Design of Dither Waveforms for Quantized Visual Signals

By J. O. LIMB

Dither signals may be added to coarsely quantized picture signals to mask undesirable contours. We show that a class of differential quantizers is equivalent to ordinary quantizers with respect to the design of dither signals. We give a design method for a number of deterministic and random dither waveforms and evaluate their visibility using a simple model of threshold vision.

I. INTRODUCTION

Television signals are invariably generated in an analog form. To obtain the advantages of digital transmission, it is necessary to quantize the signal in some way. In ordinary quantization the output levels of the quantizer are uniformly spaced throughout the range of the input signal; in the absence of any coding it would require six bits to send a signal quantized to 64 levels. In practice, at least 64 levels are required to produce a high quality picture.

A strong incentive to reduce the number of levels is that it would reduce the number of bits transmitted. For example, if the quantizer step size is doubled, the number of levels can be halved and the bit rate of the source can be reduced from six to five bits per sample. If this is done, the picture quality is degraded, but primarily for only one type of picture material, those areas in which the luminance changes slowly. These areas will be referred to as low-detail areas. The degradation takes the form of curved lines which look very much like contour lines on a map; thus this type of degradation is referred to as contouring.* The problem, then, is to eliminate the objectionable effect of contouring, which occurs only in the low-detail part of the picture, without using a larger number of levels.

---

* For example, see Fig. 3b of Ref. 1 for a differentially quantized picture showing contouring.

An effect similar to increasing the number of levels can be achieved by adding a dither signal to the true input signal. The dither signal produces rapid switching between the quantizer levels on either side of the true input signal. This switching is arranged so that the time one spends at a level depends on how close the true input signal is to that level. Thus in Fig. 1a when the signal lies midway between two levels, it oscillates between the two levels, spending equal time at each. In Fig. 1b the input lies a quarter of the distance up from the lower level; consequently, the required switched waveform should be down for three samples and up for one. The output waveform obtained when a dither signal is added to the input will be referred to as the chopped waveform or chopping pattern.

One could ask why such a strategy should be any good. While it is true that on the average the output signal will have the same amplitude as the input, it now has an additional error component which could degrade the signal further. Thus it is necessary to compare the visibility of the chopped waveform with the visibility of the contours that would otherwise be seen. Visibility is used here in the subjective sense of how easy is it to see an object. An objective visibility scale can be constructed using a fairly well defined subjective point on the visibility scale, that is, threshold, the point at which an object just becomes (or just ceases to become) visible. If the objective measure of the amplitude of a stimulus at threshold level is large, the stimulus has low visibility; conversely, the smaller the amplitude of the stimulus, the greater the visibility.

For signals near threshold, the visual system acts like a low-pass filter so that the chopped waveforms with the highest frequency components will be attenuated most and hence will be the least visible. Thus the pattern of Fig. 1a will be less visible since its repetition frequency is twice that of the pattern of Fig. 1b. In choosing suitable chopped waveforms we attempt to select those signals which have the least visibility.

The chopping patterns can be random or deterministic. Figure 1c shows a typical sample of a random pattern for an input halfway between the two quantizer levels (the same input amplitude as in Fig. 1a). The probability of each sample being high or low is one-half and is independent of previous samples. Since there is a finite chance that a given segment of the random sequence contains frequency components lower than those of the waveform of Fig. 1a, the random sequence of Fig. 1c is more visible than the deterministic pattern of Fig. 1a.

Let us now consider the problem of designing a dither signal which will produce an optimum chopping pattern at various levels. Goodall first observed that by adding a small amount of noise to the input, contouring was almost eliminated at the expense of a small increase in the granularity (or noisiness) in the picture.[2] Roberts examined the problem quantitatively and showed that in order to produce a random chopping pattern, which always averaged out to the same amplitude as the input, the probability density function (PDF) of the noise should be rectangular with a maximum amplitude of plus and minus half a quantizing interval.[3] He further showed that if one subtracted at the quantizer output the same noise that was added at the input, the root mean square error in the output signal is halved (if one forgets the correction term for the quantizing intervals at the end of the range). Limb considered the visibility of the granulation in the quantizer output.[4] Using a simple model of the visual process, it was shown that the visibility of granulation resulting from independent random samples with a rectangular probability density function of the correct amplitude is zero when the input equals a quantizer output level, and reaches a maximum when the input is midway between levels. Further, by introducing negative correlation between samples, the visibility can be reduced by about 50 percent.
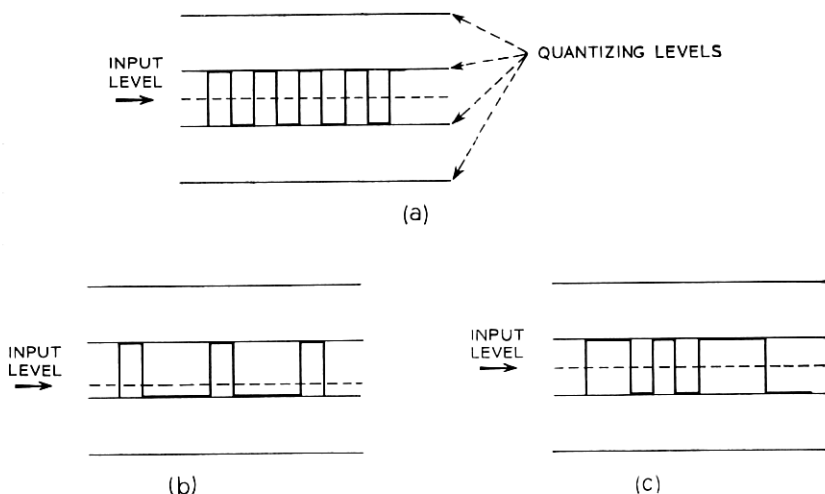

Fig. 1 — Chopping patterns for (a) input half way between levels, (b) input quarter way between levels, and (c) random pattern with input half way between levels.

