Digital Differential Quantizer for Television

By J. O. LIMB and F. W. MOUNTS

(Manuscript received January 23, 1969)

Correct tracking between the transmitter and receiver is difficult to maintain when a long integrator time constant is used. We describe a differential quantizer which has a digital integrator; this integrator enables perfect tracking to be achieved at the output of the integrator without any adjustments.

The differential quantizer gives high quality pictures when seven, eight, and nine quantizer output levels are used. We present a scheme for transmitting the nine-level signal at the rate of three bits per picture element.

Picture quality is improved significantly by adding low-amplitude dither patterns to the input signal to mask contours. The coder is more susceptible to transmission errors than coders having an analogue integrator with a short time constant; we discuss two methods for reducing the susceptibility.

I. INTRODUCTION

Differential quantization is well suited to visual signals for two reasons.[1] First, the quantizer acts like a predictive encoder, taking advantage of the large amount of correlation between adjacent elements of a picture to obtain a good prediction of the amplitude of the point being quantized.[2] Thus, the differential quantizer makes use of some of the statistical redundancy in the source. Second, the quantization can be partially matched to the changing sensitivity of vision.[3] To understand this, remember that the sensitivity of the visual system to small differences in luminance decreases markedly at boundaries between light and dark areas. The signal that is applied to the quantization stage of a differential quantizer is very nearly equal to the change in amplitude between adjacent elements. Thus, by quantizing small amplitude samples finely and large amplitude samples more coarsely a picture can be obtained which is partially matched to visual requirements.

Figure 1a is a block diagram of a differential quantizer; it differs in

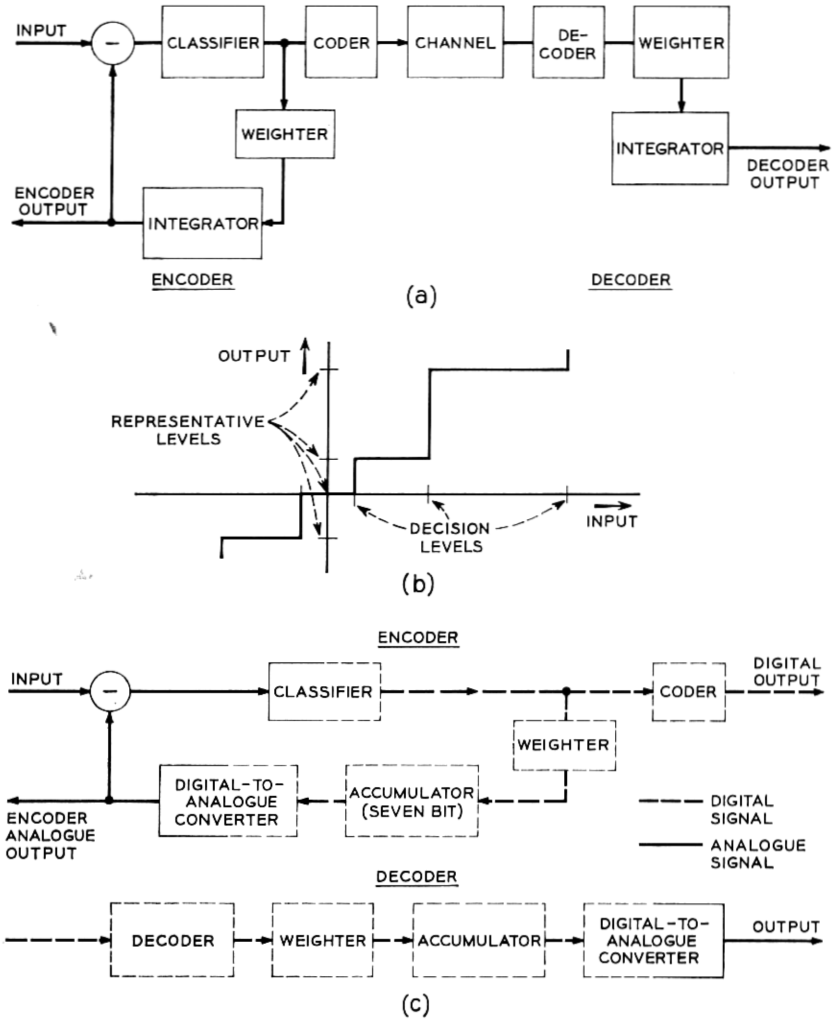

Fig. 1 — (a) Differential quantizer (DPCM coder-decoder). (b) Transfer characteristic of a quantizer. (c) Digital differential quantizer.

one small aspect from the usual representation. The quantizer section of the differential quantizer is considered as two separate parts. The first part, the classifier, contains the decision levels which divide the input signal range into a number of intervals (see Fig. 1b). Thus, the signal at the output of the classifier is digital and just denotes the interval in which the signal sample occurred. The signal at this stage is encoded for transmission to the receiver. The second part, the weighter, assigns an amplitude or weight (representative value) to each section that can have either a digital or analogue value. The integrator is usually an analogue device in which case the weights are generated as analogue values.

