Exhibit G

Dr. Stoffa is a member of the Society of Exploration Geophysicists, the American Geophysical Union, Sigma Xi, European Association of Exploration Geophysicists and the European Association for Signal Processing.

Peter Buhl was born in New Jersey on November 12, 1939. He received the B.S. degree in engineering physics from Lehigh University, Bethlehem, PA, in 1966, and the M. Phil. degree from Columbia University, New York, NY, in 1974.

In 1962 he joined Lamont-Doherty Geological Observatory, Columbia University, to participate in research cruises; he became a Graduate Research Assistant in 1966 and a Research Scientist in 1974. At present he is a Senior Staff Associate with Lamont-Doherty. He is a co-developer of the Lamont-Doherty Geological Observatory Multichannel Seismic Processing System and multiple ship-multichannel seismic experiments. His areas of interest include seismic computer processing, including frequency domain and cepstrum techniques, and wave propagation.

Mr. Buhl is a member and former officer of the ARRAY Processors Users Group. He is also a member of the Society of Exploration Geophysicists, the American Geophysical Union, the European Association of Exploration Geophysicists, and the American Association for the Advancement of Science.

Vector Quantization of Speech and Speech-Like Waveforms

HÜSEYIN ABUT, SENIOR MEMBER, IEEE, ROBERT M. GRAY, FELLOW, IEEE, AND GUILLERMO REBOLLEDO, MEMBER, IEEE

Abstract—An algorithm for the design of vector quantizers that are locally optimum in the sense of minimizing an average quantitative distortion measure is used to design 1 and 2 bit/sample vector quantizers for both real sampled speech and a simulated speech-like autoregressive random process. Both weighted and unweighted squared-error distortion measures are considered.

Several comparisons are made and discussed based on the average distortions of the vector quantization schemes. The results for the simulated speech are compared to mathematical performance bounds from information theory to provide an indication of how nearly globally optimal vector quantization is for such highly correlated sources. A comparison of the results for the real speech and the simulated speech provides a quantitative measure of the accuracy of such models and, hence, of the applicability of information theory bounds and code designs based on probabilistic models. The signal-to-quantization-noise ratios of vector quantizers designed to minimize squared-error distortion are compared to those of several popular speech waveform coding systems of similar rates.

Manuscript received October 14, 1980; revised August 1, 1981 and November 5, 1981. This work was supported in part by the U.S. Air Force Office of Scientific Research under Contract F49620-79-C-0058 and by the Consejo Nacional de Ciencia y Tecnologia of Mexico.

H. Abut was with the Speech Laboratory, National Semiconductor Corporation, Santa Clara, CA 95053. He is now with the Department of Electrical and Computer Engineering, San Diego State University, San Diego, CA 92182.

R. M. Gray is with the Information Systems Laboratory and the Department of Electrical Engineering, Stanford University, Stanford, CA 94305.

G. Rebolledo was with the Information Systems Laboratory, Stanford University, Stanford, CA 94305, and Time and Space Processing, Inc., Santa Clara, CA 95053. He is now with the Division de Estudios Superiores, Facultad de Ingeniera, University of Mexico, Mexico City, Mexico.

I. Introduction

Waveform coding speech compression systems attempt to digitally communicate a good reproduction of the actual speech waveform. Voice coding systems attempt to estimate and communicate a linear model of the speech production process rather than the specific waveform. Voice coding systems permit communication at lower rates, but waveform coders typically provide better quality and more robustness against speaker variations, multiple speakers, and background noise. It is still important, however, to use the minimum possible bit rate and, hence, the smallest possible bandwidth or storage capacity for a given fidelity.

Waveform coders range from simple fixed high rate scalar quantizers (PCM's) and predictive quantizers (DPCM's), to fairly complex low rate adaptive subband coders, transform coders, DPCM, and tree encoding systems typically operating in the range of 8 kbits/s (kilobits per second)–15 kbits/s. A variety of such systems is discussed, for example, in [1]–[10].

