

Ask The Applications Engineer—15


by Oli Josefsson

USING SIGMA-DELTA CONVERTERS—PART 1

Q: I'd like to use sigma-delta A/Ds but have some questions because they seem markedly different from what I've been using. To start with, what issues do I need to consider when designing my antialiasing filter?

A: A major benefit of oversampling converters is that the filtering required to prevent aliases can be quite simple. To understand why this is the case and what the filter constraints are, let's look at the basic digital signal processing that takes place in such a converter. For the purpose of anti-alias filter design we can think of a sigma-delta converter as a conventional high-resolution converter, sampling at a rate much faster than the Nyquist sampling rate, followed by a digital decimator/filter; the fact that the input into the digital decimator is 1-bit serial with a noise-shaping transfer function is irrelevant.

The input signal is sampled at F_{ms} , the modulator input sampling rate, which is much faster than twice the maximum input signal frequency (the Nyquist rate). The figure shows what the frequency response of a decimation filter may look like; frequency components between f_b and $F_{ms}-f_b$ are greatly attenuated. Thus, the digital filter can be used to filter out all energy from the converter within $[0, F_{ms}-f_b]$ that does not fall within the bandwidth of interest $[0, f_b]$. However, the converter can not distinguish between signals appearing at the input that are in the range $[0, \pm f_b]$ and those in the ranges, $[kF_{ms} \pm f_b]$, where k is an integer. Any signals (or noise) in those ranges get aliased down to the bandwidth of interest $[0, f_b]$ via the sampling process; the decimation filter, which works only on the digitized samples, cannot be of any help attenuating these signals.



Thus it is the input noise energy in these bands $[kF_{ms} \pm f_b]$ that must be removed by the antialiasing filter before the input signal is sampled by the converter.

Q: So if I were to use the AD1877, which has a dynamic range of 90 dB, the antialiasing filter will need attenuation well above 90 dB at $F_{ms} - f_b$ (≈ 3 MHz)?

A: Not quite. You are assuming that the A/D has full-scale input at frequencies close to the modulator sampling rate; this is simply not the case in most systems. The only signal input of concern for aliasing is normally just noise from sensors and circuitry preceding the converter. The noise is usually low enough for a simple RC filter to suffice as an antialias filter.

Q: How do I make sure that a one-pole RC filter will suffice for my application—and establish the time constant of the filter?

A: Your application will typically specify a maximum allowable attenuation of an input signal that falls within the bandwidth of interest. This in turn puts a minimum on the 3-dB point of the RC filter. Let's take a look at an example using the AD1877 to illustrate this point further and to show how one might verify that a single-pole filter will provide enough filtering.

Let's assume that we have an application where the bandwidth of interest is 0 to 20 kHz, and signals in this range must not be attenuated more than 0.1 dB, or a ratio of 0.9886 [$\text{dB} = 20 \log_{10}(\text{ratio})$ for voltage and $10 \log_{10}(\text{ratio})$ for power]. From the formula for attenuation of a single-pole filter,

$$\text{ratio} = \frac{1}{\sqrt{1 + (2\pi fRC)^2}} > 0.99 \text{ at } f = 20 \text{ kHz}$$

$$RC \leq \frac{\sqrt{1 - (\text{ratio})^2}}{(2\pi f)(\text{ratio})} \approx 1.21 \times 10^{-6} \text{ s}$$

Choosing $RC = 1.0 \mu\text{s}$, to allow for component tolerances, the -3-dB frequency will be 159 kHz. We can now calculate the attenuation the filter will provide in the frequency bands, $kF_{ms} \pm f_b$, that alias down to the baseband. Assuming that the AD1877 has a modulator sampling rate of 3.072 MHz (and output sampling rate of 48 kHz), the first frequency band occurs at 3.052 MHz to 3.092 MHz. The attenuation of the RC filter at these frequencies is approximately 25.7 dB (about 0.052) over the whole band. Over the second band (6.124 MHz to 6.164 MHz), the attenuation is 31.8 dB (0.026). We know that the noise in these two bands (and all higher bands up the scale) that escapes through the filter to the A/D input will be aliased down to the baseband and get added as root sum-of-the-squares (rss) of their rms values,

i.e., $\sqrt{n_1^2 + n_2^2 + \dots + n_i^2}$. For values given in dB, the formulas shown the Appendix can provide results directly in dB, avoiding the intermediate step of computing the ratios.

For white noise, the noise spectral density is constant as a function of frequency, and each frequency range has the same bandwidth, so each band contributes an equal amount of noise to the input of the filter. We can therefore find the effective attenuation of the RC filter by adding the attenuation of the different frequency bands in rss fashion. The noise contribution from the first two bands, for example, is the same as the contribution from a single frequency band with attenuation of

$$\sqrt{0.052^2 + 0.026^2} = 0.058, \text{ or } 24.7 \text{ dB, compared with } 25.7 \text{ dB for the first band.}$$

How many bands do we need to consider when calculating the total aliased noise? For this case, the rss sums of the first 3, 4, 5, and 6 bands are, respectively, $-24.2, -24.0, -23.9, -23.8$ dB. The first band is therefore quite dominant; its attenuation is within 2 dB of the attenuation for all bands. It is usually sufficient to take only the first band into account unless the noise is exceptionally large or has a non-white spectrum; in addition, the A/D itself, though fast, has limited bandwidth; it tends to reject high-order bands.