In practice, it is difficult to resynthesize a high-quality picture at the decoder. The longer the time constant of the encoder and decoder integrators, the greater the precision required in implementation to prevent the received picture from differing from the sent picture (referred to as mistracking). A short integrator time constant, on the other hand, makes mistracking less of a problem but introduces effects similar to noise. These, generally, are not too serious if the inner pair of levels alone is used to make the correction for integrator leak. Mistracking can stem from three sources:

(*i*) The representative levels at the encoder and decoder can be mismatched. While the setting of the largest pair of levels is not quite as critical since it is seldom used, the smallest levels must be adjusted quite accurately, especially when the integrator time constant is long.

(*ii*) The frequency response of the integrators can differ.

(*iii*) An analogue component of the signal can bypass the classifier stage (hence analogue breakthrough) and feed through the analogue weighter into the integrator. The analogue breakthrough is generally prevented from reaching the integrator in the decoder by digital regeneration in the signal path, and so a mismatch between the encoder and decoder can occur. The problem can be overcome by carefully gating the digital output of the classifier stage to remove any vestige of analogue signal.

In an attempt to overcome decoder mistracking and still have a long integrator time constant, it was decided to perform the operations of weighting and integrating digitally. This should ensure exact tracking of the decoder under all conditions except for the obvious cases of either a digital circuit malfunction or an error occurring during transmission.

Figure 1c is the block diagram of the digital differential quantizer. The analogue parts are distinguished from the digital parts. An extra block is required since it is necessary to convert the output of the accumulator (digital integrator) to an analogue quantity prior to subtracting.

The digital encoder-decoder completely eliminates mistracking; thus, the picture at the receiver cannot be distinguished from the quantized picture at the transmitter. High precision analogue weighters and integrators are now replaced by digital circuits and, further, the classifier design need not be precise. Indeed, analogue breakthrough, unless it is large enough to switch a digital circuit, is ineffectual and small variations in the position of the decision levels have negligible effect on picture quality. Requirements on the digital-to-analogue (D-A) converter are not very strict. Nonlinearity in the characteristics of the D-A converters has the effect of producing a change in the gamma of the signal, of which the eye is not very critical. In fact, changes in the D-A converter produce similar effects to changes in the D-A converter of an ordinary PCM system.

In quantizers used for ordinary PCM encoding the appearance of contour lines in low-detail areas sets the lower limit on the number of quantizing levels that can be used. For the differential quantizer a similar effect occurs. As the weight assigned to the smallest pair of representative levels is increased, contour lines become more visible. Addition of random and pseudorandom noise improves picture quality for ordinary quantizers having less than about 100 levels.[4,5] A theoretical study of the use of dither signals with differential quantizers suggests that even at a close viewing distance dither should prove effective in improving picture quality for a given number of levels.[6]

This paper first describes the digital differential quantizer and the results obtained with it. The reduction in the visibility of contours by adding specially designed dither signals in both the horizontal and vertical directions is then explored. The problem of overcoming the effects of channel errors is also briefly discussed.

II. DESCRIPTION OF SYSTEM

2.1 *Analogue Section*

See Fig. 1c. The subtractor is an emitter-coupled pair circuit that is ac coupled at both the input and feedback terminals. The classifier comprises a set of eight threshold circuits connected in parallel. Their threshold levels can be adjusted independently to give the desired partition of the input. For example, setting two decision levels to the same value reduces the number of intervals by one.

The classifier design is simplified by making the positive and negative stages identical. This is made possible by feeding them separately with signals of opposite polarity; such signals are generated by the emitter-coupled pair of the subtractor. Sampling is inherent in the operation of the decision circuits. Narrow sample pulses with a base width of approximately 20 ns are amplitude modulated by the signal. The decision as to whether the sample exceeds the threshold is then made using a high speed flip-flop. When the flip-flop is set, the signal is considered to have exceeded threshold. A reset pulse is applied to all flip-flops prior to the occurrence of the next sample pulse. The classifier successfully uses only emitter-coupled integrated circuit logic elements (for essentially analogue operations) to obtain fast decisions (< 30ns) and good stability.

