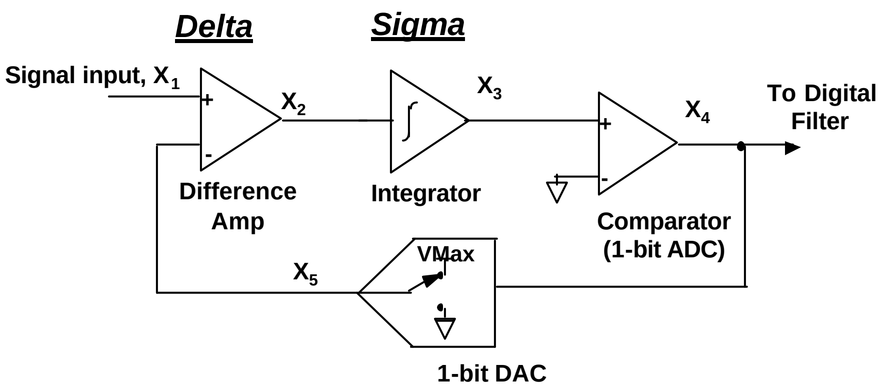

The delta-sigma modulator consists of a differential amplifier, an integrator, a comparator, and a 1-bit DAC. The input signal is subtracted from the 1-bit DAC signal, and the remainder is applied to the integrator input. When the system reaches steady state operation the integrator output signal is the sum of all the error voltages, and the integrator acting as a low-pass filter has lowered the noise content. One integrator yields –6dB noise suppression, so sometime several integrators are used to decrease the quantization noise in the passband. The integrator output signal is quantized using a 1-bit ADC (the comparator). The comparator output is a bit stream whose density of digital 1s is proportional to the ratio of the input signal to the reference signal. The DAC converts the comparator signal into a digital waveform which is compared with the input signal.

Delta-Sigma Modulator

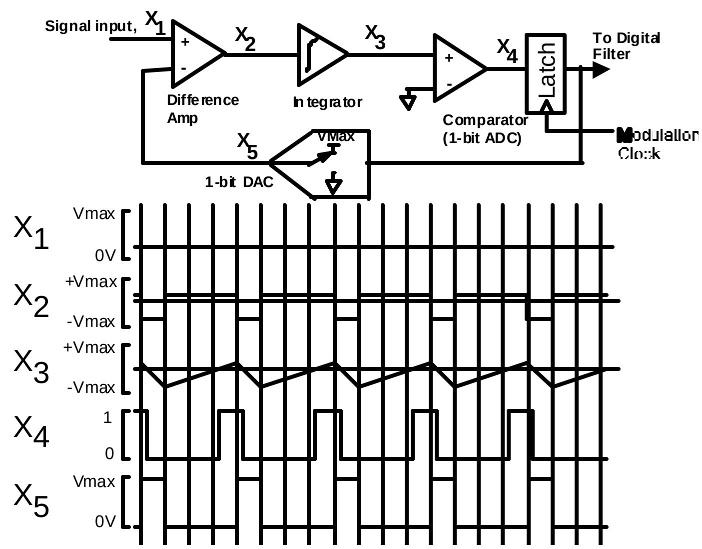

Now, let us review the waveforms found in a simple delta-sigma modulator. The input signal, X1, is at 1/4 scale. The input signal minus the DAC output signal is a pulse train with one period low and three periods high. The integrator falls for one period and rises for three periods. The only way the integrator output can fall for three periods is if that output is below ground for three periods, and looking at the comparator output, (X4) we see that this has to be the case to achieve steady state operation. The comparator output is the serial bit stream that contains the data conversion.

Each of the vertical lines represent where the comparator output is latched by the modulation clock. To analyze the operation, it is best to start with the output and see how it interacts with the input. The input voltage is 1/4 of Vmax range. We start with the output high, since the DAC follows the output it has an output of Vmax.

The initial difference amp has Vmax/4 and Vmax which creates an output of -3/4 Vmax. As can be seen, negative voltage causes the integrator to have a strong negative slope.

At the next clock, since X3 is negative, the output of X4 is zero. This is latched into the latch which causes the DAC to now output zero volts.

At zero volts on the DAC, the difference is only a positive 1/4 volt. As can be observed, it takes several cycles before it crosses the comparison threshold. You will notice that the integration cycles continue to ramp positive until the next clock cycle, which then latches a one into the output and we are back where we started.

Characteristics of Delta-Sigma

- Integration shapes the noise
- Errors that could reduce effectiveness:
  - Input offset voltage
  - Offset drift
  - Input bias
  - Sampling effects

The integrators shape the noise by spreading the quantization noise across the whole bandwidth. The shaped noise is pushed out from the data frequency, so a filter can be used to eliminate most of the shaped noise. There are several error sources that can reduce the overall effectiveness of the system. Most signals require some amplification and level shifting before they can fill the input span of an ADC. Sometimes the amplifier is internal to the ADC, and sometimes external amplifiers are used. In either case the errors introduced by the amplifier are input offset voltage, input offset voltage drift, input bias current (multiplies by an input resistor to make voltage noise), or sampling noise. The amplifier errors must be eliminated by adjustments or minimized to obtain accurate ADC conversions.

One of the benefits of integration is that high frequency noise both adds and subtracts from the low frequency or DC signal. Therefore it gives a result that is mainly influenced by the average low frequency portion of the input. This has the effect of greatly reducing the noise in the passband.

However there are errors which must be controlled.

1. Offset error would add a consistent imbalance into the integration.

2. Offset can be compensated, but only if the offset stays constant. A drift in offset could not be separated from movement of the signal.

3. Bias currents could cause offsets in the source impedance.

4. Sampling effects such as aliasing must be controlled.

Low Frequency Noise

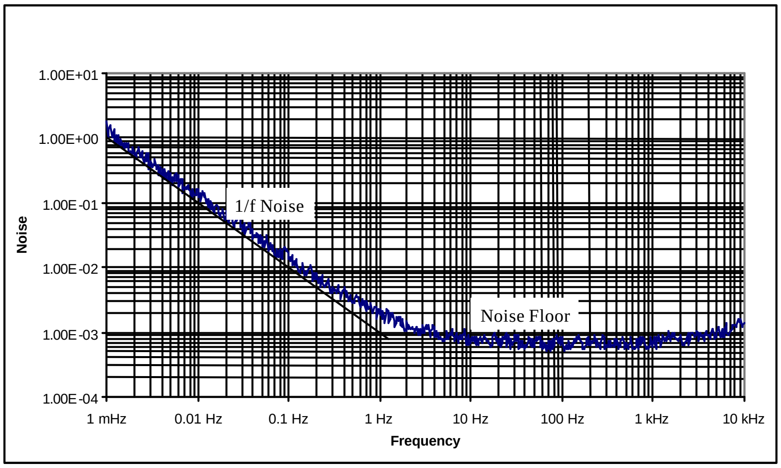

The noise curve of a typical semiconductor amplifier has two distinct regions called the 1/f noise and the noise floor. The majority of delta-sigma ADC applications are low frequency applications; hence; the 1/f noise often becomes dominant. The 1/f noise can be controlled by the proper amplifier selection.

If you look at the spectrum of noise you will get a chart that looks something like this. The noise of the device will set the noise floor with an upturn at the high frequencies. But at low frequencies we have what is known as 1/f noise. This shows the fallacy of the concept of a DC voltage. You will notice that the closer we get to DC, the bigger the noise becomes. You can usually think of this 1/f noise as drift. It is a very low frequency phenomena. One method to solve drift is through the use of a chopper input which we will examine in the next slide.

Chopper-Stabilized

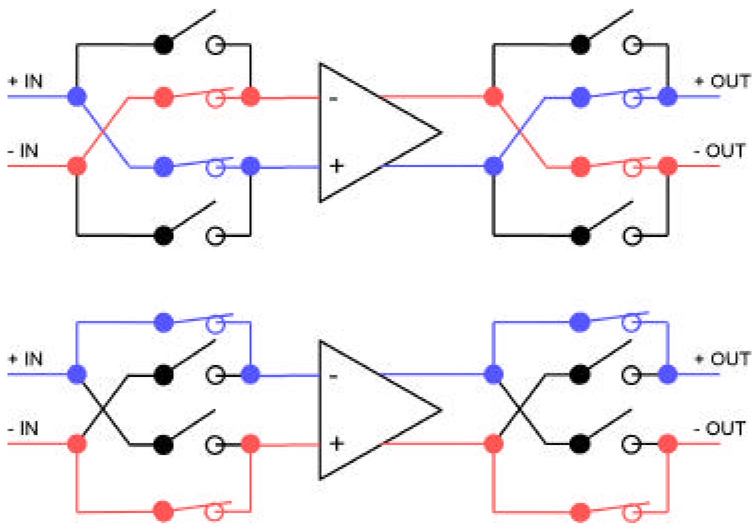

There are several methods for implementing a chopper-stabilized input. This is probably the easiest to visualize, but not the simplest to build.

Imagine that you have an offset in this differential op amp that adds 1 mV to the +input. In the top diagram that 1 mV appears on the +OUT signal.

