

Optimum Quantizer Design Using a Fixed-Point Algorithm

By A. N. NETRAVALI and R. SAIGAL

(Manuscript received June 10, 1976)

A fixed-point algorithm has been used to obtain the parameters (i.e., decision and representative levels) of an "optimum" quantizer that minimizes a quite general distortion measure, subject to an entropy constraint on its output. Construction of the algorithm starts with a point-to-set mapping whose fixed point satisfies the well-known Karush-Kuhn-Tucker conditions necessary for a local extremum. A computer program is then used to determine a fixed point of this mapping. Several examples are solved, and correspondence with the existing results in the literature is pointed out. Finally, as conjectured, the growth of the computations as a function of dimensionality n (n : number of representative levels) is found to be of the form $a \cdot n^b$ where a is a positive constant and $1.5 \leq b \leq 2.0$.

I. INTRODUCTION

Simple quantization¹⁻³ has been and continues to be a popular method of digitizing analog signals. The relative ease with which quantizers can be implemented in hardware and their near optimum performance has made them withstand the challenge from several new coding schemes.⁴⁻⁶ Universal use of quantizers has naturally spurred a significant activity in optimizing their performance, some of which is summarized in the next few paragraphs. Our objective in this paper is to show how the problem of obtaining the parameters of an optimum quantizer can be converted to the problem of obtaining fixed points of a suitably constructed mapping and then to use a fixed-point algorithm to solve the problem numerically.

Quantizers have been optimized based on several criteria. In order to discuss these in relation to the problem considered in this paper, we describe the basic quantizer equations. Given a scalar random variable T with probability density $p(t)$, a quantizer Q is a map $Q(t) = y_i$ whenever $x_i \leq t < x_{i+1}$, where x_i , $i = 1, \dots, N+1$ and y_i , $i = 1, \dots, N$ are the decision and representative levels of the quantizer, respectively. The performance of the quantizer is judged generally in terms of two quantities:

the distortion

$$D = \sum_{i=1}^N \int_{x_i}^{x_{i+1}} g(t - y_i) \times f(t) dt, \quad (1)$$

and the entropy

$$\mathcal{E} = - \sum_{i=1}^N (\log_2 p_i) \times p_i, \quad (2)$$

where g is a nonnegative function and f is a nonnegative weighting function that weights the quantization noise and

$$p_i = \int_{x_i}^{x_{i+1}} p(t) dt.$$

Optimum quantizers choose their parameters $\{x_i\}$, $i = 2, \dots, N$ and $\{y_i\}$, $i = 1, \dots, N$ (given the end points x_1, x_{N+1}) to optimize a certain combination of D and $\mathcal{E}$.

Most quantization literature uses the weighting function f to be the same as the probability function p , although in some applications⁷⁻⁹ a different weighting function performs better. Most of the earlier work is concerned with minimizing D for a given number of levels. Panter and Dite¹⁰ have used $g(\cdot) = |\cdot|^r$ ($r > 0$) and obtained an approximate optimum quantizer as one in which each of the quantizing intervals $[x_i, x_{i+1}]$ makes an equal contribution to the integral of $|t - y_i|^r$. This allowed them to choose the quantizer parameters for large N . Lloyd¹¹ and Max¹² have developed an algorithm for $r = 2$, which corresponds to minimizing the mean square error. Bruce¹³ has used dynamic programming to solve the same problem in slightly more generality by taking a general function $g(\cdot)$. Simpler suboptimal algorithms and bounds on the performance of the quantizers have been obtained by Roe,¹⁴ Algazi,¹⁵ and Zador.¹⁶

Representation of the quantizer output by a variable length code allows reduction of the average bit rate of the quantizer when p_i varies with i . Use of Huffman code¹⁷ makes the average bit rate approach the entropy of the quantizer output. Thus, the problem of designing an optimum¹⁸⁻¹⁹ quantizer can be reformulated as that of obtaining the decision and representative levels to minimize D subject to a constraint on the entropy. Gobllick and Holsinger have considered this problem for uniform quantizers and have concluded that for gaussian density, for $r = 2$, and for the same distortion, the entropy of the output of the uniform quantizer is higher than the theoretical lower bound based on the rate distortion theory by about $\frac{1}{4}$ bit. Uniform quantizers are also good in an asymptotic sense, since they are optimum for a large number of levels.²¹ Moreover, for Laplacian densities, as shown by Berger,²⁰

uniform quantizers are optimum for any value of entropy. A different type of distortion measure has been considered by Elias.²²