2.2 *Digital Section*

Because a parallel classifier is used, all threshold circuits with thresholds less than the input signal value are triggered. For example, if the signal to the classifier exceeds level number three, then an output occurs on level number three and also on levels number two and number one. The outputs of the threshold circuits are combined logically to give a sign bit and four other binary signals—one for each level. A particular signal takes the value "one" when the input to the classifier falls in the interval associated with that level. This code allows no more than one of the four outputs to be a "one" for any sample. A zero on all four bits denotes the zero interval of the classifier. This code change is not essential, but it allows the weights associated with each interval to be controlled independently, which is convenient in an experimental coder.

Each output is connected to a word generator which generates a binary number specifying the amplitude of the representative level. Notice that this method of weighting means that only symmetrical weighter configurations can be investigated because the magnitude of the levels is generated independently of the sign. The experimental arrangement allowed any set of digital weights to be wired on a small plugboard.

The accumulator uses a seven bit adder that sets the precision with which the weights can be assigned (Fig. 2). The contents of the adder for the previous sample are fed back into the adder together with the new difference signal. For a zero difference signal it can be seen that the same seven bit amplitude signal would circulate through the adder and delay stage without change. The operation of the accumulator can be expressed as $y_n = y_{n-1} + x_{n-1}$, where x_n and y_n would be the values of the input and output, respectively, of the accumulator for the nth sample.

Under certain conditions the adder could overflow or underflow, say for a large peak in the input video signal. This is prevented by the

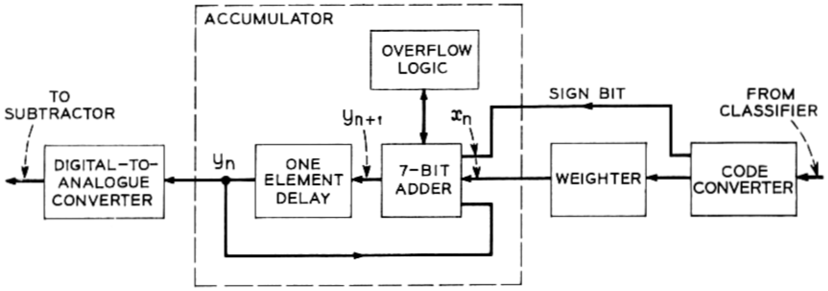

Fig. 2 — Digital section of differential quantizer.

addition of overflow logic. The circuit holds the adder output at level 127 if overflow occurs or at level zero if the adder underflows, that is,

$$y_n = 127 \quad \text{if} \quad y_{n-1} + x_{n-1} \geqq 127$$

and

$$y_n = 0 \qquad \text{if} \quad y_{n-1} + x_{n-1} \leqq 0.$$

The overflow logic effectively clamps the accumulator when it is underdriven; we make use of this action to fix the dc level in the coder. The delay period of one element is realized with a clocked flip-flop.

The digital-to-analogue converter uses the ladder method of conversion. The seven-bit converter was built using selected 1 per cent resistors and has a settling time of less than 50 ns. Resampling for display purposes was not considered necessary; within the loop the classifier resamples.

2.3 *Alignment*

Setting the input and output levels of a quantizer accurately is normally a tedious problem requiring precise equipment. In this section we describe a self-alignment technique using the digital-to-analogue converter of the differential quantizer; the level adjustment problem then becomes quite trivial.

Since the representative levels of the weighter are assigned digitally, they can be set exactly with the wired plugboard. The decision levels, however, are analogue levels, and they need to be set up accurately both in relation to themselves and to the representative levels. To do this, the feedback loop is broken between the classifier, and the weighter and the classifier levels are set up digitally in the weighter at twice the desired value. The sign bit to the accumulator is alternated each sample period so that the output of the digital-to-analogue converter is a square wave of twice the desired amplitude. The input video signal is disconnected so that only the square wave is coupled to the classifier. Since the signal is ac coupled, the excursions from the mean are of the right value to set both the positive and negative decision levels. The control on the position of the threshold is adjusted until the decision stage just triggers. Nothing more complex than a voltmeter is required to make this adjustment.

III. INVESTIGATIONS

3.1 *Ideal Weighter-Integrator*

Probably the most significant difference between previous analogue implementations of the differential quantizer and the present digital implementation is the fact that the weights are known exactly and the integrator accumulates exactly until it is reset at the end of a line. Thus, the weights of the larger levels can be set at values which are exact multiples of the smallest (or inner) levels. For example, if the inner pair of levels was set at 3/128ths and the other three pairs of levels were set at 6/128ths, 12/128ths, and 21/128ths, this would be called a multiple setting of the weighter. A multiple setting is a necessary condition for producing clear contour patterns that are the same as one gets with ordinary quantization. Figure 3b, which shows a picture processed by the digital differential quantizer, illustrates these contour patterns; for comparison Fig. 3a shows the original signal. The picture was chosen because of its flat background which has the effect of emphasizing contours. A coarse quantizer setting with only seven levels is used to make them more visible.

