# Resolution Below the Least Significant Bit in Digital Systems with Dither*

**JOHN VANDERKOOY AND STANLEY P. LIPSHITZ****

*University of Waterloo, Waterloo, Ontario N2L 3G1, Canada*

The work was motivated by some common misunderstandings about digital systems. It is commonly believed that small signals or signal details are lost if they are smaller than the quantizing step. Expanding on previous arguments, it is shown that this is not true when the signal to be quantized contains a wide-band noise dither with an amplitude of approximately the step size. The introduction traces the use of dither from video quantization through to its use in audio. Quantization error is studied in some detail and the effect of dither is analyzed theoretically and experimentally. By examples of quantized signals it is shown that the dither effectively turns signal distortion into low-level wide-band noise by linearizing the averaged quantizer staircase function, which is as perceived by the ear.

## 0 HISTORICAL BACKGROUND

The basic aims of digitization in an analog system are 1) to increase the quality of the recorded signal and 2) to prevent degradation of the signal during the record–replay process. Digitization is usually implemented by sampling the signal repetitively at an adequate rate, and quantizing the resulting analog signal into finite steps. Although the sampling of the signal is a radical process, when properly applied there is theoretically no degradation for a band-limited signal, whereas the quantization process must cause a reduction in signal quality. Utilizing smaller quantization intervals will reduce the loss in quality, but cost and/or technological considerations place limits on how small the quantization intervals can be made.

Available digital audio systems use 14- or 16-bit binary numbers to represent the signal levels in each sample, using a simple linear scheme. There has been a great deal of criticism leveled at digital recordings (and much praise as well) relating to either the sampling or the quantization processes. Typical criticisms relate to concern over the fact that time sampling ignores the signal waveform except at the sampling points, and hence that details will be lost or degraded. It is now recognized by most of the popular audio press that sampling does not cause a loss of information on band-limited signals as long as the sampling frequency is more than twice the bandwidth of the input signal. Quantization is a different matter. The effect creates distortions of the signal which may be perceived as harmonic or intermodulation distortion, wideband noise, or even more gross misrepresentations of the signal. In addition, the quantizer or analog-to-digital converter (ADC) often has a measure of nonuniformity or even nonmonotonicity in its input–output relationship. The digital signal is a set of numbers, and a digital-to-analog converter (DAC) is needed to rederive the audio signal. In this paper nonuniformity or nonmonotonicity is a property of the ADC–DAC pair being used. For convenience we prefer to think of the DAC as perfect, and lump all irregularities into the ADC. This is roughly true in practice with actual digital circuits. We should point out that in a real digital system there are other degradations that may depend on inaccuracies in the sampler or the ADC, or even those that depend on their joint inaccuracies. In this paper we ignore all such features and idealize the ADC to a perfect quantizer with finite quantization steps. Specifically we focus on the effect of adding to the audio signal a dither signal which is intended to ameliorate the effects of the ADC quantization steps. An early paper by Bennett [1] dis-

* Presented at the 72nd Convention of the Audio Engineering Society, Anaheim, California, 1982 October 23–27; revised 1983 December 7.

** Authors are members of the Guelph-Waterloo Program for Graduate Work in Physics (GWP)$^2$.

cussed the spectra of quantized signals, giving relationships for the quantization noise, and suggesting that even for weak signals there is usually enough residual noise in a system to whiten the quantizing noise and mask the relation between it and the signal. This represents a kind of unintentional but beneficial dither.

Dither signals were first used to reduce the effects of quantization in pulse-code modulation (PCM) video systems. In 1951 Goodall [2] showed that in reproducing television pictures using high-speed ADCs, the contour effects created by the quantized steps in intensity are easily visible in a 5-bit representation having 32 levels of intensity. Goodall noticed that the contours could be masked by an rms random noise level 40 dB below the peak-to-peak input signal level. Although the picture then was more "noisy," most observers agreed that it was preferable for a system with a small number of bits.

The use of noise in masking contour effects was further studied by Roberts [3]. He states that while with normal quantization it takes 6 or 7 bits per sample to obtain good pictures, when pseudorandom noise is added, it takes only 3 or 4 bits for acceptable pictures. There are several new ideas found in this work. First there is the concept of adding noise to the ADC input signal, and subtracting the same noise after the reconversion to an analog signal. Second there is an explicit statement that such a process breaks up the regular steps of digital coding into something resembling additive noise, and the channel becomes very similar to an analog channel. It is apparent that the addition–subtraction scheme generates an average output which is just the mean value of the input signal, while not increasing the rms level of the noise at all. With dither the quantization error, which is signal dependent, is turned into wide-band noise which is uncorrelated with the signal.