The problem we consider is that of obtaining the parameters of quantizers such that D is minimized for a given constraint on the entropy. Although the approach taken here is suitable for a general distortion measure of eq. (1), we consider only the case of $g(\cdot) = (\cdot)^2$, mainly to compare our results to those in the literature. In the next section, we present the necessary conditions that the optimum quantizer must satisfy for a local extremum. Then, in Section III, we construct a point-to-set mapping such that its fixed point satisfies the necessary conditions for a local extremum of our problem. A description of the algorithm is then presented for completeness. In Section IV, we present the results of use of this algorithm for uniform, Laplacian, and gaussian densities. The distortion-entropy curves are presented for each case. We also present a surprising observation on the growth of computations as a function of dimensionality (i.e., the number of quantizer parameters to be optimized).

II. FORMULATION OF THE PROBLEM AND NECESSARY CONDITIONS

Using $g(\cdot) = (\cdot)^2$, the distortion of eq. (1) becomes

$$D = \sum_{j=1}^N \int_{x_j}^{x_{j+1}} (t - y_j)^2 f(t) dt. \quad (3)$$

Then the problem is to obtain $\{x_j\}$, $j = 2, \dots, N$, $\{y_j\}$, $j = 1, \dots, N$ such that they minimize D subject to $\mathcal{E} \leq K$, for a given N . The necessary conditions from the Karush-Kuhn-Tucker theory²³ are that there exists a $\lambda \geq 0$ such that

$$\nabla D(x) + \lambda \nabla \mathcal{E}(x) = 0, \quad (4)$$

where x is a vector of quantizer parameters and ∇ denotes the gradient. For the parameters $\{y_j\}$, since $\mathcal{E}$ is independent of $\{y_j\}$, (4) becomes

$$y_j = \int_{x_j}^{x_{j+1}} t f(t) dt / \int_{x_j}^{x_{j+1}} f(t) dt, \quad j = 1, \dots, N. \quad (5)$$

This implies that the representative levels can be obtained explicitly by knowing the decision levels and therefore they do not add to the dimensionality of the problem. Also, the other necessary conditions are

$$\mathcal{E} \leq K$$

and

$$\lambda(\mathcal{E} - K) = 0. \quad (6)$$

III. FIXED-POINT APPROACH

In this section, we formulate the quantization problem as a fixed-point problem and give a general description of the algorithm that

solves this problem. This algorithm is based on the theory of complementary pivoting.²⁴

Given a point-to-set mapping Γ [i.e., to each point x in R^n it associates a subset $\Gamma(x)$ of R^n], a fixed point of such a mapping is a point x such that $x \in \Gamma(x)$. We show that the problem of finding the parameters of the optimal quantizer can be formulated as a problem of finding a fixed point of a certain point-to-set mapping.

3.1 Fixed-point formulation

Let ∇D and $\nabla \mathcal{E}$ be the gradient vectors of the distortion D and entropy $\mathcal{E}$, respectively. Then, consider the following point-to-set mapping:

$$\Gamma(x) = \begin{cases} x - \{\nabla D(x)\} & \mathcal{E}(x) < K \\ x - \text{hull}\{\nabla D(x), \nabla \mathcal{E}(x)\} & \mathcal{E}(x) = K, \\ x - \{\nabla \mathcal{E}(x)\} & \mathcal{E}(x) > K \end{cases} \quad (7)$$

where $\text{hull}\{E\}$ is the smallest convex set containing E ; i.e., the convex hull of E , and $x - A = \{x - y : y \in A\}$ for a set A in R^n . Note that the mapping as defined is upper semicontinuous* (u.s.c.) and the set $\Gamma(x)$ is convex for each x . As we subsequently see, these properties are needed if the algorithm is to find a fixed point of Γ .