Now that the attenuation is in hand, we can consider the noise magnitude itself: Let's be conservative (by about 50%) and take the effective filter attenuation to be 20 dB (i.e., 0.1 V/V). To be able to calculate the maximum allowed noise spectral density when using a single pole filter, an estimate should be made of the maximum performance degradation that aliased noise can contribute. From the dynamic specs of the AD1877 we find that the total noise power internal to the converter is 90 dB below (32 ppm of) full-scale input. If the whole system is to be within, say, 0.5 dB of this spec, the total aliased noise power can't exceed the rss difference between -90 dB and -89.5 dB or -99.1 dB (11.1×10^{-6}). Using this information, and the fact that the input scale of the AD1877 is 3V p-p, we find that aliased noise must not exceed $3/(2\sqrt{2}) \text{ V} \times 11.1 \times 10^{-6} = 11.8 \mu\text{V rms}$. If all this noise were assumed lumped in a single aliased band, and noting that rms noise = noise spectral density (N.S.D.) $\times \sqrt{\text{BW}}$,

$$\text{N.S.D.} < \frac{11.8 \mu\text{V}}{\sqrt{3.092 \text{ MHz} - 3.052 \text{ MHz}}} = 59 \text{ nV}/\sqrt{\text{Hz}}$$

This is the maximum post-filter spectral density allowed. To find the maximum prefilter spectral density (MPSD), with the effective filter attenuation of 20 dB (i.e., $\times 0.1$) established previously, M.P.S.D. = $10 \times 59 \text{ nV}/\sqrt{\text{Hz}} = 0.59 \mu\text{V}/\sqrt{\text{Hz}}$.



Clearly your system has to be pretty noisy in the 3-6-9-12-MHz regions in order for a simple RC filter not to suffice; however, as always, one must be careful of ambient rf pickup.

Q: As I understand it, the noise floor of sigma-delta converters may exhibit some irregularities. Any thoughts on that?

A: Most sigma-delta converters exhibit some spikes in the noise floor, called idle tones. In general, these spikes have low energy, not enough to substantially affect the S/N of the converter. Despite that, however, many applications cannot tolerate spikes in the frequency spectrum that extend much beyond the white noise floor. In audio applications, the human ear, for example, does an excellent job of detecting tones in the absence of large input signals even though the tones are well below the integrated (0-20-kHz) noise of the system.

There are two sources of idle tones. Their most common cause is voltage-reference modulation. To understand this mechanism a basic understanding of sigma-delta converters is needed. Here is a one minute crash course on sigma-delta converters (to probe further please consult[1]).

As the block diagram shows, a basic sigma-delta A/D converter consists of an oversampling modulator, followed by a digital filter and a decimator. The modulator output swings between two states (high and low, or 0 and 1, or +1 and -1), and the average output is proportional to the magnitude of the input signal. Since the modulator output always swings full-scale (1 bit), it will have large quantization errors. The modulator, however, is constructed so as to confine most of the quantization noise to the portion of the spectrum beyond f_b , the bandwidth of interest.

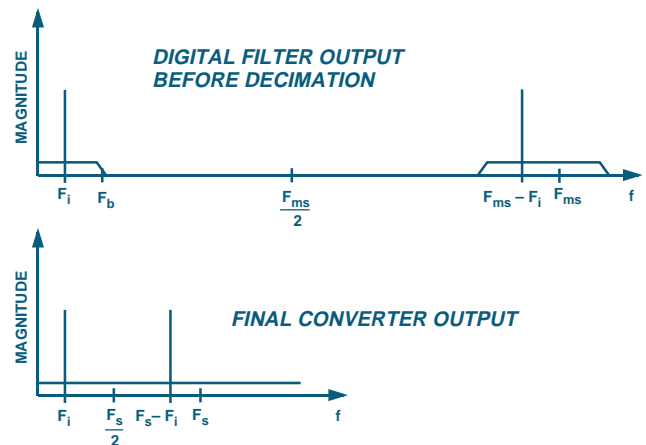


As shown, the spectral “sticks” (single frequencies) at f_i and $F_{ms} - f_i$ correspond to an input signal, while the shaded area shows how the quantization noise has been pushed (shaped) beyond the bandwidth of interest, f_b .

The digital filter, which is often an n -tap FIR filter, takes the high-speed low-resolution (1-bit) modulator output and performs a weighted average of n modulator outputs in a manner dictated by the desired filter characteristics. The output of the filter is a high-resolution word, which becomes the A/D output. The digital filter is designed to filter out “everything” between f_b and $F_{ms} - f_b$, where F_{ms} is the sampling rate of the modulator. Cleaning out all the noise in between f_b and $F_{ms} - f_b$ makes it possible to reduce the sampling rate to values between F_{ms} and $2f_b$ without causing any spectra to overlap (i.e., aliasing).


Conceptually, reducing the sample rate, i.e., *decimation*, can be thought of as only sending every d th digital filter output to the A/D output, where d is the decimation factor. This will bring the spectral images close together, as shown in the figure, which makes the output look like an output from a non oversampled converter. The upper figure shows the output of the modulator after digital filtering but prior to decimation. The lower figure shows the spectral output after decimation—the final A/D output.

In real converters, digital filtering and decimation are intimately combined for economy in design and manufacture. Thus, the terms “digital filter” and “decimator” are used interchangeably to describe the digital circuitry processing the modulator output to produce the output of the converter.



O.K., now back to “idle tones”. Let's start by looking at the output of the modulator when a dc signal is applied to the input. For an exact mid-scale dc input level, the output of the modulator is equally likely to be high (1) or low (0), in other words,

the pulse density is 0.5, very likely to result in bitstream patterns like 010101. These regular patterns mean that the output spectrum will have a spike at $F_{ms}/2$ (upper figure). If the dc input now moves somewhat off midscale, the modulator output bit pattern will change accordingly. The spectrum of the modulator output will now show spikes at $F_{ms}/2 - \partial F$ and $F_{ms}/2 + \partial F$, with ∂F proportional to the dc change from midscale (lower figure).



With effective digital filtering, how can such tones possibly find their way down to baseband? The answer is via the *voltage reference*. The digital output is a measure of the ratio of the analog input to the voltage reference. An x% change in the magnitude of the voltage reference will result in a -x% change in the magnitude of the digital output word. Voltage-reference change will, in effect, amplitude modulate the A/D output. Now, we have clocks internal to the converter, and possibly also externally, running at $F_{ms}/2$. If small amounts of these clock pulses get coupled onto the voltage reference line, they will change it slightly and, in effect, modulate the tones at $F_{ms}/2 - \partial F$ and $F_{ms}/2 + \partial F$. One of the difference frequencies created by this modulation is at ∂F , and it is clearly in the bandwidth of interest. Nonlinearities may also create tones at multiples of ∂F .