By the early 1960s the use of the term dither was widespread. It meant a second signal added to the quantizer input and then subtracted after the quantizing operation. Schuchman [4] studied the effect of dither on quantization noise. In general, quantization results in a signal which is the sum of two terms, the information signal and a noise that is a function of the information signal. He gives conditions for the dither signal which make the quantizer noise statistically independent of the information signal, which had been shown by Widrow [5] to represent the minimum loss of statistical data due to quantizing the input signal. The optimum dither, as given by Schuchman, is noise having a uniform probability density function the width of which is the quantizing step, with each sample statistically independent of the others.

Another interesting aspect of quantization noise was discussed by Spang and Schultheiss [6], who showed that by feedback around the ADC the noise spectrum could be shaped to reduce the noise in any given frequency region, although this did increase the total noise.

Jayant and Rabiner [7] analyzed the application of dither to the quantization of speech signals, and concluded that dither is very significant and beneficial if the number of bits per sample is less than 5 or 6. They also showed that for effective dithering using pseudorandom noise, the step size of the noise did not need to be less than about a fourth of the quantizer step size.

Blesser investigated quantization noise without mention of dither in an earlier paper [8], but in his comprehensive examination [9] of digital audio he clearly indicates the beneficial effects of dither and points out that the average value of the quantized signal can move continuously between two levels. It seems clear from his descriptions that for digital audio the concept of dither does not necessarily imply the addition–subtraction scheme of Roberts [3], but simply a noise added to the ADC input to eliminate digital artifacts.

We feel that the audio community in general does not yet understand the nature of the quantization error in digital systems, and in particular the beneficial effects of adding an appropriate amount of dither. We shall show that the dither really does remove the "digital" aspects of quantization error, leaving an equivalent analog signal with high resolution and some benign wide-band noise.

## 1 QUANTIZATION ERROR

It is important to emphasize the fact that time sampling does not generate noise. Consider an input signal as shown in Fig. 1(a), which is sampled to produce Fig. 1(b). When this sampled output is processed by a linear-phase low-pass filter which allows only the original band-limited range of signals to pass, and assuming that the sampling frequency exceeds the Nyquist rate, then a smooth signal is again obtained which is precisely the same as the original input signal, except for possible time delay, with no added noise. In a sampled audio system (having no quantization error), the signal is

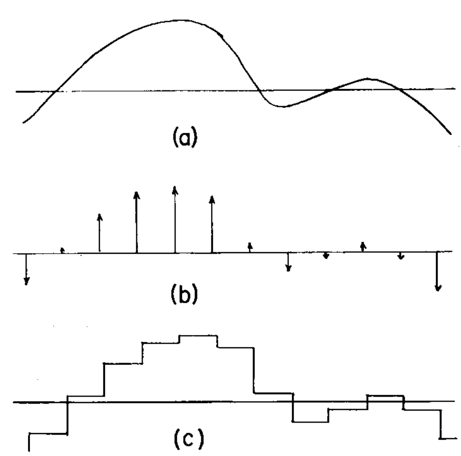


Fig. 1. Effects of time sampling alone, with no quantization error. The signal in (a) is sampled at regular intervals to produce the sampling pulses in (b). Such a signal is difficult to process, and usually the samples are held between sampling points (c). Although (c) differs from (a) by an error signal, this error does not produce noise, but only a small change in baseband frequency response.

sampled at regular intervals, but the reconstructed output is held at the particular value until the next sample. This sample-and-hold procedure is shown in Fig. 1(c). The difference between Fig. 1(c) and (a) is often misinterpreted as representing the presence of noise due to "error." But the sample-and-hold output of Fig. 1(c) is simply a convolution of the sampled output of Fig. 1(b) by a square pulse of width one sampling interval in time. This will give rise to a slight frequency response error (called aperture error) but not noise. As long as the quantization is infinitely fine grained, no noise will be introduced by time sampling.