In this paper we look at the problem of applying dither to differential quantization as opposed to ordinary quantization. The approach is the same as with ordinary quantization; design a dither waveform which, when added to the input, produces the required chopping pattern at the output (see Fig. 2a). All the components of the differential quantizer are assumed to be ideal. The chopped waveforms produced by the differential quantizer will depend on how the levels of the quantizer within the loop are positioned close to the zero level. Two commonly used configurations are (*i*) a decision, or input, level placed at zero (Fig. 2b), and (*ii*) a representative, or output, level placed at zero (Fig. 2c). We consider only the second configuration (however, see Section VII). We show that under fairly general conditions a differential quantizer, containing a quantizer stage with a representative level at zero, behaves the same as an ordinary quantizer (with equal level spacing) with respect to dither. We design a set of dither signals, both random and deterministic, which produces chopping patterns with a low visibility. The visibility of the chopping patterns is calcu-
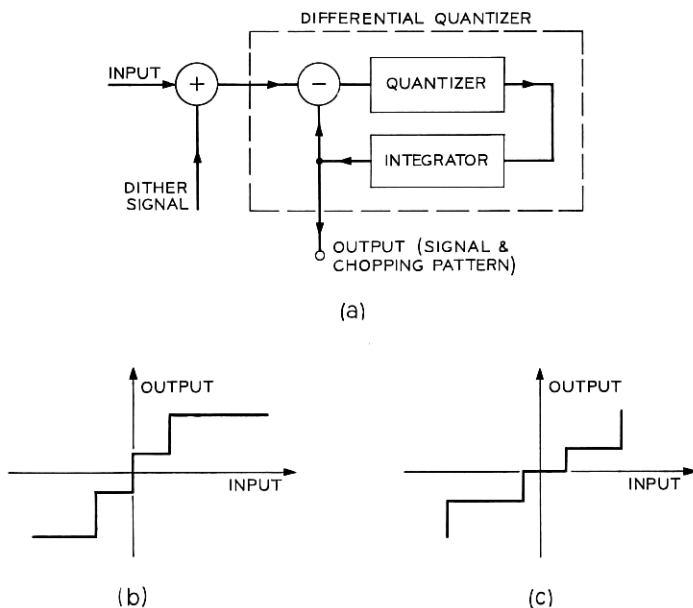

Fig. 2 — (a) Dither applied to differential quantization. (b) Transfer characteristic of quantizer with decision level at zero. (c) Transfer characteristic of quantizer with representative level at zero.

lated, enabling a comparison to be made between random and deterministic dither. It is anticipated that two-dimensional dither signals will reduce the visibility of contours by a factor of six.

If one decides to subtract the dither signal from the averaging pattern as Roberts did, the rules for generating the best dither waveforms have to be rederived. When this is done, it is found that subtracting dither signals is barely superior to not subtracting them.

II. DIFFERENTIAL QUANTIZER—QUANTIZER CHARACTERISTIC

A quantizer may be divided into two sections, the classifier which divides the signal into a number of ranges according to the position of its decision levels, D_i, and the weighter which assigns a value to each range according to the settings of the representative levels R_i. Figure 3a shows the quantizer characteristic as it is generally drawn. An alternative representation is given in Fig. 3b, where the vertical dashes represent the decision levels, and the crosses represent the representative levels. This representation is more convenient since we are concerned with the positions of the representative levels relative to the positions of the decision levels.

The input level to the classifier, in the absence of the dither signal, is denoted by Δ and expressed as a fraction of r, the distance from R_0 to R_1 (Fig. 4). Since we are considering slowly changing input signals, Δ will always lie in the range R_{-1} to R_1.

We assume that the quantizer has a representative level at zero as Fig. 4 shows. The only levels that affect the design of the dither signal are the two decision levels, D_{-1} and D_1, lying closest to zero and the adjacent representative levels R_{-1} and R_1. Furthermore, we assume that R_{-1}, D_{-1}, R_0, D_1, and R_1 are equally spaced. This is probably the most useful configuration since it satisfies Max's first condition for minimizing error, that is, the decision levels should lie midway between the corresponding representative levels.[5] In addition, $R_1 = 2D_1$ which is on the stability boundary and hence corresponds to the maximum setting of R_1 if limit cycles are not to occur.[6]

III. DESIGN OF DITHER SIGNAL

When rectangular random noise is used as a dither signal, the chopped waveform has a granular appearance and the visibility of this granulation depends on the amount of correlation in the waveform. For example, when the correlation is positive, the waveform is
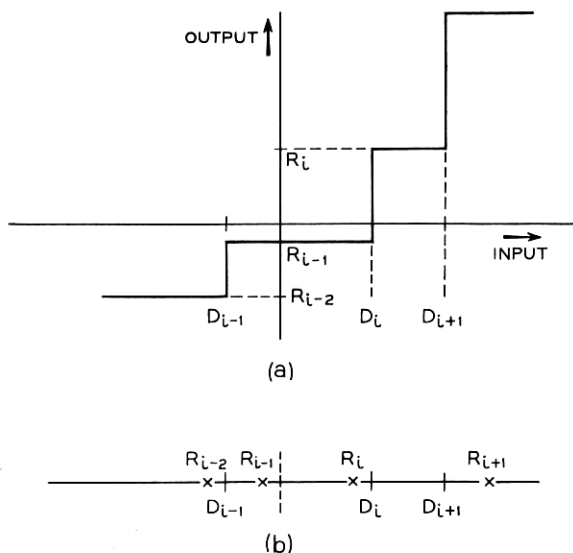
(a)

(b)

Fig. 3 — Quantizer characteristic (a) usual representation and (b) alternative representation.

more likely to contain large runs of 0's and 1's (if we use 0 and 1 to denote the two quantizer levels between which the output is chopping); if the waveform is negatively correlated, a 1 is more likely to follow a 0 and the waveform will switch back and forth more rapidly. Notice that the visibility of a perturbation is approximately proportional to the area when the area is small. Consequently, long runs of 1's or 0's are much more visible than the negatively correlated waveform containing a higher probability of short runs.

If we restrict the chopped waveform to be described completely by a second order probability density function, there are limits on the


Fig. 4 — Quantizer characteristic—configuration with representative level at zero.

amount of negative correlation that can be achieved.[4] This restricted chopped waveform with maximum negative correlation can be generated with a dither signal having the second order probability density function shown in Fig. 5. Here x_1 and x_2 represent adjacent samples of the dither signal. The negative correlation produces a sharp minimum in the visibility of granulation in the waveform when the input to the classifier is close to D_1 (or D_{-1}), that is, when $\Delta = 0.5$. The dashed curve is for uncorrelated noise and is shown for comparison.

The dither noise can also be represented as shown in Fig. 6a, which better illustrates the time series nature of the process. For example, a sample occurring at random in the top half of the amplitude range will, for the next sample, occur in the lower half. The nature of the second order probability density function ensures that the random sample oscillates between the upper and lower half of the range. This type of dither will be referred to as two-step random dither.