Most people would agree that given a rate constraint, the general goal of such digitization schemes is to minimize the distortion between the original and reconstructed speech, but there is no universally accepted definition of distortion. There are, in fact, two fundamentally different approaches to defining and measuring distortion. The first is a subjective or qualitative notion based on relative distortions perceived by listeners in listening tests. While subjective distortion is clearly of fundamental importance in the evaluation of a voice communication system, it is of little help in the actual design and

0096-3518/82/0600-0423$00.75 © 1982 IEEE

analysis of a specific system since no analytical techniques exist for the design of a system that is optimized in the sense of minimizing subjective distortion.

The second approach uses a quantitative distortion measure such as the ubiquitous squared-error distortion, weighted versions thereof, or more complicated spectral distortion measures. Such quantitative measures may result from an approximate quantification of a subjective measure or they may be the result of a mathematical development as in the Itakura-Saito distortion [37]. The advantages of a quantitative measure are numerous. 1) It provides a means for the quantitative comparison of different systems. 2) Given a suitably tractable distortion measure, a system can be designed to provide a minimum or, at least, a local minimum of the average or maximum distortion. 3) Using techniques from information theory, one may be able to characterize the optimal performance for a given class of data compression systems on a specified source, where optimal here means yielding the minimum average distortion subject to the system constraints.

In short, a quantitative distortion measure permits the application of tools and techniques of communication and information theory to the design, analysis, and comparison of data compression schemes.

An additional justification for considering a quantitative distortion measure in the design and analysis of data compression systems is that virtually all such systems perform the actual digitization or quantization step using a minimum distortion or nearest neighbor rule to select a digital representation of an analog quantity. This quantization may be applied directly to the waveform, to an error sequence within a feedback loop; to a companded, transformed, or filtered set of samples; or to some combination. The quantization operation, however, invariably involves a minimum distortion selection. Thus, even though a particular class of systems may be selected based on its subjective quality, it is of both theoretical and practical interest to develop the properties of such systems when they are optimized. The resulting performance provides a benchmark for the comparison of similar systems.

When trying to characterize the optimal performance achievable for a given class of data compression systems, it is often of interest to consider the most general possible class. This provides a more fundamental benchmark for comparison with existing systems, and an indication of what additional improvement might be achievable using less structured (and possibly more complex) systems. The original Shannon mathematical model for a data compression system is that of a block source code—a mapping of consecutive nonoverlapping blocks or frames of source data into consecutive nonoverlapping blocks of channel symbols such as binary numbers. Block source codes can be simply viewed as vector or multidimensional quantizers, and include such systems as transform coders, traditional scalar quantization (PCM), permutation codes, and any other schemes that process parsed information. Information theory has also treated sliding-block or sliding-window codes that operate as possibly nonlinear filters (e.g., DPCM) and both classes have been shown to have an optimal achievable performance (in the limit of long block length or constraint length) given by Shannon's distortion rate function (see, e.g., [11], [2]). Here we focus on the design performance of block codes or vector quantizers. The design of trellis waveform coders for speech using an algorithm similar to the vector quantizer design algorithm treated here is considered in [38].

A general goal might be to find (if possible) the optimum vector quantizer of a given rate and block length (dimension) when used on a particular source. Such a result would indicate achievable performance improvements over special cases such as transform coders. Even if one possesses a complete and accurate probabilistic model of the source, however, such theoretical characterizations do not usually exist. Rate distortion theory bounds are achievable only in the limit of long block lengths. Approximations to the optimal performance may be found for a fixed block length using asymptotic quantizer performance theory if the rate is large enough (see, e.g., [12]-[14]). Since theory can characterize only asymptotic optimal performance values (large block length or large rate), the best available means of finding the performance limits of a class of codes of nonasymptotic rate and block length remains the actual design and test of good codes incorporating optimization techniques and a minimum of structural constraints.

An obvious advantage of popular special cases is the fact that their structural constraints usually make simple and efficient implementation possible. For example, most implementations of vector quantizers consist of a simple (often linear) transformation followed by separate scalar quantizers for the transformed samples [15]-[16], [4], [8]. Such codes are suboptimal in a Shannon sense because of the requirement of scalar quantization, but they are reasonable to build, work fairly well, and are field-proven. It remains of interest, however, to determine the potential improvement available by using less structured vector quantizers and whether the improvement merits the possibly increased complexity.