If there is a finite quantization error, as in real systems, then each flat sample-and-hold interval in Fig. 1(c) will be in error because of the coarseness of the amplitude quantization. Fig. 2(a) shows a portion of a signal segment, with vertical lines representing the instants of sampling. Along the vertical axis the possible quantized levels are shown. At the sampling points the signal will in general not fall precisely at quantized levels. The convention we adopt is that the level chosen by the ADC will be the one that is closest to the signal at the time of sampling. This means that the error between the actual signal and its ADC representation as a finite-bit-length number extends over a range of ±½ of a quantizing interval. The error at each time sampling point will be held until the ADC makes its "guess" at the next sampling point. Fig. 2(a) also shows the sampled-and-held signal that would be made by a quantizer of infinite precision [which looks like Fig. 1(c)] and the actual sequence of quantized levels as produced by a real ADC. The area between the two is crosshatched. We note that the former contains no noise error by the arguments given for Fig. 1. However, the actual sampled and quantized signal differs from the former, and the resultant crosshatched error is shown

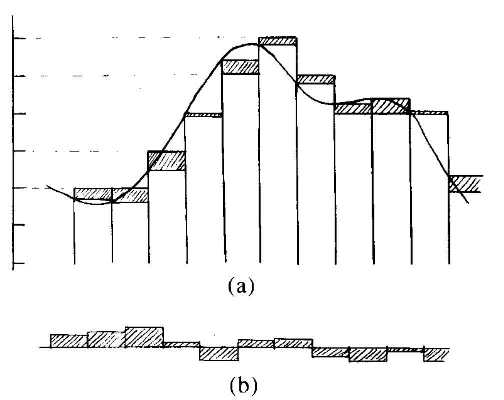

Fig. 2. Effects of time sampling using finite amplitude quantization steps. The scale on the left in (a) shows the possible ADC steps, and the time sampling intervals are represented by vertical lines. The signal segment is shown represented by two sets of horizontal lines between sampling points, one being the actual ADC chosen level, and the other the level of the signal at the sampling instant. The crosshatched areas represent the quantization error (b). At each sampling instant the ADC chooses the quantization level closest to the signal waveform, thus the error in (b) ranges over ±½ a quantizing interval. For high signal levels, adjacent error samples are uncorrelated.

in Fig. 2(b). It is this error that represents the quantization noise. As the quantization becomes more precise (more bits), this error will decrease.

The error waveform can be thought of as a *real* signal, not band-limited, which is added to the input of an *infinite* precision quantizer along with the signal. For simple input signals the related error will produce anharmonic signals in the baseband, equivalent to aliasing. On a 14-bit system such effects are audible with sinusoidal input signals. For somewhat more complex signals the error spectrum can fluctuate, resulting in a modulation of the perceived noise.

At adequately high levels with complex signals, the error from sample to sample will be statistically independent and uniformly distributed over a range $-\Delta/2$ to $+\Delta/2$, where $\Delta$ is the quantizing interval. Thus the probability density $p(v)$ of the error signal is given by

$$p(v) = \begin{cases} \dfrac{1}{\Delta}, & -\dfrac{\Delta}{2} < v < \dfrac{\Delta}{2} \\ 0, & |v| > \dfrac{\Delta}{2} \end{cases} . \tag{1}$$

The rms noise $v_n$ can now be computed[1] as

$$v_n = \left[ \int_{-\infty}^{\infty} v^2 p(v)\, dv \right]^{1/2} = \left[ \int_{-\Delta/2}^{\Delta/2} v^2 \frac{1}{\Delta}\, dv \right]^{1/2}$$

$$= \left[ \frac{\Delta^2}{12} \right]^{1/2} = \frac{\Delta}{2\sqrt{3}} . \tag{2}$$

If the digital system has $B$ bits, one being a sign bit, the peak signal levels are $\pm 2^{B-1}\Delta$, and the largest rms sinusoidal signal level that can be represented without clipping is

$$V_{\mathrm{rms(max)}} = \frac{2^{B-1}\Delta}{\sqrt{2}} . \tag{3}$$

Thus the signal-to-noise power ratio in decibels is

$$\frac{S}{N} = 10 \log \left\{ \frac{V^2_{\mathrm{rms(max)}}}{\Delta^2/12} \right\} = 20 \log\{\sqrt{3/2}\, 2^B\}$$

$$= 1.76 + 6.02B . \tag{4}$$

A perfect 16-bit system has 98 dB signal-to-noise ratio.

What is the quantization noise spectrum, assuming that the signals are complex and that the error from sample to sample is statistically random? To answer this, we compute first the autocorrelation function of the noise, defined by

$$C(\tau) = \lim_{T\to\infty} \frac{1}{T} \int_{-T/2}^{T/2} v(t)v(t + \tau)\, dt . \tag{5}$$

[1] Not all the error falls inside the signal band, but we shall ignore this point in evaluating the signal-to-noise ratio.