We now show that a fixed point of this mapping satisfies the necessary conditions of Section 2.

Theorem: Let $x \in \Gamma(x)$. Then, if $\mathcal{E}(x) \leq K$, x satisfies the necessary conditions of Section 2. Otherwise, x is a local minimizer of $\mathcal{E}(x)$.

Proof: We construct the required λ and show that (6) is satisfied. Since $x \in \Gamma(x)$ and $\mathcal{E}(x) \leq K$, we have two cases:

Case (i): $\mathcal{E}(x) < K$. Let $\lambda = 0$ and, since $0 \in \{\nabla D(x)\}$, $\nabla D(x) + \lambda \nabla \mathcal{E}(x) = 0$, satisfying (6). Note that $\lambda[\mathcal{E}(x) - K] = 0$.

Case (ii): $\mathcal{E}(x) = K$. Then, as $0 \in \text{hull}\{\nabla D(x), \nabla \mathcal{E}(x)\}$, there exist $\lambda_1 + \lambda_2 = 1$, $\lambda_1 \geq 0$, $\lambda_2 \geq 0$ such that

$$\lambda_1 \nabla D(x) + \lambda_2 \nabla \mathcal{E}(x) = 0. \quad (8)$$

Now, in case $\lambda_1 \neq 0$, letting $\lambda = \lambda_2/\lambda_1 \geq 0$, (4) is satisfied, and $\lambda[\mathcal{E}(x) - K] = 0$. In the contrary case, a constraint qualification would be violated.

In case $x \in \Gamma(x)$ and $\mathcal{E}(x) > K$, then, since $0 \in \{\nabla \mathcal{E}(x)\}$, $\nabla \mathcal{E}(x) = 0$ and we have a local minimizer of $\mathcal{E}(x)$. If $\mathcal{E}(x)$ were a convex function, our problem has no feasible solution [i.e., an x such that

* A mapping Γ is u.s.c. if, for any two sequences $\{x^k\}$, $\{y^k\}$ such that $x^k \rightarrow x$, $y^k \in \Gamma(x^k)$, and $y^k \rightarrow y$, we have $y \in \Gamma(x)$.

$\mathcal{E}(x) \leq K]$. In the contrary case, we would conclude the algorithm has failed.

3.2 Description of the algorithm

In this section we give a brief description of the algorithm that computes fixed points of point-to-set mappings. Before going into the details of the algorithm, we introduce some notation.

Given a set C in R^n , and a point-to-set mapping Γ , by $\Gamma(C)$ we represent the set $\bigcup_{x \in C} \Gamma(x)$. Also, given a one-to-one linear mapping r , we say a set C is

- (i) Γ —complete if $0 \in \text{hull} \{ \Gamma(C) \}$,
- (ii) r —complete if $0 \in \text{hull} \{ r(C) \}$, and
- (iii) $\Gamma \cup r$ —complete if $0 \in \text{hull} \{ \Gamma(C) \cup r(C) \}$.

The significance of Γ -complete sets is the following: in case Γ is u.s.c. and $\Gamma(x)$ is convex for each x , a sequence C_i , $i = 1, 2, \dots$ of Γ -complete sets whose diameter approaches 0 as i approaches ∞ , converges to a fixed point of Γ (see, for example, Refs. 25–27). The fixed-point algorithms are designed to find such a sequence of Γ -complete sets.

These algorithms work with sets C that are simplexes of appropriate dimension. (An n -dimensional simplex is a convex body obtained by taking the convex hull of $n + 1$ affinely independent points in n -space. A two-dimensional simplex is a triangle; a three-dimensional simplex is a tetrahedron.) They start with a unique r -complete simplex and generate a sequence of $\Gamma \cup r$ -complete simplexes that terminate with a Γ -complete simplex. There are essentially two basic algorithms that can be used to generate a sequence of Γ -complete simplexes of decreasing diameters. They are the *restart method* of Merrill²⁷ and the *continuous deformation method* of Eaves and Saigal.²⁶ A study of both these methods can be found in Saigal.^{28–30}

