Copyright © 1972 American Telephone and Telegraph Company
THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 51, No. 6, July-August, 1972
Printed in U.S.A.

The Application of Dither to the Quantization of Speech Signals

By N. S. JAYANT and L. R. RABINER

(Manuscript received February 15, 1972)

By adding a pseudo-random "dither" noise to a signal X that is to be quantized, and by subtracting an identical noise sequence from the quantizer output, it is possible to break up undesirable signal-dependent patterns in the quantization error sequence, without increasing the variance of the error E. The idea has been widely discussed in the context of picture coding, and it is the purpose of this paper to demonstrate application of the technique to the quantization of speech signals. Computer simulations have shown how the use of dither whitens the quantization error sequence in PCM encoding, and renders it more acceptable than signal-correlated errors of equal variance. We demonstrate, for conditions of dither and no dither, typical speech recordings, illustrative error waveforms, and data on signal-to-error correlation C, and indicate how the advantage of dithering increases monotonically with crudeness of signal quantization and becomes significant when the number of bits per sample is less than about six. While the parameter C is a simple criterion for demonstrating the effect of dither, it must be emphasized that the truly relevant criterion is the statistical independence of E and X, and not merely the decorrelation of these functions. Thus, for example, we show that for the case of a reciprocal PDF (probability density function) for X, a zero value of C can be achieved without dither. For purposes of implementation, it is desirable to employ dither noise values characterized by a discrete PDF, with a support that is equal to an integral multiple of the step-size Δ_X in the quantizer. We show that for effective dithering, the step-size Δ_N in the noise PDF need be no smaller, typically, than $\Delta_X/4$. Finally, we indicate an application of dither to the quantization of speech signals by delta modulation.

I. INTRODUCTION

Signal quantizers, in general, produce quantization error sequences that have signal-dependent patterns. The perceptibility of such patterns tends to be very small for quantizations that are fine enough to provide practically useful signal-to-error ratios; while with relatively cruder quantizations, the perceptibility of signal-dependent errors increases to a point where techniques that can make the errors independent of signal samples become very attractive, even if they do not decrease the error variance itself. Dithering[1] is precisely such a scheme. It is based on the concept of forcing the quantization error E, conditional to a given input X, to be a zero-mean random variable, rather than a deterministic function of X. The randomization of conditional error $E(X)$ is accomplished by the addition of a random dither noise sample N to the input, and quantizing $(X + N)$ instead of X. The use of a pseudo-random dither sample N permits the subtraction of N from the quantizer output $(X + N)_Q$, and this insures an error variance that is essentially no greater than that in the undithered system. Roberts[1] provided an excellent demonstration of the above concept in his pioneering paper on the use of dither for picture coding. Subsequent work on dither[2,3] has also referred to picture signals. Specifically, Limb[2] has studied application to differential quantizers, and Lippel, et al.,[3] have demonstrated the use of two-dimensional, non-random dither patterns the inherent low visibility of which makes dither subtraction from $(X + N)_Q$ irrelevant, perceptually.

The purpose of this paper is to demonstrate the utility of dithering for the quantization of speech signals. We have confined our attention to the use of a pseudo-random dither of the type Roberts[1] employed, but we have considered application to differential quantization also; specifically, to the simplest type thereof, viz, delta modulation.

Section II will describe results from a computer simulation which studied dither for uniform quantizers of the PCM type, and showed that the use of dither whitens the quantization error sequence without increasing its variance, and renders the errors more acceptable than the signal-dependent errors in the undithered system. Results are in the form of speech recordings, error waveforms, and data on the signal-to-error correlation C. These data show how the utility of dithering increases monotonically with crudeness of quantization, and becomes significant for quantizers operating with less than about six bits per sample. The parameter C is a simple criterion for our demonstration, but it is emphasized that the truly relevant criterion is the statistical independence of E and X, and not merely the decorrelation of these functions. In fact, we show that in the example of a reciprocal PDF (probability density function) for X, a zero value of C can be achieved without dither. Section III discusses how, for implementation, it is desirable to employ dither noise values characterized by a discrete PDF with a support that is equal to an integral multiple of the step-size Δ_X in the quantizer, and shows that for effective dithering, the step-size Δ_N in the noise PDF need be no smaller, typically, than $\Delta_X/4$.

Finally, in Section IV, we indicate a possible application of dither to the quantization of speech signals by delta modulation.