Since the noise is assumed random in different samples spaced by the time sampling interval $T_s$, the autocorrelation function will only be nonzero for $-T_s < \tau < T_s$. For $\tau = 0$, $C(\tau)$ measures the mean square value of $v(t)$, which by Eq. (2) is $\Delta^2/12$. Fig. 3(a) shows the autocorrelation function, described by

$$C(\tau) = \begin{cases} \dfrac{\Delta^2}{12}\left(1 - \dfrac{|\tau|}{T_s}\right), & |\tau| < T_s \\ 0, & |\tau| > T_s \end{cases} \quad . \tag{6}$$

The power spectrum of a signal is the Fourier transform of the autocorrelation function, and hence the double-sided noise power spectrum $P(\omega)$ of the quantization error shown in Fig. 3(b) is

$$\begin{aligned} P(\omega) &= \int_{-\infty}^{\infty} C(\tau) e^{-j\omega\tau}\, d\tau \\ &= \frac{\Delta^2}{12} \int_{-T_s}^{T_s} \left(1 - \frac{|\tau|}{T_s}\right) e^{-j\omega\tau}\, d\tau \\ &= \frac{\Delta^2}{12}\, T_s \frac{\sin^2(\omega T_s/2)}{(\omega T_s/2)^2} \quad . \end{aligned} \tag{7}$$

Thus it may appear that the noise spectrum falls slightly at the band edge. At the Nyquist frequency $\omega = \pi/T_s$ the noise power is about 4 dB down relative to low frequencies. However, the digital signal is held between samples, and thus the signal consists of samples convolved by a square pulse of width $T_s$ and amplitude $1/T_s$ to give unit weighting. This function is shown in Fig. 3(c). Such a convolution of the signal causes the audio spectrum to be multiplied by the Fourier transform of the square pulse, given by $[\sin(\omega T_s/2)]/(\omega T_s/2)$, which is plotted in Fig. 3(d). The audio power spectrum is thus multiplied by the square of the above, and has the same form as the noise power spectrum given by Eq. (7). Since we wish the frequency response to be flat, an equalizer is needed to restore the audio response to flatness, and in so doing the noise spectrum is also made flat. Hence quantization error has a white noise spectrum as long as our assumptions are true regarding the statistical independence of the errors in each time sample.

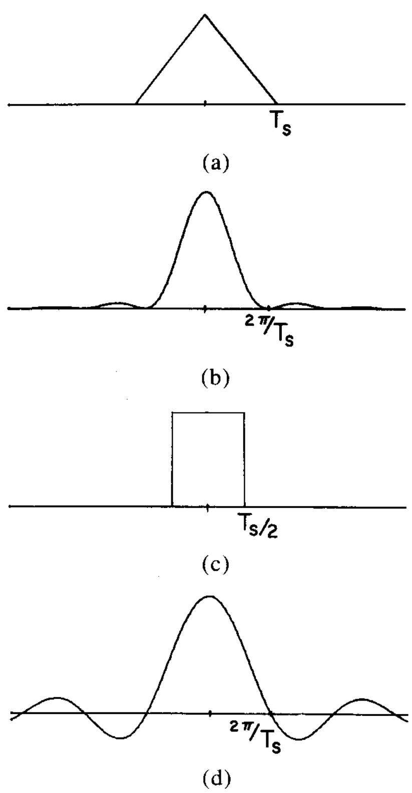


Fig. 3. Quantization error often results in white noise. (a) Autocorrelation function of the quantization error of Fig. 2(b), which results in the noise power spectrum (b). The signal shown in Fig. 2(a) is convolved by a square pulse (c) by reason of the sample-and-hold action. This convolution results in the frequency response (d) whose square is the same as the curve in (b). Hence equalization to correct the frequency response also flattens the noise spectrum.

In experiments with a 10-bit ADC we have shown that when only the least significant 2 or 3 bits are connected to the DAC, the sound produced is often similar to white noise for complex high-level signals. However, speech with periods of low levels and music with single pure tones do not produce a white-noise output. The error signal can be gritty or granular with noise modulation. For these reasons it is important to remove the structure of the quantization error so that it is just like white noise. Using more bits per sample helps, of course, but even 16-bit systems will retain such features to some degree. By adding an appropriate dither signal, all structure in the output noise can be removed.