Dither waveforms can be generated for any number of steps, although with an increase in the number of steps, the visibility of granulation will reach a minimum and then start to increase. Figure 6b shows an example of four-step dither. Notice that when the input level lies on the boundary between two steps, deterministic chopping patterns are produced. Furthermore, these patterns should have the least visibility of any chopping pattern with the same average level. In general, a low visibility pattern (LVP) has the minimum allowable cycle length (for example, cycle length of 3 at $\Delta = \frac{1}{3}$) and has the

(a)   (b)

Fig. 5 — Two-step random dither signal: (a) probability density function of correlated noise and (b) visibility of noise.
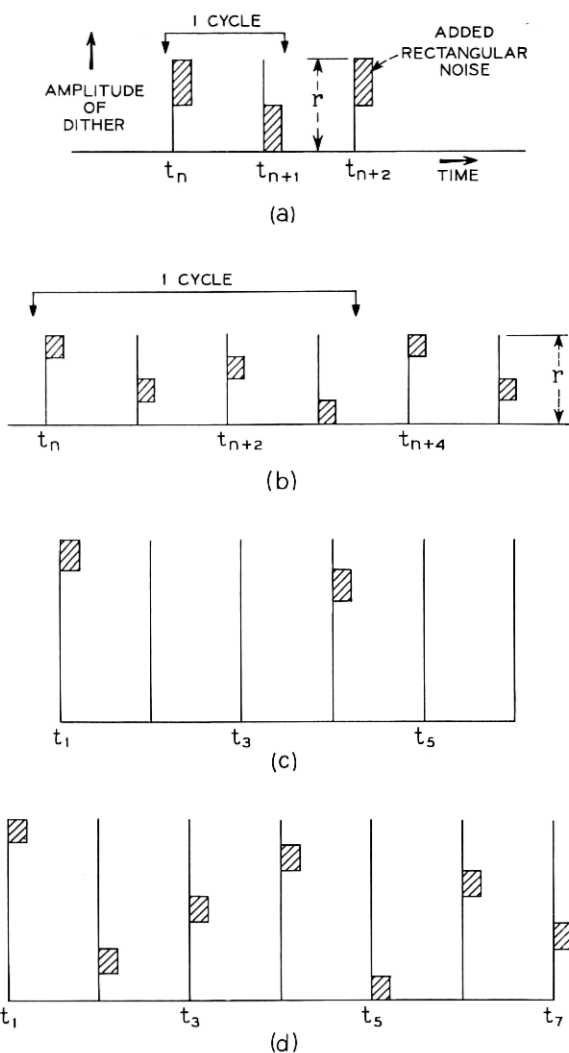
Fig. 6 — Dither signal represented in time series form: (a) two-step dither, (b) four-step dither, (c) attempt to construct six-step dither, and (d) seven-step dither.

shortest maximum run of either value (for example, 1010100 is a low visibility pattern for $\Delta = 3/7$ but 1100100 is not).

Can an n-step low-visibility chopping pattern be generated with second order noise for any value of n? The answer is no, as the following attempt to reconstruct patterns for six-step and seven-step

dither shows. In Fig. 6c the first step has been assigned arbitrarily to t_1, while the second step must then be assigned to t_4 so that when the input gives $\Delta = \frac{1}{3}$, every third sample exceeds threshold ($LVP = 1, 0, 0$). There is no sample to which the third step can be assigned to give a low visibility pattern of $(1, 0)$ as required for $\Delta = \frac{1}{2}$.

For the seven-step case (Fig. 6d), the first step is assigned to t_1, the second step may be assigned to either t_4 or t_5, both giving low visibility patterns (assume t_4). The third step, if assigned to t_6, will again give a low visibility pattern $(1, 0, 0, 1, 0, 1, 0)$. Similarly, all the other steps can be assigned to give low visibility patterns.

In the general case of n-step dither, tests for low visibility patterns can be made as follows:

Assign first step   — t_1
      second step — $t_{n/2}$ n even
and   third step   — $t_{n/4}$ n divisible by 4.

To have the least visibility, the resulting pattern after assigning the third step must not have runs of 0's differing in length by more than one, otherwise the position of a 1 could be moved to shorten the longest run. However, this would affect steps 1 and 2 which have given low visibility patterns. Thus any multiple of 4 equal to or greater than 8 will not give low visibility patterns. Again:

Assign third step $-t_{(n\pm2)/4}$ n even, not divisible by four.

By the same argument as above for $n \geq 6$ and even, low visibility patterns are not obtained. Similarly, the odd numbers can be tested. In all, low visibility patterns can be obtained for $n = 2, 3, 4, 5,$ and 7.

In the scheme considered so far, each step in the quantizing interval has been filled with rectangular noise of amplitude equal to the step height. Random patterns are produced whenever Δ lies within a step, while a deterministic pattern is generated when Δ lies exactly on the boundary between two steps. Consider changing the rectangular noise to a fixed level at the midpoint of the step. The chopping pattern will now switch from one deterministic pattern to another as Δ changes. We examine the visibility of granulation associated with both random and deterministic dither in Section IV.

Implementation of either random or deterministic dither schemes would be a simple matter requiring little additional hardware. Fig. 7 shows the output from a computer simulation of a differential quantizer with four-step deterministic dither in Fig. 7a and seven-step deterministic dither in Fig. 7b. The visibility of granularity in these two dither schemes will differ; in Section IV, calculations of visibility are made to enable the most promising schemes to be selected.

(a)

(b)

Fig. 7 — Chopping patterns resulting from deterministic dither signal for (a) four-step dither and (b) seven-step dither. The straight lines denote the inputs.

IV. VISIBILITY OF THE CHOPPED WAVEFORM

The measure of the visibility of the discrete waveform is based upon a simple model of threshold vision that has proved reasonably accurate.[4,7] Briefly, threshold vision is assumed to act like a spatial low-pass filter, and the difference in visibility between two displays (in this case the display resulting from the analog signal and the display resulting from the quantized signal) is measured by the difference between the filtered version of the two signals.* Two different measures of the difference are considered: one is the mean square, and the other is the mean modulus. The measure of visibility is denoted by $U(\Delta)$, which depends on Δ since the value of Δ determines the shape of the chopped waveform.

4.1 Deterministic Patterns

The solid-line curves in Figs. 8 and 9 show the calculated visibility of granulation for three- and four-step patterns at a viewing distance of 36 inches. The visibility is shown for only half the range of Δ,

---

* Appendix B gives more detail.