By going to a nonmultiple setting, the coded amplitude in a flat area of the picture will vary from line to line depending upon what levels were used in the previous part of the line. It might be argued that this effect could be used to mask contouring. It can, but it is not very successful as Fig. 3c shows. There is still contouring on the left side of the picture; the right side is quite streaky compared with what can be done using other methods (Fig. 3d). These methods are discussed further in Section 3.4.

Thus, for a multiple setting of the weighter the low-detail areas of the picture are quite free from random noise. Quantizing errors show up as more or less visible contours (depending on the level setting and the type of picture material) which are quite sharp if the input signal-to-noise ratio is high. Quantization error at edges and in high detail areas, on the other hand, is more random in nature.

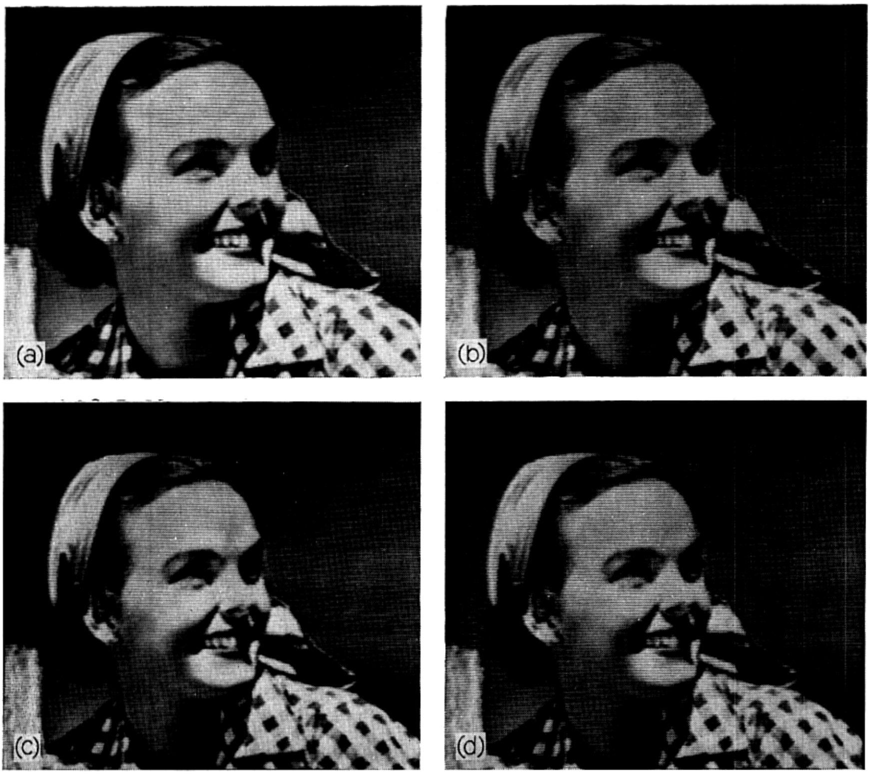

Fig. 3 — Pictures processed by digital differential quantizer: (a) original analogue signal, (b) processed picture—seven levels with multiple setting of levels, (*c*) processed picture—seven levels with nonmultiple setting, and (*d*) processed picture—same level setting as (b) with 4 × 4-step dither signal added. (The scan lines and printing screen cause moiré patterns that are not in the originals.) Glossy prints of this figure can be obtained by writing to the authors.

3.2 *Quantizer Characteristic*

We are limited to symmetrical configurations of the decision and representative levels because of the way the quantizer was designed. However, there is nothing to suggest that a nonsymmetrical setting would have any significant advantage. Many settings were tried using nine representative levels—four positive levels, four negative levels, and a zero level. The scale of Table I was found to give good results—no contouring is apparent but the skilled viewer can detect slight degradation at edges. The best setting changes only slightly with the subject matter. For contrasty pictures and graphics, edges can be improved by expanding the scale slightly. This is done by reducing the input amplitude and increasing the output amplitude by a compensating amount.

Eight-level settings of the weighter were investigated using the same configuration as with nine levels, except that the zero representative level was removed by moving the first pair of decision levels at ±1 percent to zero. The smallest output step of this configuration, however, is unchanged. Edges should also be reproduced with the same fidelity.