Q: From your explanation it seems that if I apply an ac signal to the converter I do not have to worry about idle tones?

A: Well, any ac signal generally has a dc component associated with it, which will have to be represented by the modulator output, so the explanation above still applies. But if the total dc input offset (i.e., internal converter offset plus external offsets) in your system is exactly 0, the tones will be at dc (0 Hz).


There is another source of idle tones in lower-order (<3rd-order) modulators. The order of the modulator (number of integrations) is a measure of how much quantization-noise shaping takes place. Second-order modulators can actually exhibit bit patterns that show up directly in the baseband, even if voltage-reference modulation is not occurring. This is one of the reasons why sigma-delta converters from Analog Devices that are designed for ac applications use higher-order (≥ 3) sigma-delta modulators.

Q: So what can I do to minimize the chances of idle tones interfering with my A/D conversion?

A: Follow the layout recommendations and bypassing schemes recommended by the manufacturer of the converter. This applies not only to the voltage reference, but to power supplies and grounding as well. It is the manufacturer's responsibility to minimize the voltage-reference corruption that takes place *inside* the converter, but it is up to the system designer to minimize the *external* coupling. By following those guidelines, the user should be able to reduce the coupling to a negligible level. If, despite the proper design precautions, idle tones are still an issue, there is yet another option that can be pursued. As I explained previously, frequency of the idle tones is a function of the dc input. This opens up the possibility of introducing enough dc offset on the A/D input to move the idle tones out of the bandwidth of interest to where they will be filtered out by the decimation filter. If the user does not want the dc offset to propagate through the system it can be subtracted out by the processor that handles the data from the A/D.

Q: What kind of a load does the input of sigma-delta converters present to my signal conditioning circuitry?

A: It depends on the converter. Some sigma-delta converters have a buffer at the input, in which case the input impedance is very high and loading is negligible. But in many cases the input is connected directly to the modulator of the converter. A switched-capacitor sigma-delta modulator will have a simplified equivalent circuit like that shown in the figure.



Switches S1 and S2 are controlled by the two phases of a clock to produce alternating closures. While S1 is closed, the input capacitor samples the input voltage. When S1 is opened, S2 is closed and the charge on C is dumped into the integrator, thus discharging the capacitor. The input impedance can be computed by calculating the average charge that gets drawn by C from the external circuitry. It can be shown that if C is allowed to fully charge up to the input voltage before S1 is opened then the average current into the input is the same as if there were a resistor of $1/(F_{sw}C)$ ohms connected between the input and ground, where F_{sw} is the rate at which the input capacitor is sampling the input voltage. F_{sw} is directly proportional to the frequency of the clock applied to the converter. This means that the input impedance is inversely proportional to the converter output sample rate.

Sometimes other factors, such as gain, can influence the input impedance. This is the case for the 16/24-bit AD771x family of signal conditioning A/Ds. The inputs of these converters can be programmed for gains of 1 to 128 V/V. The gain is adjusted using a patented technique that effectively increases F_{sw} (but keeps the converter output sample rate constant) and combines the charges from multiple samples. The input impedance of these converters is, for example, 2.3 M Ω when the device's external clock is 10 MHz and the input gain is 1. With input gain of 8, the input impedance is reduced to 288 k Ω .


These impedances, as noted earlier, represent the average current flow into or out of the converters. However, they are not the impedances to consider when determining the maximum allowable output impedance of the A/D driver circuitry. Instead, one needs to consider the charging time of the capacitor, C, when S1 is closed. For dc applications the driver circuit impedance has only to be low enough so that the capacitor, C, will be charged to a value within the required accuracy before S1 is opened. The impedance will be a function of how long S1 is closed (proportional to the sampling rate), the capacitance, C and C_{EXT} in parallel with the input (unless C_{EXT} >> C). The table shows allowable values of external series resistance with f_{CLKIN} = 10 MHz which will avoid gain error of 1 LSB of 20 bits—for various values of gain and external capacitance on the AD7710.

Typical external series resistance which will not introduce 20-bit gain error


Gain	External Capacitance (pF)					
	0	50	100	100	500	5000
1	145 kΩ	34.5 kΩ	20.4 kΩ	5.2 kΩ	2.8 kΩ	700 Ω
2	70.5 kΩ	16.9 kΩ	10 kΩ	2.5 kΩ	1.4 kΩ	350 Ω
4	31.8 kΩ	8.0 kΩ	4.8 kΩ	1.2 kΩ	670 Ω	170 Ω
8-128	13.4 kΩ	3.6 kΩ	2.2 kΩ	550 Ω	300 Ω	80 Ω

For ac applications, such as audio, where the modulator sample rate is around 3 MHz for 64× oversampling, the input capacitor voltage may not have enough time to settle within the accuracy indicated by the resolution of the converter before the capacitor is switched to discharging. It actually turns out that as long as the input capacitor charging follows the exponential curve of RC circuits, only the gain accuracy suffers if the input capacitor is switched away too early.

The requirement of exponential charging means that an op amp can not drive the switched capacitor input directly. When a capacitive load is switched onto the output of an op amp, the amplitude will momentarily drop. The op amp will try to correct the situation and in the process hits its slew rate limit (non linear response), which can cause the output to ring excessively. To remedy the situation, an RC filter with a short time constant can be interposed between the amplifier and the A/D input as shown in the figure. The (low) resistance isolates the amplifier from the switched capacitor, and the capacitance between the input and ground supplies or sinks most of the charge needed to charge up the switched capacitor. This ensures that the op amp will never see the transient nature of the load. This additional filter can also provide antialiasing.



For converters that have a differential input, a differential version of this circuit may be used, as shown in the figure below. Since one input is positive with respect to ground while the other is negative, one input (the negative one) needs to be supplied negative charge while the other needs to get rid of negative charge when the input capacitors are switched on line. Connecting a capacitor between the two inputs enables most of the charge that is needed by one input to be effectively supplied by the other input. This minimizes undesirable charge transfers to and from the analog ground.