II. DITHERING FOR PCM QUANTIZATION OF SPEECH

Referring to Fig. 1, the input X was speech sampled at the Nyquist rate (6 kHz) and included about 6,000 samples from a 1-second male utterance, "Have you seen Bill?". The dither noise N had a uniform PDF with a zero mean and a range equal to the step-size of the B-bit quantizer:

$$p(N) = \frac{1}{\Delta}\,;\; -\frac{\Delta}{2} < N < \frac{\Delta}{2}\,, \tag{1}$$

$$\Delta = \frac{\text{Peak-to-peak value of } X}{2^B}\,. \tag{2}$$

The uniform quantizer was described by the output-input relation

$$X_Q = \left[\frac{X}{\Delta}\right]\cdot\Delta + \frac{\Delta}{2}\,, \tag{3}$$

where the square brackets denote the "greatest integer in." The input X also has an integral part X_I and a fractional part X_F :

$$X = \left[\frac{X}{\Delta}\right]\cdot\Delta + X_F = X_I + X_F\,; \qquad 0 \leqq X_F < \Delta \tag{4}$$

and the quantization error E_0 , *without dither*, is simply the difference between $\Delta/2$ and X_F :

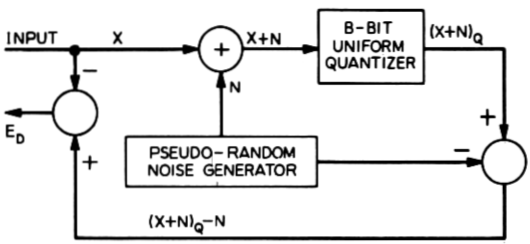

Fig. 1—PCM quantizer with dither.

$$E_0 = X_Q - X = \frac{\Delta}{2} - X_F . \tag{5}$$

The signal-independent, random dither noise N has the effect of making the quantizing error *with dither*,

$$E_D = (X + N)_Q - N - X, \tag{6}$$

statistically independent of X; and the subtraction of N from the quantizer output insures that the variance of E is no more (forgetting a correction term for the end steps of the quantizer) than that in the undithered case, which is given by the well-known expression (assuming a uniform distribution)

$$\langle E^2 \rangle = \int_{-\Delta/2}^{\Delta/2} e^2 \cdot p(e)\, de = \int_{-\Delta/2}^{\Delta/2} e^2 \cdot \frac{1}{\Delta}\, de = \frac{\Delta^2}{12} , \tag{7}$$

where $\langle \cdot \rangle$ denotes "average value of." The reader is referred to Roberts' paper[1] for a demonstration of the effect of dither on the properties of E, but we will briefly indicate here how E is statistically independent of X.

Let us rewrite eq. (6) in the form

$$E_D = (X + N)_Q - (X + N). \tag{8}$$

Referring to the example in Fig. 2b, one sees that for any given X, the dithered quantizer input $(X + N)$ has a uniform PDF of width Δ, centered around X. In general, a portion of this range (the hatched area) falls outside of the quantizer slot that included X. In view of eqs. (8) and (3), this portion is equivalent, for error calculations, to a corresponding portion (the horizontally striped area) in the quantizer slot including X. In other words, the fractional part (4) of $(X + N)$ can have any value between 0 and Δ, irrespective of the value of X. Hence, the error E_D (8) has the following X-independent distribution:

$$p(E_D/X) = \frac{1}{\Delta} \, ; \, -\frac{\Delta}{2} < E_D < \frac{\Delta}{2} \, ; \quad \text{any} \quad X \tag{9}$$

while the error E_0 [Fig. 2a and eq. (5)] is a deterministic function of X.

Figure 3 illustrates typical waveforms of E_D and E_0, the quantization errors with and without dither for three illustrative values of B. The following observations emerge:

(*i*) The perceived signal dependence of E_0 is a monotonically decreasing function of B.

(*ii*) The introduction of dither serves to decorrelate the error E_D from the input X even for the worst case of $B = 1$.

(*iii*) A broad similarity between the E_0 and E_D waveforms, attained at $B = 5$, suggests that for larger values of B, the advantage gained by dither tends to become insignificant.