The pictures obtained with eight levels are very similar to the nine-level pictures. Contours still occur with about the same visibility for coarse-level settings but now a constant amplitude consists of an oscillation between two levels separated by an amount equal to the smallest step size. This oscillation is largely removed by the filtering at the receiver and the filtering taking place in the eye.

Since there is little choice between the picture qualities of the eight-level setting and the nine-level setting, the eight-level setting would be preferred since fewer levels are required. However, the nine-level configuration of Table I can be altered to significantly reduce the visibility of contours by placing the first pair of decision levels at half the value given in Table I. One can show that this is equivalent to adding a deterministic, two amplitude dither signal (which is random in the vertical and temporal directions) in the horizontal direction (see Section 3.4 and Ref. 6). The quantizer levels can now be expanded (taking advantage of the reduction in visibility of contours) to reproduce high-detail areas more accurately, giving an overall improvement in picture quality.

In a number of analogue differential quantizers built previously by others, representative level settings of approximately 1, 3, 7, and 20 percent (for example, Ref. 2) have been found satisfactory. This setting compares with approximately 2, 6, 14, and 24 percent for the digital differential quantizer. The difference in the inner level settings

Table I—Level Positions of Quantizer*

Level		±1		±2		±3		±4	
Decision Level		±1		±4		±10		±19	
Representative Level	0		±2		±6		±14		±24

* Expressed as percent of peak-to-peak signal.

is quite large and is almost surely the result of the different integrator characteristics.

The picture quality is rather insensitive to a change in input signal amplitude. For example, an increase of 4 dB in signal level causes a slight loss of sharpness at edges with a decrease in the visibility of contours, while a decrease of 4 dB improves the edges and makes the contours a little more visible.

3.3 *Sign Predictor*

If there is to be no further coding of the digital signal (apart from assigning a constant length word to each output sample), it is more efficient to have the number of quantizer output levels equal to a power of two. Thus, the improvement in quality obtained using nine output levels would be negated if four bits instead of three had to be assigned to each picture sample.

Now we describe a simply implemented scheme (which may also be used with differential quantizers having analogue integrators) for reducing the required number of levels by one. The scheme enables the nine-level setting to be used with a channel transmitting at the rate of three bits per picture element.

The probability of having the largest level (level number four) preceded by a level of the opposite sign is small. One factor tending to reduce this probability is the smoothing provided by the normalizing filter at the input to the differential quantizer.* Consequently, the sign of an outside level can be predicted fairly accurately by assuming that it is the same as the previous sign. If the prediction is wrong, a number three level, rather than a number four level, is used and will thus have the correct sign.

Thus, a level is effectively eliminated since instead of indicating that the fourth positive or the fourth negative level has occurred (one of two possible events), it is only necessary to indicate to the decoder that the fourth level has occurred (one event) and the decoder then assigns the sign of the previous sample. The signal is modified in this way immediately after the classifier stage; this modification is best regarded as an adjunct to classification. Thus, the encoder and decoder keep in track; in the event of an outside level being preceded by a level of the opposite sign, the slope capability of the encoder and decoder is reduced. For most pictures it is difficult to detect any

---

* The input filter is 3 dB down at 670 kHz and 14 dB down at half the sampling frequency (1 MHz). The output filter is 6 dB down at 1 MHz, giving an overall attenuation of 20 dB at 1 MHz.

change in picture quality when level elimination is used with nine levels. It is simple to implement—in fact it requires only an additional flip-flop and two gates.

3.4 *Dither*

3.4.1 *Straight Quantization with Dither*

In quantizers used for straight PCM encoding, contouring becomes obvious if less than about 100 levels are used. However, small amplitude, high frequency waveforms may be added to the coarsely quantized signal to reduce the visibility of the contours. The penalty associated with adding dither (as these added waveforms will be called) is that the background noise level in the picture is increased slightly. The task in designing dither waveforms is to select that waveform which has the minimum visibility and hence disturbs the picture the least. In the past, random and pseudorandom waveforms have been investigated,[5,7] but more recent calculations have been made with deterministic waveforms (or patterns) indicating their superiority.[6]

3.4.2 *Differential Quantization With Dither*

Dither can be used to advantage in the digital differential quantizer. The quantizer levels are generally set so that contouring is not detectable; but if dither is used, the level spacing can be expanded to either enable the number of levels to be reduced (say from nine to seven) or improve the picture quality by reproducing edges more sharply.

For a certain quantizer configuration, the design of the dither waveform becomes identical to the design for ordinary quantizers.[6] Figure 4 shows this configuration, which has a representative level at zero and the first pair of representative levels set at twice the value of the first pair of decision levels. On the other hand, dither with a decision level at zero produces complex multilevel output waveforms and is not considered here.