A general vector quantizer consists of a codebook of possible reproduction vectors and a minimum distortion encoding rule. With advances in large-scale integration, cheap memory, and fast search techniques, such a structure will itself become increasingly amenable to real implementation.

In some applications, the potential gains of general vector quantization over simple techniques are known to be small. It is well known, for example, that for large rates, simple PCM performs quite close to the rate distortion bound for Gauss memoryless sources [12], [17]. Based on these rigorous arguments, intuitive developments have been advanced so that DPCM provides nearly optimal performance for large rates and Gauss and Markov sources [18]. These are, however, asymptotic results and should not discourage the investigation of the potential improvements from optimization at more interesting lower (not asymptotically large) rates. We also point out that the Gaussian assumption is a "worst case" assumption in source coding—the Gaussian source is the "least compressible" source, given the source's autocorrelation, in that theoretically more bits are required to achieve a given average distortion [11]. Hence, conclusions based on a Gaussian source may be too conservative for other sources.

An obvious difficulty with the goal of developing general optimal vector quantizers is the fact that there is no general theory for the construction of optimal source codes. For this reason, the standard design technique for vector quantizers comprises two components: 1) choose a code structure that is reasonable to implement and optimize; and 2) develop a probabilistic model of the source for use in the optimization. Almost all traditional data compression schemes use only scalar quantization and the technique of Lloyd [19] (Method II) and Max [20] applied to a model distribution to find the optimal scalar quantizers. Examples are PCM and ADPCM systems [5], [1], [3] and transform coders [4], [8], which are optimized for models and then applied to real speech.

Other model-based techniques have also been incorporated into waveform coder design and analysis. In order to use the transform coder optimization techniques of Kramer and Mathews [15] and Huang and Schultheiss [16] on real speech, the speech is assumed to have a Gaussian distribution [4], [8]. This assumption yields transformed samples that are also Gaussian and provides rate distortion functions in a useful form for determining good bit allocations for the scalar quantizers for the transformed samples.

Thus, in the above examples and in more sophisticated systems based on these examples, the code is optimized for a model of the source and not the source itself. Intuitively, if the model is accurate, the resulting theory and performance should hold approximately for real speech. There is, however, no quantitative measure of the accuracy of a model and the validity of the approximation for real speech.

An alternative to the Lloyd–Max model-based quantizer design technique that has not previously been applied to speech waveform compression is the generalization of Lloyd's Method I [19] to the design of vector quantizers based on a long training sequence of actual data [21]–[24]. Both Lloyd's Method I and Method II, in general, guarantee only locally optimal quantizers in the sense of yielding a local minimum of an average distortion, but Method I has several advantages over the Lloyd (Method II)–Max technique.

First, it requires no explicit modeling of the source. The algorithm is instead driven directly by a long training sequence, and not indirectly through a probabilistic model based on a training sequence. In particular, no assumptions of the source being Gaussian, Laplacian, or autoregressive of a particular order are required. It can be shown mathematically that under certain stationarity assumptions, the design algorithm run on a long training sequence will produce nearly the same code and performance that it would produce if run on the true (but unknown) probabilistic description [22]. No such convergence result is known for the Lloyd–Max algorithm run on a model inferred from training data. Intuitively, this convergence implies that a quantizer designed on a sufficiently rich training sequence will yield similar performance on future data—a property that is experimentally supported for speech waveform coding in a later section.

Second, the technique works for vector quantization as well as scalar quantization to yield a locally optimum quantizer of a given rate and dimension.

Third, quite general distortion measures can be used. The only requirement is the existence of a generalized centroid or center of gravity with respect to an arithmetic average.

Fourth, the design algorithm can be run either on real speech data or as a Monte Carlo design for a speech model. This permits a comparison of the model performance with rate distortion bounds and a new and, we think, informative quantitative comparison between real and simulated speech.