These observations are confirmed by speech recordings of several sentences, which compare, for values of B in the range 1 to 10, the speech output $\{X_0\}$ without dither, and the speech output $\{X_D\}$ with dither. The quantizing error in the latter has an obvious white-noisy nature, while that in the undithered case is perceived to be signal dependent, especially for crude quantizations ($B < 5$ or 6). Recordings of the respective error waveforms E_0 and E_D confirm the point; the E_0 waveforms begin to sound more and more like speech as the quantization gets coarser; and for all values of B, the signal-dependent distortion in X_0 is more degrading than the white-noise in X_D. Incidentally, we have also verified that, for all B, the use of dither has no effect on the error-variance itself, as shown by Roberts.[1]

We have compiled, as a further quantitative description of the effect

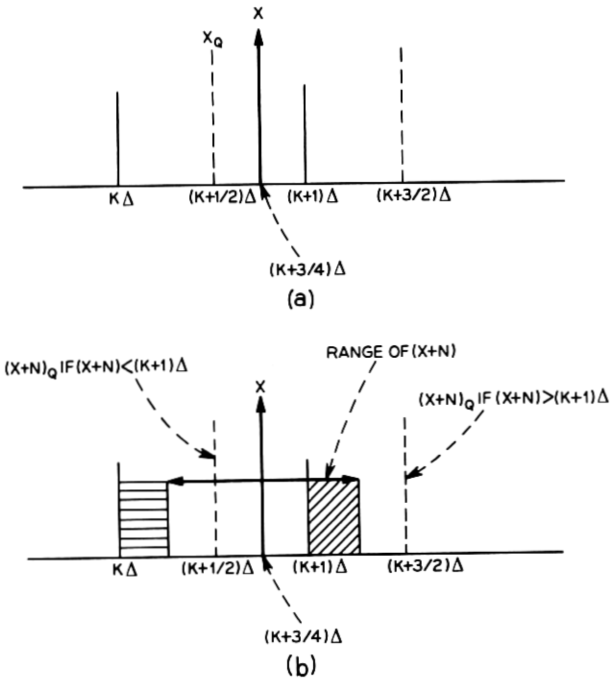

Fig. 2—Illustration of quantization error characteristics (a) without dither, (b) with dither.

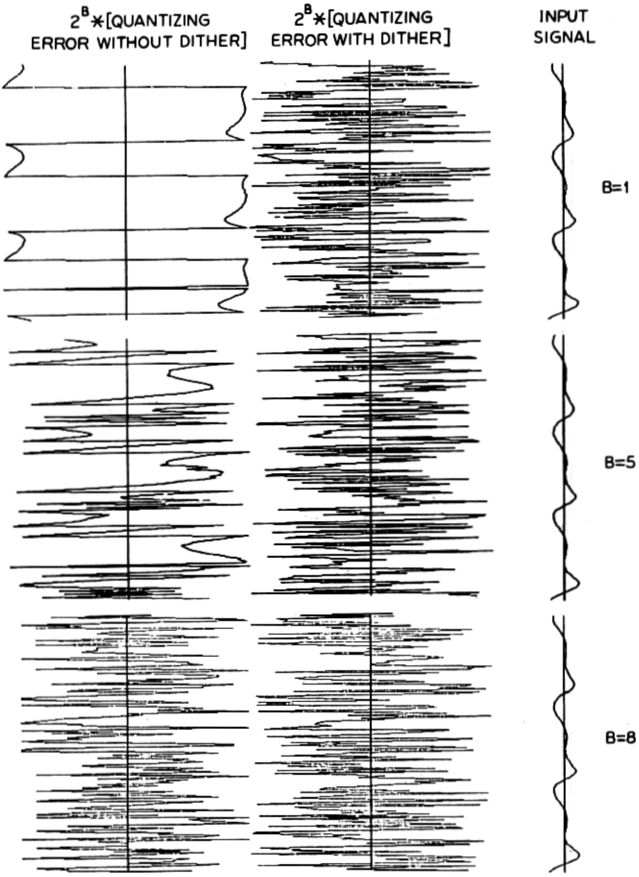

Fig. 3—Waveforms of quantization error.

of dither, values of the signal-to-error correlation (both X and E are assumed to be zero-mean functions)

$$C = \frac{\langle XE \rangle}{\sqrt{\langle X^2 \rangle \langle E^2 \rangle}}. \tag{10}$$

Figure 4 plots C_0 (without dither) and C_D (with dither) as functions of B. It is clear once again, that there is a value of B, say 6, below which the perceptibility of signal dependence in E_0 (as reflected by C_0) is large enough for the decorrelating effect of dither to be significant.