Fig. 8 — Visibility of granularity produced by three-step dither at 36 inches viewing distance (— — random; ——— deterministic).


Fig. 9 — Visibility of granularity produced by four-step dither at 36 inches viewing distance (— — random; ——— deterministic).

since $U(\Delta)$ is symmetrical about $\Delta = 0.5$. The arrows on the abscissa indicate the amplitude levels of the dither pattern. Thus until the value of Δ exceeds 0.167 in Fig. 8, no switching occurs in the output. There are minima at $\Delta = \frac{1}{3}$ and $\frac{2}{3}$ as expected for three-step dither. The curves of $U(\Delta)$ for the two criterion functions are similar in shape, the rms curve lying slightly above the mean modulus curve. In Fig. 9 the minimum at $\Delta = 0.25$ is not very large, and one would expect granulation to be more visible for patterns having a greater number of steps. Figure 10 clearly shows this for a five-step pattern which has a higher minimum than Figs. 8 and 9. By comparing the average value of U for three-, four-, and five-step patterns, four-step is just better than three-step, and both are superior to five-step patterns.

Figure 10 also gives U for a five-step dither at a viewing distance of 72 inches. The spread of the visual impulse response is now much greater in relation to the size of a picture element. In fact, $U(\Delta)$ is not very different from what would be expected with infinite smoothing by the eye. With infinite smoothing all minima would be zero and joined to the maxima at 0.1 by straight lines; that is, five equal triangles of amplitude 0.1. The similarity of U to the result expected for infinite smoothing would suggest that a pattern with a larger number of steps would be superior. Going to the maximum of seven steps
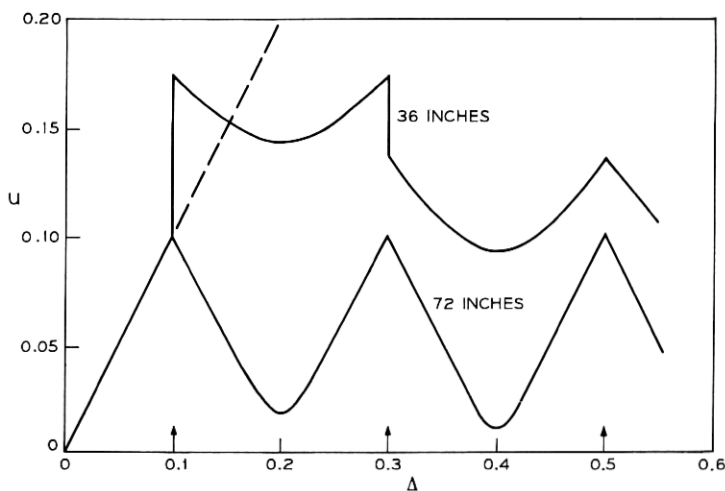

Fig. 10 — Visibility of granularity produced by a five-step dither at 36 and 72 inches viewing distance; mean square error criterion.

(Fig. 11) significantly reduces the mean value of U, and the curve is no longer similar to the curve for infinite smoothing. By comparing Fig. 11 with Fig. 9, one can see that increasing the viewing distance by a factor of two has reduced the calculated visibility of granulation by about one-half for the best pattern in each case—a not altogether surprising result.

4.2 *Random Patterns*

Random noise was added to the deterministic patterns in the manner shown in Figs. 6 and 7. Notice that the pattern generated after quantization is deterministic when the decision level lies at the junction of two steps and is the same as the pattern produced in the absence of noise. Figs. 8, 9, and 11 give the calculated visibility of random patterns for mean square and mean modulus error criteria. As required, the random curves touch the deterministic curves between steps, and in most other places the curves lie above them. Four-step dither still gives the smallest average U, $\langle U \rangle_{av}$; three-step dither is the next best.

4.3 *Deterministic versus Random Patterns*

With deterministic patterns, n-step dither results effectively in inserting $n-1$ levels in the original quantization interval. The brightness at these new levels is not constant, however, and has a variance about the true analog input value given by the minima, $U(\Delta_{min})$. As Δ changes from Δ_{min}, the variance remains unchanged but a constant error is introduced since the average value of the output no longer equals the average value of the true analog input.

With random patterns, the average value of the output always equals the average value of the input. Thus at the maxima of $U(\Delta)$, the variance of the perceived image with deterministic patterns is less than with random patterns, but there is an additional error resulting from differences in the perceived average values of the true analog input and the chopped waveform. By using a decision theory model of threshold vision, the visibility in the two situations could be compared. However, such models have not proved accurate enough to apply to this type of second order effect.

On the basis of the mean square and mean modulus criteria it appears that deterministic dither is slightly superior; but because each case has different distributions for the perceived brightness, such comparisons are risky and best wait experimental confirmation.

Fig. 11 — Visibility of granularity produced by seven-step dither at 72 inches viewing distance; random and deterministic, mean square error criterion.

V. DITHER APPLIED IN TWO AND THREE DIMENSIONS

Devising low visibility patterns in two dimensions is more difficult than in one dimension. In fact it appears that there are only two equivalent, trivial low visibility patterns. These patterns occur for two by two step interpolation; they are,

$$
\begin{array}{c} - x \rightarrow \\ \left| \begin{array}{cc} 1 & 3 \\ 4 & 2 \end{array} \right. \end{array} y \quad \text{and} \quad \begin{array}{c} - x \rightarrow \\ \left| \begin{array}{cc} 1 & 4 \\ 3 & 2 \end{array} \right. \end{array} y .
$$

For larger patterns it appears that we must settle for something less ideal. A four by four step pattern was generated by considering it to consist of four two by two patterns, which were themselves generated in the manner of a two by two pattern, as the partly completed pattern in Fig. 12a shows.

The computer program used previously for the one-dimensional case was extended to calculate the visibility of the four by four pattern. $U(\Delta)$ is shown in Fig. 13 for a viewing distance of 36 inches and a mean square error criterion. $\langle U \rangle_{av}$ has been reduced to about one-third in going from one dimension to two in this example. This pattern does not have minimum visibility. This can be seen for $\Delta = \frac{1}{4}$ where a lower visibility pattern could be obtained by the chopping pattern of Fig. 12b. This would make little difference to $\langle U \rangle_{av}$, however, since $U(\Delta)$ for $\Delta = \frac{1}{4}$ is already very small. Undoubtedly patterns approximating the ideal could be found for a larger number of steps.

