Signals From Noise: Calculating Delta-Sigma SNRs

by Dave Van Ess, Principal Application Engineer, MTS, Cypress Semiconductor

Understanding the operation of a Delta-Sigma (Δ-Σ) analog-to-digital converter (ADC) can appear to be complicated. Many articles and textbooks, when trying to explain the operation, choose to gloss over the basics and have you accept their word that they function the way they do. I refer to this as “Proof By Intimidation” (**PBI**). Actually, if broken down correctly, Delta-Sigma ADC operation can be explained in several easy to manage portions.

A Delta-Sigma ADC contains the following:

- A Delta-Sigma Modulator (**DSM**) to convert an analog signal to a digital density
- A digital filter to convert a density signal to a digital word. This digital word is sometimes called Pulse Code Modulation (**PCM**)

A simplified block diagram is shown below.

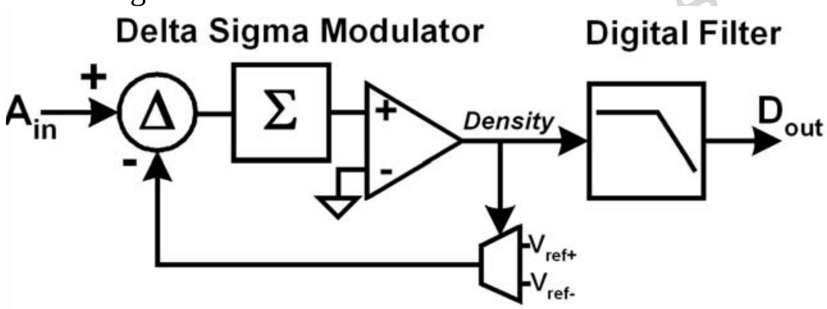

Fig. 1: ADC Comprised With Delta-Sigma Modulator And Digital Filter

A Bit about Density

Before starting to analyze the Delta-Sigma modulator, a brief explanation of digital density is warranted. Suppose you wish to supply 75% of a reference voltage. You can either design circuitry that will supply 75% of the reference all the time, or supply all of the reference 75% of the time. Either way, on the average you get 75% of the reference. An example is a PWM filtered DAC. Increase the PWM duty cycle and the filtered output gets larger. Reduce it and the output gets smaller. The figure below shows a signal value “x” that is between two references values (±½**s**. The difference between the reference values is the quantization level **s**.

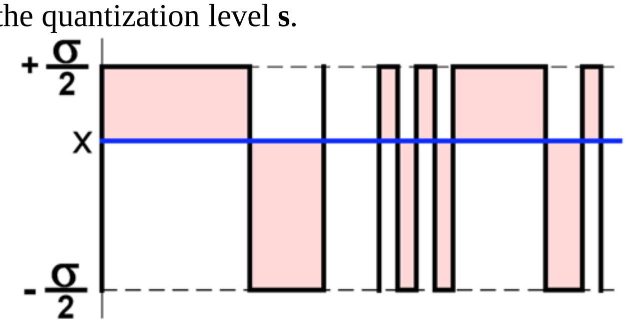

Fig. 2: Signal Value Represented By Percentages Of Two Bounding References

The density is set so the average of the digital waveform equals the average of the signal. Note that in the figure two different waveforms are shown and both average to the signal value. The first is a simple PWM waveform where its density is the “duty cycle.” The second waveform is far more complex. With density, the actual waveform shape is not important. What matters is the percentage that it is high. Given the references **±½s** and a density value the average signal is:

$$x = density \cdot \frac{+\sigma}{2} + (1 - density) \cdot \frac{-\sigma}{2} \quad ;0\% \le density \le 100\% \qquad (1)$$

In Fig. 2 the error between the density waveform and the signal is shown in red. While it averages to zero, the rms error is:

$$\varepsilon(x)_{rms} = \sqrt{density \cdot \left(\frac{+\sigma}{2} - x\right)^2 + (1 - density) \cdot \left(\frac{-\sigma}{2} - x\right)^2} = \sqrt{\tfrac{1}{4}\sigma^2 - x^2} \qquad (2)$$

The error is as little as zero at the reference boundaries and is the largest (**½s**), half way between the references. For a varying signal that is evenly distributed between the boundaries, the rms error is:

$$\varepsilon_0 = \sqrt{\frac{1}{\sigma} \int_{\frac{-\sigma}{x}}^{\frac{\sigma}{x}} \varepsilon(x)_{rms}^2 \, dx} = \frac{\sigma}{\sqrt{6}} \qquad (3)$$

This equation doesn’t state where the noise will be. The shape of density waveform determines this. What it says is that when all the quantization noise is accounted for, it will be **s/√6**. This value will be important in calculating the noise density distribution of a Delta-Sigma Modulator.