To be continued. Topics to be covered in the next installment include multiplexing, clock signals, noise, dither, averaging, spec clarifications

APPENDIX

RSS addition of logarithmic quantities: The root-sum square of two rms signals, S₁ and S₂, has an rms value of $\sqrt{S_1^2 + S_2^2}$. One often needs to calculate the rss sum of two numbers that are expressed in dB relative a given reference. To do this one has to take the antilogs, perform the rss addition, then convert the result back to dB. These three operations can be combined into one convenient formula: If D₁ and D₂ are ratios expressed in dB [negative or positive] their sum, expressed in dB, is

$$10 \log_{10} \left(10^{D_1/10} + 10^{D_2/10} \right)$$

Similarly, to find the difference between two rms quantities,

$$x = \sqrt{S_2^2 - S_1^2}$$

the result, x, expressed in dB, is

$$10 \log_{10} \left(10^{D_2/10} - 10^{D_1/10} \right)$$

References (not available from Analog Devices):

- ¹ *Oversampling Delta-Sigma Data Converters—Theory, Design, and Simulation*, edited by J.C. Candy and G.C. Temes, IEEE Press, Piscataway, NJ, 1991.
- ² J. Vanderkooy and S.P. Lipshitz, “Resolution Below the Least Significant Bit in Digital Systems with Dither,” *J. Audio Eng. Soc.*, vol. 32, pp. 106-113 (1984 Mar.); correction *ibid.*, p.889 (1984 Nov.).
- ³ A.H. Bowker and G.J. Lieberman, *Engineering Statistics*, Prentice Hall, Englewood Cliffs, NJ, 1972.

Ask The Applications Engineer—16

by Oli Josefsson

USING SIGMA-DELTA CONVERTERS—PART 2

This is a continuation of a discussion of sigma-delta converters begun in the last issue. We covered antialiasing requirements, idle tones, and loading on the signal source.

Q: What happens if my input signal is beyond the input range of the sigma-delta converter? I remember hearing something about the converter becoming unstable?

A: The modulator can become temporarily unstable if it is driven with inputs outside the recommended range. However, this instability is invisible to the user, since decimators are generally designed to simply clip the digital output and show either negative or positive full scale, just as one would expect with a conventional converter.

Q: The specifications for sigma-delta converters assume a certain input clock rate and therefore a specific sampling rate. Can I safely use the converter with a higher or lower clock frequency?

A: While the specs are measured at a particular sampling frequency, we often specify a range of input clock frequencies that the device can be operated with. This translates into a range of possible sampling rates. If you plan to go much beyond that range you can expect some performance degradation. If you sample at higher rates than specified, the internal switched-capacitor circuits may not be able to settle to the required accuracy before a new clock edge comes along. With too slow a sampling rate, capacitor leakage will degrade performance.


The digital filter characteristics of the converter (group delay, cutoff frequency, etc.) scale with sampling rate; so too do the input impedance (unless the input is buffered) and power consumption.

Q: I am planning to use a sigma-delta converter to digitize several signals by using a multiplexer at the input of the converter. Is that a problem?

A: While sigma-delta converters have a certain appeal due to their ease of antialiasing, they do not lend themselves well to applications for multiplexed ac signals. The reason for this is that the output of a sigma-delta converter is a function not only of the latest analog input but also of previous inputs. This is mostly due to the memory that the digital filter has of previous inputs, but the modulator has some memory as well. In a multiplexing application, after switching from one input to another, all information the filter has about the old input needs to be flushed out before the converter output word represents the new input.

Most decimation filters in sigma-delta converters intended for ac applications are FIR filters, principally because of their linear phase-response. For FIR filters, it is easy to calculate the time it takes to rid the filter of any information about the old input. The figure shows the structure of a FIR filter; the number of clock cycles required to clock all old data points out (i.e., the filter settling time) is equal to k , the number of taps in the filter. While data corresponding to a new input is propagating through the filter and replacing the earlier data, the output of the filter is calculated from a combination of the old data and

the new data. The AD1879, for example, an 18-bit audio A/D converter, has a 4096-tap FIR filter which, when running at 3.072 MHz, has a 1.33-ms settling time.



The effective sampling rate for sigma-delta converters in multiplexed applications is quite low because of this need to wait for the old signal to be flushed out before capturing a valid data point for the new input. Traditional converters, which convert directly, or in a small number of stages, are therefore a much better choice in applications requiring the capture of multiple ac channels.

For a multichannel dc application where time is available to wait after switching between channels, or if the application does not require frequent changes between channels, the use of a sigma-delta converter can be very feasible. In fact, Analog Devices offers 16-24-bit converters with multiplexers on the input (AD771x family) specifically for such applications.

Q: Does this also explain why sigma-delta converters are not suitable for some control applications?

A: Yes. Since delays in control loops must be minimized for stability, sigma-delta converters are not suitable for control applications where they add a relatively long time delay. However, the actual delay is predictable; in applications that involve relatively slow signals, the converter phase delay, and therefore the effect on pole and zero locations of the control loop, may be negligible. However, even if this is the case, a traditional non-oversampling converter may still be a much better choice for the application, because a sigma-delta converter would need to run at a much faster sampling rate than a traditional converter in order to have the same phase delay. This will unnecessarily burden the circuitry that processes the A/D data.

Q: Are there any other issues I should be aware of when using sigma-delta converters?

A: In addition to the general guidelines on grounding, power supply bypassing, etc., that apply to all converters, there are a couple of points worth remembering when designing with sigma-delta converters. The first issue involves their input. As mentioned earlier, some sigma-delta converters (such as the AD1877) have buffers on the input; others (such as the AD1879), without a buffer, present a switched-capacitor load, which needs periodic current transients to charge the input capacitor. It is important that the circuitry driving the converter be as close to the converter as possible to minimize the inductance in the leads between the external circuitry and the switched-capacitor node. This reduces the settling time of the input and minimizes radiation from the input to other parts of the circuit board.