In applying dither in the time dimension, care must be taken not to introduce "temporal granularity," that is, flicker. To study the visibility of flicker would require an entirely new model, accounting

for the variation in sensitivity over the retina to temporal changes in luminance. Flicker occurs when large picture areas differ in luminance periodically from frame to frame. By arranging for the average luminance of an area to change as little as possible from frame to frame, flicker can be minimized. Thus the two-dimensional pattern

|	— Distance →				
|	1	15	7	10	Frame 1
Time	12	5	14	4	Frame 2
	8	9	2	16	Frame 3
↓	13	3	11	6	Frame 4

which was built up with the help of the two by two low visibility pattern, will have an average luminance which varies at most by $\frac{1}{16}$ of a quantizing interval from frame to frame. This is not true of the sequence,

	— Distance →				
|	1	13	4	16	Frame 1
Time	9	5	12	8	Frame 2
	3	15	2	14	Frame 3
↓	11	7	10	6	Frame 4

which is simply derived from the two by two pattern and nearly identical to the pattern of Fig. 12a. Notice that if the input signal has a uniform distribution over the quantizing interval, the average luminance of each frame will be the same. For example, for $0 < \Delta < 0.25$, frames 1 and 3 have greater average luminance, while for $0.5 < \Delta < 0.75$, frames 2 and 4 have greater average luminance. Since the probability of obtaining signals that do not vary (lie within one step) over "large" areas is small for high quality pictures (which

(a)   (b)

Fig. 12 — Generation of four by four step pattern: (a) partially completed pattern (b) two dimensional chopping pattern having lower visibility than the corresponding pattern resulting from (a).

consequently have a small step size), the probability of obtaining flicker should be small. There is advantage in using the second pattern since it provides better smoothing.

Dither applied in the time domain should be more successful than in one spatial dimension at a 36-inch viewing distance since the temporal impulse response, even at high ambient illumination, probably has a greater spread; the problem of flicker, however, should be kept in mind. Another advantage of the temporal dimension is that the amount of smoothing should be independent of the viewing distance.

In comparing deterministic patterns with random patterns, temporal smoothing has often been neglected; this leads to incorrect conclusions. For example, if deterministic two-dimensional spatial dither is compared with random dither, the random pattern would provide smoothing in three dimensions since added noise components in adjacent frames are uncorrelated, and, as just shown, the improvement in smoothing provided by an additional dimension is large. A valid comparison could be made by using "frozen" noise, that is, noise that repeats from frame to frame.

Section 3.43 of Ref. 1 describes results that were obtained when two types of dither waveform were added at the input of a differential quantizer. The first pattern was a one dimensional four step waveform added vertically. The second pattern was a four by four step pattern added horizontally and vertically. Figure 3d of Ref. 1 shows the effect of adding the two dimensional dither to a picture while Figs. 6c and 6d of Ref. 1 show the effect of adding one and two dimensional dither respectively, to a low amplitude ramp waveform.

VI. RECEIVER SUBTRACTION

Roberts added pseudorandom noise having a rectangular probability density function to the signal prior to quantization, and subtracted the same noise from the signal at the receiver.[3] Neglecting end effects from the smallest and largest quantization levels, a reduction of one-half in the variance of the output signal is obtained. Roberts states that adding noise to the input and subtracting it from the output is equivalent to adding a level of noise to the signal, but that this is not the same noise as was added to the input. Since we are concerned with the exact sequence in the output signal (this will critically affect the visibility of the added noise), the relation between the added input noise and the equivalent output noise will be derived.

In Fig. 14, rectangular noise is added to the input signal of value

Fig. 13 — Visibility of granularity produced by two dimensional four by four step pattern. Viewing distance is 36 inches with mean square error criterion.

$R_n + \Delta$. All noise components which cause the input signal to exceed D_n are represented by R_{n+1}, and all components producing a combined signal less than D_n are represented by R_n. Thus, in subtracting the input noise from the quantized signal, components lying between $r/2 - \Delta$ and $r/2$ are subtracted from R_{n+1}, while the other components are subtracted from R_n. When the noise is subtracted one sees that the whole process is equivalent to adding noise of the same amplitude to the unquantized signal. The noise to be added can be obtained from the input noise by inverting separately amplitudes greater and less than $r/2 - \Delta$ as Fig. 14 shows. For example, amplitudes above D_n [such as $(r/2 - \Delta) + \Gamma$] go to $r/2 - \Gamma$ where Γ is any increment between 0 and Δ. This relationship is very useful since now we can forget the quantization and consider just the distortion of the added noise component.

If an n-step dither sequence is quantized and the original sequence subtracted, inversion occurs at every step except where Δ is less than $r/2n$. Consequently, $n - 1$ new sequences will be produced and only in special cases will the new sequences be the same as the original. A technique will be developed for rapidly estimating the new output sequences from the input sequence.

A sequence can be written as a function of time,

$$\text{Time} \qquad 1, \quad 2, \quad 3, 4, \cdots, n$$
$$\text{Amplitude} \quad A_1, A_2, \qquad \cdots, A_n$$

where A_i is an integer between 1 and n denoting the step. A sequence can also be written

$$\text{Amplitude} \quad 1, \quad 2, \quad 3, \quad \cdots, n$$
$$\text{Time} \qquad T_1, T_2, T_3, \cdots, T_n$$

Fig. 14 — Derivation of the properties of add-subtract noise patterns.

where T_i is an integer between 1 and n denoting the time slot in which the ith amplitude occurs. Conversion from the time representation to the amplitude representation can be simply accomplished. For example, at amplitude A_1 the corresponding time slot is 1; that is, T corresponding to A_1 is 1.

Consider the sequence $T_1 \ldots T_i, T_j \ldots T_n$. For Δ lying between the ith and jth steps, the new sequence is $T_i \ldots T_1, T_n \ldots T_j$. Notice that the cyclic order is reversed but otherwise unchanged. Thus if Δ changes by i steps, the amplitude sequence shifts by i steps but the order is unchanged unless Δ is less than $r/2n$, in which case the order reverses. However, since the visibility of a sequence does not change if the order is reversed, this may be neglected. A cyclic shift in the amplitude sequence must now be converted to the time sequence, since we use the time sequence to calculate visibility. A shift by one step in the amplitude representation corresponds to an addition or subtraction by one, modulo n (depending on the direction of the shift), in the time representation. Thus if the time sequence was $1, 2, \ldots, n$ (a bad sequence from the point of view of visibility), the sequence at the ith level would be $n - i + 1, \ldots 1, 2, \ldots, n - i$, which is in fact the same sequence. This particular case is one of a set of sequences that remain unchanged as Δ changes from step to step.