We now discuss an application of the algorithm. A real number $d > 0$ is chosen. Then the space $R^n \times [0, d]$ is triangulated (i.e., each point in the space lies in an $(n + 1)$ -dimensional simplex, and these simplexes overlap only on their boundaries) such that the vertices of the triangulation are only in the set $R^n \times \{d/2^k\}$, $k = 0, 1, \dots$. In addition, the diameter of each n -dimensional face of each $(n + 1)$ -dimensional simplex that lies in $R^n \times [d/2^{k+1}, d/2^k]$ is at most $d/2^k$. Now, an arbitrary starting point x_0 is chosen. We then define

$$r(x) = -x + x_0, \quad (9)$$

which is a one-to-one linear mapping.

The sequence of $\Gamma \cup r$ -complete simplexes is then generated as

follows:

Step 1: Start with an r -complete simplex in the triangulation that contains (x_0, d) . The triangulation is arranged in such a way that there is a unique such simplex, and that this simplex has exactly one vertex in $R^n \times \{d/2\}$, and $(n+1)$ vertices in $R^n \times \{d\}$. The entering vertex is the one in $R^n \times \{d/2\}$. Design the labeling function L on the vertices of the triangulation with

$$L(x, t) = \begin{cases} z - x & \text{for some } z \in \Gamma(x) \text{ if } t < d \\ x^0 - x & \text{if } t = d \end{cases} \quad (10)$$

Step 2: Find the label on the entering vertex.

Step 3: Find a new $\Gamma \cup r$ -complete simplex that includes the entering vertex, in place of some vertex of the older simplex. This is equivalent to the basic pivot operation of the simplex method.³¹

Step 4: Find the other $(n+1)$ -dimensional simplex that contains the new $\Gamma \cup r$ -complete simplex found in Step 3, and determine the entering vertex.

Step 5: If the entering vertex is outside $R^n \times \{d/2^K, d\}$, stop. The earlier $\Gamma \cup r$ -complete simplex is actually Γ -complete. Otherwise, go to Step 3.

Having found a Γ -complete simplex τ , say, whose vertices are $V^1, V^2, \dots, V^{n+1}$, where $V^i = (v^i, d_i)$, $i = 1, \dots, n+1$, we have determined points $z^i \in \Gamma(v^i)$ and a $\lambda = (\lambda_1, \dots, \lambda_{n+1}) \geq 0$ such that

$$\sum_{i=1}^{n+1} \lambda_i z^i = 0 \quad (11)$$

$$\sum_{i=1}^{n+1} \lambda_i = 1$$

has a solution. In this case, we say that the point x determined by

$$x = \sum_{i=1}^{n+1} \lambda_i v^i \quad (12)$$

is an approximate fixed point (for justification, see Ref. 26).

Since the stopping criterion at Step 5 requires that we generate a vertex in $R^n \times \{d/2^K\}$, we have generated a sequence of Γ -complete sets C_i , the last one of diameter less than $d/2^K$, and have thus found a reasonable solution.

The procedure for triangulation $R^n \times (0, d]$ generally used is called *J3* in the literature. For a more detailed description of this algorithm, the reader is referred to Ref. 32.

IV. EXAMPLES

In this section, we discuss some examples that we solved using the algorithm described in the previous section. Three of the four examples had $f(\cdot) = p(\cdot)$ corresponding to mean square quantization error as a measure of distortion. The fourth example, on the other hand, uses a different weighting of the quantization noise; it is motivated by the problem of quantizer design for simple element differential coding of picture signals.⁹ The examples are:

- (i) $f(x) = p(x) = \frac{1}{32}, \quad -16 \leq x \leq +16$
 $= 0 \quad \text{otherwise}$
- (ii) $f(x) = p(x) = \frac{1}{\alpha} e^{-\alpha|x|}, \quad -\infty < x < +\infty, \alpha = 0.1$
- (iii) $f(x) = p(x) = \frac{\exp(-u^2/2\alpha)}{\sqrt{2\pi\alpha}}, \quad -\infty < x < +\infty, \alpha = 1 \quad (13)$
- (iv) $f(x) = \frac{1}{\beta} e^{-\beta|x|}; \quad p(x) = \frac{1}{\alpha} e^{-\alpha|x|}, \quad -\infty < x < +\infty$
 $\alpha = 0.18, \beta = 0.1; \quad \text{and} \quad \alpha = 0.1, \beta = 0.065.$