In the bottom circuit the 1 mV is added to the x –OUT signal. Since it is alternately added to the positive and then the negative output, the net result is that on average the 1 mV offset doesn't appear on the output. This works particularly well for delta-sigma converters since they integrate the difference signal output.

Drift is a phenomenon where the amount of offset changes with time. The actual value of the offset doesn't matter in the chopper-stabilized circuit. Therefore, the fact that it might drift or change over time doesn't affect the result.

Sampling and Quantization

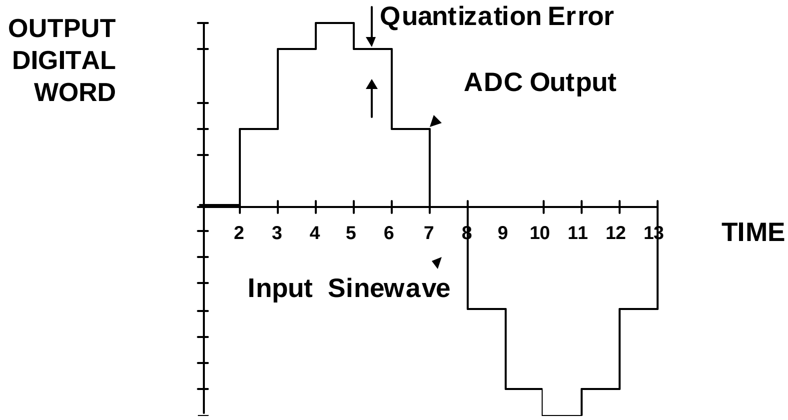

4-bit (16 level) ADC sampling a sinewave input, time domain

The time domain transfer function of a 4-bit ADC digitizing a full scale sine wave is shown here. The ADC output is drawn to correspond to the digital word. If this signal was reproduced with a DAC, the effects of sampling and quantization would be easily noticed.

Sampling means that the input waveform is only captured at certain discrete points in time. Between these points the output value is held constant. The rate at which the input is sampled is known as the sampling frequency. Nyquist theory tells us that the sampling frequency must be at least twice the information bandwidth of the input signal. Sampling faster than this minimum rate is known as oversampling. Delta-sigma ADCs take great advantage of oversampling capabilities--benefits will be shown later.

Quantization is the effect of the continuous analog input amplitude being divided into discrete levels (16 in this case). This means that most sampling points do not capture the exact input value. Rather, they capture the nearest quantized bin level. The difference between these two is the quantization error. The value of this quantization error is up to ±1/2 LSB for each sample. As you will notice, the DAC output will have more error than 1/2 LSB eventhough the A/D error is less than 1/2 LSB. This will be reduced for a higher sampling rate. If the input waveform is not correlated to the sampling frequency, the quantization errors are effectively random.

The Frequency Domain

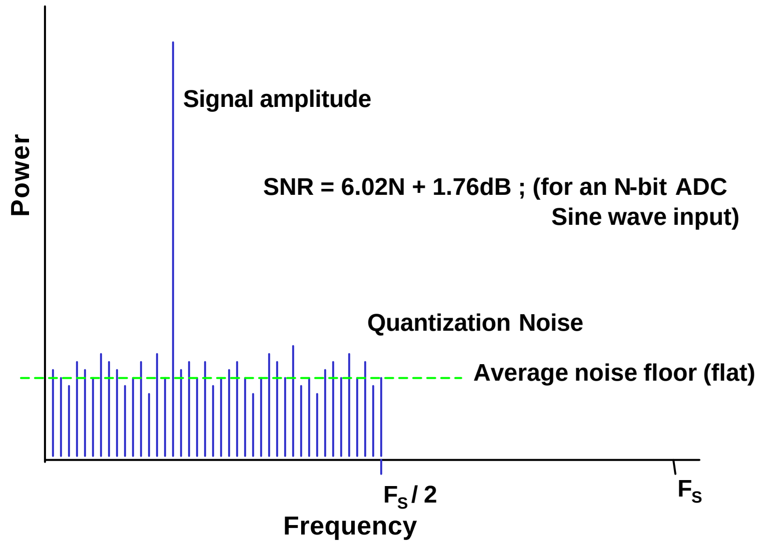

Having looked at the effects of sampling and quantization noise in the time domain, the frequency domain of quantization is examined next. This slide shows a diagram of the FFT of the waveform from the previous slide.

The signal being captured is a pure sine wave, which clearly shows up in a single frequency bin of the FFT. The FFT also shows a lot of random noise in all other bins, distributed flat across the frequency range from DC to $F_S/2$. This is quantization error and is known as quantization noise. Remember that every sample taken of the input signal has a quantization error and that the magnitude of this error is random up to ±1/2 LSB. It is this random error, present on every sample, that generates the noise floor in the FFT. The noise floor shown does not consider the inherent shortcomings of the ADC. ADCs have error sources, but in this case, we are looking at a perfect ADC.

Taking the RMS sum of all the frequency bins containing noise (i.e. all except the fundamental) and dividing this into the fundamental amplitude gives the signal-to-noise ratio (SNR). It can be shown that for an N-bit ADC, the SNR ratio is given by SNR = 6.02 N + 1.76 dB (for a full-scale sine wave input). This formula clearly shows that the obvious way to improve the SNR is to increase the number of bits in the ADC. Conversely, when the number of bits are increased, the reproduced input signal is more accurate.

Delta-sigma ADCs take another approach. This is to use a 1-bit ADC and by using the techniques of oversampling, noise shaping and filtering, improve the accuracy.

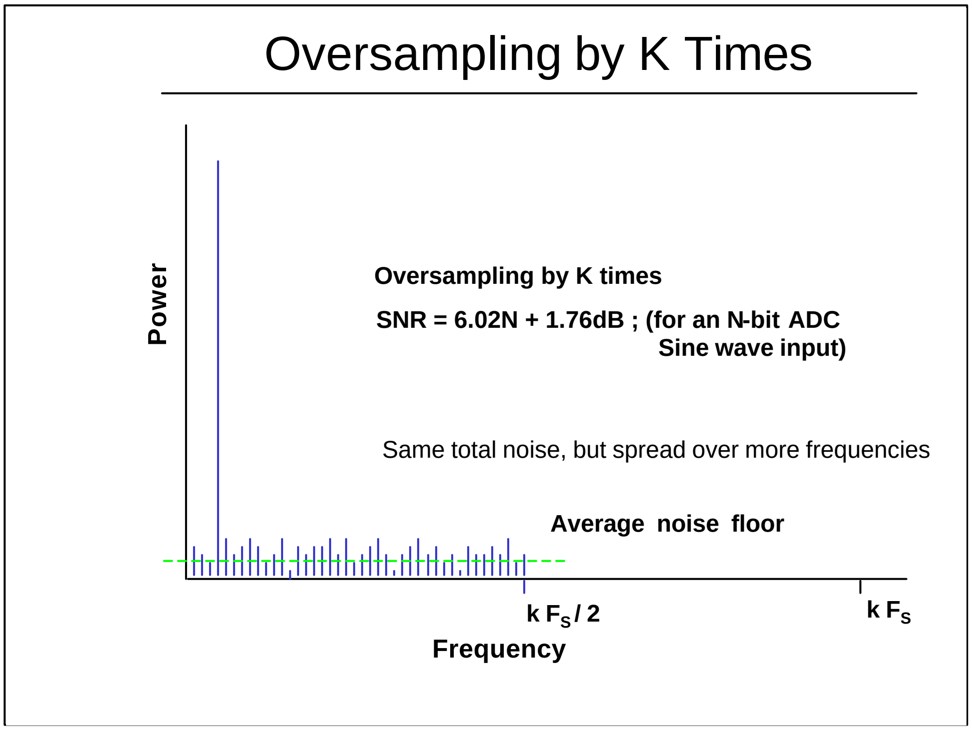

The total noise is independent of the oversampling; it stays constant. Sampling spreads the quantization noise over the sampled bandwidth. Oversampling increases the bandwidth considerably, hence it spreads the quantization noise over a wider bandwidth. The result is lower inband quantization noise, and this is the only noise that can't be filtered out.

In the previous example, the signal frequency was fairly close to the Nyquist frequency of $F_S/2$. Here the affects of oversampling can be seen. The signal frequency is the same, but the sampling frequency has been increased by an oversampling ratio of k to kF_S.

Notice how this FFT shows that the noise floor has dropped. It is important to realize that the SNR is still the same as it was before as the formula on the previous slide made no reference to sampling frequency. The total amount of noise energy is still the same. But as it has just been spread over a wider frequency range, the noise level in each frequency bin has been reduced.