6.1 *Visibility of Sequences*

There are at most $(n - 1)!/2$ different sequences that can be generated for a particular value of n where sequences are regarded as different if they have different visibilities; that is, they are not shifted

in time or reversed versions of another sequence. There is only one unique sequence for $n = 3$ and three unique sequences for $n = 4$. For $n = 4$ the three possible sequences are: (*i*) 1, 2, 3, 4; (*ii*) 1, 3, 2, 4; (*iii*) 1, 2, 4, 3. Sequence *i* produces three output sequences which are the same as itself. Sequence *ii* produces the output sequences

$$
\begin{array}{lllll}
\text{Input} & 1, 3, 2, 4 = 1, 3, 2, 4 & \text{Output} & 0 \\
& 4, 2, 1, 3 = 1, 3, 4, 2 & & 1 \\
& 3, 1, 4, 2 = 1, 3, 2, 4 & & 2 \\
& 2, 4, 3, 1 = 1, 3, 4, 2 & & 3 \\
& 1, 3, 2, 4 & &
\end{array}
$$

and *iii* produces the output sequences

$$
\begin{array}{lllll}
\text{Input} & 1, 2, 4, 3 = 1, 3, 4, 2 & \text{Output} & 0 \\
& 4, 1, 3, 2 = 1, 3, 2, 4 & & 1 \\
& 3, 4, 2, 1 = 1, 3, 4, 2 & & 2 \\
& 2, 3, 1, 4 = 1, 3, 2, 4 & & 3 \\
& 1, 2, 4, 3 & &
\end{array}
$$

These outputs are just shifted versions of one another and they should yield the same overall value of U. The output sequence 1, 3, 2, 4 provides better smoothing than 1, 3, 4, 2 which contains lower and hence more visible frequency components. This can be seen in the curves of U for the two sequences which were calculated independently of the arguments of this section (Fig. 15).

For comparison, $U(\Delta)$ is shown for the case previously considered in which the dither waveform is not subtracted from the output. The subtraction method gives a slightly lower average value of U (< 2 percent lower). The value of U for uncorrelated rectangular noise with subtraction is also shown. U is now independent of Δ. However, the problems associated with comparing random and deterministic dither schemes should be borne in mind (see Section 4.3).

6.2 *Constant Sequences*

Here is a technique for finding sequences which do not change as Δ changes from step to step (constant sequences). A step change in Δ results in an increment, modulo n, of each number in the sequence; thus, the numbers must be arranged so that the order remains unchanged after a shift. Constant sequences can be constructed simply by using a geometric method. In Fig. 16, five points are spaced equally around a circle, each point corresponding to a number in the sequence. Starting from any point, a line is drawn to another point to cor-

Fig. 15 — Visibility of granularity produced by four-step add-subtract dither pattern. Viewing distance is 36 inches (——— input sequence 1, 3, 2, 4; ——— ——— input sequence 1, 2, 4, 3; - - - - addition only 1, 3, 2, 4).

respond to a shift of the number of points cut off by the line: 1 in Fig. 16a and 2 in Fig. 16b. This second point is then shifted the same distance in the same direction. The shifting process is repeated until all points are covered, and we arrive back at the starting point if the number of points shifted is not a divisor of n (excluding 1). The number of unique constant sequences for an n-step pattern is equal to the number of nondivisor integers less than $n/2$ plus 1. Thus Figs. 16a and b represent the two constant sequences for $n = 5$. Unfortunately, for larger n, constant sequences do not have low visibility, as the five constant sequences,

$$
\begin{array}{cccccccccc}
1, & 2, 3, & 4, 5, & 6, & 7, & 8, & 9, 10, 11 \\
1, & 7, 2, & 8, 3, & 9, & 4, 10, & 5, 11, & 6 \\
1, & 5, 9, & 2, 6, 10, & 3, & 7, 11, & 4, & 8 \\
1, & 4, 7, 10, 2, & 5, & 8, 11, & 3, & 6, & 9 \\
1, & 10, 8, & 6, 4, & 2, 11, & 9, & 7, & 5, & 3
\end{array}
$$

show for $n = 11$. Probably the best sequence is the third, but this is significantly inferior to a sequence such as

$$1, \quad 11, \quad 2, \quad 10, \quad 3, \quad 9, \quad 4, \quad 8, \quad 5, \quad 6, \quad 7.$$

Figure 17 is a graph of $U(\Delta)$ for $n = 5$ for the constant sequence 1, 4, 2, 5, 3 and sequence 1, 2, 5, 4, 3. The constant sequence has an

Fig. 16 — Generation of constant sequences.

average value $\langle U \rangle_{av}$ of 0.095, which is lower than the other sequence and the low visibility pattern derived in Section 4.1 (which is also shown for comparison). The subtracted sequences for $n = 4$ (Fig. 15) give a slightly greater value of $\langle U \rangle_{av}$ (0.099) compared with $n = 5$. Notice the very low minimum at $\Delta = 0.2$ for the nonconstant sequence. The sequence producing this minimum may be calculated by subtracting one from each digit of the input sequence and is thus 1, 4, 3, 2, 5.

VII. DISCUSSION

The quantizer configuration with a decision level at zero was referred to briefly in the introduction. This configuration results in an
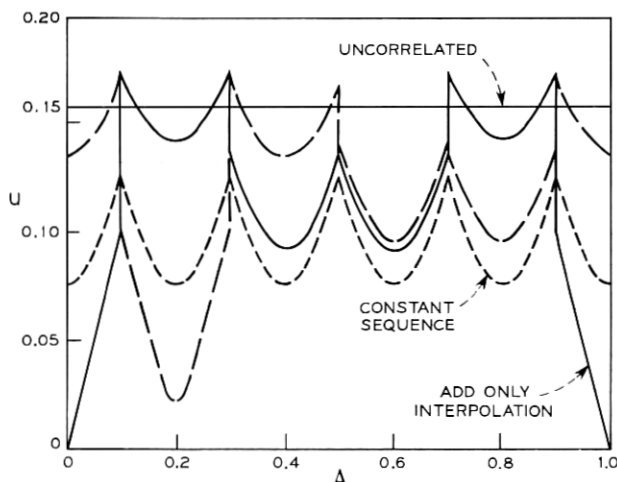

Fig. 17 — Visibility of granularity produced by five-step add-subtract dither pattern. Viewing distance is 36 inches.

even number of representative levels and has received more attention in the literature. In the absence of a dither signal the output will oscillate between R_1 and R_{-1}, but otherwise has no inherent dithering ability of its own. It will produce contours in low detail areas with much the same visibility as the quantizer configuration we have investigated. For an uncorrelated random dither signal, the switching waveform is not constrained to lie between the two adjacent quantizer levels as it is with a representative level at zero. For illustrative purposes, two switching waveforms have been generated for two different input levels assuming a random, uncorrelated dither signal (Fig. 18). Although it would be more complex to do so, one could calculate the visibility of these types of waveforms as done previously and compare the results with those just obtained. One problem is to decide upon the relative amplitudes of R_1 and R_{-1} for the two configurations.

