

Improving A/D Converter Performance Using Dither

National Semiconductor
Application Note 804
Leon Melkonian
February 1992



1.0 INTRODUCTION

Many analog-to-digital converter applications require low distortion for a very wide dynamic range of signals. Unfortunately, the distortion caused by digitizing an analog signal increases as the signal amplitude decreases, and is especially severe when the signal amplitude is of the same order as the quantizing step. In digital audio applications, for example, low-level signals occur often, sometimes alone and sometimes in the presence of larger signals. If these low-level signals are severely distorted by the quantization process, the usefulness of the system is greatly diminished.

It is, in fact, possible to reduce the distortion, and also to improve the resolution below an LSB (least significant bit), by adding noise (dither) to the signal of interest. For ideal converters, the optimum dither is white noise at a voltage level of about $\frac{1}{2}$ LSB rms. The addition of dither effectively smoothes the ADC transfer function, which normally has a staircase-like appearance. The price one pays for the benefits of reduced distortion and improved resolution is a slight reduction of the signal-to-noise ratio. There are many applications for which the benefits of dither are well worth the cost. Applications where the spectral distortion caused by the quantization of low level signals is particularly undesirable will especially benefit from applying dither. These applications include audio, acoustical instrumentation, the analysis of rotating or vibrating machinery, and the characterization of the nonlinear distortion in electronic circuits.

This application note introduces the concept of dither, its relation to quantization noise, and how it can improve ADC performance. Additionally, experimental data are presented which explicitly show the effects of dither on a 10-bit converter.

2.0 SAMPLING AND QUANTIZATION

Two processes which are inherent to the digital conversion of analog signals are sampling and quantization. It is important to review these processes and examine their effects on ADC noise.

In analog sampling, a signal is "looked at" on only certain parts of its waveform. Although most of the waveform is "thrown away", all of the information contained in the signal can be retained if the conditions of the sampling theorem are met. The sampling theorem requires that for the perfect reproduction of a signal, it is necessary to sample at a rate greater than twice the highest frequency component in the signal. Frequencies which are greater than the Nyquist frequency (half the sampling rate) will be aliased to lower frequencies, resulting in signal degradation. To prevent aliasing, the input analog signal must be processed through a low pass (anti-aliasing) filter, which allows only frequencies below the Nyquist frequency to reach the input of the ADC. With ideal sampling, no noise will be added to the signal.

Unlike sampling, quantization inherently adds noise to the signal. In an ADC, a continuous range of voltages is converted to a discrete set of output codes (Figures 1 and 2). All voltages which fall within the same quantizing step are assigned to a single output code. Thus, quantization unavoidably results in a loss of information. Adding bits to the converter improves the resolution, but the drawback is increased cost and complexity.

The difference between an analog voltage and the voltage corresponding to its digital output code (for ideal converters) is known as the quantization error. If one identifies an output code with the voltage at the center of its range, the maximum quantization error is $\pm q/2$, where q is the size of the quantizing interval (Figure 3).




FIGURE 1. Quantization without Sampling




FIGURE 2. 3-Bit ADC Transfer Curve





FIGURE 3. Quantization Error is $\pm \frac{1}{2}$ LSB for a Perfect ADC



TL/H/11320-4

FIGURE 4. Treatment of Quantizing Error as an Explicit Quantization Noise

Alternatively, one can think of the quantization error in ADCs as a quantization noise which is added to a signal entering an ADC of infinite resolution (*Figure 4*). This is the way quantization error should be treated for the purposes of discussing dither.