Notice that C_D oscillates, without any obvious structure, in the range

(−0.01, 0.01). Speech recordings mentioned earlier indicate that the ear cannot really resolve colorations corresponding to different values of C in the above range; and, in fact, we believe that a useful criterion for perceptually sufficient signal-to-error decorrelation would be an empirical requirement of the form

$$-0.01 < C < 0.01. \tag{11}$$

We should emphasize, however, that while the correlation measure C is very demonstrative, the truly relevant criterion in question is the statistical independence of E and X, and not merely the decorrelation of these quantities. As a matter of fact, C can be forced to be zero even without dither, as seen in the following example:

Assume that X has a reciprocal PDF. Let us compute, for all values of X in the Kth quantizer slot, the expected value of $X \cdot E_0$:

$$\langle XE_0 \rangle = \int_{K\Delta}^{(K+1)\Delta} \{(K + \tfrac{1}{2})\Delta - X\} \cdot X \cdot p(X)\, dX. \tag{12}$$

Obviously, if $p(X) = 1/X$, $\langle XE_0 \rangle_K$ vanishes, for all K, and, as per eq. (10), C_0 will be zero, without the use of dither! To reiterate, therefore, the idea of using dither is not merely to decorrelate E and X, but to make E statistically independent of X; which, of course, also ensures that C is zero, by definition.

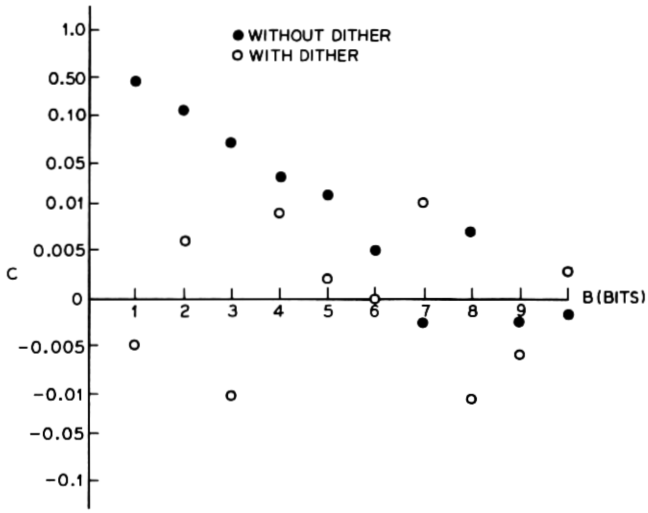

Fig. 4—Signal-to-error correlations.

III. IMPLEMENTATION

The discussion of the previous section has assumed a dither noise with a continuum of sample values uniformly distributed with a zero mean and a support equal to Δ. It is clear, however, that the error randomization or signal-smearing mechanism of Fig. 2 can work even if the support of the noise distribution is equal to an integral multiple of Δ; further, that the noise distribution can be discrete, allowing only a finite number of equiprobable, equally spaced, values; and that for successful dithering, the step-size Δ_N in the noise generator (spacing between consecutive allowed values of N) must be much smaller than the step-size of the signal quantizer itself.

$$\Delta_N \ll \Delta. \tag{13}$$

The above description of the dither noise turns out to be an important one for purposes of practical implementation. Computer simulations were carried out to demonstrate the condition (13) above. Results appear in Fig. 5 which plots the signal-error correlation C_D as a function of (Δ_N/Δ) for different values of B. (The speech input used here was different from that of Fig. 4, but this is immaterial.) Recall now, from (11), that a criterion for perceptually sufficient signal-to-error decorre-

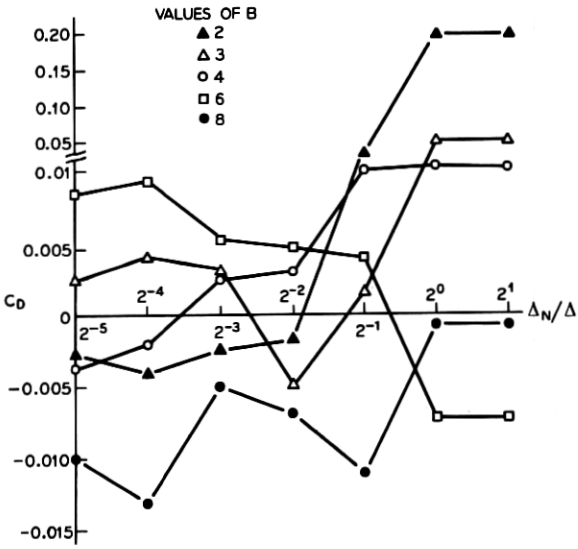

Fig. 5—The effect of discrete dither noise.