The Digital Filter

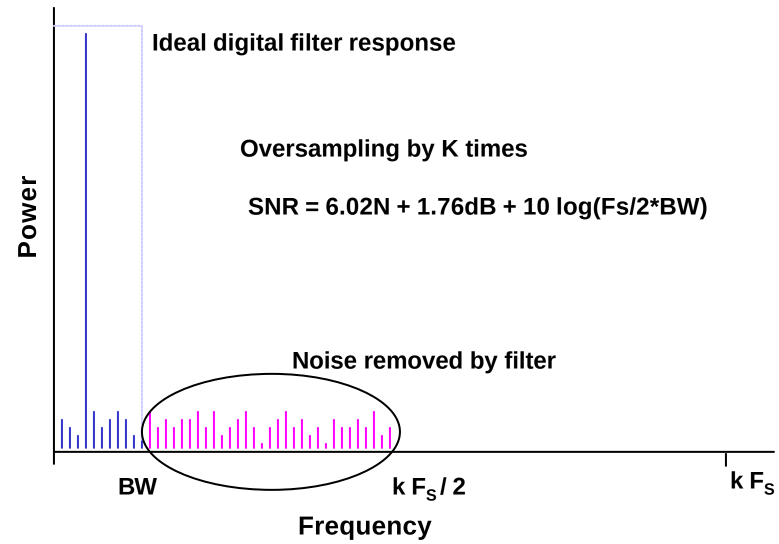

In the previous slide it was shown how oversampling spreads the noise over a wider bandwidth and hence reduces the level of the noise floor.

Delta-sigma converters make use of this effect by following the 1-bit ADC with a digital filter. The effect of this filter is to restrict the noise bandwidth. Since most of the noise cannot now pass through the digital filter, the RMS noise (i.e. the RMS sum of the noise in those frequency bins that can pass) is reduced.

This technique of spreading the noise over a wide frequency range, then filtering out most of the noise, is how a delta-sigma converter achieves a wide dynamic range from a low resolution ADC.

Noise-Shaped Spectrum

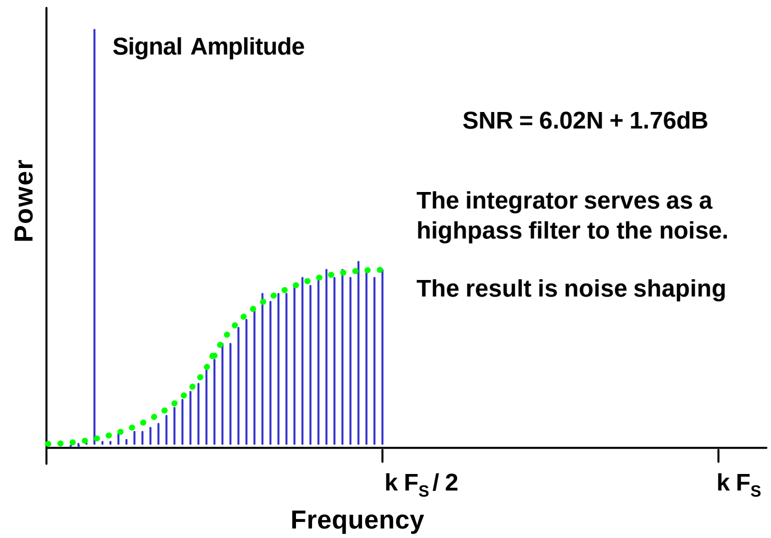

The effect of the integrator in the delta-sigma modulator is shown here. The noise spectrum can be seen to rise as the frequency increases. Note that the total noise power i.e. the RMS sum of all the frequency bins has not changed. There is no less total noise than in the case of simply oversampling, but the distribution of the noise has changed.

Filtering the Shaped Noise

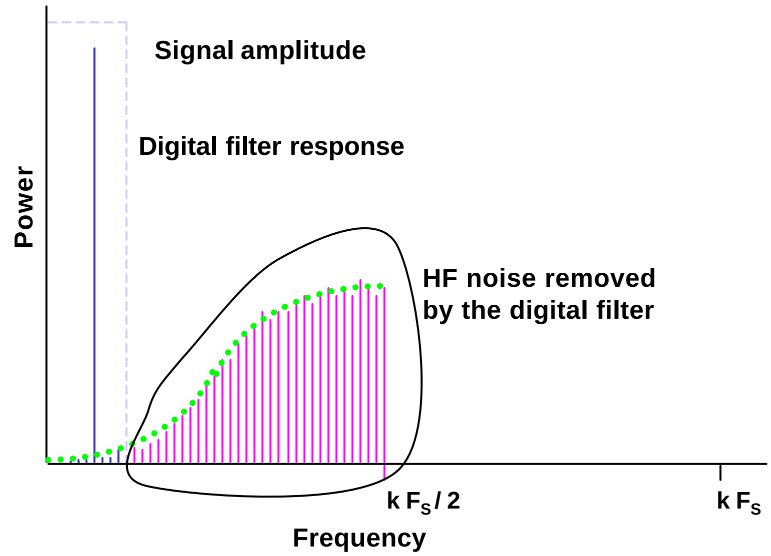

A digital filter is now applied to the noise shaped delta-sigma modulator. More noise is removed than in the simple oversampling case. The delta-sigma modulator just described is a first-order system and gives a 9-dB improvement in SNR for every doubling in sample rate. Compare this with a mere 3 dB achieved by oversampling alone. Using this architecture, it is now possible to implement a high-accuracy ADC from a single-bit delta-sigma modulator with a practical oversampling ratio.

The Delta-Sigma Modulator

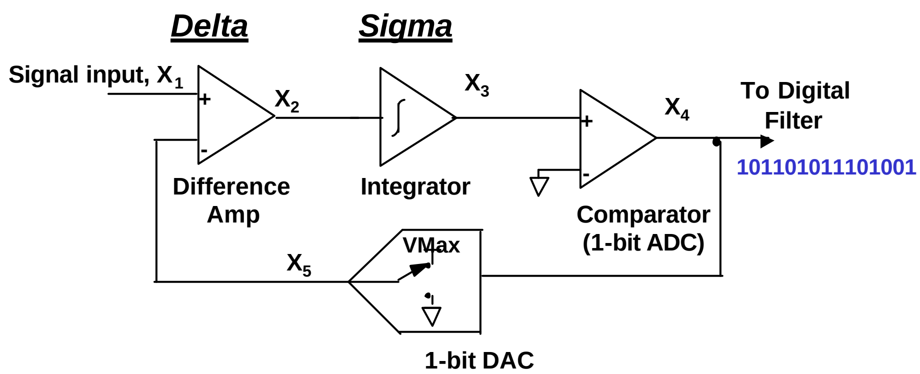

This returns us to our earlier view of the elements of the delta-sigma modulator. The digital values from the modulator go into a digital filter to be filtered and decimated into the high resolution value. The pattern of ones and zeros is a representation of the analog input voltage.

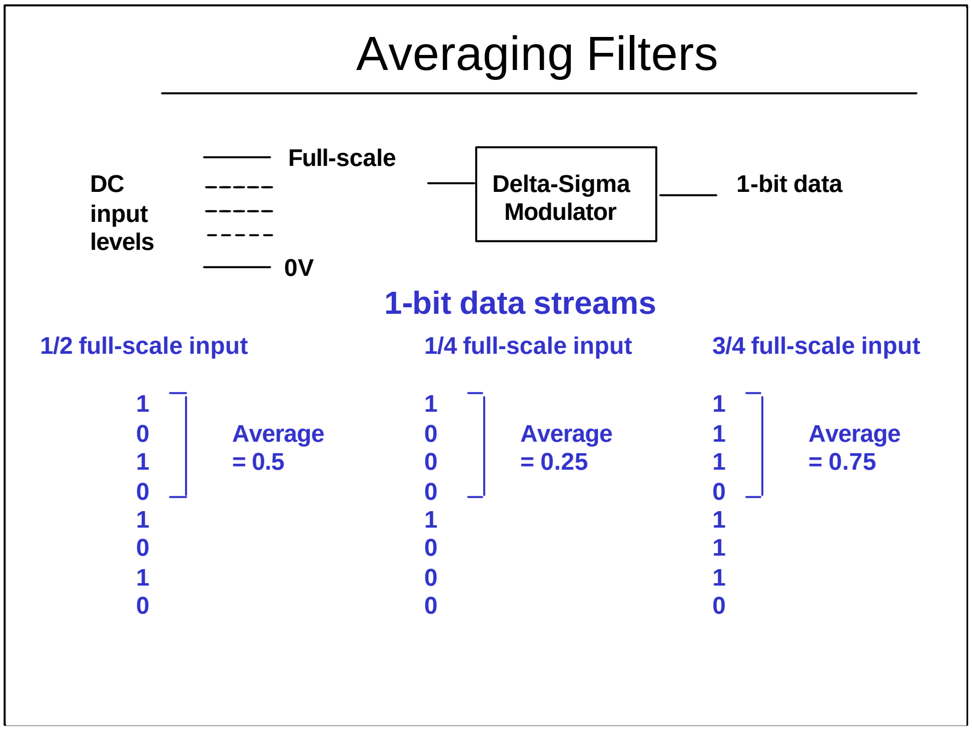

In order to see the action of the digital filter, first consider the case of a DC input level. The slide shows the values present on the 1-bit data stream for three values of the input level. As is clearly shown, the average value of the1-bit data represents the value of the input signal. In this case, 4 samples of the 1-bit data were averaged, but any number of samples could have been chosen. Higher numbers of sample averages improves the accuracy of the result. The filter used in a delta-sigma ADC is similar to a averaging filter. But as we will see in the next slide, averaging by itself will not work.