The presence of quantization noise limits the signal-to-noise (S/N) ratio attainable by an ADC. For large amplitude, complex signals, the quantization error from sample to sample will be statistically independent and uniformly distributed over the quantizing interval. The probability density $P(v)$ for the quantization error v will be given by

$$P(v) = \begin{cases} \frac{1}{q}, & |v| < \frac{q}{2} \\ 0, & |v| > \frac{q}{2} \end{cases}$$

The rms noise v_{noise} is then

$$\begin{aligned} v_{\text{noise}} &= \left[\int_{-\infty}^{\infty} v^2 P(v) dv \right]^{1/2} \\ &= \left[\int_{-q/2}^{q/2} v^2 \frac{1}{q} dv \right]^{1/2} = \left[\frac{q^2}{12} \right]^{1/2} \end{aligned}$$

For an n -bit ADC, the largest sinusoidal signal that can be converted without clipping has a peak-to-peak amplitude of $2^n \cdot q$. This corresponds to an rms level of

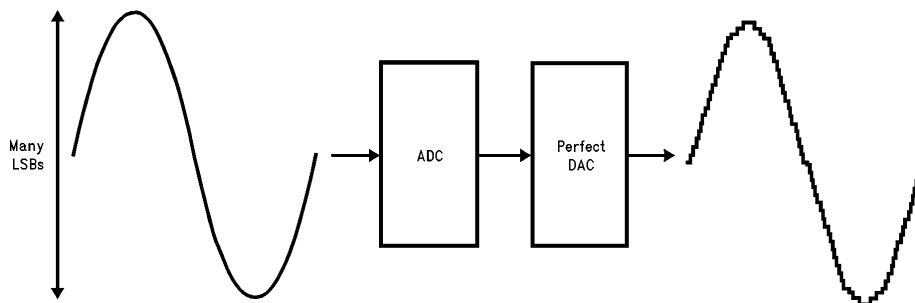
$$v_{\text{signal}} = \frac{2^{n-1}q}{\sqrt{2}}$$

The signal-to-noise ratio in decibels is then

$$\begin{aligned} \frac{S}{N} &= 20 \log \left\{ \frac{v_{\text{signal}}}{v_{\text{noise}}} \right\} = 20 \log \left\{ \frac{2^{n-1} q/\sqrt{2}}{q/\sqrt{12}} \right\} \\ &= 1.76 + 6.02 n \end{aligned}$$

Thus, a 10-bit converter will have a maximum signal-to-noise ratio (for sinusoidal signals) of 62 dB.

For large amplitude, complex signals, it can be shown that the quantization noise will be white (Appendix). This arises from the fact that for this type of signal, the quantization error for a large number of samples is uniformly distributed over the quantizing interval. For low level or simple (in terms of frequency content) signals, the quantization error cannot be treated as white noise added to the input signal. For example, let us look at a large amplitude, sinusoidal signal. (This is a simple signal, because it is composed of only one frequency.) If we digitize this with an ADC and then reconstruct it with a perfect DAC, we obtain what appears to be a sine wave with steps in it (*Figure 5*). These steps represent the various quantization levels in the ADC. If we were to do a spectral analysis of this reconstructed signal, we would see in addition to a peak at the input frequency, peaks at harmonics of the input frequency. The process of digitization has effectively added frequency components to the signal.



TL/H/11320-5

FIGURE 5. A/D Converters Inherently Distort Signals

Distortion is greatest in small amplitude, simple signals. For example the digitization of a sinusoidal signal with a peak-to-peak amplitude of 1 LSB will involve just two digital codes; the reconstructed waveform looks like a square wave (Figure 6). In this case, the quantization error is principally in the form of odd-order harmonic distortion.


Hence, for large, complex signals, the quantization noise is white. For small, simple signals, the quantization noise, because it is mainly in the form of harmonic distortion, effectively adds frequency components to the signal. Spectral analysis of such a signal would give erroneous results. A spectrum of a large, complex signal, on the other hand, would only show the input signal, with a flat noise floor.

3.0 DITHER

To ameliorate the negative effects of quantization, early workers added analog white noise to the ADC input signal. In 1951 Goodall¹ noticed that the addition of dither to signals masked the contour effects in video systems. In 1960 Widrow² determined that the signal loss due to quantization is minimized if the quantization error is independent of the signal. Schuchman³ determined the forms of dither signals which yield a quantization error which is independent of the signal. He found that for ideal converters, the optimum dither is $\frac{1}{3}$ LSB rms of white noise. Vanderkooy and Lipshitz⁴ showed that with dither, the resolution of an ADC can be improved to below an LSB.