Delta-Sigma Modulator

A Delta-Sigma Modulator must have:

- A difference (**delta**) circuit
- An accumulate (**sigma**) circuit
- A quantization (**modulation**) circuit consisting of an ADC and a DAC

The modulator in Fig. 1 meets these requirements as the comparator acts as a single bit ADC and the multiplexed references make up a single bit DAC. The waveforms for a simulation, in Excel, are shown in Fig. 3, overleaf.

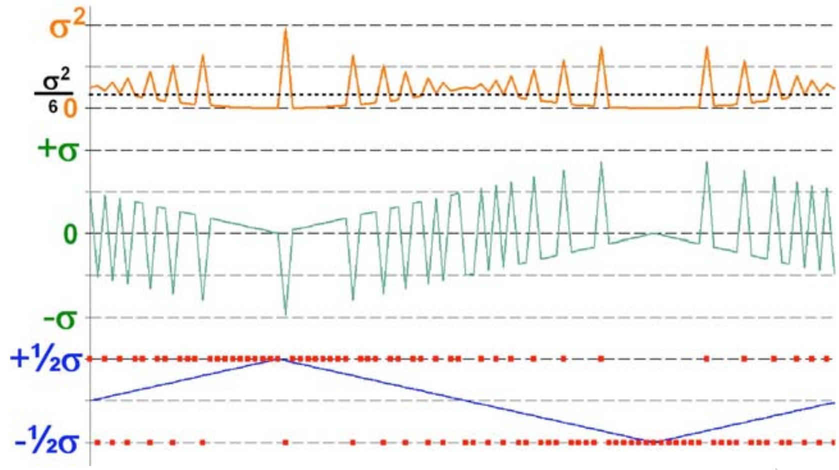

Fig. 3: Waveforms For Signal Stage Delta-Sigma Modulator

The blue signal is the input and the red is its quantized output. The difference between them is the quantization error that is displayed in green. Note that it has a range of **±s**. The square of this error is shown in orange and has a range of 0 to $\mathbf{s^2}$. The mean of this squared error signal is $\mathbf{s^2/6}$. Taking the square root of this mean results in an rms value of **s/√6**. This is the expected value for the quantization noise of a density signal. Again, it doesn't say where the noise is but that when all totaled it equals **s/√6**. To figure out the shape of the noise, the quantization noise model (see Fig. 4) is used.

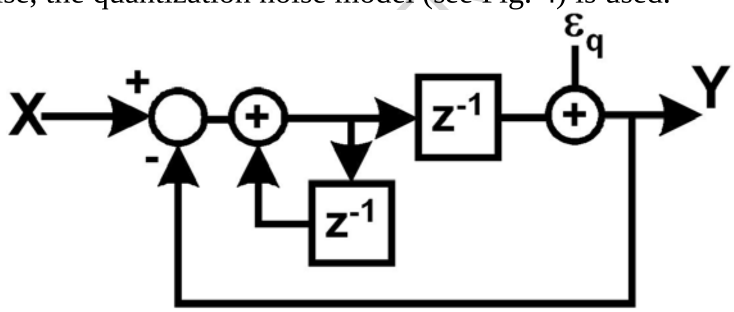

Fig. 4: Delta-Sigma Modulator Quantization Noise Model

To calculate the shape of the noise we now assume the input is adequately busy so the quantization error looks like white noise. Equation shows the output to be a function of the input and the quantization noise:

$$y = x \cdot z^{-1} + \varepsilon_q \cdot \left(1 - z^{-1}\right) \qquad (4)$$

Separating out just the noise component results in:

$$\eta(f) = \varepsilon_q \cdot 2 \cdot \sin(\pi \cdot f / f_s) \quad \int_0^{\frac{f_s}{2}} \eta(f)^2 df = \frac{\sigma^2}{6} \qquad (5)$$

Now, all noise will be found within the Nyquist Region (0 to ½f_s). Integrating the squared signal over this spectrum allows the value of the noise density (e) to be calculated. It has already been shown, both mathematically and empirically, that the total noise must add up to s/√6.