Oversampling and Signal-to-Noise

- Oversampling and averaging improves SNR
  - Each oversample by a factor of 4 gives a 6-dB (1-bit) improvement in SNR
- For a 1-bit ADC for oversampling and averaging there are:
  - 2-bits for 4x oversampling (4^1)
  - 3-bits for 16x oversampling (4^2)
  - 4-bits for 64x oversampling (4^3)
  - .. ..
  - .. ..
  - 24-bits for 4^{23}x oversampling !!!!! (70,368,744,177,664)
- At a 40-kHz modulation rate it takes over 55 years to average 4^{23} bits
- Clearly 24-bits from a 1-bit ADC is not realizable just by oversampling

Just how much can be gained by oversampling and filtering? Each oversample by a factor of 4 increases the SNR by 6 dB. Put another way, each 6 dB increase in dynamic range is equivalent to gaining1-bit of resolution. Hence, each time we oversample by 4 times we gain 1-bit of resolution.

Using a 1-bit ADC and oversampling by 4 times, 2-bits of h can be achieved. Oversampling by 16 times achieves 3-bits, by 64 times achieves 4-bits, etc. To achieve 24-bits of resolution from a 1-bit ADC by oversampling alone would require an oversampling ratio of 4^{23} times. Clearly, not a practical idea.

Delta-sigma ADCs overcome this limitation by using the technique of noise shaping to gain more than 6 dB of dynamic range for each factor of 4 times oversampling as will be shown.

Before leaving the subject of oversampling, it is appropriate to point out that any ADC can be improved by oversampling and filtering, the simplest filter being averaging. If a little noise is added (only needed if there isn't already enough noise) to the input of a 12-bit ADC, and for each reading point an average of 4 readings are taken, 13-bits of resolution will be achieved. Averaging 64 readings will achieve a resolution of 16-bits. Indeed, 16-bit ADCs are limited by internal noise and rarely achieve 16-bits of effective accuracy. If 16-bits of accuracy is required, then averaging multiple readings is used to achieve this precision.

One problem with averaging is that there is a point a diminishing returns. Additional averaging only works if the signal doesn't have any drift. The Allan Variance can be used to show that there is an optimum number of averages. Continuing to average after the optimum value will lead to a reduction in accuracy instead of an increase in accuracy.

Limits of Averaging

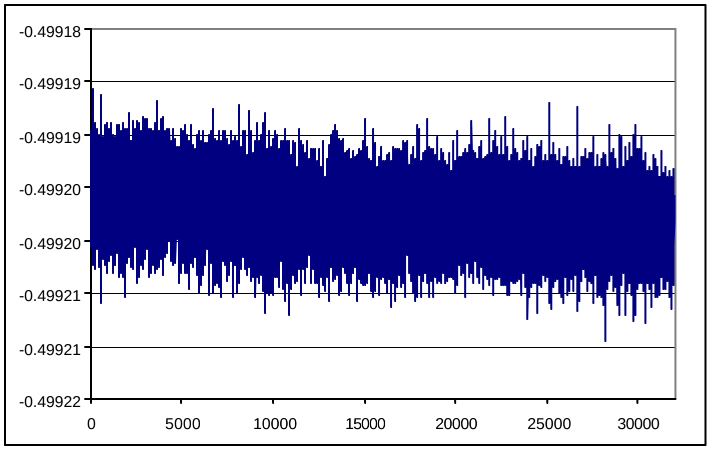

ADS1210

-5-V Input 60-Hz Rate

~9 minutes of Data

This data illustrates one problem with additional averaging. There is a limit to the benefit of averaging. If it takes a long time to acquire the necessary number of averages, the data might drift during that time. For a specific type of data and drift, there is a maximum number of averages which are beneficial. The Allan Variance calculates this maximum by taking multiple sets of longer and longer averages and then measures the resultant noise from each set.

Allan Variance

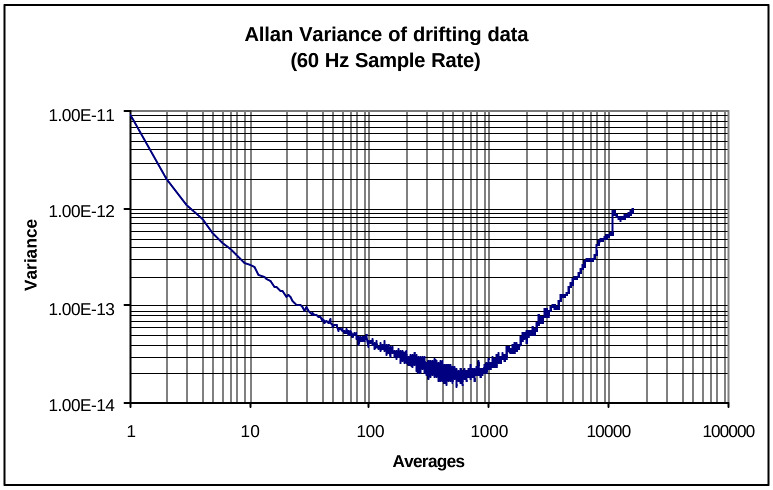

Allan Variance is a method that breaks up the data into sub data packets and averages those packets. It then computes the variance of the result as the size of the packets are changed. The results are plotted on a log-log plot where it become immediately obvious what the optimum number of averages are for that set of data.

As can be seen by this set of data, there is no benefit to average more than about 500 samples. Averaging more than 500 points starts to be influenced by the data drift and actually gives you a noise performance that is worse.

Low Drift Data

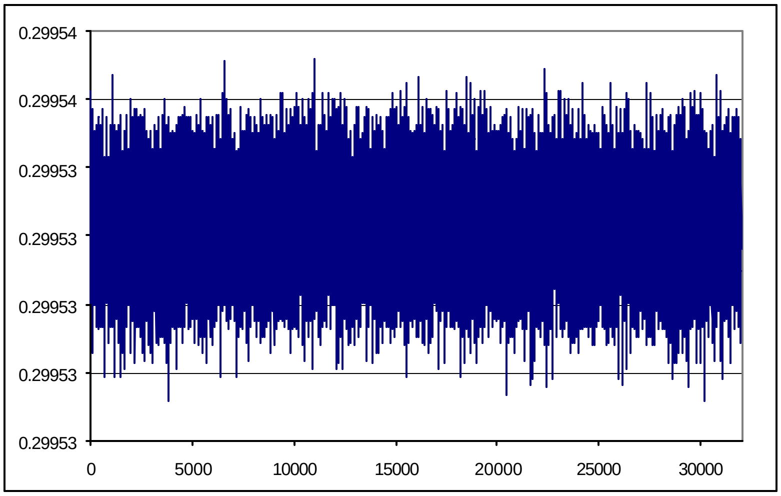

3-V Input

Here is a similar set of data. Although the random noise is similar in amplitude there isn’t any observable drift. If we try that same Allan Variance technique we would expect that this data would be benefited by a higher number of averages.

Low Drift Allan Variance

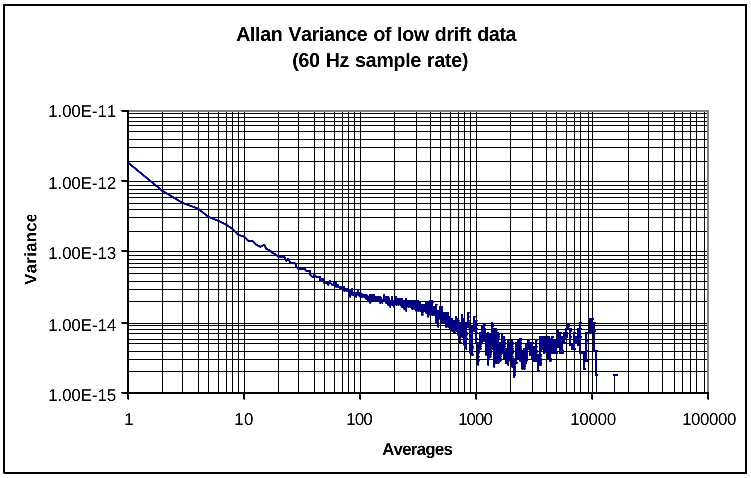

3-V Input

This data shows that up to 2000 averages are beneficial. The slight hump in the plot at about 300 averages might be indicative of another phenomenon which can also be a problem. There is a slight periodicity or "tone" to the data. With some early delta-sigma converters, this "tone" was significant and objectionable. When a tone is present in the data, the averaging could actually give poorer performance than no averaging at all.

Manufacturers have taken different approaches to solving input tones. The tones are the largest with an input of zero volts DC. Therefore, some manufacturers have added an offset to move the tone voltage away from DC. But that only moves the problem, it doesn't solve it. All of the delta-sigma converters from TI, which have been introduced since 1999, have solved this problem.

