

**Taking the Mystery out of the Infamous Formula,
"SNR = 6.02N + 1.76dB," and Why You Should Care**

by Walt Kester

INTRODUCTION

You don't have to deal with ADCs or DACs for long before running across this often quoted formula for the theoretical signal-to-noise ratio (SNR) of a converter. Rather than blindly accepting it on face value, a fundamental knowledge of its origin is important, because the formula encompasses some subtleties which if not understood can lead to significant misinterpretation of both data sheet specifications and converter performance. Remember that this formula represents the theoretical performance of a perfect N-bit ADC. You can compare the actual ADC SNR with the theoretical SNR and get an idea of how the ADC stacks up.

This tutorial first derives the theoretical quantization noise of an N-bit analog-to-digital converter (ADC). Once the rms quantization noise voltage is known, the theoretical signal-to-noise ratio (SNR) is computed. The effects of oversampling on the SNR are also analyzed.

QUANTIZATION NOISE MODEL

The maximum error an ideal converter makes when digitizing a signal is $\pm\frac{1}{2}$ LSB as shown in the transfer function of an ideal N-bit ADC (Figure 1). The quantization error for any ac signal which spans more than a few LSBs can be approximated by an uncorrelated sawtooth waveform having a peak-to-peak amplitude of q, the weight of an LSB. Another way to view this approximation is that the actual quantization error is equally probable to occur at any point within the range $\pm\frac{1}{2}$ q. Although this analysis is not precise, it is accurate enough for most applications.




Figure 1: Ideal N-bit ADC Quantization Noise

W. R. Bennett of Bell Laboratories analyzed the actual spectrum of quantization noise in his classic 1948 paper (Reference 1). With the simplifying assumptions previously mentioned, his detailed mathematical analysis simplifies to that of Figure 1. Other significant papers and books on converter noise followed Bennett's classic publication (References 2-6).

The quantization error as a function of time is shown in more detail in Figure 2. Again, a simple sawtooth waveform provides a sufficiently accurate model for analysis. The equation of the sawtooth error is given by

$$e(t) = st, -q/2s < t < +q/2s. \quad \text{Eq. 1}$$

The mean-square value of $e(t)$ can be written:

$$\overline{e^2(t)} = \frac{s}{q} \int_{-q/2s}^{+q/2s} (st)^2 dt. \quad \text{Eq. 2}$$

Performing the simple integration and simplifying,

$$\overline{e^2(t)} = \frac{q^2}{12}. \quad \text{Eq. 3}$$

The root-mean-square quantization error is therefore

$$\text{rms quantization noise} = \sqrt{\overline{e^2(t)}} = \frac{q}{\sqrt{12}}. \quad \text{Eq. 4}$$




Figure 2: Quantization Noise as a Function of Time

The sawtooth error waveform produces harmonics which extend well past the Nyquist bandwidth of dc to $f_s/2$. However, all these higher order harmonics must fold (alias) back into the Nyquist bandwidth and sum together to produce an rms noise equal to $q/\sqrt{12}$.

As Bennett points out (Reference 1), the quantization noise is approximately Gaussian and spread more or less uniformly over the Nyquist bandwidth dc to $f_s/2$. The underlying assumption

here is that the quantization noise is uncorrelated to the input signal. Under certain conditions where the sampling clock and the signal are harmonically related, the quantization noise becomes correlated, and the energy is concentrated in the harmonics of the signal—however, the rms value remains approximately $q/\sqrt{12}$. The theoretical signal-to-noise ratio can now be calculated assuming a full-scale input sine wave:

$$\text{Input FS Sinewave} = v(t) = \frac{q2^N}{2} \sin(2\pi ft). \quad \text{Eq. 5}$$

The rms value of the input signal is therefore

$$\text{rms value of FS input} = \frac{q2^N}{2\sqrt{2}}. \quad \text{Eq. 6}$$

The rms signal-to-noise ratio for an ideal N-bit converter is therefore

$$\text{SNR} = 20\log_{10} \frac{\text{rms value of FS input}}{\text{rms value of quantization noise}} \quad \text{Eq. 7}$$

$$\text{SNR} = 20\log_{10} \left[\frac{q2^N / 2\sqrt{2}}{q / \sqrt{12}} \right] = 20\log_{10} 2^N + 20\log_{10} \sqrt{\frac{3}{2}} \quad \text{Eq. 8}$$

$$\text{SNR} = 6.02N + 1.76\text{dB}, \quad \text{over the dc to } f_s/2 \text{ bandwidth.} \quad \text{Eq. 9}$$

Bennett's paper shows that although the actual spectrum of the quantization noise is quite complex to analyze, the simplified analysis which leads to Eq. 9 is accurate enough for most purposes. However, it is important to emphasize again that the rms quantization noise is measured over the full Nyquist bandwidth, dc to $f_s/2$.