Due to symmetry of functions $f(\cdot)$ and $p(\cdot)$, the optimum quantizers are symmetric and, for simplicity therefore, quantizers were con-




Fig. 1—Quantizer performance for uniform density. Minimum mean square error (MMSE) is plotted against entropy for a fixed number of levels (N). Only odd-level quantizers are considered. For each fixed number of levels, MMSE decreases with entropy up to a certain point, after which there is no further decrease in mean square error.




Fig. 2—Quantizer performance for Laplacian density.

strained to be symmetric. Also, without loss of any generality, only quantizers having odd numbers of levels were considered. In each case, several problems were solved by varying the entropy constraint and the number of levels. The number of levels were varied from 3 to 21, and the entropy constraint was varied from 1.0 bit to the largest possible bits using a particular number of levels.

Results of these simulations are given in Figs. 1 through 5. In these




Fig. 3—Quantizer performance for gaussian density.




Fig. 4—Quantizer performance for Laplacian density and exponential weighting. Quantization noise is weighted by $1/|\beta| \exp(-\beta|x|)$, whereas the probability density is taken to be $1/|\alpha| \exp(-\alpha|x|)$. Such situations arise in quantization of the prediction errors in predictive coding of the television signals: $\alpha = 0.1, \beta = 0.065$.

figures the distortion is plotted logarithmically on y -axis and the entropy is plotted linearly on x -axis in bits. Alternate solid and broken lines are shown for different values of quantizer levels. For a given number of levels, the minimum distortion decreases approximately exponentially with respect to the entropy up to a certain point and




Fig. 5—Quantizer performance for Laplacian density and exponential weighting: $\alpha = 0.18, \beta = 0.1$.

then the entropy constraint is not operative any longer, and consequently the distortion remains a constant. These are indeed Lloyd-Max^{11,12} quantizers that minimize the distortion for a given number of levels with no constraint on the entropy. The distortion versus entropy curves are lower bounded by the following functions:

$$\begin{aligned}
 \text{Example (1)} \quad D &= 85.3 \exp(-1.39E) \\
 \text{Example (2)} \quad D &= 196.18 \exp(-1.32E) \\
 \text{Example (3)} \quad D &= 1.40 \exp(-1.39E) \\
 \text{Example (4a)} \quad D &= 62.50 \exp(-1.31E) \\
 \text{Example (4b)} \quad D &= 176.71 \exp(-1.24E).
 \end{aligned} \tag{14}$$

In the case of uniform densities, the optimum quantizer is non-uniform whenever the entropy constraint is operative, but when the entropy constraint is too large and inoperative, the optimum quantizers are uniform. Laplacian densities, on the other hand, always have uniform quantizers as the optimum quantizers. This has been shown by Berger.²⁰ In the case of gaussian density, the optimum quantizers were not uniform; however, a comparison of our results with those given by Goblick and Holsinger¹⁸ indicates that, although nonuniform quantizers perform better than uniform quantizers, the differences in the performance of the two are somewhat small. This conclusion has also been reached by Wood¹⁹ and Berger.²⁰ The case of an exponential weighting function falling slower than the probability density function arises in quantization of the prediction error in a simple element differential coding of picture signals. In this case, the density of the prediction error is approximately Laplacian, whereas the perceptual visibility^{9,33} of the quantization noise may be approximated by an exponential function decaying somewhat slower than the probability density. The distortion-entropy curves for this case show larger improvement (that is, for a given entropy the distortion decreases much more than in the previous examples) as the number of levels is increased. Also, the optimum quantizers are nonuniform. Improvement in their performance over that of the uniform quantizers is more significant than in the previous examples. It is interesting to note that our algorithm can solve Lloyd-Max problem trivially by setting the entropy constraint to a very high value. This algorithm was also used in other applications related to adaptive quantization³⁴ of picture signals. The problems in this case were such that they had uniform (constant) weighting functions and two-sided exponentials as the density functions. The resulting quantizers had interesting structure and were used quite successfully.