This paper has two distinct but strongly related goals. One is to use the vector quantizer design algorithm of [21] to design, simulate, and evaluate speech waveform coders using vector quantization of 1 and 2 bits/sample and short block lengths of 8 and less. Both a squared error and a weighted squared error are considered. The resulting performance is compared to that of several popular speech waveform coders of comparable rates including optimized PCM and DPCM, transform coding, subband coding, and ADPCM.

The other goal is to apply the same vector quantization algorithm to common probabilistic models of the speech rather than to the speech itself. The models selected are tenth-order autoregressive sources with matching correlations. This is of interest since information-theoretic bounds on optimal performance can be computed for such mathematical models and not for real speech. Thus, the model—which bears at least a resemblance to real speech—provides a means of comparing the performance of the designed vector quantizers with the theoretical optimum, and hence provides an indication of the actual (global) optimality of the algorithm. Such comparisons have not previously been made for waveform coders operating on such highly correlated sources.

While it is generally acknowledged that models have provided a useful tool in the design of speech waveform coders, it is also generally accepted that conclusions regarding the performance using real speech cannot safely be made based only on the models. We plot results of separate experiments on real speech and simulated speech on common graphs to quantify the differences between real speech and mathematical models of speech. The results point out significant differences in both the code structure and the resulting performance.

This paper is concerned with the average quantitative distortion of vector quantizers designed to minimize average distortion. In speech waveform coding systems, however, subjective distortion is also important. For this reason, the digitized speech produced in our experiments was synthesized and listened to by a variety of audiences. We briefly describe these informal subjective tests and their support of the quantitative results.

II. Vector Quantization

Let $X = (X_0, \cdots, X_{k-1})^t$ be a k-dimensional real-valued random vector described by a cumulative distribution function $F(x) = \Pr(X_i \leqslant x_i, i = 0, \cdots, k-1)$. When the dimension k is not clear from context, we will denote X as X^k. We assume for simplicity that X has zero mean, that is, $EX = 0$ where E denotes expectation with respect to F. A vector quantizer (block quantizer, k-dimensional quantizer, block source code) consists of a reproduction alphabet or codebook $\mathcal{C} = \{y_i; i = 1, \cdots, N\}$ of N vectors together with a mapping q that assigns to each x a codeword $y_i \in \mathcal{C}$. For waveform coding, the reproduction vectors or codewords y_i should be similar to the source vectors x, that is, the codewords are also k-dimensional real-valued vectors. It is convenient to define subsets S_i, $i = 1, \cdots, N$, of k-dimensional Euclidean space by $S_i = \{x: q(x) = y_i\}$, that is, S_i, collects together all input vectors mapping into the ith codeword. The vector quantizer is then completely described by the codebook $\mathcal{C}$ and the collection of disjoint sets $\mathcal{S} = \{S_i; i = 1, \cdots, N\}$.

To measure the performance of a vector quantizer, we assume a distortion measure $d(x, \hat{x})$ that assigns a nonnegative distortion value to reproducing x as $\hat{x}$. We here focus on weighted-squared error distrotion measures of the form

$$d(x, \hat{x}) = (x - \hat{x})^t B(x - \hat{x})$$
$$= \sum_{i=0}^{k-1} \sum_{j=0}^{k-1} (x_i - \hat{x}_i) B_{i,j} (x_j - \hat{x}_j) \tag{1}$$

where $B = \{B_{i,j}; i = 0, \cdots, k\ 1; j = 0, \cdots, k - 1\}$ is a symmetric, positive definite matrix. In spite of the existence of several more sophisticated and more subjectively meaningful distortion measures for speech coding systems, such as spectral deviation and the Itakura-Saito distortion measure (see, e.g., [23]), various forms of weighted-squared error remain the most popular distortion measure for waveform coding systems because of their simplicity, tractability in analysis and design, and computability in applications.