## 2 EFFECT OF DITHER

To elucidate the effect of quantizing a signal with added dither we analyze several situations with and without dither. Consider a low-level low-frequency sine wave applied to the quantizer as shown in Fig. 4. In Fig. 4(a) the sine wave is centered exactly at a quantization step, and the output is a square wave of amplitude $\Delta$, the quantizing interval. In Fig. 4(b) the sine wave is registered between two steps, and the output is the dc level between these two steps. Such distortion of a low-level signal is severe.

Now consider the same signal with an added wideband noise having peak-to-peak value somewhat less than two quantizing intervals. Fig. 5 shows the same two situations as before but with an added dither signal. Although there are differences in the two outputs, it is clear that the noise in each case is similar, and that there is a duty-cycle modulation which with short-term averaging would recover a relatively undistorted sine wave in each case. The noise still has a nearly white spectrum, being similar to random binary noise, so audibly the result is like a sine wave with added noise.

Fig. 6 shows the effect of adding dither to a low-level 1-kHz sinusoidal signal with a peak-to-peak excursion of about 1 least significant bit (LSB). The upper trace shows the DAC output waveform with no added dither into the ADC. Note that the situation is similar to that indicated in Fig. 4(a). The second trace in Fig.

6 shows the output when a Gaussian dither signal is added with an rms value of 1 LSB. Occasionally the output spikes above and below the two levels shown, but for neatness we have chosen a signal segment in which this does not happen. The third trace shows the result of averaging 32 traces such as the second one, and the fourth and final trace shows the result when 960 traces are time averaged. To obtain these graphs we switched the ADC to produce only a few levels, so that real analog noise was not a problem. In each case the signal averager was triggered from the oscillator used to produce the 1-LSB peak-to-peak sinusoidal signal. Note how the averaging restores a virtually undistorted sine wave even though the signal from the ADC–DAC was basically a binary waveform. Audibly a signal such as the second trace sounds like a sine wave with a fair amount of added noise, whereas the undithered first trace has the harsh sound of a square wave.

In simple experiments it is relatively easy to hear 2-kHz sine waves when they have levels about 10 dB below accompanying wide-band noise. Suppose that the critical bands of the human ear are about one-third-octave wide as regards noise bandwidth. Then we might surmise that if a sinusoidal signal has a larger rms value than the noise in the appropriate critical band, it will be audible. A one-third-octave filter centered at 2 kHz (a frequency near most acute hearing) has a noise bandwidth of about 500 Hz. If wide-band noise from low frequencies to 20 kHz is used as a representation of quantization noise, then the rms value of the noise falling in a 500-Hz band will be smaller by a factor $\sqrt{500/20\,000} = 0.158$, or 16 dB. Thus we might expect a spectrally pure signal such as a sine wave to be audible even though it is 16 dB below the wide-band noise. The example of Fig. 6 shows that dither is very effective even though it is comparable in level to the quantization error. Although the dither noise has a digital *visual* character, it sounds and acts just like ordinary analog noise. Hence there has been only a little addition to the noise, in exchange for the ability to resolve signals having levels well below the quantizing level with low distortion. An alternative way of looking at the ear's ability to resolve narrow-band signals below the noise is to consider the averaging properties of the basilar membrane filter. Modeled as a one-third-octave filter, it has a $Q$ of about 4, and its impulse response would be extended over at least that many oscillations. The quantization error, now given a white noise character by the dither, tends to average out, while narrow-band signals are heard without distortion.

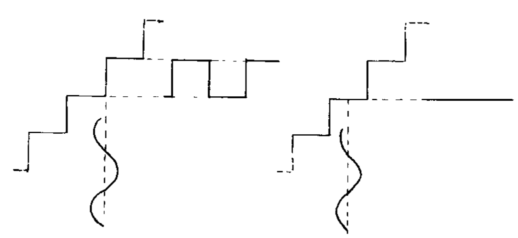

Fig. 4. Effect of the ADC steps on small signals with no applied dither. (a) The sine-wave input signal gives a large clipped output. (b) No output.

Fig. 5. Effects of dither on the situations of Fig. 4. The dither noise causes the ADC to make transitions, and the sine wave is not lost, being preserved in the duty-cycle modulation created by the added dither noise.

To analyze the effect of dither more precisely, we use the probability density of the noise $p(v)$ to compute the average output of the quantizer with input voltage $V_{in}$. Suppose that the quantizer produces a staircase output $q$ given by

$$q = g(V_{in}) \quad . \tag{8}$$

Then the average of $q$ is given by

$$\bar{q} = \int_{-\infty}^{\infty} p(v) g(V_{in} + v) \, dv \tag{9}$$

in which all possible dither noise values, properly weighted, are added to the input voltage $V_{in}$. We note that the noise voltage basically smears the quantizer input–output relation, Eq. (8); $\bar{q}$ is essentially $q$ convolved with the noise probability density.