Solving for e_q results in:

$$\varepsilon_q = \frac{1}{\sqrt{f_s}} \frac{\sigma}{\sqrt{6}} \tag{6}$$

Substituting this into Equation 5 results in:

$$\eta(f) = \frac{1}{\sqrt{f_s}} \frac{\sigma}{\sqrt{6}} \cdot 2 \cdot \sin(\pi \cdot f / f_s) \tag{7}$$

The spectral density for this modulator with a sample rate of 1Msample/s and the quantization level (s) of 1 V is shown in Fig. 5.

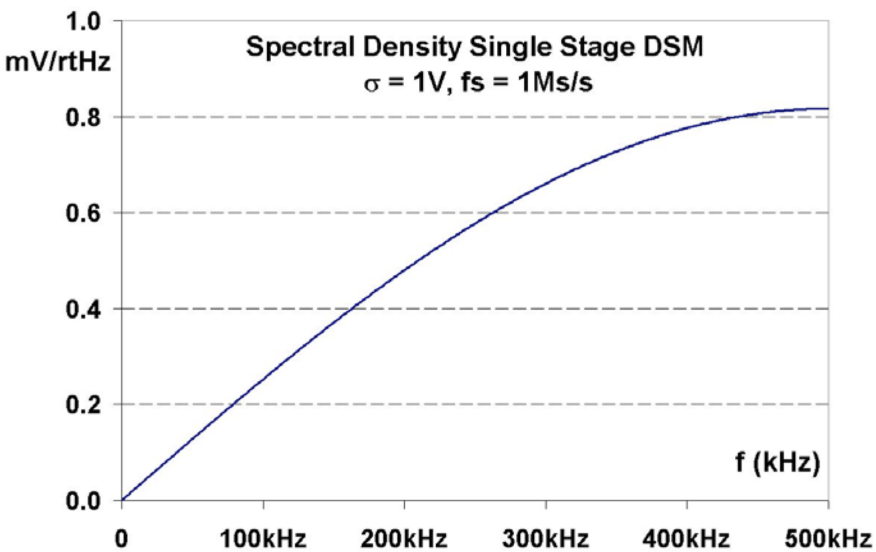

Fig. 5: Spectral Density For Single-Stage DSM

Multi-Stage Delta-Sigma Modulator

A higher-order modulator can be built by adding a difference circuit and integrator to the input of a lower-order modulator. The input to this new differentiator is the input signal and the fed back output of the lower-order modulator. A quantization module for a multi-order modulator is shown in Fig. 6.

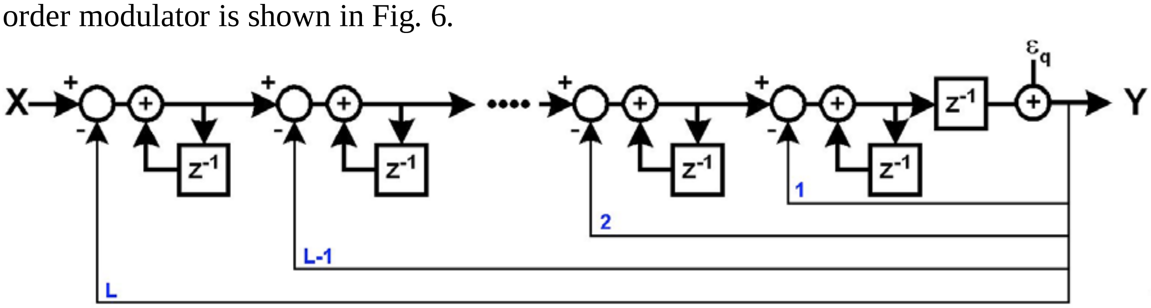

Fig. 6: Lth-Order Delta-Sigma Modulator Quantization Noise Model

With the addition of a new integrator also comes a new feedback loop. As shown above, this modulator has “L” feedback loops. Equation 8 shows the output to be a function of the input and the quantization noise.:

$$y = x \cdot z^{-1} + \varepsilon_q \cdot \left(1 - z^{-1}\right)^L \tag{8}$$

Showing this to be true for all values of “L” requires proof by deduction. For L = 1 the equation is:

$$y = x \cdot z^{-1} + \varepsilon_q \cdot \left(1 - z^{-1}\right) \tag{9}$$

Earlier this was shown to be true. If the response for an Mth-order modulator is:

$$y = x \cdot z^{-1} + \varepsilon_q \cdot \left(1 - z^{-1}\right)^M \tag{10}$$

Making an (M+1)th-order modulator requires that an integrator and feedback path be added to the input. Doing so results in:

$$y = \frac{(x - y)}{1 - z^{-1}} \cdot z^{-1} + \varepsilon_q \cdot \left(1 - z^{-1}\right)^M \quad y = x \cdot z^{-1} + \varepsilon_q \left(1 - z^{-1}\right)^{M+1} \tag{11}$$

So if the equation holds for a particular-order modulator then it also holds for the next higher-order modulator. Since it was shown to hold for a 1st-order modulator, it follows that it holds for all higher-order modulators.

To calculate the shape of the noise, we now assume the input is adequately busy so that the quantization error looks like white noise. Equation 12 shows the output to be a function of the input and the quantization noise.

$$y = x \cdot z^{-1} + \varepsilon_q \cdot \left(1 - z^{-1}\right)^L \tag{12}$$

Separating out just the noise component results in:

$$\eta(f) = \varepsilon_q \cdot \left(2 \cdot \sin(\pi \cdot f / f_s)\right)^L \quad \int_0^{\frac{f_s}{2}} \eta(f)^2 \, df = \frac{\sigma^2}{6} \tag{13}$$

Now all noise will be found within the Nyquist Region (0 to ½ f_s). Integrating the squared signal over this spectrum allows the value of the noise density (e) to be calculated. Using:

$$\int_0^{\frac{f_s}{2}} \left(\left(2 \sin(\pi \cdot f / f_s)\right)^L\right)^2 = \frac{L!}{\sqrt{(2 \cdot L)!}} \cdot \frac{f_s}{2} \tag{14}$$

Integrating the noise over this region results in:

$$\varepsilon_q = \sqrt{\frac{1}{f_s}} \cdot \sqrt{\frac{\sigma}{3}} \cdot \frac{L!}{\sqrt{(2 \cdot L)!}} \tag{15}$$

Substituting this value into Equation 13 results in Equation 16 for the shape and magnitude to the quantization noise:

$$\eta(f) = \sqrt{\frac{1}{f_s}} \cdot \frac{\sigma}{\sqrt{3}} \cdot \frac{L!}{\sqrt{(2 \cdot L)!}} \cdot \left(2 \cdot \sin(\pi \cdot f / f_s)\right)^L \tag{16}$$

The spectral density for this modulator with a sample rate of 1Msample/s and the quantization level (s) of 1 V is shown Fig. 7.

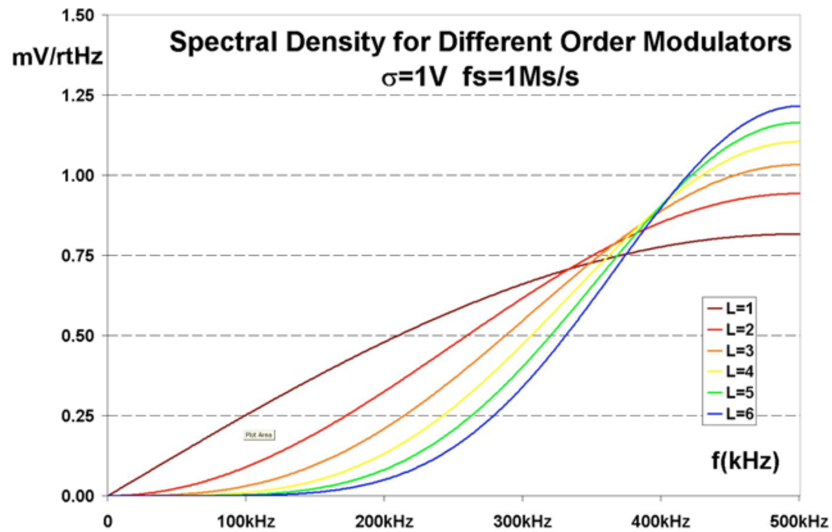

Fig. 7: Spectral Density For Multi-Order Stage DSMs

(In every case the integrated noise over the whole Nyquist bandwidth is **s/√6**)