The mathematical results to be developed are valid for a general weighting matrix B. The actual quantizer design and comparison will, however, focus on two special cases—the traditional squared error with B an identity matrix

$$d_1(x, \hat{x}) = \|x - \hat{x}\|^2 \triangleq (x - \hat{x})^t (x - \hat{x}) \tag{2}$$

and the Mahalanobis distortion measure defined by the weighting $B = R^{-1}$ where $R = E\{(X - EX)(X - EX)^t\}$ is the covariance matrix of the k-dimensional vector X

$$d_2(x, \hat{x}) = (x - \hat{x})^t R^{-1} (x - \hat{x}). \tag{3}$$

We require that R be positive definite so that R^{-1} is well defined. The Mahalanobis distortion has been successfully used in a variety of pattern classification problems (see, e.g., [25]-[29]), including speech classification problems [30], and hence seems a potentially useful distortion measure for vector quantization of speech, which can be viewed as a nearest neighbor pattern classification system [21].

Since R is symmetric and positive definite, so is R^{-1} and we can form a triangular decomposition of the form $R^{-1} = G^t G$ where G is a lower triangular matrix. Thus,

$$d_2(x, \hat{x}) = \|Gx - G\hat{x}\|^2 \tag{4}$$

where Gx is a "whitened" vector since its covariance matrix is the identity matrix. Minimizing the Mahalanobis distortion is equivalent to minimizing the squared-error distortion between the "residual" vectors Gx and $G\hat{x}$ formed by passing x and $\hat{x}$ through a causal time-varying filter described by G.

If the random vector X is quantized to obtain $q(X)$ and an error vector $e = X - q(X)$, then the expected Mahalanobis distortion can be expressed as

$$Ed_2(X, q(X)) = E\{e^t R^{-1} e\}$$
$$= \operatorname{tr}\{R^{-1} R_e\} \tag{5}$$

where "tr" denotes the trace operation and $R_e = E\{(X - q(X))(X - q(X))^t\}$ is the covariance matrix of the error signal. Comparison with

$$Ed_1(X, q(X)) = E(e^t e) = \operatorname{tr}\{R_e\} \tag{6}$$

points out that the Mahalanobis distortion couples the components of the error vector, that is, the distortion depends on the cross correlations of the errors as well as on their power.

Given a distortion measure, the performance of a vector quantizer q on a random vector X can be measured by the average or expected distortion per sample

$$\bar{D}(q) = k^{-1} Ed(X, q(X))$$
$$= k^{-1} \sum_{i=1}^{N} E\{d(X, \hat{y}_i) | X \in S_i\} \Pr\{X \in S_i\}. \tag{7}$$

The rate of a quantizer is defined by $R = \log_2 N$ bits/block or $r = R/k$ bits/sample.

A vector quantizer is implemented on a discrete-time random process $\{X_i\}_{i=0}^{\infty}$ by successively quantizing nonoverlapping vectors $X_i = (X_{ik}, X_{ik+1}, \cdots, X_{ik+k-1})$ as $q(X_i)$. If the process $\{X_i\}_{i=0}^{\infty}$ is k-stationary and k-ergodic (i.e., the vector process $\{X_i\}_{i=0}^{\infty}$ is stationary and ergodic), then the ergodic theorem implies that

$$\lim_{L \to \infty} L^{-1} \sum_{i=0}^{L-1} d(X_i, q(X_i)) = Ed(X, q(X)) \tag{8}$$

with probability one, and hence the expected distortion yields the long-term time average distortion. Hence, we define a quantizer q to be *optimum* for a given number of levels N and dimension k if it minimizes the expected distortion and, hence also minimizes the long-term sample average distortion over all quantizers of N levels and dimension k.

The goal of vector quantizer design is to produce an optimum vector quantizer. Except for a few special cases, however, the best one can hope for is a locally optimum quantizer, that is, a q such that $\bar{D}(q)$ is a local minimum. In particular, subject to certain technical assumptions, the vector quantizer design algorithm of [21] can be used to design locally optimum quantizers for quite general distributions F and distortion measures d [21]-[24], [37].