VIII. SUMMARY AND CONCLUSIONS

The design of dither signals for ordinary quantizers is the same as the design for differential quantizers for quantizer characteristics of specific types. The requirements for equivalence are that the char-
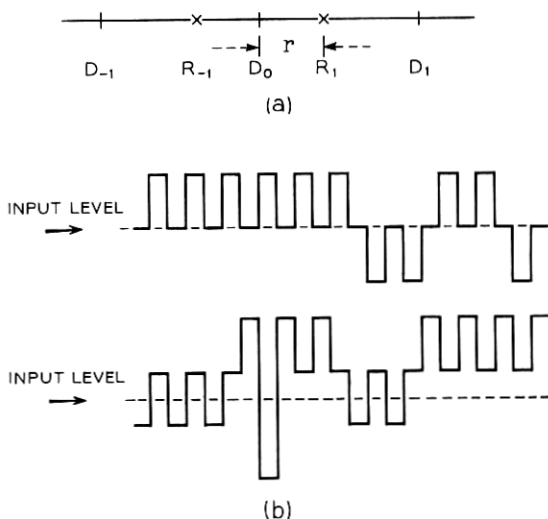

Fig. 18 — (a) Quantizer characteristic with decision level at zero. (b) Chopping pattern for input at $\Delta = 0$. (c) Chopping pattern for input at $\Delta = 1/2\ R_1$.

acteristic have a representative level at zero and uniform spacing of the adjacent pair of representative and decision levels.

Dither patterns may be three dimensional in design, varying horizontally, vertically, and from frame-to-frame. A pattern that varies in only one dimension can be generated having two, three, four, five, or seven amplitude levels (and no others), such that the visibility of the added pattern is a minimum for each level of the pattern.

We predict that a deterministic four-level pattern will give minimum visibility or granularity for *Picturephone®* visual telephone viewed at 36 inches. The use of this one dimensional dither signal should reduce the visibility of contours by a factor of about two when compared with a picture with no dither.

At 72 inches viewing distance (or say 36 inches with twice the sampling frequency) seven level dither should be used.

Four-level dither applied in two dimensions should reduce the visibility of contours by a factor of six compared with a picture having no dither. A further significant reduction should occur when dither is applied to the temporal dimension as well.

The dither signal may be subtracted from the received signal to further reduce the visibility of the added waveform. But the rules for determining the best patterns are different. For four-level dither the best addition-subtraction patterns give results that are only marginally better than the best patterns when they are not subtracted from the receiver.

APPENDIX A

Equivalence of Dither for Quantization

The method of proof is to show that for a Markov dither pattern (having an arbitrary conditional probability density function) the conditional probability of the switching pattern being at either level, given the previous value of the dither signal, is the same for both ordinary and differential quantizers.

Assume a Markov dither pattern described by the transition probability density function $P(x_i/x_{i-1})$ where $|x_i| \leq r/2$. The pattern could be either deterministic or random. For ordinary quantization, the probability of obtaining level R_n and level R_{n+1} for an input analog amplitude of $R_n + \Delta$ (see Fig. 19b) is

$$\Pr\{R_n/x_{i-1}\} = \int_{(-r/2)}^{(r/2)-\Delta} P(x_i/x_{i-1}) \, dx_i = I_1 \tag{1}$$

Fig. 19 — Equivalence of dither for ordinary and differential quantization—definition of terms—(a) conditional probability density function, (b) ordinary quantization, and (c) differential quantization.

and

$$\Pr \{R_{n+1}/x_{i-1}\} = \int_{(r/2)-\Delta}^{r/2} P(x_i/x_{i-1}) \, dx_i = I_2 , \qquad (2)$$

where $\Pr\{a/b\}$ is the probability of event a occurring given that event b has occurred. Thus the probability of obtaining levels R_n and R_{n+1} is conditional only on x_{i-1}.

For differential quantization, feedback occurs from the previous sample value, and it becomes necessary to distinguish between the output of the quantizer (primed) and the output of the complete differential quantizer (unprimed). Again, for an analog input of $R_n + \Delta$, assuming equal spacing of $R_{-1}, D_{-1}, R_0, D_1,$ and R_1 we have

$$\Pr\{R_n/x_{i-1} \; ; x_{i-1} < (r/2) - \Delta\}$$

$$= \int_{-(r/2)}^{(r/2)-\Delta} P(x_i/x_{i-1} \; ; x_{i-1} < (r/2) - \Delta) \, dx_i \tag{3}$$

$$= \Pr\{x_{i-1} < (r/2) - \Delta\} \int_{-(r/2)}^{(r/2)-\Delta} P(x_i/x_{i-1}) \, dx_i$$

$$= \Pr\{x_{i-1} < (r/2) - \Delta\}\{I_1\}. \tag{4}$$

Now

$$\Pr\{R_n/x_{i-1} \; ; x_{i-1} > (r/2) - \Delta\}$$

$$= \Pr\{R'_{-1}/x_{i-1} - r; x_{i-1} > (r/2) - \Delta\}$$

$$= \Pr\{x_{i-1} > (r/2) - \Delta\} \Pr\{R'_{-1}/x_{i-1} - r\}. \tag{5}$$

But since

$$\Pr\{R'_{-1}/x_{i-1} - r\} = \Pr\{R'_0/x_{i-1}\} = \Pr\{R_n/x_{i-1}\}, \tag{6}$$

$$\Pr\{R_n/x_{i-1} \; ; x_{i-1} > (r/2) - \Delta\} = \Pr\{x_{i-1} > (r/2) - \Delta\}\{I_1\}.$$

Thus from equations (3) and (6) one can see that $\Pr\{R_n/x_i\}$ is independent of the previous state of the differential quantizer and equal to the value obtained for the ordinary quantizer. By a similar argument, $\Pr\{R_{n+1}/x_{i-1}\}$ can be equated for the two quantizers.