We now show how it is possible to reduce harmonic distortion and improve resolution in the A/D conversion of signals by adding dither. In Figure 6 we show a low level sinusoidal signal centered on a quantization step. The peak-to-peak amplitude of this signal is 1 LSB. When this signal is input to an ADC, it will be represented by only two codes. If the sine wave is centered on the threshold between the two codes, the digital output will represent a square wave. Any offset will change the duty cycle, but the digitized signal will always take the form of a series of pulses at the same frequency as the input. This is obviously a very poor representation of a sine wave.

The addition of dither will cause the quantizer to toggle between the two (and possibly additional) states more frequently (Figure 7). Sub-LSB information is preserved in the percentage of time spent between levels. With time averaging the resolution can be increased significantly beyond an LSB. What has been accomplished by adding dither is an effective linearization of the ADC transfer curve. A power spectrum of the output would show that the harmonic distortion arising from the quantization process has been significantly reduced, as compared to the case with no applied dither. What one pays for a reduction in total harmonic distortion (THD) and improved resolution is a slightly degraded signal-to-noise ratio and, if one uses time averaging, an increase in the effective conversion time.



TL/H/11320-6

FIGURE 6. Quantization of Low Level Signals Results in Severe Error



TL/H/11320-7

FIGURE 7. With dither, sub-LSB information can be extracted from percentage of time spent between the two levels.

Another way of looking at dither is by looking at how it affects the ADC transfer curve. It fundamentally affects its meaning. With no dither, each analog input voltage is assigned one and only one code. Thus, there is no difference in the output for voltages located on the same "step" of the ADC's "staircase" transfer curve. With dither, each analog input voltage is assigned a probability distribution for being in one of several digital codes. Now, different voltages within the same "step" of the original ADC transfer function are assigned different probability distributions. Thus, one can see how the resolution of an ADC can be improved to below an LSB.

To illustrate the effects of dither on sinusoidal signals, a 10-bit A/D converter was connected as shown in Figure 8. The source of the dither is a pseudo-random noise generator. The dither must be attenuated so that at the ADC input, it has an rms level of $\frac{1}{3}$ LSB, which is given by:

$$V_{\text{dither}} = \frac{1}{3} \frac{V_{\text{REF}}}{2^n}$$

For a 10-bit converter with $V_{\text{REF}} = 5\text{V}$, we have $V_{\text{dither}} = 1.6\text{ mV}$. The signal and dither are filtered through an 8th order Chebyshev low pass filter with a 20 kHz corner frequency (Figure 9). The signal is also AC coupled to remove any DC offset introduced by the op amps. The controllable DC offset (Figure 8) gives the ADC a unipolar input and allows one to center the signal with respect to the quantization levels. The A/D converter is the ADC1061, which is a 10-bit, 1.8 μs conversion time multi-step converter which operates from a 5V supply.

For applications in the audio frequency range, one can use the MM5437 pseudo-random noise generator chip as the noise source (Figure 10). To convert from a pseudo-random digital waveform to white noise, one must low-pass filter the MM5437 output. This can conveniently be done through the anti-aliasing filter. Prior to filtering, the MM5437 output must be attenuated to a level which gives $\frac{1}{3}$ LSB of dither at the ADC input.