FREQUENCY SPECTRUM OF QUANTIZATION NOISE

In many applications, the actual signal of interest occupies a smaller bandwidth, BW, which is less than the Nyquist bandwidth (see Figure 3). If digital filtering is used to filter out noise components outside the bandwidth BW, then a correction factor (called *process gain*) must be included in the equation to account for the resulting increase in SNR as shown in Eq. 10.

$$\text{SNR} = 6.02N + 1.76 \text{ dB} + 10\log_{10} \frac{f_s}{2 \cdot \text{BW}}, \quad \text{over the bandwidth BW.} \quad \text{Eq. 10}$$

The process of sampling a signal at a rate which is greater than twice its bandwidth is referred to as *oversampling*. Oversampling in conjunction with quantization noise shaping and digital filtering are the key concepts in sigma-delta converters, although oversampling can be used with any ADC architecture.

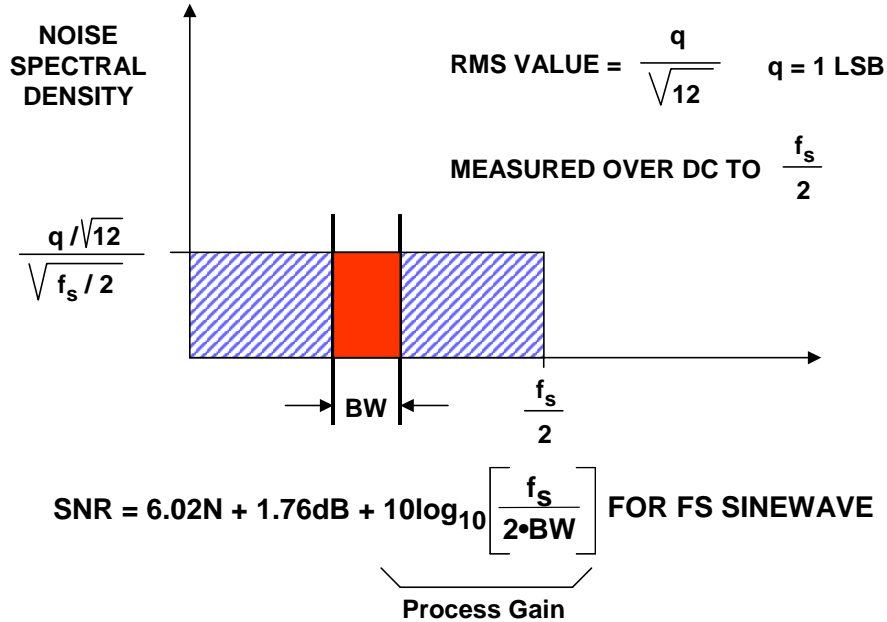


Figure 3: Quantization Noise Spectrum Showing Process Gain

The significance of process gain can be seen from the following example. In many digital basestations or other wideband receivers the signal bandwidth is composed of many individual channels, and a single ADC is used to digitize the entire bandwidth. For instance, the analog cellular radio system (AMPS) in the U.S. consists of 416 30-kHz wide channels, occupying a bandwidth of approximately 12.5 MHz. Assume a 65-MSPS sampling frequency, and that digital filtering is used to separate the individual 30-kHz channels. The process gain due to oversampling for these conditions is given by:

$$\text{Process Gain} = 10\log_{10} \frac{f_s}{2 \cdot \text{BW}} = 10\log_{10} \frac{65 \times 10^6}{2 \times 30 \times 10^3} = 30.3 \text{ dB}. \quad \text{Eq. 11}$$

The process gain is added to the ADC SNR specification to yield the SNR in the 30-kHz bandwidth. In the above example, if the ADC SNR specification is 65 dB (dc to $f_s/2$), then it is increased to 95.3 dB in the 30-kHz channel bandwidth (after appropriate digital filtering).

Figure 4 shows an application which combines oversampling and undersampling. The signal of interest has a bandwidth BW and is centered around a carrier frequency f_c . The sampling frequency can be much less than f_c and is chosen such that the signal of interest is centered in its Nyquist zone. Analog and digital filtering removes the noise outside the signal bandwidth of interest, and therefore results in process gain per Eq. 10.