4.1 Computational effort as a function of n

The increase in computational effort as a function of dimension, n , of the problem is important in the study of algorithms. In the case of fixed-point algorithms, Saigal²⁸ had speculated that different triangulations would have different effects on this growth and he was the first to propose a measure to describe it. For the triangulation employed in our experimentation, his measure predicted the growth rate of the number of iterations as n^2 . Subsequently, Todd³⁵ refined his measure to predict an "average" growth rate of the iterations as $n^{\frac{1}{2}}$. The measure of Saigal, in some sense, predicts the "worst case" behavior.

The computational experiments in Section IV were ideally suited to test the theoretical predictions of Refs. 28 and 35, since the dimension of the problem was increased in a regular manner, the starting points were chosen in a regular way, and the problems of dimensions varying between 1 and 10 were solved. A number of results for various entropy values were plotted on the log-log paper. A representative plot is given in Fig. 6. It is seen that the experimental points lie on a straight line. The slope of these lines for different cases was a function of the entropy constraint and the probability density used and varied from 1.55 to 1.88, which is between 1.5 predicted by Todd³⁵ and 2 predicted by Saigal.²⁸

Thus, we can conclude, with a high degree of certainty, that the number of iterations of the algorithm to solve a problem of dimension




Fig. 6—Growth of computations vs dimensionality. Number of representative levels N is two times the dimensionality n plus 1. Straight line drawn is the minimum mean square error fit to the observations shown by Δ .

n would require an^b for some b between 1.5 and 2.0. Since each iteration requires $O(n^2)$ multiplications and at most one evaluation of the function, the number of function evaluations is bounded by $O(n^2)$ and multiplications by $O(n^4)$.

V. CONCLUSIONS

A fixed-point formulation has been developed to minimize the distortion, using a fairly general distortion measure, with respect to parameters of a quantizer under an entropy constraint on the quantized output. A point-to-set mapping is first developed whose fixed point satisfies the necessary conditions for a local extremum. Then a computer program is developed to compute its fixed points. Several examples are solved to show the usefulness of the algorithm. Finally, the rate of growth of the computations used by the algorithm as a function of the dimensionality of the problem is also discussed.

VI. ACKNOWLEDGMENTS

The authors would like to thank Barry Haskell for many inspiring discussions throughout the course of this work and Ivan Cermak and Ron Olsen for the use of the CDC-6600 computer.

REFERENCES

1. Clavier, Panter, and Greig, "PCM Distortion Analysis," *Elec. Eng.*, **66** (1947), pp. 1110-1122.
2. B. M. Oliver, J. R. Pierce, and C. E. Shannon, "The Philosophy of PCM," *Proc. IRE*, **36** (November 1948), pp. 1324-1331.
3. B. Widrow, "A Study of Rough Amplitude Quantization by Means of the Nyquist Sampling Theorem," *Trans. IRE Prof. Group Circuit Theory*, **CT-3** (1956), pp. 266-276.
4. T. Berger, F. Jelinek, and J. K. Wolf, "Permutation Codes for Sources," *IEEE Trans. Inform. Theory*, **IT-18** (1972), pp. 160-169.
5. J. Anderson, "Algorithms for Tree Source Coding with a Fidelity Criteria," Ph.D. dissertation, Cornell University, Ithaca, N. Y., 1972.
6. R. J. Dick, "Tree Coding for Gaussian Sources," Ph.D. dissertation, Cornell University, Ithaca, N. Y., 1973.
7. J. O. Limb, "Source-Receiver Encoding of Television Signals," *Proc. IEEE*, **55**, No. 3 (March 1967), pp. 364-380.
8. F. W. Mounts, A. N. Netravali, and B. Prasada, "Design of Quantizers for Real-Time Hadamard Transform Coding of Pictures," unpublished work.
9. A. N. Netravali, "On Quantizers for DPCM Coding of Picture Signals," unpublished work.
10. P. F. Panter and W. Dite, "Quantization Distortion in Pulse-Count Modulation with Nonuniform Spacing of Levels," *Proc. IRE*, **39** (1951), pp. 44-48.
11. S. P. Lloyd, "Least Squares Quantization in Pulse Count Modulations," unpublished work.
12. J. Max, "Quantizing for Minimum Distortion," *IEEE Trans. Inform. Theory*, **IT-8** (1960), pp. 7-12.
13. J. D. Bruce, "Optimum Quantization," Research Laboratory of Electronics, M.I.T., Tech. Rept. 429, March 1965.
14. G. M. Roe, "Quantizing for Minimum Distortion," *IEEE Trans. Inform. Theory* (correspondence), **IT-10** (1964), pp. 384-385.
15. V. Algazi, "Useful Approximations to Optimum Quantization," *IEEE Trans. Commun. Technol.*, **COM-14** (1966), pp. 297-301.