Signal-to-Noise Ratio For An Ideal Filter

With the noise pushed towards the Nyquist limit, filtering the data with a low-pass digital filter greatly reduces the noise seen in the result. Equation 17 gives the total quantization noise for an ideal filter, as a function of the cutoff frequency (**f$_c$**) and DSM order (**L**):

$$\eta_c = \sqrt{\int_0^{f_c} \eta(f)^2\,df} = \sqrt{\frac{1}{f_s}}\cdot\frac{\sigma}{\sqrt{3}}\cdot\frac{L!}{\sqrt{(2\cdot L)}}\cdot 2^L \sqrt{\int_0^{f_c}\sin(\pi\cdot f/f_s)^{2L}\,df} \qquad (17)$$

For a cutoff frequency much lower than the sample frequency, Equation 18 is a reasonable approximation:

$$\eta_c \approx \sqrt{\frac{1}{f_s}}\cdot\frac{\sigma}{\sqrt{3}}\cdot\frac{L!}{\sqrt{(2L)}}\cdot\left(\frac{2\pi}{f_s}\right)^L \sqrt{\int_0^{f_c}(f_c)^{2L}\,df} = \frac{\sigma}{\sqrt{6\pi}}\cdot\frac{L!}{\sqrt{2L+1}}\cdot\left(\frac{2\pi f_c}{f_s}\right)^{\frac{2L+1}{2}} \quad ; f_c << f_s \quad (18)$$

The maximum peak-to-peak amplitude for an input sinusoid is s. Its rms value is s/√8. Given this the signal-to-noise ratio is:

$$SNR \approx \frac{1}{\sqrt{8}}\cdot\sqrt{6\pi}\,\frac{\sqrt{(2L+1)}}{L!}\cdot\left(\frac{OSR}{2\pi}\right)^{\frac{2L+1}{2}} \quad ; OSR = \frac{f_s}{f_c} \qquad (19)$$

Note that the cutoff frequency term has been replaced with an over-sample ratio (**OSR**). To express SNR in dB, the following approximation is used for the factorial term:

$$20\log\left(\frac{\sqrt{(2L+1)}}{L!}\right) \approx 6.56\cdot L + 1.73 \qquad (20)$$

It has accuracy better than a third of a dB for modulation orders between one and six.

Using this approximation, the SNR shown in Equation 19 can be expressed in dB. With an OSR of $\mathbf{2^n}$ the SNR is:

$$SNR_{dB} = 3.01 \cdot n \cdot (2L + 1) - 9.36L - 2.76 \tag{21}$$

Table 1 gives the SNR and equivalent number of bit (ENOB) for modulation order and over sample value.

		OSR (2^n)				
		2^4	**2^5**	**2^6**	**2^7**	**2^8**
Modulator Order (L)	**1**	24 dB 3¾ bits	33 dB 5¼ bits	42 dB 6¾ bits	51 dB 8¼ bits	60 dB 9¾ bits
	2	39 dB 6¼ bits	54 dB 8¾ bits	69 dB 11¼ bits	84 dB 13¾ bits	99 dB 16¼ bits
	3	53 dB 8 ¾bits	75 dB 12¼ bits	96 dB 15¾ bits	117 dB 19¼ bits	138 dB 22¾ bits
	4	68 dB 11¼ bits	95 dB 15¾ bits	122 dB 20¼ bits	149 dB 24¾ bits	177 dB 29½ bits
	5	83 dB 13½ bits	116 dB 19 bits	149 dB 24½ bits	182 dB 30 bits	215 dB 35½ bits
	6	99 dB 16 bits	137 dB 22½ bits	176 dB 29 bits	215 dB 35½ bits	254 dB 42 bits

Table 1: SNR And ENOB As Function Of Modulator Order And OSR

Note that for a doubling of the OSR (halving of the cutoff frequency), the resolution increases by L+ ½ bits. Now in no way does this say that you could slap together a sixth-order modulator and /256 filter to build a 40-bit ADC. There are many other noise sources that will reduce the overall SNR. This only gives the quantization noise contribution to SNR.

Signal-to-Noise Ratio For A SINC^K Filter

A SINC^K is a digital low pass filter that collects the output of the DSM at some higher sample frequency ($\mathbf{f_s}$) and uses it to generate a higher resolution signal at some lower decimated output frequency ($\mathbf{f_d}$). Its wider use is the result of ease of construction, as it does not require the use of a multiplier to generate weighted samples to accumulate. It is constructed with a series of digital accumulators operating at the sample frequency followed by difference circuits operating at the lower decimation frequency. A block diagram is shown in Fig. 8, overleaf.