In many cases such as speech, however, there is no generally accepted distribution F describing x. If one does possess a training sequence $\{x_i; i = 0, \cdots, L - 1\}$ of data from the source and L is large, then the expectations can be computed based on the sample distribution provided by the data instead of the "true"—but unknown—underlying distribution. If the vector source were stationary and ergodic, the resulting sample averages would be good approximations to the "true" averages. In particular, from (8) one should be able to use the design algorithm to minimize the sample distortion

$$k\bar{D}(q) = L^{-1} \sum_{i=0}^{L-1} d(x_i, q(x_i)) \tag{9}$$

and thereby approximately minimize the true expectation. This approach is made precise in [22] and implies, in particular, that if L is sufficiently large, then a quantizer yielding a given sample distortion on a training sequence will yield nearly the same time average distortion on future data produced by the same source. Furthermore, this fact justifies the use of the algorithm on sample distortions even when F is known, as this is equivalent to a Monte Carlo design for a known source. For this reason, we henceforth focus on sample averages of the form (9).

In order to state the design algorithm, we require the notion of the centroid or center of gravity of a set S of $|S|$ k-dimensional real-valued vectors with respect to the distortion measure d of (1). The centroid of S, cen (S), is defined as the vector $\boldsymbol{u}$ that minimizes the sample average

$$\Delta_S(\boldsymbol{u}) = \sum_{\boldsymbol{x} \in S} d(\boldsymbol{x}, \boldsymbol{u}).$$

It is shown in the Appendix that regardless of the specific positive definite matrix $\boldsymbol{B}$, the centroid is given by

$$\text{cen}\,(S) = |S|^{-1} \sum_{\boldsymbol{x} \in S} \boldsymbol{x}, \tag{10}$$

the usual Euclidean center of gravity of the set S of vectors.

Two versions of the vector quantizer algorithm are considered. The difference is in the manner of generating initial guesses for larger codes from good small codes. Usually, the initial guess for a vector of dimension k is formed as a product code of a dimension $k - 1$ code with a scalar code; thus, higher dimension codes are constructed from good smaller dimension codes. For comparison, some of the codes were designed using the splitting technique of [21] (see [24] for a flowchart), which designs codes for a fixed dimension code by designing codes of rate 0, 1, 2, $\cdots$ bits/block by splitting each word in a rate R code into two close words to form a rate $R + 1$ code. As shall be seen, the two techniques occasionally give distinct results, indicating that genuinely distinct local minima are possible.

The product code form of the design algorithm is next described. It should be pointed out that for the weighted squares distortion being considered, the iteration of Steps 2 and 3 is equivalent to Diday and Simon's Algorithm α for cluster analysis [29].

A. Design Algorithm—(Product Code Initialization)

Step 0—(Initialization): Given are a training sequence of sample data (real numbers) $T = \{x_i; i = 0, \cdots, J - 1\}$, an integer rate r bits/sample, a maximum dimension (vector size) K, a distortion threshold $\epsilon > 0$, and a symmetric positive-definite $K \times K$ weighting matrix $\boldsymbol{B} = \{B(i,j);\ i,j = 0, 1, \cdots, K - 1\}$. Let $\boldsymbol{B}_k$ denote the $k \times k$ submatrices $\{B(i,j);\ i,j = 0, 1, \cdots, k - 1\}$, $k = 1, \cdots, K$ (which are also positive definite), and define for $k = 1, 2, \cdots, K$ the distortion measures $d_k(\boldsymbol{x}, \hat{\boldsymbol{x}}) = (\boldsymbol{x} - \hat{\boldsymbol{x}})^t \boldsymbol{B}_k (\boldsymbol{x} - \hat{\boldsymbol{x}})$ where $\boldsymbol{x} = \boldsymbol{x}^k$. Set $k = 1$.