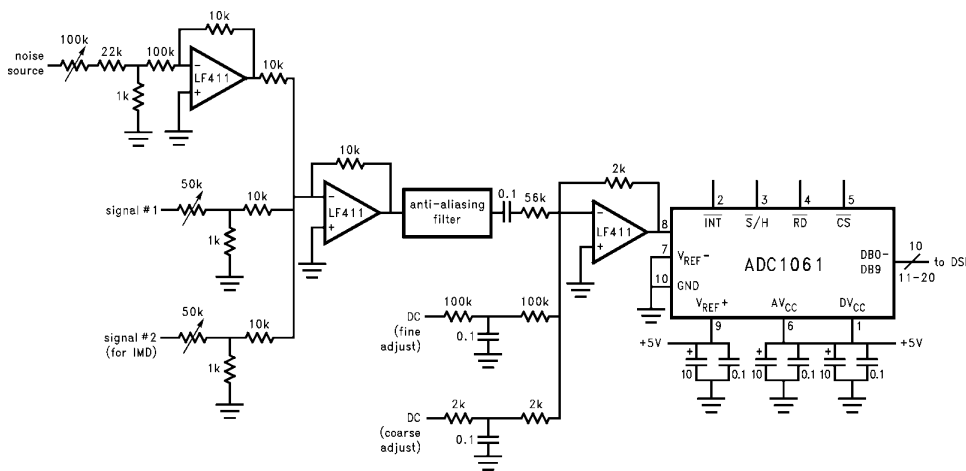


FIGURE 8. Setup for Investigating the Effects of Dither

TL/H/11320-8




FIGURE 9. 8th Order Chebyshev Anti-Aliasing Filter

TL/H/11320-9





FIGURE 10. The MM5437 Digital Noise Generator is a convenient source of dither for applications in the audio frequency range.

TL/H/11320-10

Figure 11a shows the spectrum, averaged 25 times, of a digitized 1 kHz signal which has a peak-to-peak amplitude of 1 LSB. Note the high relative amplitudes of the odd harmonics. With the addition of $\frac{1}{3}$ LSB rms of wide-band dither, the only harmonic to be seen is a small fraction of the third (Figure 11b). Notice that the noise floor has increased relative to the case in Figure 11a. Figure 12 shows the spectrum of a $\frac{1}{4}$ LSB signal with $\frac{1}{3}$ LSB dither applied. Without dither, this signal could not be detected at all. In Figure 13 are displayed spectra for larger signals, with a peak-to-peak amplitude of 5 LSB.

Dither is very effective in reducing harmonic distortion for signal levels up to about 10 LSB. Beyond that, one gets into the transition region between low level signals, where the quantization produces harmonic distortion, and large signals, where the quantization process produces white noise. Dither is also effective in ameliorating intermodulation distortion. Figure 14a shows the effects of quantization on 1 LSB amplitude signals with frequencies of 600 Hz and 1 kHz. When $\frac{1}{3}$ LSB of dither is applied, all the distortion products, both harmonic and intermodulation disappear.



(a) (b)
FIGURE 11. a) Spectrum of a 1 kHz, 1 LSB Signal. b) With dither, the harmonics have been eliminated. These spectra were averaged 25 times.






FIGURE 12. Without dither, this $\frac{1}{4}$ LSB, 1 kHz signal could not be detected at all. The spectrum was averaged 50 times.



(a) (b)
FIGURE 13. Effects of dither on larger signals (5 LSB). a) Without dither. b) With dither. These spectra were averaged 25 times.



(a) (b)
FIGURE 14. a) Two 1 LSB signals at 600 Hz and 1 kHz show both harmonic and intermodulation distortion. b) With dither, only the fundamentals remain.

As one goes to higher resolution A/D converters, the inherent analog noise becomes a significant fraction of the optimum dither. For example, for a 12-bit converter operating from a 5V reference, $\frac{1}{3}$ LSB corresponds to $400 \mu\text{V}$. If the inherent noise level is a few hundred microvolts but not close to $400 \mu\text{V}$, without any added dither the analog noise will cause a degraded signal-to-noise ratio without the benefits of a reduced harmonic distortion. Hence, for 12-bit or higher converters, such as the ADC12441 and ADC12451, dither is especially useful in improving performance.

4.0. ALTERNATIVE DITHERING SCHEMES

Before we end the discussion of dither, we should mention other dithering schemes. Depending upon the application and the degree of nonlinearity of the ADC, other types of dithering may be more appropriate than the white noise dither which we have discussed. The white noise dither is best for AC signals which cover a wide bandwidth, and for an ADC whose transfer function is close to ideal.

If the ADC has significant differential nonlinearity, a better type of dither is narrow band noise centered at the Nyquist frequency. An amplitude of several LSB, on the order of 4 to 5 LSB, will yield the best results. The amplitude of the dither needs to be higher because the spacing between codes will be higher due to the significant differential nonlinearity.