lation is the requirement that C lie within the range $(-0.01, 0.01)$. It is hence apparent from Fig. 5 that dither is quite ineffective for relatively fine quantization ($B > 6$) while, for coarse quantization, a safe requirement for achieving the error-smearing effect of dither can be expressed typically in the form

$$\frac{\Delta_N}{\Delta} \ngtr \frac{1}{4}. \tag{14}$$

It is interesting that (14) is also borne out by previously mentioned literature on picture-coding.[2,3]

Finally, the observation that C_D does not decrease monotonically with Δ/Δ_N is very interesting, especially since it can be shown that there are other important properties of discrete dither which are indeed monotonically related to Δ/Δ_N. For example, it can be shown that the PDF of the quantization error E_D, with discrete dither, is given simply by eq. (9) with a correction term that is inversely proportional to Δ/Δ_N.

IV. THE USE OF DITHER IN DELTA MODULATION

Limb[2] has mentioned the applicability of dither to differential quantizers for picture coding. The differential quantizer that we will discuss here for speech, is a simple one-bit differential quantizer, or a delta modulator (DM); and the dither noise that we will consider is the simple pseudo-random noise considered by Roberts[1] and discussed in Sections II and III.

Figure 6 is a block diagram of a simple delta modulator,[4] which builds a staircase approximation Y to a band-limited input X on the basis of the equation

$$Y_r = Y_{r-1} + \delta_r \operatorname{sgn} (X_r - Y_{r-1}). \tag{15}$$

In other words, each increment in Y follows the direction of the difference between the current value of X_r and the latest staircase approximation to it. With a linear delta modulator (LDM), the step-size δ_r is time-invariant and is tailored to the slope statistics of the input for optimal encoding:[5]

$$\delta_r = \delta_{\text{OPT}}, \tag{16}$$

while in adaptive delta modulation (ADM), δ_r is allowed to follow the slope variations in the input.[5,6] As a result, encoding errors in ADM not only exhibit a smaller variance than in LDM (for a given sampling rate), but are also less dependent on the input signal. The dependence of

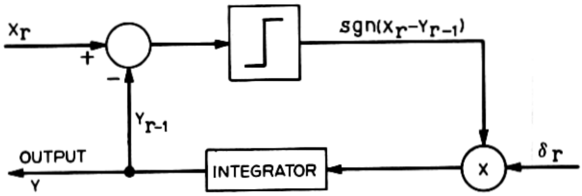

Fig. 6—Block diagram of a linear delta modulator.

encoding error on the input in LDM takes the very specific form of "slope-overload" distortion encountered in the encoding of relatively steep segments of the input. The dither experiment to be mentioned was thus motivated by LDM; and indeed, it proved to be less applicable for ADM.

We ought to emphasize, before proceeding, that for useful encoding, the sampling rate in LDM is at least an order of magnitude times greater than the Nyquist frequency of the input, and the perceptually relevant part of the encoding error in DM is the "in-band" noise as obtained by low-pass-filtering the high-frequency DM noise to the frequency band of the band-limited input. This error is shown as E_D in Fig. 7; here, the input was the utterance, "This is a recording of delta modulated speech," about 3 seconds long, and band-limited to 3.3 kHz. It was sampled for LDM at 60 kHz. The optimum step-size δ_{OPT} was determined in an earlier simulation using signal-to-in-band-error-ratio as a criterion of

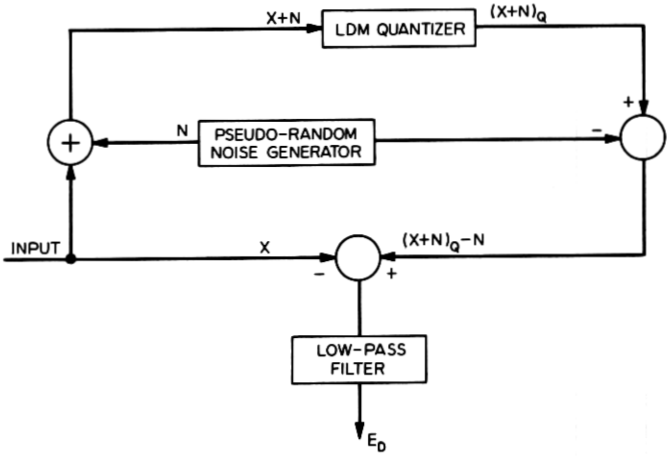

Fig. 7—LDM quantizer with dither.

good encoding. The pseudo-random dither was a zero-mean, uniformly distributed quantity (also sampled at 60 kHz), and had a variable range δ_N:

$$p(N) = \frac{1}{\delta_N}\,; \frac{-\delta_N}{2} < N < \frac{\delta_N}{2}. \tag{17}$$

The low-pass filter was a recursive 3-pole filter with a nominal 3.3-kHz cutoff.