Digital Filter and Decimation Filter

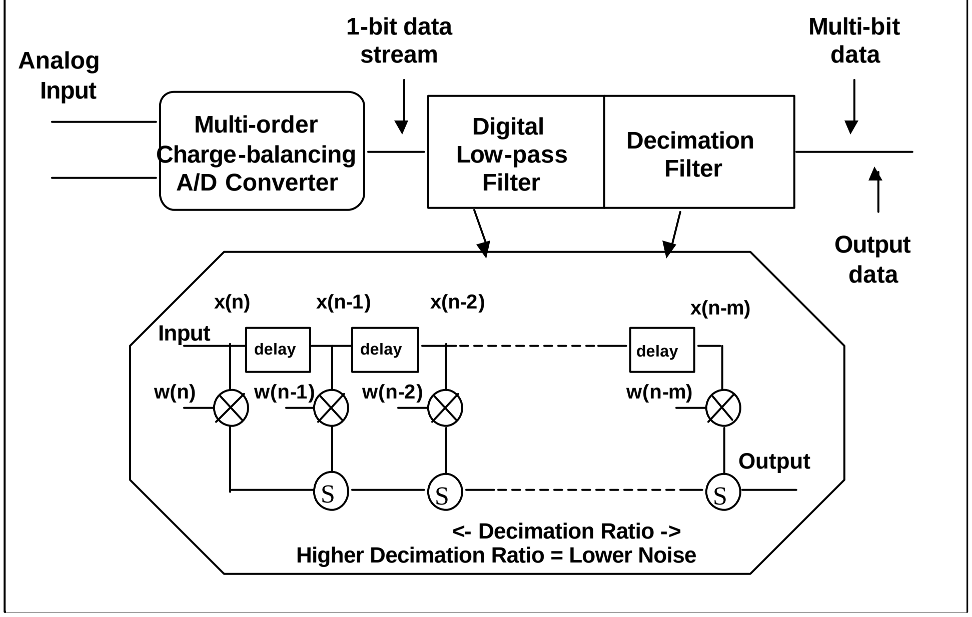

Sometimes a decimation filter is represented as just taking fewer of the output samples, but to achieve a good understanding requires a more accurate presentation. This gives us a representation of a digital filter. The performance parameters of the filter such as frequency roll-off, ripple, slope of attenuation, group delay, etc. are determined by the weights chosen for w(n) to w(n-m). All input samples are multiplied by the weights and then summed to create the output. The number of elements included in the filter is directly related to the decimation ratio. Higher decimation ratios require more elements in the filter and, therefore, yield lower noise in the result.

Delta-Sigma Tradeoffs

- For Best Noise
  - More filtering
    - Higher order filtering will increase latency
  - To keep data rate increase the modulator speed
- For Best Data Rate
  - Less filtering and a high modulator speed
    - But less filtering adds more noise
- Why Not Run the Highest Modulator Speed?
  - High modulator speed means more power and larger die size

The best SNR is achieved with high-order brick wall filters. High-order filters create a longer data pipeline that has to be filled before the output is meaningful. Thus, high-order filtering means high latency and low data rates. The best data rate is achieved with less filtering and high modulator speed. Modulator speed becomes the limiting factor because it determines the power and eventually the cost of the ADC (increased die size). Some of these tradeoffs will be made in the integrated circuit to simplify the operation.

ENOB (Effective Number of Bits)

- Measured signal-to-noise ratio = RMSSIGNAL/RMSNOISE
- SNR is usually specified in dB = 20 * log (ratio)
- SNR = 20 log(signal/noise) = 20 log (1/ppm noise) = 6.02 N

- ENOB = N = 20 log (1/ppm) / 6.02
- For example if noise = 0.4 ppm
- ENOB = 20 log (1/0.4e-6)/6.02 = 21.3 bits

- 2ENOB = Full Scale/Noise = 224/ σ (A/D codes)
- log2(2ENOB) = log2(FS/σ)
- ENOB = log2(224) – log2(σ codes)
- ENOB = 24 – log2(codes)

ENOB is different from the calculated SNR, because it deals with measured quantities rather than theoretical equations. Noise can be specified as ppm of full scale or bits.

SNR = 6.02N + 1.76dB ; (for an N-bit ADC sine wave input)

SNR = 6.02N + 1.76dB + 10 log(Fs/2*BW) for oversampling

Delta-sigma has oversampling and noise shaping. For simplicity, the increase in SNR is measured in bits; therefore, SNR = 6.02 N

Effective Number of Bits gives us the number of bits of resolution we can expect from the observed Signal-to-Noise ratio. Since ENOB gives us a digital representation it is best understood using base 2 math. With this approach, the ENOB is simply computed from the total number of bits available minus the log to base 2 of the number of codes of noise.

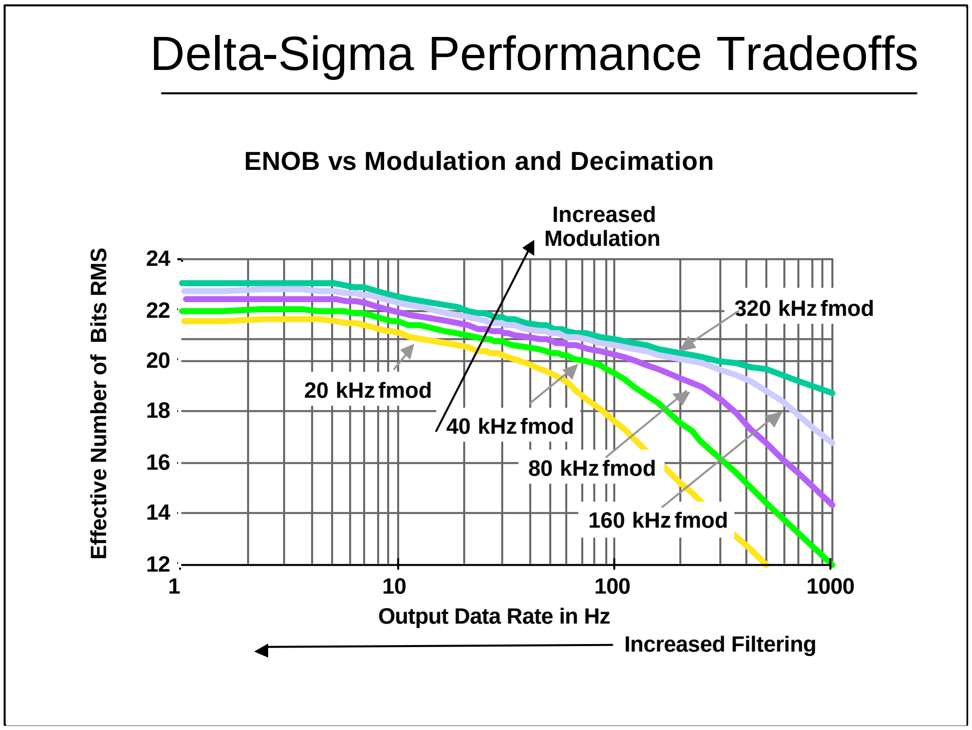

The above diagram illustrates the Effective Number of Bits or ENOB available from either the ADS1210 or the ADS1211 for any given oversampling rate. For the same output data rate, a higher sampling rate provides more samples which are used to produce each output.

The price to be paid for oversampling a signal is conversion time or converter power dissipation. In the case of the ADS1210/11, using the turbo mode to oversample by a factor of two does not double the conversion time. Rather, the digital filter used to average the digital equivalent of the data word is run at a higher or faster rate. The effect on the converter is an increase in power dissipation.

ADS1216 ENOB vs PGA

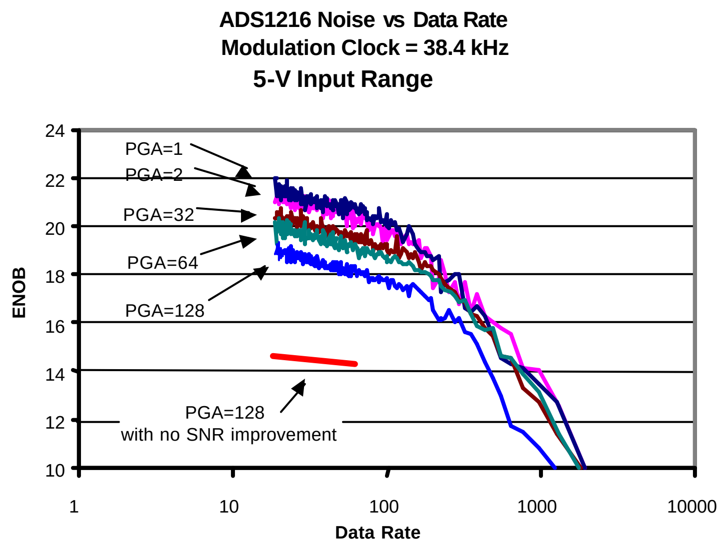

Programmable amplifiers are effective at increasing the ENOB, because they help fill the converter's input span. If there were no advantage with increased gains, then our ENOB for a gain of 128 would be where the red line is located—at about 14.5 bits.

Care should be taken when examining this type of data. If the ENOB is reduced by 1-bit when the PGA is doubled, then there is not net gain.