A variation of the white noise dither is the subtractive white noise dithering technique. In this technique, the digital output of the pseudo-random noise generator is subtracted from the digitized result of the ADC. This has the effect of reducing the noise penalty of dithering. References 5 and 6 discuss these other schemes in more detail.

5.0. CONCLUSIONS

This application note has attempted to introduce to the reader the concept of dither and how it affects A/D converter performance. The addition of $\frac{1}{3}$ LSB of wide-band noise to the input of an ADC has several beneficial effects. The dither will significantly reduce the harmonic and intermodulation distortion of small signals caused by the quantization process. With time averaging, resolution can be improved to levels significantly below an LSB. Fields such as audio where the spectral distortion caused by the quantization of low level signals is particularly undesirable will especially benefit from applying dither.

APPENDIX

We will determine the spectrum of quantization noise, assuming that the quantization error from sample to sample is statistically independent. This condition will be met for large amplitude, complex signals. The autocorrelation function of the noise, $R_{VV}(x)$, is defined by

$$R_{VV}(x) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{T/2} v(t) v(t+x) dt$$

Since the quantization error from sample to sample is independent, the autocorrelation function must be zero for $|x| > t_s$, where t_s is the time between samples. If we take into account the fact that the ADC maintains its output until the result of the next conversion is available, then it is legitimate

to speak of times which are between sample times. The autocorrelation function is then

$$R_{VV}(x) = \begin{cases} \frac{q^2}{12} \left(1 - \frac{|x|}{t_s}\right), & |x| < t_s \\ 0, & |x| > t_s \end{cases}$$

The value of $R_{VV}(x)$ at $x = 0$ corresponds to the mean square value of the quantization noise, $q^2/12$. The power spectrum of a signal is the Fourier transform of its autocorrelation function, hence for the power spectrum of the noise we get

$$\begin{aligned} P(\omega) &= \int_{-\infty}^{\infty} R_{VV}(x) e^{-j\omega x} dx \\ &= \frac{q^2}{12} \int_{-t_s}^{t_s} \left(1 - \frac{|x|}{t_s}\right) e^{-j\omega x} dx \\ &= \frac{q^2}{12} t_s \frac{\sin^2(\omega t_s/2)}{(\omega t_s/2)^2} \end{aligned}$$

The corresponding linear spectrum is

$$L(\omega) = \frac{q}{\sqrt{12}} \sqrt{t_s} \frac{\sin(\omega t_s/2)}{(\omega t_s/2)}$$

This spectrum actually represents the convolution of the quantization noise, $g(t)$, with a square pulse, $h(t)$, of width t_s and amplitude $1/t_s$.

$$\begin{aligned} L(\omega) &= F\{g(t) * h(t)\} = G(\omega) H(\omega), \\ \text{where } G(\omega) &= F\{g(t)\}, H(\omega) = F\{h(t)\} \end{aligned}$$

The square pulse arises from the fact that the ADC output is held fixed between samples. The Fourier transform of the square pulse is $[\sin(\omega t_s/2)]/(\omega t_s/2)$. Hence, the power spectrum of the quantization noise itself is

$$P_N(\omega) = \frac{q^2}{12} t_s$$

Since this is not frequency dependent, we have shown that the spectrum of quantization noise is white.

We can also determine the spectrum of the quantization noise by explicitly taking the sampled nature of an ADC into account. In this case, the output of the quantizer is defined only when the signal is sampled. The autocorrelation function of a discrete function is defined by a sequence of N terms, each of which is given by

$$\begin{aligned} R_{VV}(mt_s) &= \lim_{N \rightarrow \infty} \frac{1}{N+1} \sum_{n=-\frac{N}{2}}^{\frac{N}{2}} v(n t_s) v(n t_s + m t_s), \\ m &= -\frac{N}{2}, \dots, -1, 0, 1, \dots, \frac{N}{2} \end{aligned}$$