APPENDIX B

Calculation of Visibility of Dither Signals

B.1 *Model of Vision*

Figure 20 shows a simple model used to describe the visibility of small amplitude signals.[7] $I(x, y, t)$ represents the spatial and temporal luminance pattern incident at the eye. The filter $\lambda(x, y, t)$ accounts for spread of the signal in space and time caused by the optics, the receptors, and subsequent neural processing. The amplitude of the hypothetical signal $E(x, y, t)$ is proportional to the observed visibility of the display. Thus the difference in visibility between two


Fig. 20 — Model of threshold vision.

displays can be measured by evaluating the average of some function of the difference between the value E resulting from one display, and the value of E resulting from the other.

We wish to know how well the discrete waveform with added dither approximates the analog signal in flat areas where contours are most bothersome. Thus it is reasonably accurate to represent the analog signal by a constant amplitude E_a, and the measure of the visibility of the discrete waveform is

$$U(\Delta) = E\{f[E_a - E_\Delta(x)]\},$$

where $E\{\cdot\}$ denotes the expected value and, as before, Δ denotes the position of the input within the quantizing interval. $E(x)$ varies with Δ since the chopping pattern $I(x)$ changes as Δ is varied. In a number of cases, $U(\Delta)$ has been evaluated for two different f functions, the square and the modulus.

B.2 *Visibility of Waveform*

The method of evaluating the visibility differs from that used previously.[4] Earlier, $E(x)$ was calculated for every possible input combination occurring in a signal segment of the length of the significant part of the impulse response. The probability density function of $E(x)$ was then calculated by weighting each output by the probability that the corresponding input occurred. From the probability density function the error can be calculated for the required criterion function.

The method now used is to first calculate a combined impulse response for the reconstruction filter and visual filter: this is then convolved with the input signal to obtain an output from which a measure of the granularity is derived. This technique is fast and accurate for deterministic signals which repeat after a short length, but slower if accurate results are required for random inputs. Fortunately, most of the signals investigated were deterministic.

Denoting the impulse response of the low-pass filter by $h_1(x)$ and the visual spatial filter (for example, in the horizontal dimension) by $h_2(x)$, then the combined impulse response is given by[8]

$$\lambda(x) = \int_{-\infty}^{\infty} h_1(y)h_2(x - y)\, dy.$$

This integral was evaluated for

$$h_1(x) = \frac{1}{x'}\left(\sin\frac{\pi x}{x'}\right)\bigg/\frac{\pi x}{x'}$$

and

$$h_2(x) = \pi^{-1/2} \exp\left[-0.0833\left(\frac{Ax}{x'}\right)^2\right],$$

where x' is the spatial Nyquist interval and A, which depends upon the viewing distance, is the width of a picture element in minutes of arc; $h_2(x)$ is the same impulse response as used previously.[4]

Figure 21 shows the combined impulse response of the normalizing low-pass filter and the visual system for Mod. II *Picturephone*® visual
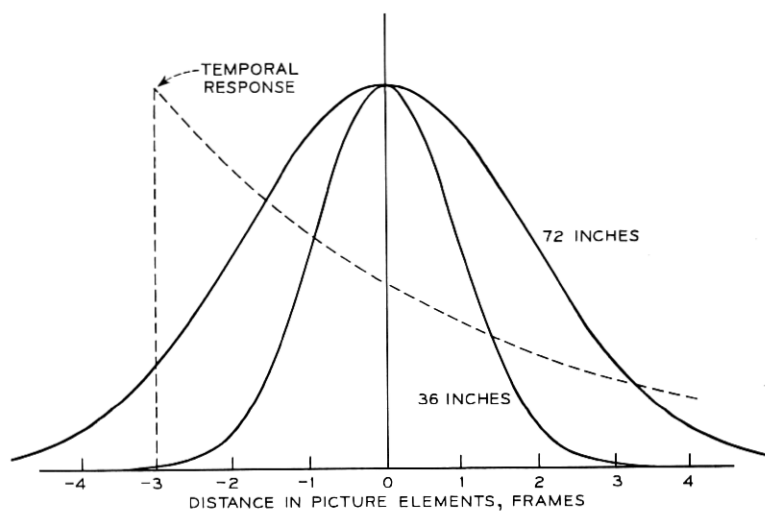

Fig. 21 — Combined spatial impulse response in one direction at viewing distances of 36 and 72 inches.

telephone viewed at 36 inches. The corresponding impulse response for a viewing distance of 72 inches (or alternatively, for 36 inches at twice the sampling frequency) is also shown and agrees to three decimal places with the impulse response of the visual system itself. In other words, the visibility of threshold detail is almost completely unaffected by the horizontal resolution of the display at 72-inch viewing distance (resolution limited by eye).

For a given input $I(x)$ the output is (using the convolution theorem),

$$E(x) = \int_{-\infty}^{\infty} I(y)\lambda(y - x)\,dy.$$

The limits of the integral can be reduced so as to integrate over only those values of $(y - x)$ for which λ is significantly greater than zero (in practice, greater than 0.1 percent). For random inputs, simulations were run for between 300 and 900 samples.

REFERENCES

1. Limb, J. O., and Mounts, F. W., "Digital Differential Quantization," B.S.T.J., this issue, pp. 2583–2599.
2. Goodall, W. W., "Television by Pulse Code Modulation," B.S.T.J., *30*, No. 1 (January 1951), pp. 33–49.
3. Roberts, L. G., "Picture Coding Using Pseudo-Random Noise," IRE Trans., *IT-8*, No. 2 (February 1962), pp. 145–154.
4. Limb, J. O., "Coarse Quantization of Visual Signals," Australian Telecommunication Res., *1*, Nos. 1 and 2 (November 1967), pp. 32–42.
5. Max, J., "Quantizing for Minimum Distortion," IEEE Trans. Inform. Theory, *IT-6*, No. 1 (March 1960), pp. 7–12.
6. Graham, R. E., "Predictive Quantizing of Television Signals," IRE Wescon Conv. Rec. *2*, part 4 (1958), pp. 147–157.
7. Budrikis, Z. L., "Visual Thresholds and Visibility of Random Noise in TV," Proc. Inst. Radio Eng. (Australian) *22*, No. 12 (December 1961), pp. 751–759.
8. Mason, S. J., and Zimmerman, N. J., *Electronic Circuits, Signals, and Systems*, New York: John Wiley, 1960, p. 327.