Fig. 7 shows the effect of adding various amounts of wide-band Gaussian dither to a quantizer input. The curve is produced by applying a low-frequency ramp signal to a coarse quantizer. The quantizer consists of a 10-bit ADC in which only the three most significant bits are used, giving only eight levels with a quantizing interval of about 0.35 V. The start of the ramp triggers

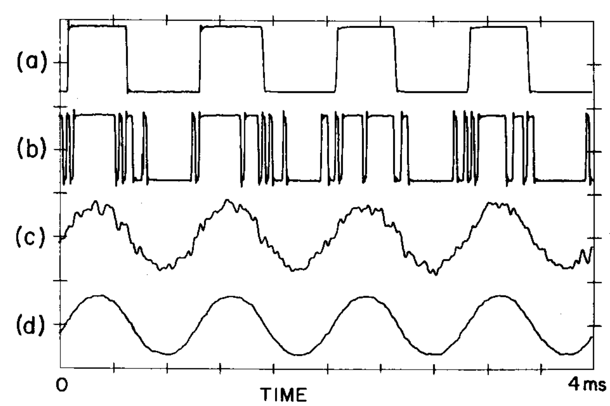


Fig. 6. Effects of dither when a signal is averaged. (a) ADC/DAC output when a 1-kHz sine wave is applied to the ADC with an amplitude of about ½ LSB. (b) Wide-band dither noise of about 1 LSB rms has been added to the sine wave, and the ADC rapidly toggles between two quantization levels, especially when the input signal swings through zero. (c), (d) Results of signal averaging respectively 32 and 960 traces such as shown in (b). Note that with adequate averaging, a virtually undistorted input signal can be recovered.

a signal averager and the quantizer output is thus repetitively averaged, while an appropriate wide-band Gaussian dither signal is added to the quantizer input as well. Note that with no dither added, the quantizer output $\bar{q}$ is related to the input voltage $V_{in}$ like a staircase function. As the dither is increased, the staircase edges are rounded, as would be predicted by Eq. (9), and the slope of the flat portions becomes significantly greater than zero when the rms dither is about one-half of the quantizing step size. When the rms dither is about equal to the quantizing interval, the staircase has been smeared to nearly a straight line. With this amount of applied Gaussian dither, the audible distortion of low-level signals will be reasonably small. For this amount of dither the adjacent sample errors are essentially statistically independent.

To illustrate the reduction in distortion of signals when dither is applied, the spectra of the signals can be measured. The upper trace of Fig. 8 shows the spectrum of the signal appearing in the upper trace of Fig. 6. The harmonics of the 1-kHz fundamental are all significant, since the signal is basically a square wave with somewhat less than 50% duty cycle.[2] The lower trace of Fig. 8 shows 16 averaged output spectra when Gaussian dither of 1 LSB rms is applied, corresponding to the signal in the second trace in Fig. 6. Note that the only harmonic remaining is the third. The fundamental amplitude was actually reduced with dither applied, since the ADC output no longer is simply clipping the sine-wave input. By adding the dither, distortion (harmonic in this case) has been turned into low-level white noise.

Fig. 9 shows a similar reduction in intermodulation distortion when dither is applied. There are two overlaid spectrum traces. The two largest peaks represent the two sinusoidal signals of approximately 600 Hz and 1 kHz, each of amplitude approximately 1 LSB. Without dither there is a myriad of sum and difference tones of all orders. When about 1 LSB of wide-band dither is added, all the intermodulation products disappear to be replaced by low-level wide-band noise, which is a far preferable situation.

If dither noise with a rectangular probability density is used, then when its peak-to-peak amplitude is exactly equal to a quantization interval, the staircase relationhsip between $\bar{q}$ and $V_{in}$ becomes perfectly linear. There is a small residual noise modulation which does not exceed ½ LSB. The resulting output noise has a triangular probability density function, but since adjacent samples are not statistically correlated, it will still appear audibly as white noise. Due to the complexity of producing this form of noise at precisely the right level, dither noise would normally be chosen to be Gaussian noise. In a practical audio system the dither noise may not be specifically incorporated, but inherent analog noise in the electronics or in the signal source itself can act as dither.