Figure 4: Undersampling and Oversampling Combined Results in Process Gain

CORRELATION BETWEEN QUANTIZATION NOISE AND INPUT SIGNAL YIELDS MISLEADING RESULTS

Although the rms value of the noise is accurately approximated by $q/\sqrt{12}$, its frequency domain content may be highly correlated to the ac-input signal under certain conditions. For instance, there is greater correlation for low amplitude periodic signals than for large amplitude random signals. Quite often, the assumption is made that the theoretical quantization noise appears as white noise, spread uniformly over the Nyquist bandwidth dc to $f_s/2$. Unfortunately, this is not true in all cases. In the case of strong correlation, the quantization noise appears concentrated at the various harmonics of the input signal, just where you don't want them.

In most practical applications, the input to the ADC is a band of frequencies (always summed with some unavoidable system noise), so the quantization noise tends to be random. In spectral analysis applications (or in performing FFTs on ADCs using spectrally pure sinewaves as inputs, however, the correlation between the quantization noise and the signal depends upon the ratio of the sampling frequency to the input signal.

This is demonstrated in Figure 5, where the output of an ideal 12-bit ADC is analyzed using a 4096-point FFT. In the left-hand FFT plot (A), the ratio of the sampling frequency (80.000 MSPS) to the input frequency (2.000 MHz) was chosen to be exactly 40, and the worst harmonic is about 77 dB below the fundamental. The right hand diagram (B) shows the effects of slightly offsetting the input frequency to 2.111 MHz, showing a relatively random noise spectrum, where the SFDR is now about 93 dBc and is limited by the spikes in the noise floor of the FFT. In both cases, the rms value of all the noise components is approximately $q/\sqrt{12}$ (yielding a theoretical SNR of 74 dB) but in the first case, the noise is concentrated at harmonics of the fundamental because of the correlation.

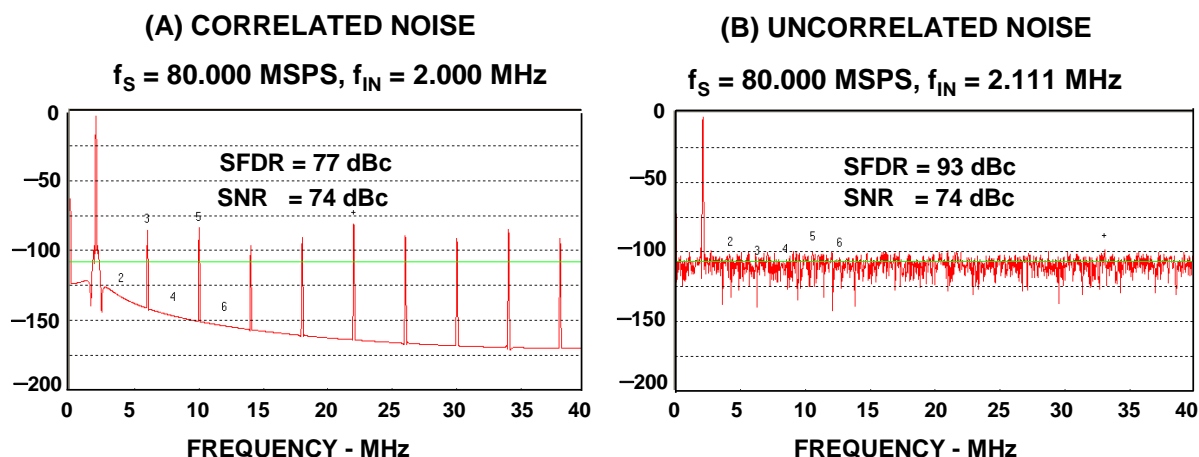


Figure 5: Effect of Ratio of Sampling Clock to Input Frequency on Quantization Noise Frequency Spectrum for Ideal 12-bit ADC, 4096-Point FFT. (A) Correlated Noise, (B) Uncorrelated Noise

Note that this variation in the apparent harmonic distortion of the ADC is an artifact of the sampling process caused by the correlation of the quantization error with the input frequency. In a practical ADC application, the quantization error generally appears as random noise because of the random nature of the wideband input signal and the additional fact that there is a usually a small amount of system noise which acts as a *dither* signal to further randomize the quantization error spectrum.

It is important to understand the above point, because single-tone sinewave FFT testing of ADCs is one of the universally accepted methods of performance evaluation. In order to accurately measure the harmonic distortion of an ADC, steps must be taken to ensure that the test setup truly measures the ADC distortion, not the artifacts due to quantization noise correlation. This is done by properly choosing the frequency ratio and sometimes by summing a small amount of noise (dither) with the input signal. The exact same precautions apply to measuring DAC distortion with an analog spectrum analyzer.