where N samples are taken spaced by an interval t_s . Since the quantization error from sample to sample is independent,

dent, the only nonzero term in the sequence will be the one for $m = 0$. This term represents the mean square value of the quantization noise, $q^2/12$. Hence, the autocorrelation function is the sequence

$$R_{VV}(m t_s) = \{ \dots 0, 0, q^2/12, 0, 0, \dots \}$$

The power spectrum of the quantization noise, which is the Fourier transform of its autocorrelation function, is given by the sequence of N terms, each term of which is given by

$$P(k\Omega) = \sum_{m=-\frac{N}{2}}^{\frac{N}{2}} R_{VV}(m t_s) e^{-j\Omega t_s k m},$$

$$k = -\frac{N}{2}, \dots, -1, 0, 1, \dots, \frac{N}{2}$$

where $\Omega = 2\pi/Nt_s$. Since the only nonzero value of $R_{VV}(m t_s)$ in the sequence is for $m = 0$, the equation above reduces to

$$P(k\Omega) = R_{VV}(0) = q^2/12$$

Since the expression for $P(k\Omega)$ is independent of k , every term in the power spectrum will be identical; therefore, the quantization noise will be white.

Hence, we have shown, using two different approaches, that the quantization noise for a large amplitude, complex signal will be white noise.

REFERENCES

1. W.M. Goodall, "Television by Pulse Code Modulation," *Bell Sys. Tech. J.*, vol. 30, pp. 33-49 (1951).
2. B. Widrow, "Statistical Analysis of Amplitude-Quantized Sampled-Data Systems," *Trans. AIEE (Applications and Industry)*, vol. 79, pp. 555-568 (1960).
3. L. Schuchman, "Dither Signals and Their Effect on Quantization Noise," *IEEE Trans. Commun. Technol.*, vol. COM-12, pp. 162-165 (1964).
4. 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.).
5. B. Blesser and B. Locanthi, "The Application of Narrow-Band Dither Operating at the Nyquist Frequency in Digital Systems to Provide Improved Signal-to-Noise Ratio over Conventional Dithering," *J. Audio Eng. Soc.*, vol. 35, pp. 446-454 (1987 June).
6. J. Vanderkooy and S.P. Lipshitz, "Dither in Digital Audio," *J. Audio Eng. Soc.*, vol. 35, pp. 966-974 (1987 Dec.).

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.



National Semiconductor Corporation
 2900 Semiconductor Drive
 P.O. Box 58090
 Santa Clara, CA 95052-8090
 Tel: 1(800) 272-9959
 TWX: (910) 339-9240

National Semiconductor GmbH
 Livny-Gargan-Str. 10
 D-82256 Fürstenfeldbruck
 Germany
 Tel: (81-41) 35-0
 Telex: 527849
 Fax: (81-41) 35-1

National Semiconductor Japan Ltd.
 Sumitomo Chemical
 Engineering Center
 Bldg. 7F
 1-7-1, Nakase, Mihama-Ku
 Chiba-City,
 Ciba Prefecture 261
 Tel: (043) 299-2300
 Fax: (043) 299-2500

National Semiconductor Hong Kong Ltd.
 13th Floor, Straight Block,
 Ocean Centre, 5 Canton Rd.
 Tsimshatsui, Kowloon
 Hong Kong
 Tel: (852) 2737-1600
 Fax: (852) 2736-9960

National Semicondutores Do Brazil Ltda.
 Rue Deputado Lacorda Franco
 120-3A
 Sao Paulo-SP
 Brazil 05418-000
 Tel: (55-11) 212-5066
 Telex: 391-1131931 NSBR BR
 Fax: (55-11) 212-1181

National Semiconductor (Australia) Pty, Ltd.
 Building 16
 Business Park Drive
 Monash Business Park
 Nottingham, Melbourne
 Victoria 3168 Australia
 Tel: (3) 558-9999
 Fax: (3) 558-9998

National does not assume any responsibility for use of any circuitry described, no circuit patent licenses are implied and National reserves the right at any time without notice to change said circuitry and specifications.