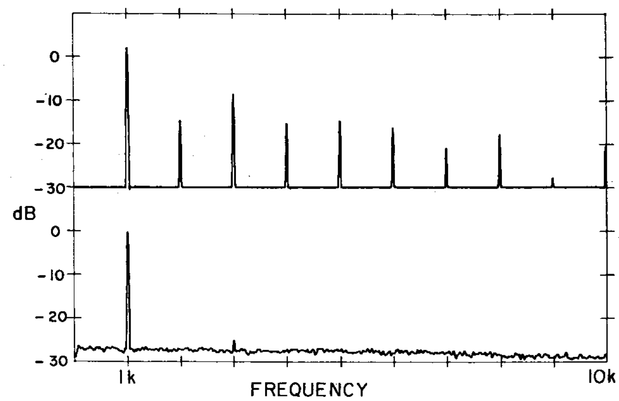


Fig. 8. Dither reduces harmonic distortion in a digital system. The upper trace shows the spectrum of the signal shown in Fig. 6(a), consisting of a 1-kHz sine wave with an amplitude of about ½ LSB. Note all the harmonics produced by the staircase action of the ADC. The lower trace shows the resultant spectrum when about 1 LSB rms of dither is applied. Only a small amount of third harmonic remains, along with wide-band noise. Although the wide-band noise has about the same total power as the signal, the high $Q$ filtering action of the basilar membrane allows the ear to clearly resolve sine waves well below the noise, in much the same way that the spectrum analyzer can.

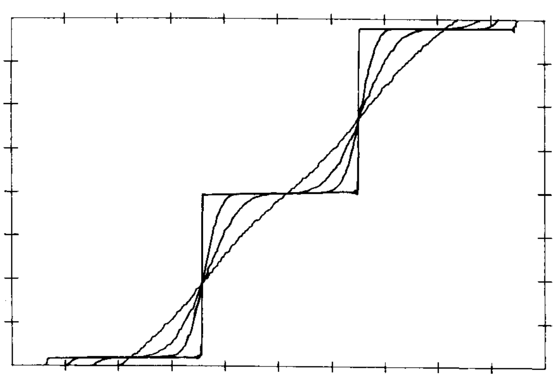

Fig. 7. Effects of dither on the time-averaged input–output characteristic of a digital system. When no dither is applied, the sharp staircase results. As Gaussian dither is applied equivalent to ¼, ½, and 1 LSB rms, the staircase smears until finally an almost linear relationship is obtained.

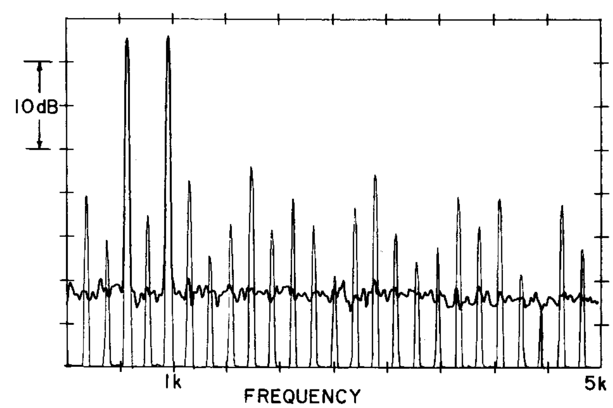


Fig. 9. Intermodulation distortion reduction by the addition of dither. Signals of approximately 600 Hz and 1 kHz, each having a level of about 1 LSB, produce numerous intermodulation products. Overlapped with the undithered spectrum is one resulting from the same two signals with 1 LSB of dither. All intermodulation products have been converted to wide-band noise.

[2] Note that for this input signal the quantization error is principally harmonic distortion rather than broad-band noise.

An insufficient amount of dither in an audio system will give rise to various deleterious effects. Low-level signals will be severely distorted or even altogether lost. The inherent noise of the system will vary depending on whether or not the input voltage is straddling a converter transition. Low-frequency low-level signals or dc drifts can create modulation effects in the perceived noise. For all of these reasons it appears wise to ensure that an adequate dither exists, the only penalty being a slight increase in the noise, which is likely to be less audible than the effects outlined above.

We have made measurements on several EIAJ standard 14-bit digital recorders to obtain data similar to those of Fig. 7, and we have found that there was insufficient dither to produce a fairly straight input–output relationship. The measurements show that the available wide-band noise acts as a partial dither with an rms value of about ½ LSB. Thus there is the possibility of some of the effects mentioned above. Perhaps the manufacturers feel that adding adequate dither does not leave the digital machine with significant noise advantage over an analog machine. The EIAJ machines also use 10 dB of pre- and deemphasis for the high frequencies to reduce the noise further. Even so we find that the output noise spectrum does not fall at high frequencies, which may be due to other irregularities in the ADC or to analog feedthrough from the digital circuits. On very-low-level sinusoidal input signals the various predicted quantization distortions and noise modulation effects can indeed be heard on these machines, although their audibility on program is subject to doubt. A switchable 14/16-bit consumer system displayed adequate dither in 16-bit mode and no low-level distortion. We have not had sufficient experience with professional systems to draw any general conclusions about their low-level performance.