Digital Filter Response

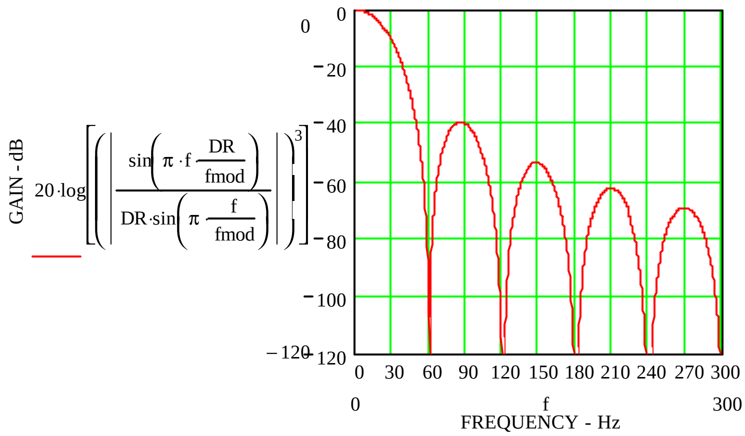

This graph gives the frequency response for a sinc^3 filter. For this example the data rate is 60 Hz. Notice that at 90 Hz the filter is only down 40 dB. If there are frequencies of 90 Hz in the input, they will be aliased at 30 Hz with an attenuation of 40 dB. Any anti-aliasing filters will have to take this into account.

Delta-Sigma Latency

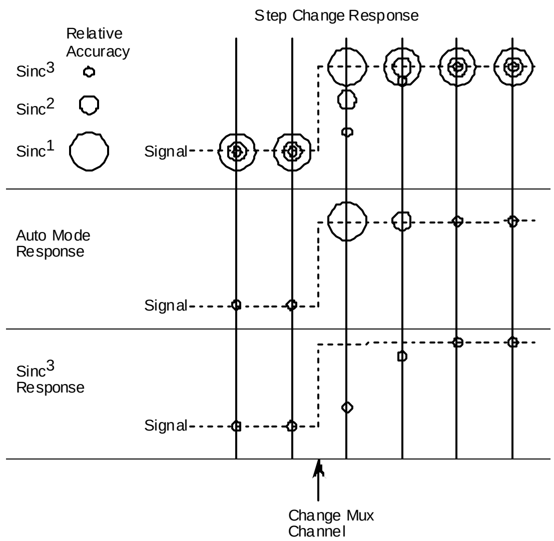

Here we see how the output data responds to an abrupt change on the inputs. Such a change could either be the signal voltage making a quick change or changing the input multiplexer to another channel. The top trace is a composite representation of the Fast Response (labeled $Sinc^1$), $Sinc^2$ and $Sinc^3$ responses. The size of the circle represents how precise the measurement is. As can be seen the Fast Settling filter settles fast, but isn't as precise as the $Sinc^3$ which takes three cycles to settle, but is more precise. In auto mode the system starts with the Fast Settling filter for two cycles (the chart only shows one), then it outputs the result from the $Sinc^2$ filter and finally the result from the $Sinc^3$ filter. This has the advantage of settling fast and gives you as much precision as is available.

ADS1216 Block Diagram

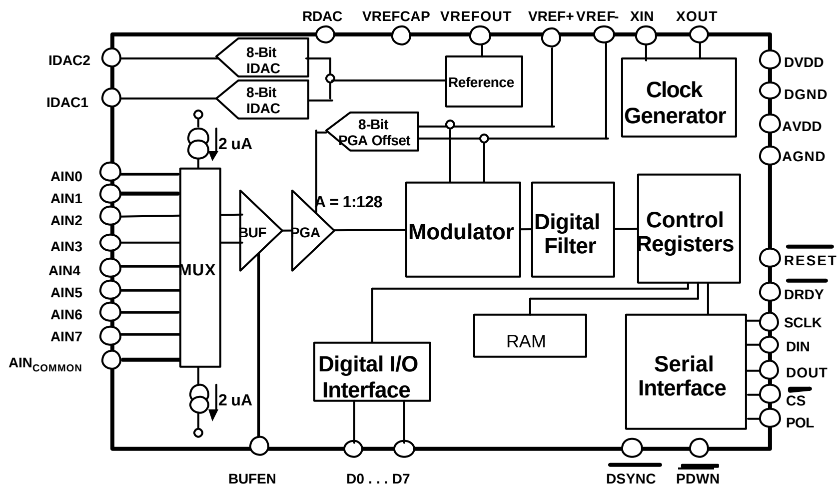

The ADS1216 is used for the next applications circuit. In addition to the analog inputs, which can differentially measure between any two analog inputs, there are two current output DACs and an on-chip reference.

Highly Integrated Measurement System

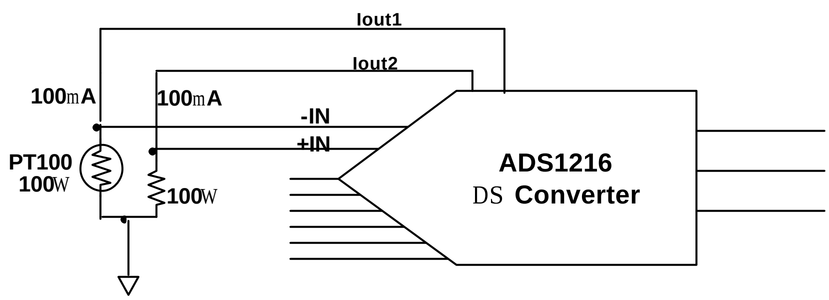

◆ Sources of Error
- Resistor tolerance
- Matching of current sources

The current DACs provide the excitation required by the PT100 RTD. The same current is sourced through a reference resistor of 100 ohms. The differential voltage is measured by the ADS1216. The internal PGA is used to extend resolution of the measurement.

Improved Measurement System

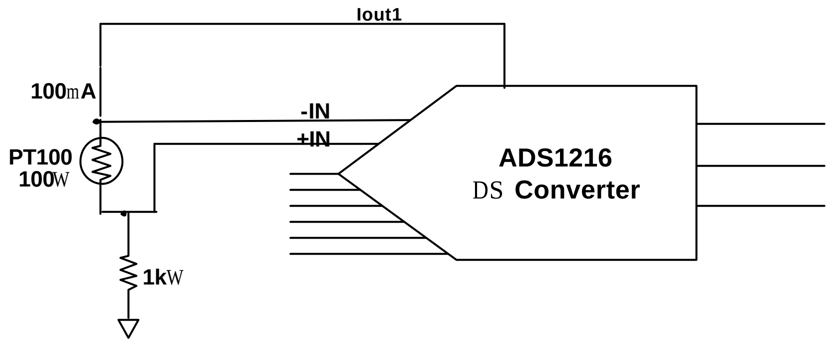

◆Sources of Error

- Requires accurate current source

The ADS1216 with 24-bits resolution provides plenty of dynamic range to measure the voltage directly across the PT100. This eliminates two sources of error: the second current source and the accuracy of the reference resistor.

Ratiometric Measurement System

- Advantages
  - Does not need accurate current source
  - PGA can be changed to suit temperature

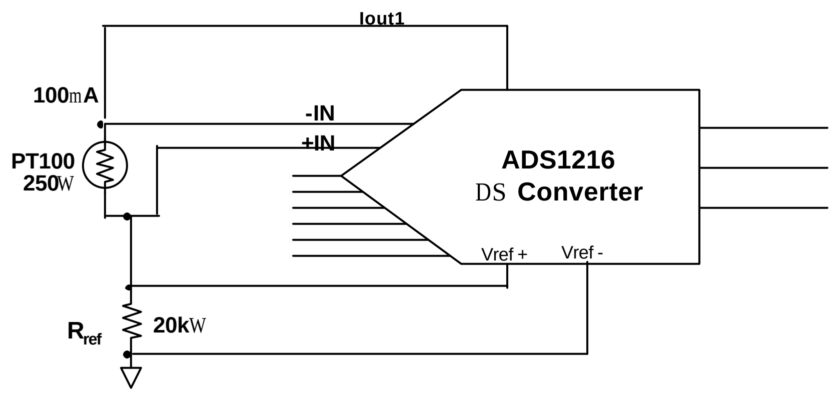

The current source can be removed as a source of error by setting up a ratiometric system. By using the current through the PT100 to create the reference voltage with R_{ref}, the actual current doesn't matter. Any change in the current is compensated by a corresponding change in the reference voltage.