Step 1—(Initialization for Dimension k): If $k = 1$, let $\mathcal{C}_0(k)$ be an initial scalar code of rate r having 2^r entries. (This can be obtained by the splitting technique of [21] or by an intuitive guess such as a uniform quantizer on some range.) If $k > 1$, then form the product code $\mathcal{C}_0(k) = \mathcal{C}(k - 1) \times \mathcal{C}(1)$ where $\mathcal{C}(k - 1)$ is the final dimension $k - 1$ codebook of Step 3. In other words, each word in $\boldsymbol{y}^k$ in $\mathcal{C}_0(k)$ has the form $(\boldsymbol{u}^{k-1}, z)$ where $\boldsymbol{u}^{k-1} \in \mathcal{C}(k - 1)$ and $z \in \mathcal{C}(1)$. Set $m = 0$, $D_{-1} = \infty$, and $L = J/k$ (which we assume an integer for simplicity). Let T_k denote the k-dimensional vector sequence $\{\boldsymbol{x}_i; i = 1, \cdots, L\}$ where $\boldsymbol{x}_i = \boldsymbol{x}_i^k = (x_{ki}, \cdots, x_{k(i+1)-1})^t$ is formed by blocking the training sequence.

Step 2: Given $\mathcal{C}_m(k) = \{\boldsymbol{y}_i; i = 1, \cdots, 2^{rk}\}$, group all vectors $\boldsymbol{x}_i$ in the training sequence according to the minimum distortion codeword by forming the sets

$$S_i = \{\boldsymbol{x}_j\colon\ d(\boldsymbol{x}_j, \boldsymbol{y}_i) \leqslant d(\boldsymbol{x}_j, \boldsymbol{y}_l), \quad \text{all } l \neq i\}$$

with some arbitrary tie-breaking rule. Compute the distortion resulting when using the quantizer q described by the codebook $\mathcal{C}_m(k)$ and the sets $\mathcal{S} = \{S_i\}$

$$D_m = L^{-1} \sum_{i=0}^{L-1} d(\boldsymbol{x}_i, q(\boldsymbol{x}_i)) = L^{-1} \sum_{i=0}^{L-1} \min_{\boldsymbol{y} \in \mathcal{C}_m(k)} d(\boldsymbol{x}_i, \boldsymbol{y})$$

$$= L^{-1} \sum_{i=1}^{2^{rk}} \sum_{\boldsymbol{x} \in S_i} d(\boldsymbol{x}, \boldsymbol{y}_i).$$

Step 3: If $(D_{m-1} - D_m)/D_m > \epsilon$, then form a new codebook

$$\mathcal{C}_{m+1}(k) = \{\text{cen}\,(S_i); i = 1, \cdots, 2^{rk}\}.$$

(This must reduce distortion or leave it unchanged.) Set $m = m + 1$ and go to Step 2.

If $(D_{m-1} - D_m)/D_m \leqslant \epsilon$, then we are done for dimension k with final codebook $\mathcal{C}(k) = \mathcal{C}_m(k)$.

Step 4: If $k < K$, then set $k = k + 1$ and go to Step 1. If $k = K$, then halt with final codebook $\mathcal{C}(K)$.

III. Simulated Speech

A common probabilistic model for speech that has proved useful in the design of waveform coding systems for real speech is a finite-order autoregressive or all-pole process with a driving process and second-order characteristics estimated from real speech (see, e.g., [4], [6], [8]–[10]). Numerous driving (innovations) densities have been proposed, including Gaussian, Laplacian, gamma, and K_0 (see, e.g., the above references and [31]).

Such autoregressive processes are stationary and ergodic if the autoregressive filter is stable. A standard bromide in the speech literature asserts that speech is only locally or quasi-stationary and "hence" not stationary. Local stationarity and global (usual) stationarity are not, however, incompatible. One can have a model exhibiting both characteristics and have the local stationary properties as the fine structure. (Such models are treated mathematically in Fontana *et al.* [32].) Thus, there is no conflict in considering such autoregressive models as consistent but coarse models of real speech. Such a process models the global behavior of speech, that is, the average over the local properties.