## 3 CONCLUSIONS

This paper has been an attempt to aid in the understanding of quantization error, which in many cases produces benign wide-band noise. There are signal conditions for which this is not true, and we have shown how dither signals can linearize the cogging of the digital system so that the net effect is that of having a truly analog system of low distortion for low-level signals, with only a small and probably inconsequential addition to the noise.

## 4 ACKNOWLEDGMENT

This work was supported by the Natural Sciences and Engineering Research Council of Canada. The authors wish to thank Ivan Frola for his help in constructing the digital audio test circuit, and Jonathan Halliday for helpful comments on noise modulation.

## 5 REFERENCES

[1] W. R. Bennett, "Spectra of Quantized Signals," *Bell Sys. Tech. J.*, vol. 27, pp. 446–472 (1948).

[2] W. M. Goodall, "Television by Pulse Code Modulation," *Bell Sys. Tech. J.*, vol. 30, pp. 33–49 (1951).

[3] L. G. Roberts, "Picture Coding Using Pseudo-Random Noise," *IRE Trans. Inform. Theory*, vol. IT-8, pp. 145–154 (1962).

[4] L. Schuchman, "Dither Signals and Their Effect on Quantization Noise," *IEEE Trans. Commun. Technol.*, vol. COM-12, pp. 162–165 (1964).

[5] B. Widrow, "Statistical Analysis of Amplitude-Quantized Sampled-Data Systems," *Trans. AIEE (Applications and Industry)*, vol. 79, pp. 555–568 (1960).

[6] H. A. Spang and P. M. Schultheiss, "Reduction of Quantizing Noise by Use of Feedback," *IRE Trans. Commun. Sys.*, vol. CS-10, pp. 373–380 (1962).

[7] N. S. Jayant and L. R. Rabiner, "The Application of Dither to the Quantization of Speech Signals," *Bell Sys. Tech. J.*, vol. 51, pp. 1293–1304 (1972).

[8] B. A. Blesser, "An Investigation of Quantization Noise," *J. Audio Eng. Soc. (Project Notes)*, vol. 22, pp. 20–22 (1974 Jan./Feb.).

[9] B. A. Blesser, "Digitization of Audio: A Comprehensive Examination of Theory, Implementation, and Current Practice," *J. Audio Eng. Soc.*, vol. 26, pp. 739–771 (1978 Oct).

## THE AUTHORS

J. Vanderkooy

S. P. Lipshitz





John Vanderkooy was educated in Hamilton, Ontario, where he received an undergraduate degree in engineering physics, and later earned a Ph.D. in physics from McMaster University in 1967, studying low-temperature magnetic properties of metals. He spent two years on a postdoctoral fellowship in Cambridge, England, in similar research. Since then he has been in the physics department at the University of Waterloo, presently as an associate professor. Although he has published papers dealing with low-temperature metals research, he has always had a strong interest in electronics and high fidelity. He has contributed articles on oscillators and compandors, and his current interests are directed toward loudspeakers, acoustics, and digital signal processing.

Dr. Vanderkooy is a member of the AES and has published a number of papers in audio over the years, mainly in collaboration with his colleague Stanley Lipshitz. A main topic of their research has been a study of the phase, polar, and bandsplitting aspects of loudspeaker crossover networks.

•

Stanley P. Lipshitz was born and educated in South Africa, where he received a Ph.D. in mathematics from the University of the Witwatersrand in Johannesburg. He has been on the faculty of the University of Waterloo in Waterloo, Ontario, since 1970, and is currently an associate professor in both the departments of applied mathematics and physics. He is a fellow of the Audio Engineering Society, as well as a member of the American Mathematical Society, the IEEE and the Acoustical Society of America. He and his colleague John Vanderkooy form the core of the Audio Research Group at the University of Waterloo. His current research interests lie particularly in the areas of loudspeaker design and analysis and digital signal processing for audio. He is vice president of the Kitchener-Waterloo Chamber Music Society and records about forty of its concerts each season for broadcast on the University's FM radio station.