Another issue has to do with interference from clock signals affecting the A/D conversion. As I noted earlier, the digital decimation filter can't provide any filtering of signals whose frequencies are close to multiples of the modulator sampling rate. To be precise, the passbands are $[kF_{ms} \pm f_b]s$ where k is an integer, F_{ms} is the modulator sampling rate, and f_b is the decimator cutoff frequency.

Besides the consequences for anti-aliasing discussed earlier, the decimator cutoff frequencies have a bearing on the selection of clock frequencies for devices that operate in the same system as the converter. These frequency bands (i.e., the passbands) embody the converter's greatest vulnerability to interference (inductive or capacitive coupling, power supply noise, etc.), because any signals in these frequency bands that manage to get into the modulator will not be subjected to attenuation in the filter. Therefore one is wise to avoid using clock frequencies that fall in these bands to minimize the possibility of interfering with the conversion—unless they are synchronous with the converter clock.

QUESTIONS ON NOISE IN CONVERTERS

Q: I recently evaluated a dual-supply A/D converter; one of the tests I did was to ground the input and look at the output codes on a LED register. To my big surprise I got a range of output codes instead of a single code output as I expected?

A: The cause is *circuit noise*. When the dc input is at the transition between two output codes, just a little circuit noise in even the finest dc converters will ensure that two codes will appear at the output. This is a fact of life in the converter world. In many instances, as in your case, the internal noise may be large enough to cause several output codes to appear. Consider, for example, a converter with peak-to-peak noise of just over 2 LSB. When the input of this converter is grounded, or a clean dc source is connected to the input, we will always see three—and sometimes even four—codes appear at the output. The circuit noise prevents the voltage being sampled from being confined to a voltage bin that corresponds to one digital code. Any external noise on the A/D input (including a noisy signal), on the power supplies, or on the control lines will add to the internal circuit noise—and possibly result in more bits toggling.

Q: Is there a way to determine how many codes I can expect to appear when I apply a dc signal to a converter?

A: It would not be hard in the ideal case where you knew the noise distribution, the exact size of the codes where the dc input is at and where within a code quantum the input lies (in the center, on the edge of two codes, etc.). But in reality you don't have this information. However, knowing some of the ac specifications (S/N, dynamic range, etc.) of the converter, you can make an *estimate*. From these specs you can find the magnitude of the rms converter noise relative to full scale. The noise will in all likelihood have a Gaussian amplitude distribution, so the standard deviation (sd) of the distribution equals the rms value. This also means that the codes that appear will not have equal probability of occurring. Using the fact that 99.7% of a Gaussian distribution occurs within ± 3 standard deviations from the mean, we can estimate the peak-to-peak noise voltage at six times the standard deviation.

If N_{rms} is the rms value of the converter noise and V_{LSB} is the size of the LSB in volts ($= V_{span}/2^b$, where b is the number of bits in the output word) the peak to peak noise in terms of LSBs, N_B , is

$$N_B = \frac{6 \times N_{rms}}{V_{LSB}} = \frac{6 \times 2^b \times N_{rms}}{V_{span}}$$

If the signal-to-noise ratio of a converter expresses noise power relative to full scale, rms signal ($V_{span}/(2\sqrt{2})$), we have

$$N_B = \frac{3}{\sqrt{2}} \times 2^b \times 10^{-SNR/20}$$

How many *codes* show up at the output depends where the mean of the input, i.e., the dc input value, is with respect to code transitions. If the mean is close to the boundary between two output codes, more codes are likely to appear than if the mean is half way between two output codes. It can easily be shown that N_C , the number of codes appearing for a particular value of N_B , is either $INT(N_B)+1$ or $INT(N_B)+2$, depending on the dc input value [$INT(N_B)$ is the integer portion of N_B]. And don't be surprised to see even more codes from the less-probable noise amplitudes $>\pm 3$ standard deviations.

How many *bits* will N_C cause to toggle on the output? The number of bits needed to represent N_C codes is

$$INT\left(\frac{\log N_C}{\log 2} + 0.5\right)$$

We can, however, see many more bits toggle, since the number of bits toggling is a function of the actual value of the converter's dc input. Consider, for example, that a one-code transition from an output word of -1 to 0 on a 2s-complement-coded converter involves inverting all the output bits.


Lets look at an example using the AD1879, an 18-bit sigma-delta converter with dynamic range of 103 dB. From the definition of dynamic range we have

$$103 = 20 \log \frac{S}{N_{rms}}$$

From the AD1879 data sheet, we find that the rms value of a full-scale input signal, S , is $6/\sqrt{2}$ V rms. This allows us to solve for N_{rms} which turns out to be $30 \mu V$. We next find the LSB size by dividing the full input range by the number of possible output codes:

$$V_{LSB} = \frac{12}{2^{18}} = 45.8 \mu V$$

Thus N_B is 3.9. We can therefore expect either 4 or 5 different codes to appear at the AD1879 output when the input is grounded (ground corresponds to a midscale input for the AD1879).





One can take this estimation one step further: If the standard deviation (the rms value) of a Gaussian distribution and the mean (the mean of the noise is 0 in this case) are known, one can use standard tables for the Gaussian distribution to calculate what percent of the time the noise will fall into a voltage interval corresponding to a specific output code. A histogram can be estimated, showing the distribution of codes at the output. Also the process can be reversed: a histogram showing the distribution of noise codes at a given value of dc output permits one to estimate the S/N ratio for a converter.