Dither may be applied in two ways; the design of the waveform depends on the way it is applied. In the first way it is just added at the input. In the second way, besides adding the waveform at the input it is subtracted from the output. For random uncorrelated waveforms Roberts has shown that for random dither the addition-subtraction technique is superior to addition alone[5]—the variance of the output waveform is reduced by one half.

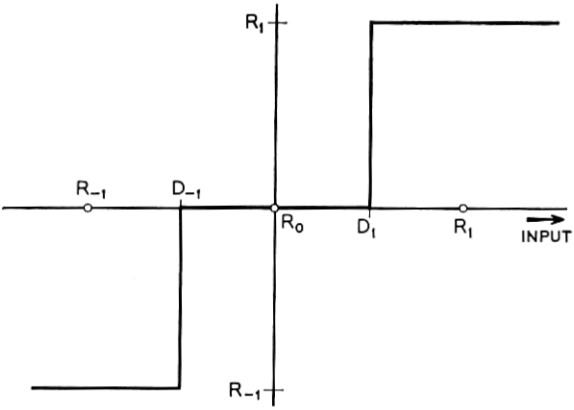

Fig. 4 — Configuration of quantizer with representative level at zero.

However, for a deterministic waveform it has been predicted that there is little difference in the visibility of waveforms designed for addition and waveforms designed for addition-subtraction.[6] This prediction has been tested using a four-amplitude (four-step) pattern with the sequence 1, 3, 2, 4 applied to one dimension, the vertical dimension (Fig. 5). Notice, that although the added waveform is written 1, 3, 2, 4, it has a mean of zero and the levels are positioned uniformly within the quantizing interval as shown in the figure. Theoretically, this is the best one-dimensional, four-step dither waveform for the addition method. The design technique for the addition-subtraction waveform is different from the design for the addition waveform and, in fact, there are two waveforms that have minimum visibility; they are sequence 1, 3, 2, 4 (as for the addition method) and sequence 1, 2, 4, 3.

3.4.3 *Result with Dither*

We now describe the results obtained by adding the waveforms to two particularly sensitive types of display. The first display is a ramp applied in the horizontal direction (Fig. 6); the second is a picture with large flat areas (Fig. 3). For the addition method, the waveform introduces a small amount of frame flicker as a result of interlace since the component added to the first and third lines (field one) is not equal to the component added to the second and fourth lines (field two). Although the add-subtract method does appear to give a smoother looking display, the difference is very slight.

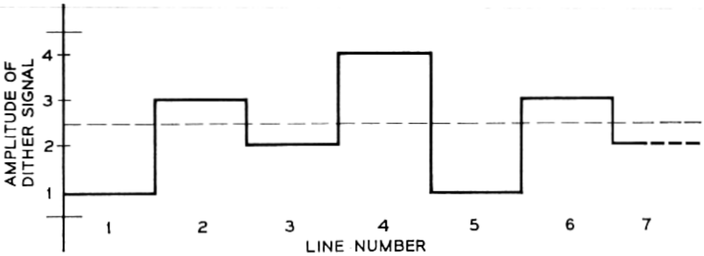

Fig. 5 — Four-step dither waveform.