16. P. Zador, "Development and Evaluation of Procedures for Quantizing Multivariate Distribution," Ph.D. dissertation, Stanford University, Stanford, Cal., 1964.
17. D. Huffman, "A Method for the Construction of Minimum Redundancy Codes," Proc. IRE, 40 (September 1952), pp. 1098-1101.
18. T. J. Goblick, Jr. and J. L. Holsinger, "Analog Source Digitization: A Comparison of Theory and Practice," IEEE Trans. Inform. Theory, IT-13 (April 1967), pp. 323-326.
19. R. C. Wood, "Optimum Quantization," IEEE Trans. Inform. Theory, IT-15 (March 1969), pp. 248-252.
20. T. Berger, "Optimum Quantizers and Permutation Codes," IEEE Trans. Inform. Theory, IT-18 (November 1972), pp. 759-765.
21. H. Gish and J. N. Pierce, "Asymptotically Efficient Quantizing," IEEE Trans. Inform. Theory, IT-14 (September 1968), pp. 676-683.
22. P. Elias, "Bounds on Performance of Optimum Quantizers," IEEE Trans. Inform. Theory, IT-16 (March 1970), pp. 172-184.
23. W. I. Zangwill, *Nonlinear Programming*, Englewood Cliffs, N. J.: Prentice Hall, 1969.
24. C. E. Lemke, "Bimatrix Equilibrium Points and Mathematical Programming," Management Science, 11 (1965), pp. 681-689.
25. B. C. Eaves, "Homotopies for Computation of Fixed Points," Math. Prog., 3, No. 1 (1972), pp. 1-22.
26. B. C. Eaves and R. Saigal, "Homotopies for Computation of Fixed Points on Unbounded Regions," Math. Prog., 3, No. 2 (1972), pp. 225-237.
27. O. H. Merrill, "Applications and Extensions of an Algorithm that Computes Fixed Points of Certain Upper Semi-Continuous Point to Set Mappings," Ph.D. dissertation, University of Michigan, Ann Arbor, 1972.
28. R. Saigal, "Investigations into the Efficiency of Fixed Point Algorithms," to appear in *Fixed Points—Algorithms and Applications*, S. Karamardian, ed, New York: Academic Press.
29. R. Saigal, "Paths Generated by Fixed Point Algorithms," unpublished work.
30. R. Saigal, "On the Convergence Rate of Algorithms for Solving Equations that are Based on Complementary Pivoting," unpublished work.
31. G. B. Dantzig, "Linear Programming and Extensions," Princeton: Princeton University Press, 1963.
32. R. Saigal, "Fixed Point Computing Methods," *Encyclopedia of Computer Science and Technology*, M. Dekker, Inc., 1976.
33. J. C. Candy and R. H. Bosworth, "Methods for Designing Differential Quantizers Based on Subjective Evaluations of Edge Busyness," B.S.T.J., 51, No. 7 (September 1972), pp. 1495-1516.
34. A. N. Netravali and B. Prasada, "Adaptive Quantization of Picture Signals Based on Spatial Masking," unpublished work.
35. M. J. Todd, "On Triangulations for Computing Fixed Points," Tech. Report No. 249, Department of Operations Research, Cornell University, Ithaca, New York, March 1975.