To make all this real, let's continue our example involving the AD1879. Consider two cases, one where the input lies midway between two output codes and one when the input is on the transition between two codes. From the calculations above, we found that the standard deviation (sd) of the noise (the rms value) was 30 μV. The size of one LSB in terms of sd is

$$\frac{45.78 \mu\text{V}}{30.0 \mu\text{V}} = 1.524$$

In the case where the dc input is midway between code transitions, as shown below, it is clear that any noise that falls within -0.5 LSBs to +0.5 LSBs from the input will result in the correct code at the A/D output. This corresponds to the noise being confined to a range of (-0.5 × 1.524) sd to (+0.5 × 1.524) sd from its mean (0). From standard tables one can find that the noise will fall in this range 55.4% of the time. If the noise falls within 0.5 LSBs to 1.5 LSBs, the output will be one code too high. Again from standard tables one can find that this will occur 21.2% of the time. Continuing in this manner one can calculate the whole histogram showing the distribution of output codes.



The upper figure shows an actual measurement where the dc input happened to be -25 LSBs. Five output codes, ranging from -27 to -23, appeared. 1024 measurements were taken and the percentage distribution of each code is shown on top of each column. The calculated distribution is listed in brackets on top of each column. As can be seen, the experimental results agree well with the calculated values. The lower figure shows a case where the dc input is close to the boundary between two codes. By following a similar procedure, one can calculate how the histogram should look. Again the experimental and calculated values are in excellent agreement. Note that the

actual applied dc input is slightly above the border between the two codes, whereas the calculations assume it is exactly on the border.


The biggest weakness of this estimating technique is the fact that in conventional converters the code width (the amount the dc input has to be increased to increase the digital output by one bit) varies from code to code. This means that if the dc input is in an area where codes are narrow, we can expect more bits to be toggling than in an area where the codes are wide. This method also assumes that the circuit noise within the converter stays constant, whether the applied signal is ac or dc. This is not exactly true in many cases.

The estimate will probably be more accurate when used with sigma-delta converters (except for “dead bands”), because neither of the two factors mentioned above is an issue in such converters.


Q: Ah, now I understand why there are multiple codes at the output. But why not discard the bits that toggle and only bring out the bits that stay steady, since the others are really indeterminate? Isn't that the real resolution of the converter?

A: Many converters are designed for ac or dynamic applications where THD (total harmonic distortion) and THD+N (total harmonic distortion+noise) are the most important specs. The design therefore focuses on minimizing harmonic distortion for high- and low-level input signals, while keeping the noise to acceptable levels. As it turns out, these requirements somewhat contradict the requirements for a good dc converter, which is optimized for precision conversion of slow moving signals where harmonic distortion is not an issue. It is actually desirable to have some noise (called dither) superimposed on the input signal to minimize distortion at very low input signal levels; dither can also be used to improve dc accuracy where repeated measurements can be made.

To understand how this may be, let's start by looking at quantization noise. The output of an ideal A/D converter has finite accuracy because of the finite number of bits available to represent the input voltage. Each one of the 2^b quanta represents with one single value all values in the analog range from -0.5 LSB to +0.5 LSB of its nominal input value. The A/D output can therefore be thought of as a discrete version of the analog input plus an error signal (quantization noise). When a large and varying input signal (dozens, hundreds, or thousands of LSBs in amplitude) is applied to a converter, the quantization noise has very little correlation with the input signal. It is, in other words, approximately white noise. The figure shows the quantization noise of a perfect A/D converter at various instants of time when the input signal is a sinusoid of about 100 LSBs in amplitude.



When the A/D input is very low in amplitude, so that the amplitude does not change more than a fraction of a LSB between samples, the samples stay in the same quantum, and are therefore constant for a few sample periods. This is depicted in the figure below, which shows a sinusoidal input signal that has an amplitude of only 1.5 LSBs, the A/D output and the quantization noise. Note that the quantization error follows the input waveform exactly while the samples are staying constant. The longer the samples stay constant, the more the quantization noise looks like the input waveform, i.e., the correlation between the input signal and the quantization noise increases. While the rms of the quantization error may not have changed, the quantization error will take on a non-uniform spectral shape. In fact, the correlated quantization noise shows up as harmonics in the A/D spectrum.

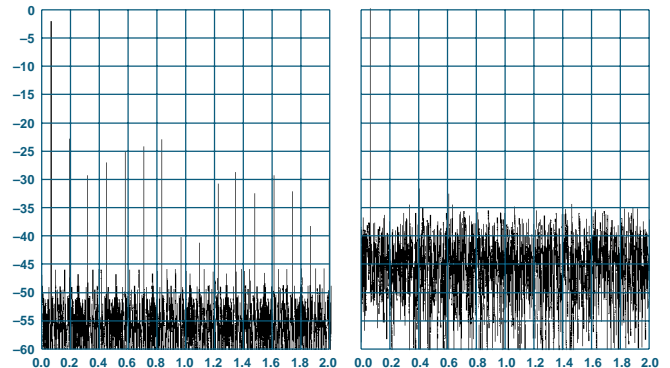


Another way to look at this phenomenon is to consider the case when the (sinusoidal) input signal is only around 1 LSB in size and the digital output resembles a square wave. Square waves are rich in harmonics! The harmonics, or noise modulation products, are very objectionable in many converter applications, especially audio.

To get around this problem, a technique called dithering is used to trade correlated quantization noise for white noise, which is less offensive to the human ear than correlated noise. Dithering is done by using circuit elements to *add* random noise to the input signal. While this will result in an increase of the total converter noise, the added noise breaks up the simple square wave patterns in the output code. The quantization error will not be a function of the input signal but of the instantaneous value of the dither noise. Thus the dither decorrelates the quantization noise and the input signal. The size of the dither

signal is often about 1/3 LSB rms (2 LSBs peak-to-peak if the noise is Gaussian). Clearly, this will result in a converter that will have more than two codes at the output when the input is grounded. We saw an example earlier involving the AD1879 which had either four or five codes appear on the output depending on the dc input level.

The figure below shows the simulated output of an A/D converter with an undithered low level input signal. The quantization noise is a function of the input signal magnitude at the sample instant. This correlation between the quantization noise and the input signal shows up as a cluster of harmonically related sticks in the A/D output spectrum. Note that the magnitude scale in the figure is referenced to the input signal (not full scale input).