Maximum temp = 400 ºC, (PT100 = 250 ohms)

Using PGA of 2, maximum voltage is 1.6 V

Set reference voltage to 2 V, R_{ref} = 20 k

Therefore, maximum voltage for –IN = 3.6 V

8 Differential Ratiometric Measurements

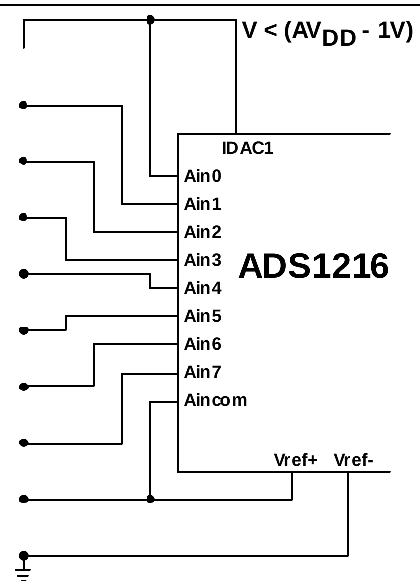

Here is the same circuit, but now we are measuring 8 RTDs with one current source. Using the flexibility of ADS1216's multiplexing capability, voltage measurement between any two inputs can be accomplished.

This circuit was tested with 2-k thermistors and a 10-k reference sense resistor. IDAC1 was set to 49.43 µA, 99.26 µA, and 148.59 µA. For all current settings, the voltage readings from the thermistors stayed the same. The voltage at the IDAC1 pin with the current of 148.59 µA would have been 3.86 V, which is close to the maximum voltage of 4.0 V.

Temperature Measurement

- Thermistor
  - High-voltage output
  - Highly sensitive
  - Highly nonlinear
  - Requires current source
  - Susceptible to self-heating
  - Resolution on the magnitude of <10 mV/ºC

A thermistor provides a good example for testing the data acquisition requirements. The voltage output with a thermistor is much larger, but the response is nonlinear.

Measurement Requirements

- Let's Make the Design Easy
  - Thermistor provides the largest signal change as temperature changes.
  - Does require processing to convert from voltage to ºC
    - Our Specification:

1ºC Accuracy

over Industrial Temperature Range (-40ºC to +85ºC)

We'll try microcontrollers with a 12-bit A/D

(125ºC)/ 2^{12} = 0.03ºC resolution

For this design, the accuracy goal is 1ºC over the temperature range of –40 ºC to +85 ºC. Since the signal level from a thermistor is high, this should be an easy requirement to meet.

Microcontroller A/D Noise Performance

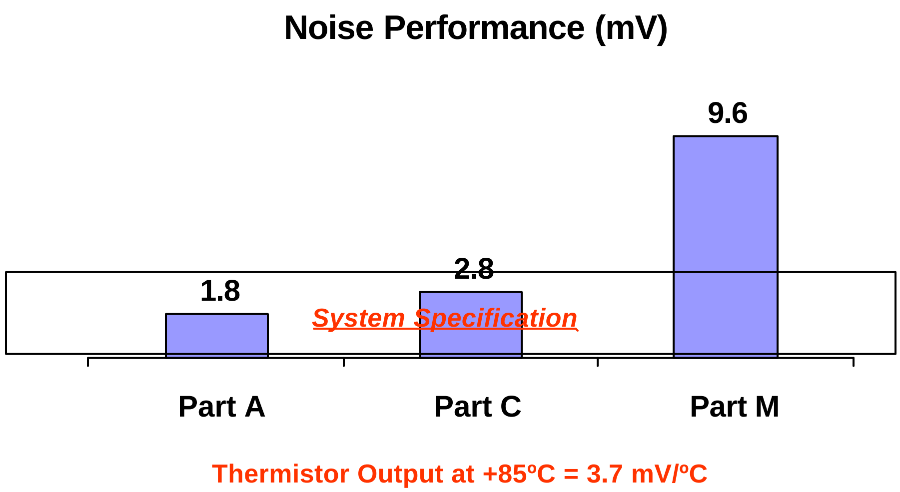

This experiment compares three microcontrollers with 12-bit A/D converters. At 85ºC the thermistor has the smallest voltage change of 3.7 mV/ ºC. The noise performance for these three devices is graphed to show how they match the requirement.

Error Sources

- What Are Other Limitations?
  - Offset error
  - Offset drift
  - Gain error
  - Gain drift
  - V_{ref} accuracy
  - V_{ref} drift
  - INL
  - DNL

Besides noise, there are many other error sources that will have an effect on temperature measurement. There are A/D errors such as offset, and gain and drift in the offset and gain. Also, linearity performance and reference voltage accuracy will add additional errors.

Error Analysis

What's the Total Error in mV?

	Root Sum Squared (RSS)	Worst Case
Part A	51.8 mV	72.6 mV
Part C	60.6 mV	77.0 mV
Part M	104.8 mV	128.5 mV

Goal: 3.7 mV

Adding all of the error sources gleaned from the specifications yields a total error (Root Sum method or Worst Case summation) much larger than the 3.7-mV requirement.

The two different methods provide a different approach to determine the expected result. The Worst Case method adds all of the maximum specifications to get a worst case overall specification. The Root Sum method assumes that all specifications will not be a maximum at the same time. It takes the square root of the summation of all the worst-case specifications squared.

Root Sum Method = square root (par12 + par22 + par32 + . . . parN2)

Worst Case = par1 + par2 + par3 + . . . parN

Thermistor Measurement

- System Specification: 1ºC
- The Resolution Achievable (using RSS)
  - Part A
    - Resolution of 14ºC
  - Part C
    - Resolution of 16ºC
  - Part M
    - Resolution of 28ºC
- So does this rule out 12-bit converters?

In terms of temperature accuracy, these error sources lead to as much as 38ºC of error.

It might seem amazing to realize that we started with a 12-bit A/D, which has a resolution of about 0.03ºC. But when we add all the other sources of error, we find that we aren't even close to our goal of 1ºC.

Error Reduction

- Can Use Microcontroller Functions For...
  - Offset and gain calibration
  - V_{ref} calibration
  - Noise reduction
    - Averaging doesn't improve noise performance significantly for this application
  - Linearization of INL

Many of these errors can be compensated for by using special algorithms or calibration techniques.

Offset and Gain require to ability to switch calibration voltages (0 V and full-scale) into the input and do a conversion.

V_{ref} calibration uses a known temperature or voltage to compute errors in the V_{ref}. This is usually done over temperature so that it can be corrected at several temperatures.

Averaging only helps if the noise is a significant portion of the error. Except in one case, the noise was below the desired specification.

Linearization is a process of factory calibration where a calibrated voltage ramp is applied and the errors across the range are stored. These values are then used to correct the results toward the ideal values.

Best Case Results

- System Specification: 1ºC
- How Does This Improve Performance? (using RSS)
  - Part A
    - Resolution of 4ºC
  - Part C
    - Resolution of 2ºC
  - Part M
    - Resolution of 4ºC
- We're closer, but still not there.

This result is obtained by assuming that some of the errors can be corrected or compensated. This is still not small enough to give us 1ºC at a temperature accuracy at 85ºC.

Assumes Perfect Adjustment for Offset, Gain, V_{ref}, INL

Implementation Issues

- Overhead for offset and gain calibration
- Assumes offset and gain drift correction
- Requires in-system V_{ref} calibration
- Overhead for linearization of INL
  - How good can it be? How many points?
- Reduces system resources for other functions
- Impractical to filter line frequency noise

The requirements to correct for several of the error sources has an impact on the overall system. Many of these methods are time consuming, are not perfect in their result, and use valuable system resources. There must be a better way.

Alternative Solutions

- 24-bit with µC
- 16-bit SAR with external µC
- MSC1210

The initial requirement for 1ºC of accuracy would not seem to need 16 bits or more of resolution. But, as we observed, the 12-bit solutions didn't give us much margin for error. Here are a few possible higher resolution approaches.

Why use 16-bit and 24-bit solutions?

1. 14-bits might be O.K., but 16-bits are more available.

2. 24-bit delta-sigma converters are low-cost.

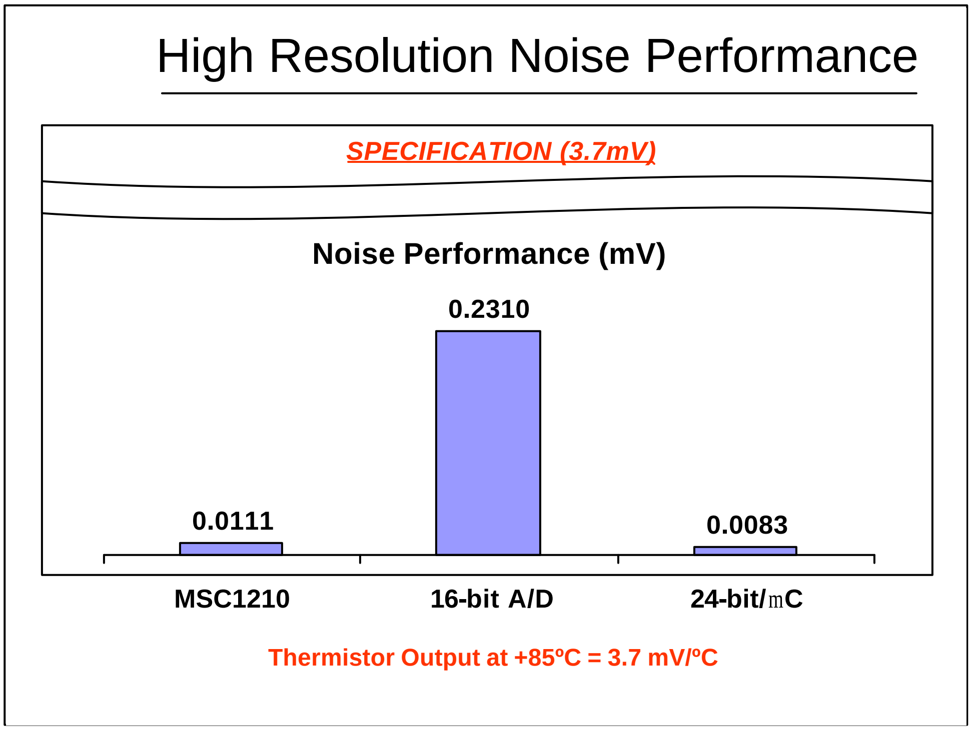

As can be seen, the noise of these approaches is significantly below the specified requirement.