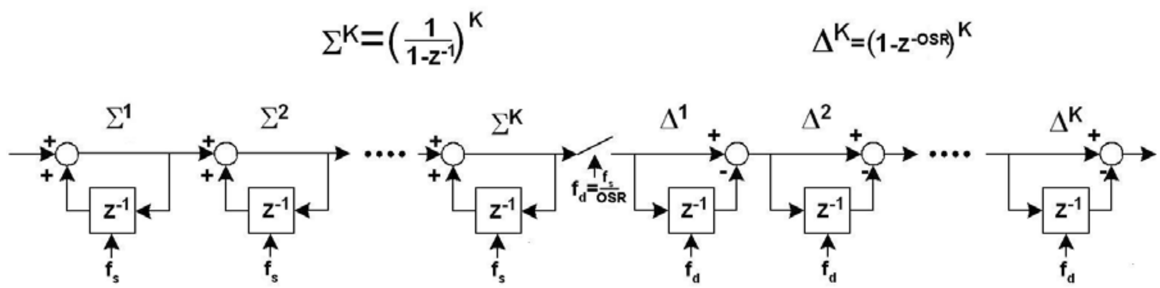

Fig. 8: Block Diagram For SINC^K Decimation Filter

The over-sample ratio (**OSR**) is the ratio of these two frequencies. The transfer function, expressed in z transforms is:

$$H(z) = \left(\frac{1}{M}\cdot\frac{1-z^{-OSR}}{1-z^{-1}}\right)^K \quad OSR = \frac{f_s}{f_d} \tag{22}$$

Using Equation 22, the frequency response is:

$$H(f) = \left(\frac{\sin(OSR\cdot\pi\ f/f_s)}{OSR\cdot\sin(\pi\cdot f/f_s)}\right)^K = \left(\frac{\mathrm{sinc}(\pi\ f/f_d)}{\mathrm{sinc}(\pi\ f/f_s)}\right)^K \tag{23}$$

From the equations it is apparent why it is called a "SINC" filter.

To calculate the total noise, the spectral noise of the DSM, from Equation 16, must be multiplied by the frequency response of the filter from Equation 23. The result is:

$$\eta(f) = \sqrt{\frac{1}{f_s}}\cdot\frac{\sigma}{\sqrt{3}}\cdot\frac{L!}{\sqrt{(2\cdot L)!}}\cdot(2\cdot\sin(\pi\cdot f/f_s))^L\cdot\left(\frac{\sin(OSR\cdot\pi\cdot f/f_s)}{OSR\cdot\sin(\pi\cdot f/f_s)}\right)^K\cdot \tag{24}$$

This filter is not ideal and does not remove the entire signal above some cutoff frequency. To calculate the total noise, it must be integrated over the whole Nyquist region. Doing so results in:

$$\eta_0(f) = \sqrt{\frac{1}{f_s}}\cdot\frac{\sigma}{\sqrt{3}}\cdot\frac{L!}{\sqrt{(2\cdot L)!}}\cdot\frac{2^L}{OSR^K}\cdot\sqrt{\int_0^{\frac{fs}{2}}\sin(\pi\cdot f/f_s)^{2(L-K)}\cdot\sin(OSR\cdot\pi\cdot f/f_s)^{2K}\,df} \tag{25}$$

I am not going to even try to integrate this function. So the total noise depends on the modulator order (L), decimator order (K) and over-sample ratio. The maximum input sinusoid is still s/√8. This works out to three variables and unfortunately tables in print best work for only two. So Table 2, overleaf, gives the signal-to-noise ratio and equivalent number of bits for an OSR of 16.