The right-hand figure shows the A/D output after a dither signal that is 4 dB above the quantization noise floor is added to the input. In this case the quantization noise depends on the magnitude of the dither signal at the instant when a sample is taken. Since the value of the dither doesn't depend on the input signal, the quantization noise becomes uncorrelated to the input and the harmonics in the A/D spectrum are eliminated, but at the cost of an overall increase in the noise floor.

Instead of actually adding noise to the A/D input, dithering can be accomplished by using the thermal noise of the converter as the dither signal and calculating enough output bits to ensure a decorrelated quantization noise.

Though I have used A/D converters in my examples, the idea of using dither also applies to D/A converters as well. Dither is applied to D/A converters by adding the output of a digital noise generator to the digital word sent to the D/A.

Q: But in dc applications, I want to make an accurate measurement each time and may not be able to tolerate the uncertainty of having a few LSBs of error in a particular measurement.

A: If you need n -bit dc accuracy in each conversion and you have problems finding a suitable n -bit converter, you have two options. One is to use an $(n+2)$ -bit converter and simply ignore the two LSBs. However, if your hardware has the capability (and time) to do some signal processing, you can enhance the resolution of a noisy (dithered) dc converter and, in fact, get more than n -bit accuracy out of an n -bit converter if the accuracy is limited by noise.

To understand why this may be so, think of an ideal n -bit converter. For a particular value of dc input, you will get one digital code at the output. However, you do not know where the input lies within the code quantum (i.e., in the middle,

close to the upper transition, etc.). That may be sufficiently accurate for your application, but if you add noise to the input of the converter—so that several codes can appear at the output—you will find that the code distribution contains information to place the dc value of the input more exactly.

In the earlier examples involving the AD1879, we saw how the code distribution looks when the input is in the vicinity of a code transition; the two most-frequent output codes are the ones on either side of the transition. Their average is therefore a good estimate of where the input lies. In fact, taking the average of a lot of conversions, while the input stays put, is an excellent way of enhancing the resolution of the converter. One has to be careful, when processing the converter output, to allow the output word length to grow without introducing roundoff errors. Otherwise one actually injects unwanted noise—called *requantization noise*—into the final output. Note that filtering out the noise is only just that; it will have no effect on other error sources of the converter, such as integral and differential nonlinearity.

This concept of resolution enhancement is an interesting one and is not restricted to the dc domain. One can actually trade resolution for bandwidth in the ac domain and combine the outputs of several converters or to construct a more-accurate output. The basic principle is that signal repetitions (which are self-correlated) add linearly, while repetitions of random noise produce root-square increases. Thus, a fourfold increase in number of samples increases S/N by 6 dB.

Q: You mentioned a couple of converter ac specifications above. I am somewhat confused about how S/N, THD+N, THD, S/THD, S/THD+N, and dynamic range are measured on A/D and D/A converters and how they relate to each other. Can you shed any light on this?

A: Your confusion is quite understandable. There is unfortunately no industry standard on exactly how these quantities are measured and therefore, what exactly they mean. Sometimes manufacturers are guilty of choosing the definition that portrays their part favorably.

Most often data sheets include a note on the testing conditions and how the different specs were calculated. The best advice I can give is to read these very carefully. By simple calculations you can often convert a specification for one part to a number that allows a fair comparison to a specification for another part. Most specifications are not expressed in absolute units, but as relative measurements or ratios. Noise, for example, is not specified in rms volts, but as SNR, or the ratio between signal power and noise power under particular test conditions. These ratios are usually expressed in decibels, dB, and occasionally as percentages (%). A power ratio, x , expressed in bels, is defined as $\log_{10}x$; multiply by 10 if expressed in decibels (one tenth of a bel): $10 \log_{10}x$. SNR is therefore equal to $10 \log_{10}$ (signal power/noise power) dB. Evaluated in terms of rms voltage quantities, $SNR = 20 \log_{10}(V_{signal}/V_{noise})$.

Armed with this knowledge, let's see whether we can make sense out of the multiple specifications you mentioned above (many of which are redundant). Those specifications seek to describe how the imperfections of the converter affect the characteristics of an ac signal that gets processed by the converter. For dc applications, a listing of the magnitude of

the actual imperfections suffices, but these can only suggest ac performance. For example, integral nonlinearity is a major factor in determining large-signal distortion (along with glitch energy for D/As) while differential nonlinearity governs small-signal distortion. To accurately determine the ac performance, at least two types of tests are performed in the case of A/Ds. The tests are as follows:

a) Full-scale sine

A sinusoidal signal approaching full-scale is applied to the converter. The signal is large enough so that converter's imperfections cause significant harmonic components to occur at multiples of the input signal frequency. The harmonics will show up in the output spectrum, along with noise. A common performance measure is the relative magnitude of the harmonic components, usually expressed in dB. Relative to what? Two possibilities are the applied input signal and the full scale of the converter (which in most cases is different from the applied input signal). Referring the harmonics to full scale will clearly yield a lower (more attractive) number than referring them to the rms value of the actual input signal. This reference issue causes a lot of confusion when dynamic specifications are evaluated, because there is no universally accepted standard for what each performance measure should be referred to. The best advice I can give you is: never assume anything; read manufacturers' data sheets very carefully.

Sometimes the magnitudes of the individual harmonics are specified, but most often only the total harmonic distortion (THD) is specified. The THD measures the total power of the harmonics and is found by adding the individual harmonics in rss fashion. The formula then for THD when referred to the input signal is

$$20 \log_{10} \left[\frac{\sqrt{\sum_{i=2}^m H^2(i)_{rms}}}{S} \right] \text{ or } 10 \log_{10} \left[\frac{\sum_{i=2}^m H^2(i)_{rms}}{S^2} \right]$$

where $H(i)_{rms}$ refers to the rms value of i th harmonic component and S to the rms value of the input signal. Usually, harmonics 2 through 5 are sufficient. Note that the input-frequency, or *fundamental*, component is the first harmonic. To refer any harmonic to full scale, add x dB to the formula above, where x is the magnitude of the input signal relative to full scale. This simple conversion formula can be applied to other specifications, but take care to observe proper polarity of the log quantities.