Total Errors

What's the Total Error in mV?

	Root Sum Squared (RSS)	Worst Case
MSC1210 Total Error (mV)	5.5	7.4
16-bit A/D with µC (mV)	4.3	7.5
24-bit/µC Total Error (mV)	20.0	28.2

Goal: 3.7mV

When we add up all the error sources, we can still have a problem easily meeting our requirement of 3.7-mV resolution per degree.

Error Correction — Again

- Can Use Microcontroller Functions For...
  - Offset and gain calibration
  - Vref calibration
  - Noise reduction
    - Averaging doesn't improve noise performance significantly for this application (because the noise is already very low).
  - Linearization of INL

Although noise isn't the limiting factor, we still have a need for calibration and linearization to achieve our required goal of 1ºC of accuracy.

Offset and Gain require to ability to switch calibration voltages (0 V and full-scale) into the input and do a conversion.

V_{ref} calibration uses a known temperature or voltage to compute errors in the V_{ref}. This is usually done over temperature so that it can be corrected at several temperatures.

Averaging only helps if the noise is a significant portion of the error. Except in one case, the noise was below the desired specification.

Linearization is a process of factory calibration where a calibrated voltage ramp is applied and the errors across the range are stored. These values are then used to correct the results toward the ideal values.

Corrected Total Error

What's the Corrected Total Error in mV?

	Root Sum Squared (RSS)	Worst Case
MSC1210 Total Error (mV)	2.3	2.4
16-Bit A/D with µC (mV)	3.1	3.4
24-bit/µC Total Error (mV)	15.6	15.7

Goal: 3.7mV

By using the various corrections possible, we now can meet our requirements with a couple solutions. The third solution could be improved with a better reference.

Notice that although the last solution starts with a low-noise, 24-bit converter, it doesn't meet the overall requirements. That is because the internal reference voltage isn't very good. It has significant drift which can't be corrected with a production calibration.

Optimized Performance

- MSC1210
  - Accuracy of 1ºC
- 16-bit SAR with µC
  - Accuracy of 1ºC
- A 24-bit mixed signal part
  - Accuracy of 5ºC

With two of these approaches we can achieve our goal of 1ºC. You will notice that it is not sufficient to increase the resolution. The third method had a 24-bit A/D, but because of the performance of the reference, it achieved the same result.

This Assumes Perfect Adjustment for Offset, Gain, V_{ref}, INL

Drawbacks of 16-Bit A/D

- Overhead for offset and gain calibration
- Assumes offset and gain drift correction
- Requires in-system Vref calibration
- Overhead for linearization of INL
  - How good can it be?
- Reduces system resources for other functions
- Impractical to filter line frequency noise
- $40 solution

Although the 16-bit A/D could achieve our required goal, it is fairly expensive.

To achieve our required accuracy we need:

1. To perform calibration operations for Offset and Gain. These need to be repeated often enough to compensate for drift.

2. The V_{ref} has to be calibrated in production and correction factors applied as the temperature changes.

3. The points need to be calibrated over the range of the converter and then corrections applied to A/D results to improve linearity.

4. The conversion frequency might not make it convenient to provide line frequency filtering. You would have to average multiple samples over the 60-Hz period (16.66 ms) to get some line frequency rejection.

Advantages of MSC1210

- Low overhead for offset and gain calibration
- Requires in-system V_{ref} calibration
- No overhead for linearization of INL
  - It's good!
- Allows system resources for other functions
- Inherent filtering of line frequency noise
- $8 solution

Here is a total integrated solution that gives you good performance. This can all be done for a lower cost.

1. Offset and Gain calibration is built-in.

2. The internal V_{ref} is very good without any calibration—0.2% accuracy, 5 ppm/ºC drift. But without V_{ref} calibration, the achievable accuracy is about 1.3ºC. This can be solved with either an external reference or V_{ref} calibration.

Thermistor Circuit

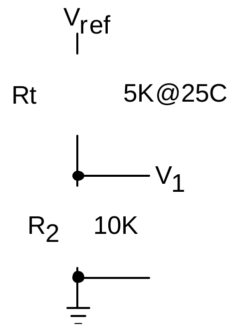

$$R_t = \frac{V}{I} = \frac{(V_{ref} - V_1)}{V_1/R_2}$$

$$= \frac{(V_{ref} - V_1)10K}{V_1}$$

$$= (V_{ref}/V_1 - 1)10K$$

The temperature of the thermistor is directly related to the resistance. We determine the resistance by measuring the voltage across the thermistor and the current through it.

System Schematic

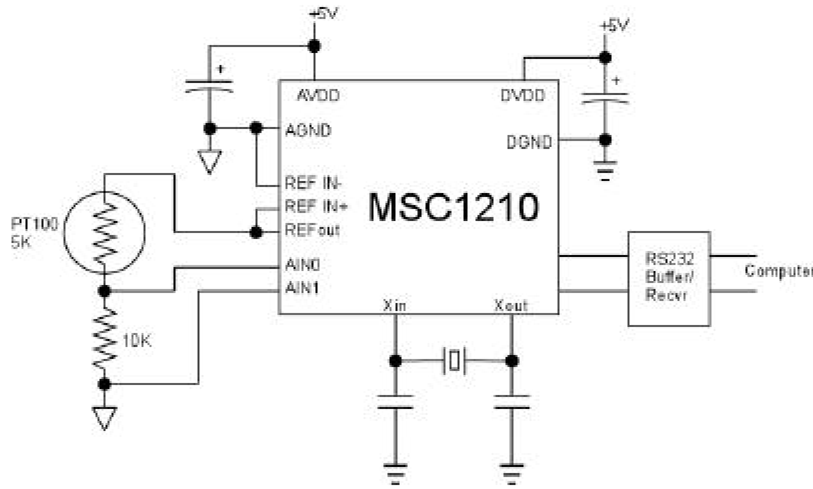

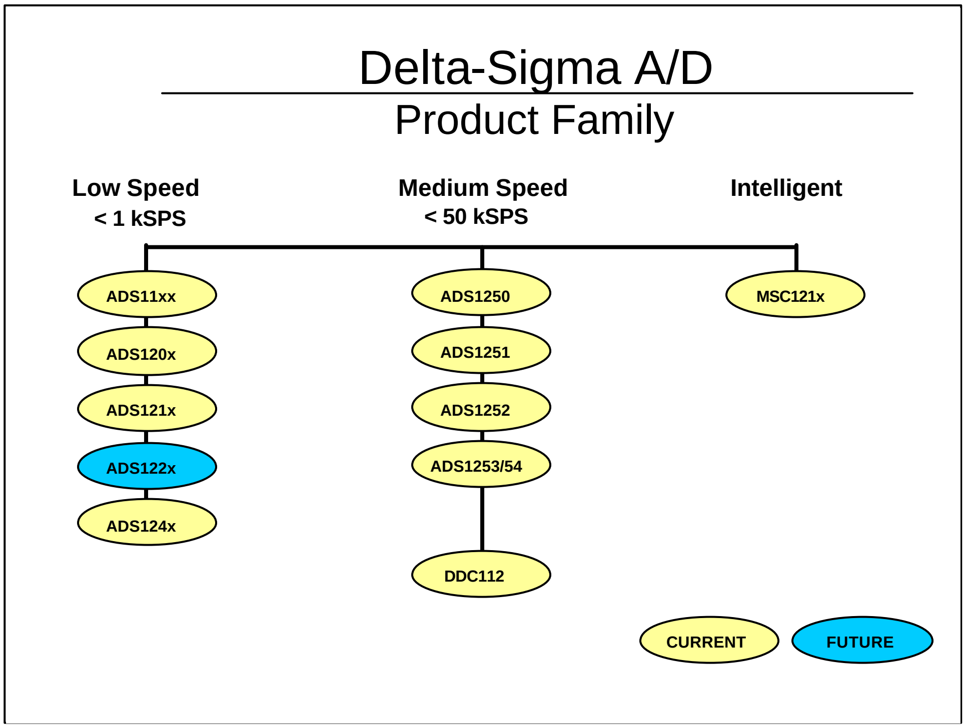

Not only are we expanding the number of products, we are bridging sampling rates and providing varying levels of integration.

Mailing Address:

Texas Instruments
Post Office Box 655303
Dallas, Texas 75265