Modulator Order (L) \ Decimator Order (K)	1	2	3	4	5	6
1	22.8 dB 3½ bits	34.9 dB 5½ bits	37.6 dB 6 bits	39.6 dB 6¼ bits	41.0 dB 6½ bits	42.1 dB 6¾ bits
2	24.6 dB 3¾ bits	46.9 dB 7½ bits	60.0 dB 9½ bits	62.5 dB 10 bits	64.7 dB 10½ bits	66.6 dB 10¾ bits
3	25.1 dB 3¾ bits	49.1 dB 7¾ bits	71.0 dB 11½ bits	83.0 dB 13½ bits	86.8 dB 14¼ bits	89.4 dB 14½ bits
4	25.3 dB 4 bits	49.8 dB 8 bits	73.4 dB 12 bits	95.1 dB 15½ bits	107.1 dB 17½ bits	111.1 dB 18¼ bits
5	25.4 dB 4 bits	50.1 dB 8 bits	74.0 dB 12 bits	97.6 dB 16 bits	119.2 dB 19½ bits	131.2 dB 21½ bits
6	25.5 dB 4 bits	50.3 dB 8 bits	74.6 dB 12 bits	98.5 dB 16 bits	121.8 dB 20 bits	143.3 dB 23½ bits

Table 2: SNR/ENOB As Function Of L And K, OSR = 16

For any particular row, the most significant increase in resolution is reached for a decimation order that is one larger than the modulation order (K = L+1). Table 3 is for the same parameters except the OSR is set to 32.

Modulator Order (L) \ Decimator Order (K)	1	2	3	4	5	6
1	28.9 dB 4 ½ bits D 1	43.9 dB 7 bits D 1½	46.9 dB 7½ bits D 1½	48.7 dB 7¾ bits D 1½	50.0 dB 8 bits D 1½	51 dB 8¼ bits D 1½
2	30.6 dB 4 ¾ bits D 1	59.0 dB 9½ bits D 2	74.0 dB 12 bits D 2½	77.5 dB 12½ bits D 2½	79.8 dB 13 bits D 2½	81.6 dB 13¼ bits D 2½
3	31.1 dB 4 ¾ bits D 1	61.2 dB 9¾ bits D 2	89.1 dB 14½ bits D 3	104.1 dB 17 bits D 3½	107.9 dB 17½ bits D 3½	110.5 dB 18 bits D 3½
4	31.3 dB 5 bits D 1	61.8 dB 10 bits D 2	91.5 dB 15 bits D 3	119.2 dB 19½ bits D 4	134.2 dB 22 bits D 4½	138.2 dB 22¾ bits D 4½
5	31.4 dB 5bits D 1	62.2 dB 10 bits D 2	92.3 dB 15 bits D 3	121.7 dB 20 bits D 4	149.3 dB 24½ bits D 5	164.3 dB 27 bits D 5 ½
6	31.5 dB 5 bits D 1	62.4 dB 10 bit D 2	92.7 dB 15 bits D 3	122.56 dB 20 bits D 4	151.9 dB 25 bits D 5	179.4 dB 29½ bits D 6

Table 3: SNR/ENOB As Function Of L And K, OSR = 32

Also included in the table is the difference in increased resolution for a doubling of the OSR. The increase of resolution, when L ≤ K, is K bits. When K > L, the increase limits to L + ½ bits.

For an ADC with a decimator order one larger than the modulation order, the resolution, in bits, is:

$$ENOB = \frac{\log(OSR)}{\log(2)} \cdot \left(L + \tfrac{1}{2}\right) - \tfrac{1}{2} \tag{26}$$

Again, in no way does this say that you could slap together a third-order modulator, fourth-order decimator, and a 64 OSR and expect to get a real 20½ bits. There are many other noise sources that will reduce the overall SNR. It only tells you the quantization noise contribution to SNR.

Defining the cutoff frequency as the point where the signal is 3 dB down, the following equations define it as a function of the decimator output frequency and decimator order:

$$\left(\frac{\sin\left(\pi \frac{f_c}{f_d}\right)}{OSR \cdot \sin\left(\pi \frac{f_c}{f_d \cdot OSR}\right)}\right)^K \approx \left(1 - \frac{1}{6}\left(\pi \frac{f_c}{f_d}\right)^2\right)^K = \frac{1}{\sqrt{2}} \qquad f_c = \frac{f_d}{\pi} \cdot \sqrt{6 \cdot \left(1 - \left(\frac{1}{\sqrt{2}}\right)^{\frac{1}{K}}\right)} \tag{27}$$

For a third-order decimator with an output frequency of 10 kHz, the cutoff frequency is 2.58 kHz.

Knowing how a DSM alters the spectral density of noise and knowing your digital filters transfer function, it possible to calculate an ADC’s total quantization noise. With this valuation it is straightforward to determine its SNR and effective number of bits.