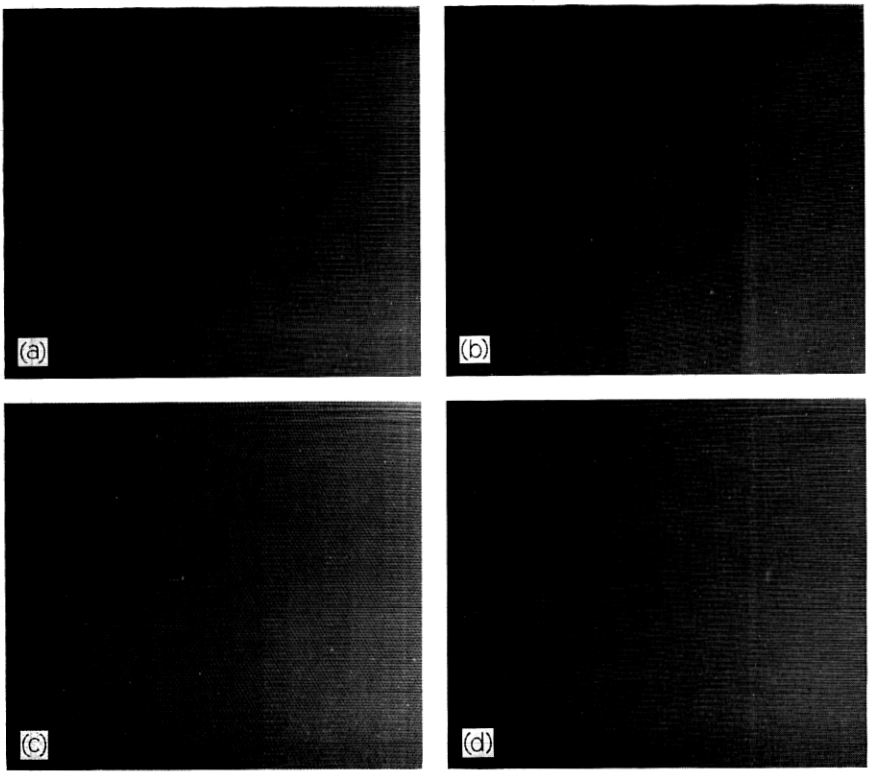

Fig. 6 — Ramp signal processed by digital differential quantizer: (a) Original analogue signal, (b) processed signal—no dither, (c) processed signal—dither added in vertical direction with four-step pattern, and (d) processed signal—dither added in vertical and horizontal directions with 4 × 4-step pattern. (The scan lines and printing screen cause moiré patterns that are not in the originals.) Glossy prints of this figure can be obtained by writing to the authors.

The 1, 2, 4, 3 pattern, as expected, gave the same results as the 1, 3, 2, 4 pattern for the addition-subtraction method. For the addition method, frame flicker was eliminated with a 1, 2, 4, 3 pattern but the output waveform was more visible since the dither waveform is less suitable.

A two dimensional 4 × 4 step dither waveform was generated for the horizontal and vertical dimensions. Table II shows the pattern. With reference to Table II, the sequence 1, 9, 3, 11, 1, 9, 3, 11 . . . was added to the first line; 14, 6, 16, 8, 14, 6, 16, 8 . . . was added to the second line, and so on. Thus 1, 9, 3, 11 . . . would be added again to the fifth line.

Since we have interlace, line one is in a different field than line two. The average contribution to each field can be found by adding the numbers in each line. The first and third lines contribute to the first field and have a total of $24 + 28 = 52$. The second and fourth lines contribute to field two and have a total of $44 + 40 = 84$. Hence, the average contribution to field one is not equal to the average contribution to field two and a small amount of 30 Hz flicker results.

The pattern in Table III does average out over a frame and hence would produce no flicker but was not investigated at this time. Figure 3d shows the improvement obtained by using the dither signal of Table II, compared with not using dither (Figure 3b). The photographs indicate fairly accurately the improvement due to dither, since this particular dither waveform does not rely on time averaging.

An attempt was made to quantitatively assess the improvement resulting from dither. Theoretical predictions were that four-step interpolation would reduce the visibility of contours by 7.1 dB and 4 × 4-step dither by 16.8 db.[6] The method used to test these figures was to add dither to the quantized ramp display (Fig. 6b). A subject attenuated the displayed picture signal until the visibility of the display without dither was equal to the display with dither. The amount of attenuation was then recorded. The viewing distance was

TABLE II—TWO DIMENSIONAL 4 × 4-STEP DITHER WAVEFORM WITH UNBALANCED FIELD CONTRIBUTIONS

	Horizontal →				Line Totals	
Vertical ↓	1	9	3	11	Field 1	24
	14	6	16	8	Field 2	44
	4	12	2	10	Field 1	28
	15	7	13	5	Field 2	40

36 inches. For the four-step pattern (Fig. 6c), four technical subjects gave an average value of 9 dB with a spread of 4 dB, probably reflecting the difficulty of making a match in the presence of flicker. For the 4 × 4-step pattern (Fig. 6d), three subjects gave an average attenuation of 14.3 dB with a spread of 1 dB. These measurements support the theoretical figures of 7.1 dB and 16.8 dB.

A quantitative comparison with picture material would be more difficult and has not been attempted. The method above is not suitable because attenuating the amplitude of the output signal for the picture without dither produces displays that are quite different in appearance. Qualitatively, the reduction in visibility of contours in low-detail areas of the picture is quite dramatic.

3.5 *Further Coding*

When subsequent coding is permitted, the whole approach to the design of the encoder changes. For example, applying dither in the horizontal direction would lead to a less efficient Huffman code (as would a short integrator time constant) while dither applied vertically or temporally would have little effect.

3.6 *Transmission Errors*

A chief disadvantage of the digital accumulator is that an error in transmission will affect subsequent picture elements until the accumulator is reset. Thus, each error will produce a horizontal streak which starts at the point where the error is made and persists to the right edge of the picture where the accumulator is reset. With an analogue integrator the length of an error streak is commensurate with the time constant of the integrator (time constants as short as six picture elements have been used); however, as mentioned previously, a short time constant has other disadvantages.