Figure 8 plots the signal-to-error ratio SNR as a function of dither range δ_N. One notices a peak SNR advantage at the value

$$\delta_N \sim \delta_{\text{OPT}}\,. \tag{18}$$

In the figure, the case of $\delta_N = 0$ represents the case of no dither.

The SNR advantage due to dither in LDM is significant in that there is no parallel result in PCM. In the quantizers reported by earlier workers, and in Section II, the role of a pseudo-random dither was merely to smear the error sequence, without changing its variance; and this left the SNR unaltered in spite of dither. The same preservation of error variance is expected to hold in LDM, but only with reference to the unfiltered (high-frequency) encoding error. Once again, the role of dither is to smear or whiten this error sequence. But since the unfiltered error in LDM is expected to have considerable low-frequency components (recall the signal-dependent slope-overload distortion), the whitening of this error has the effect of decreasing the error variance within the signal band, hence, the increase in signal-to-error ratio.

A comparison of LDM speech, without dither and with an optimal dither (18), reflects the 2-dB SNR advantage in Fig. 8, and, more obviously, a desirable whitening of the encoding error.

Without going into details, it should be mentioned that the advantage

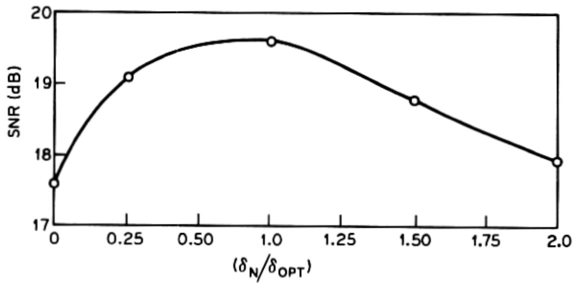

Fig. 8—LDM performance with dither.

of dither was considerably less in evidence when sampling rates of less than 60 kHz were employed, or when the delta modulation was adaptive. This is probably due, in the first case, to a lesser out-of-band noise rejection factor, and, in the second, to the fact that ADM starts out with much lesser signal-to-error dependencies than LDM, and hence, has less to gain from the dither technique.

V. CONCLUSION

The concept of an error-whitening dither noise, utilized so far generally for picture quantization, has been shown to be applicable to the coding of speech signals via PCM and LDM. The demonstrated advantages of dither have considerable practical significance at values of B (bits-per-sample) in the range 4 to 6, for PCM; and typically, for 60-kHz sampling in LDM. The qualities of speech encoding in the two cases are comparable, but they both fall short of toll-quality. However, they still represent a quality range that is obviously quite usable; and the error-whitening property of dither appears to be a very efficient way of enhancing the acceptability of speech in this quality range.

VI. ACKNOWLEDGMENTS

The programming assistance rendered by Mrs. K. Shipley is greatly appreciated. We would also like to thank J. L. Flanagan for suggesting the idea that dithering might be successfully applied to speech.

REFERENCES

1. Roberts, L. G., "Picture Coding Using Pseudo-Random Noise," IRE Transactions on Information Theory, *IT-8*, No. 2 (February 1962), pp. 145–154.
2. Limb, J. O., "Design of Dithered Waveforms For Quantized Visual Signals," B.S.T.J., *48*, No. 7 (September 1969), pp. 2555–2582.
3. Lippel, Z., Kurland, M., and Marsh, A. H., "Ordered Dither Patterns For Coarse Quantization of Pictures," Proc. IEEE (Letters), (March 1971), pp. 429–431.
4. De Yager, F., "Delta Modulation, A Method of PCM Transmission Using a 1-Unit Code," Phillips Res. Rep., No. 6 (December 1952), pp. 442–466.
5. Abate, J. E., "Linear and Adaptive Delta Modulation," Proc. IEEE, *55*, No. 3 (March 1967), pp. 298–307.
6. Jayant, N. S., "Adaptive Delta Modulation With a One-Bit Memory," B.S.T.J., *48*, No. 3 (March 1970), pp. 321–342.