Nowadays, clear distinction is usually made between total harmonic distortion plus noise (*THD+N*) and *THD*. This has not always been the case. *THD+N* includes not only the harmonics that are generated in the conversion, but also the noise. The formula for *THD+N* when referred to the input signal is:

$$20 \log_{10} \left[\frac{\sqrt{N^2_{rms} + \sum_{i=2}^m H^2(i)_{rms}}}{S} \right]$$

or

$$10 \log_{10} \left[\frac{N^2_{rms} + \sum_{i=2}^m H^2(i)_{rms}}{S^2} \right]$$

where N_{rms} is the rms value of the integrated noise in the bandwidth specified for the measurement.

Another commonly used specification is signal to noise-plus-distortion ($S/[N+D]$, or $S/[THD+N]$), also called *sinad*. This is essentially the inverse of $THD+N$, when referred to the signal; its dB number is the same, but with opposite polarity.

Another performance measure describing the test results is the signal to noise ratio, S/N or SNR , which is a measure of the relative noise power, most useful for estimating response to small signals in the absence of harmonics. If S/N is not specified, but THD and $THD+N$ are provided, relative to the input signal, THD can be rss-subtracted from $THD+N$ to obtain the noise to signal ratio [$= 1/(S/N)$]. If the numbers are given in dB, the rss subtraction formula for logarithmic quantities in the Appendix can be used as follows

$$SNR = -10 \log_{10} \left(10^{(THD+N)/10} - 10^{THD/10} \right)$$

to yield the input signal power relative to noise power expressed in dB.

b) Low-level sine

The second test usually performed is to apply a sinusoidal signal well below full scale to the converter (usually -60 dB). At this input level, sigma-delta converters usually exhibit negligible nonlinearities, so only noise (no harmonic components) appears in the spectrum. At this level, $S/N = S/N+D = -THD+N = -THD$, when all are referred to the same level. As a result, one specification indicating the noise level suffices to describe the result of this test. This specification called *dynamic range* (inversely, dynamic-range *distortion*), specifies the magnitude of the integrated noise (and harmonics if they exist) over a specific bandwidth relative to full scale, when a -60-dB input signal is applied to the converter.

Conventional (i.e. not sigma-delta) converters can exhibit harmonics in their output spectrum even with low-level input signals because all the codes may not have equal width (differential nonlinearity). In some such instances, the S/N , which ignores harmonics, measured with a -60-dB input signal, is different from dynamic range.

Frequently one sees $THD+N$ at -60-dB and dynamic range specified for the same converter. These really are, as explained above, redundant since they only differ in the reference used. The only twist on dynamic range is that sometimes, when audio converters are specified, a filter that mimics the frequency response of the human ear is applied to the converter output. This processing of the converter output is called A-weighting (because an A-weighting filter is used); it will effectively decrease the noise floor, and therefore increase the signal-to-noise ratio, if the noise is white.

Everything discussed above applies to both A/D and D/A converters, with the possible exception of signal to noise ratio. Sometimes (particularly for audio D/A converters) S/N is a measure of how “quiet” the D/A output is when zero (midscale) code is sent to the converter. Under these conditions, the S/N expresses the analog noise power at the D/A output relative to full scale output.

It’s important to note that the performance measures above are affected by: *bandwidth* of the measurement, the *sampling frequency*, and the input *signal frequency*. For a fair comparison of two converters, one has to make sure that these test conditions are similar for both.

Image Filtering Question

Q: I intend to use Analog’s AD1800 family of audio D/A converters for a digital audio playback application. I understand that using an interpolator ahead of the D/A will make it easier to filter the D/A output, assuming I want to get rid of all the images at the D/A output. But is it really necessary to filter the output, since all the images will be above the audible range as long as sampling is at >40 kHz?

A: Good question. The audio equipment (audio amplifiers, equalizers, power amplifiers, etc.) that may eventually receive the output of your D/As are typically built to handle 20-Hz to 20-kHz signals. Since they are not intended to respond at frequencies much beyond 20 kHz—and in effect themselves function as filters—they may not have the necessary slew rate and gain to handle incoming signals from an unfiltered D/A output having significant energy well above 20 kHz. With their slew-rate and gain limitations, the amplifiers are driven into nonlinear regions, generating distortion. These distortion products are not limited to high frequencies but can affect the 20-Hz to 20-kHz range as well. Attenuating the high frequency signals at the DAC will therefore reduce the possibility of distortion. CD players often include filters steep enough to reduce the total out-of-band energy to >80 dB below full scale.

APPENDIX

RSS addition of logarithmic quantities: The root-square sum of two rms signals, S_1 and S_2 , has an rms value of $\sqrt{S_1^2 + S_2^2}$. One often needs to calculate the rss sum of two numbers that are expressed in dB relative a given reference. To do this one has to take the antilogs, perform the rss addition, then convert the result back to dB. These three operations can be combined into one convenient formula: If D_1 and D_2 are ratios expressed in dB, their sum, expressed in dB, is

$$10 \log_{10} \left(10^{D_1/10} + 10^{D_2/10} \right)$$

Similarly, to find the difference between two rms quantities,

$$x = \sqrt{S_2^2 - S_1^2}$$

the result, x , expressed in dB, is

$$10 \log_{10} \left(10^{D_2/10} - 10^{D_1/10} \right) \quad \blacksquare$$

References:

[1] *Oversampling Delta-Sigma Data Converters—Theory, Design, and Simulation*, edited by J.C. Candy and G.C. Temes, IEEE Press, Piscataway, NJ, 1991.
 [2] J. Vanderkooy and S.P. Lipshitz, “Resolution Below the Least Significant Bit in Digital Systems with Dither,” *J.Audio Eng. Soc.*, vol. 32, pp. 106-113 (1984 Mar.); correction *ibid.*, p.889 (1984 Nov.).
 [3] A.H. Bowker and G.J. Lieberman, *Engineering Statistics*, Prentice Hall, Englewood Cliffs, NJ, 1972.