The length of an error streak in the digital implementation can be shortened by updating the accumulator during the line. For example,

TABLE III—TWO DIMENSIONAL 4 × 4-STEP DITHER WAVEFORM WITH BALANCED FIELD CONTRIBUTIONS

	Horizontal →				Line Totals	
	1	14	3	16	Field 1	34
Vertical	10	5	12	7	Field 2	34
↓	4	15	2	13	Field 1	34
	11	8	9	6	Field 2	34

a full seven-bit PCM signal may be transmitted halfway through each line. The accumulator at the receiver could be updated to this value thus truncating any error which may have occurred. Of course, this process could be repeated more often (with a consequent reduction in information transmission efficiency) for a high transmission error rate. This method of reducing the effect of errors is analogous in a way with shortening the time constant of an analogue integrator without the same disadvantage.

The precision of digital integration leads to another method for reducing the effect of errors. Assume that overflow or underflow of the transmit accumulator is inhibited. Then, when the coder at the transmitter is reset to a pre-assigned value at the end of a block of data (say a line), the decoder at the receiver should recover to the same value. If it does not, transmission errors have occurred. Errors that would escape detection in this way are self-correcting errors, for example, level "a" is received as level "b" followed by level "b" received as level "a". In practice, the probability of an error being of this type would be small.

When a block of data is detected as being in error it can be replaced by an estimate of that block. Now since there is a large amount of unexploited redundancy in a television signal, a reasonable estimate of the line can be made. For example, the previous line could be used or the next and previous lines could be averaged. This technique would probably be satisfactory down to error rates where the probability of obtaining errors in adjacent blocks becomes significant. If the block length was one quarter of a line, error rates of one per line or approximately one in 10^3 might still give a reasonable picture. The degradation would appear as a slight loss in vertical resolution. This proposal has the disadvantage that at least two lines of storage would be required at the receiver.

IV. SUMMARY

We were able to construct a rugged, adjustment-free quantizer with low precision components by using digital techniques at certain points in the path of a differential quantizer. High quality pictures are obtained by using either seven, eight, or nine quantizer output levels. The quality appears somewhat different from the pictures obtained with analogue implementations. The pictures appear less noisy. This is attributed to two facts: (*i*) the values of the quantizer output levels are assigned digitally allowing the larger levels to be set at an exact multiple of the smallest level; (*ii*) integration is also performed digitally, resulting in a virtually infinite time constant.

A simple technique was described whereby the number of quantizer output levels required to be transmitted can be reduced by one. Thus, a nine-level picture, which gives an improvement in picture quality over an eight-level picture can be transmitted at the rate of three bits per picture element.

By expanding the quantizing scale (spacing the levels further apart), edges were reproduced more sharply; however, contouring (exactly as encountered in straight quantization) becomes obvious in low-detail areas of the picture. By adding specially designed dither waveforms to the input signal, contours were "washed out" at the expense of a very small increase in background noise. Changing the input amplitude by ±4 dB produced little change in picture quality.

Because of the long effective time constant of the digital integrator, transmission errors are more visible than in differential quantizers employing a short integrator time constant. Two methods for reducing the visibility of such errors were discussed.

In a visual communication system, a coder must be reliable and have a long adjustment-free life under a wide variety of environmental conditions. Further, a decoder must be capable of working with every encoder in the system. The digital differential quantizer is ideally suited to such a situation.

V. ACKNOWLEDGMENT

We thank Miss M. Arthur for her assistance in constructing and testing the equipment.

REFERENCES

1. Cutler, C. C., "Differential Quantization of Communication Signals," Patent No. 2,605,361, applied for June 29, 1950, issued July 1952.
2. Graham, R. E., "Predictive Quantizing of Television Signals," IRE Wescon Conv. Rec., *2*, part 4 (1958), pp. 147–157.
3. Limb, J. O., "Source-Receiver Encoding of Television Signals," Proc. IEEE *55*, No. 3 (March 1967), pp. 364–379.
4. Goodall, W. W., "Television by Pulse Code Modulation," B.S.T.J., *30*, No. 1 (January 1951), pp. 33–49.
5. Roberts, L. G., "Picture Coding using Pseudo Random Noise," IRE Trans. on Inform. Theory, *IT-8*, No. 2 (February 1962), pp. 145–154.
6. Limb, J. O., "Design of Dither Waveforms for Quantized Visual Signals," this issue, pp. 2555–2582.
7. Limb, J. O., "Coarse Quantization of Television Signals," Australian Telecommunications Research *1*, Nos. 1 and 2 (November 1967), pp. 32–42.