SNR, PROCESS GAIN, AND FFT NOISE FLOOR RELATIONSHIPS

Figure 6 shows the FFT output for an ideal 12-bit ADC. Note that the average value of the noise floor of the FFT is approximately 107 dB below full-scale, but the theoretical SNR of a 12-bit ADC is 74 dB. The FFT noise floor is *not* the SNR of the ADC, because the FFT acts like an analog spectrum analyzer with a bandwidth of f_s/M , where M is the number of points in the FFT. The theoretical FFT noise floor is therefore $10\log_{10}(M/2)$ dB below the quantization noise floor due to the *processing gain* of the FFT.

In the case of an ideal 12-bit ADC with an SNR of 74 dB, a 4096-point FFT would result in a processing gain of $10\log_{10}(4096/2) = 33$ dB, thereby resulting in an overall FFT noise floor of $74 + 33 = 107$ dBc. In fact, the FFT noise floor can be reduced even further by going to larger and larger FFTs; just as an analog spectrum analyzer's noise floor can be reduced by narrowing

the bandwidth. When testing ADCs using FFTs, it is therefore important to ensure that the FFT size is large enough so that the distortion products can be distinguished from the FFT noise floor itself. Averaging a number of FFTs does not further reduce the noise floor, it simply reduces the variations between the individual noise spectral component amplitudes.

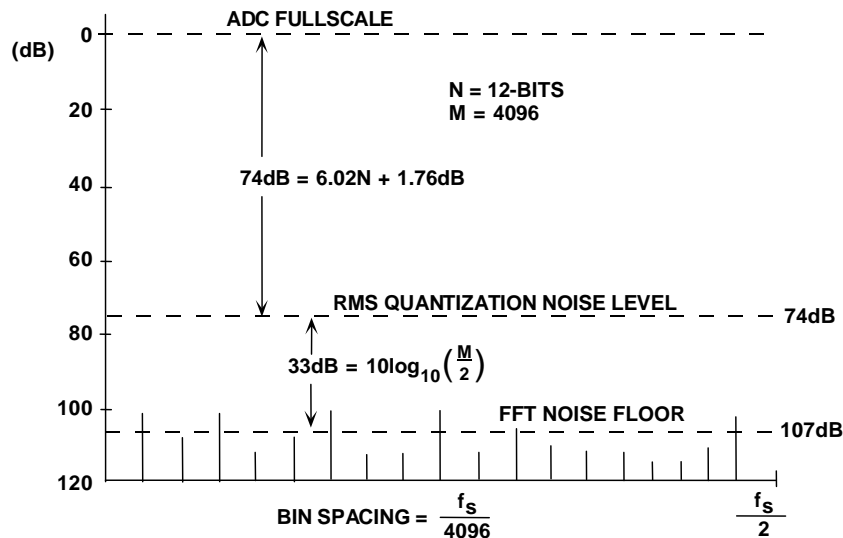


Figure 6: Noise Floor for an Ideal 12-bit ADC Using 4096-point FFT

REFERENCES

1. W. R. Bennett, "Spectra of Quantized Signals," *Bell System Technical Journal*, Vol. 27, July 1948, pp. 446-471.
2. B. M. Oliver, J. R. Pierce, and C. E. Shannon, "The Philosophy of PCM," *Proceedings IRE*, Vol. 36, November 1948, pp. 1324-1331.
3. W. R. Bennett, "Noise in PCM Systems," *Bell Labs Record*, Vol. 26, December 1948, pp. 495-499.
4. H. S. Black and J. O. Edson, "Pulse Code Modulation," *AIEE Transactions*, Vol. 66, 1947, pp. 895-899.
5. H. S. Black, "Pulse Code Modulation," *Bell Labs Record*, Vol. 25, July 1947, pp. 265-269.
6. K. W. Cattermole, *Principles of Pulse Code Modulation*, American Elsevier Publishing Company, Inc., 1969, New York NY, ISBN 444-19747-8.
7. Walt Kester, *Analog-Digital Conversion*, Analog Devices, 2004, ISBN 0-916550-27-3, Chapter 2. Also Available as *The Data Conversion Handbook*, Elsevier/Newnes, 2005, ISBN 0-7506-7841-0, Chapter 2.
8. Hank Zumbahlen, *Basic Linear Design*, Analog Devices, 2006, ISBN: 0-915550-28-1. Also available as *Linear Circuit Design Handbook*, Elsevier-Newnes, 2008, ISBN-10: 0750687037, ISBN-13: 978-0750687034. Chapter 5.

Copyright 2009, Analog Devices, Inc. All rights reserved. Analog Devices assumes no responsibility for customer product design or the use or application of customers' products or for any infringements of patents or rights of others which may result from Analog Devices assistance. All trademarks and logos are property of their respective holders. Information furnished by Analog Devices applications and development tools engineers is believed to be accurate and reliable, however no responsibility is assumed by Analog Devices regarding technical accuracy and topicality of the content provided in Analog Devices Tutorials.