An Mth-order autoregressive or all-pole model for speech is a random process $\{X_n\}_{n=-\infty}^{\infty}$ specified by the difference equation

$$X_n = -\sum_{k=1}^{M} a_k X_{n-k} + \sigma Z_n$$

where $\{Z_n\}_{n=-\infty}^{\infty}$ is an independent, identically distributed sequence of random variables described by a probability density function f_Z with $EZ_n = 0$, $E(Z_n^2) = 1$, and where the a_k are chosen so that the filter with z-transform $1/A(z) = (1 + \sum_{k=1}^{M} a_k z^{-k})^{-1}$ is stable (and hence, $\{X_n\}$ is stationary and ergodic). The parameter σ^2 is the squared gain or one-step prediction error of the process $\{X_n\}$.

If one forms an estimate of the autocorrelation of speech as $r(k)$; $k = 0, 1, \cdots, M$ and chooses σ and the a_k in the model so that $r(k) = EX_nX_{n-k}$, then the resulting process is called the matching autoregressive source (MARS) [6]. If $r(k)$ is given by a sample autocorrelation, then from the correlation matching property of linear prediction (see, e.g., Geronimus [33, p. 12] or Markel and Gray [34]), $\{X_n\}$ is also a linear prediction model for speech. If $r(k)$ is estimated based on a short (~20 ms) speech frame, then the model is exactly the LPC model for that speech frame as used in LPC speech compression systems. We here take $1/A(Z)$ as the zero rate Itakura-Saito centroid of [23], that is, the optimum zero rate LPC model based on a training sequence of 5000 frames of 128 samples each (640 000 samples) normal speech spoken by five male speakers with a sampling rate of 6.5 kHz. This is equivalent to the LPC model for the average of the frame sample autocorrelations. (See [23] for details.) The autocorrelation and inverse filter coefficients are given in Table I. We shall use the same training sequence and portions thereof for direct vector quantizer design.

The driving process pdf $f_Z(z)$ was chosen to approximate the residual marginal pdf produced for the model filter $A(Z)/\sigma$, that is, the real speech is passed through the filter A/σ selected above, and $f_Z(z)$ is chosen to approximate the output. Two pdf's yielding good fits [5], [31], [10] are the unit variance K_0 pdf defined by $f_Z(z) = \pi^{-1}K_0(z)$ where K_0 is a zero-order modified Bessel function, and the unit variance Laplace (or double exponential) density $f_Z(z) = (\sqrt{2})^{-1} \exp(-\sqrt{2}|z|)$. Both pdf's provide good fits to the data and both will be seen to yield similar results. For comparison purposes, we will also evaluate the performance bound for a unit variance Gaussian pdf $f_Z(z) = (2_\pi)^{-1/2} \exp(-z^2/2)$ which provides a "worst case" bound for all sources having the same autocorrelation, that is, the largest possible minimum achievable average distortion.

IV. Performance Bounds for Speech-Like Sources

Given a probabilistic source model, a distortion measure d, a dimension or block length k, and a rate r bits/sample ($R = kr$ bits/block), the optimal achievable performance is defined by

$$\rho_k(r) = \min_q \overline{D}(q),$$

$$\overline{D}(q) = k^{-1}Ed(X, q(X)) \tag{11}$$

where the minimum is taken over all quantizers q with rate r bits/sample and the k^{-1} normalizes the distortion to a per-sample basis. Information-theoretic techniques can be used to obtain lower bounds to $\rho_k(r)$ for the models considered here.

The first bound is given by the generalization of Yamada *et al.* [14] of the Shannon lower bound to the distortion-rate function (see, e.g., [11]). For the distortion measure $(x - \hat{x})^t B(x - \hat{x})$ this bound is given from [14, eq. (52) and (6c)] as

TABLE I
Sample Autocorrelation Coefficients and Filter Parameters

Lag k	r(k)/r(0)	a_k
0	1.0	1.0
1	0.400940	-0.35573
2	-0.047447	0.29671
3	-0.167337	0.09716
4	-0.182098	0.13798
5	-0.198497	0.10661
6	-0.263764	0.18392
7	-0.25939	0.17867
8	-0.129546	0.07244
9	-0.010753	0.11004
10	0.104571	0.03551

r(0) = 9662.57
σ^2 = 6561.25

$$\rho_k(r) \geq k^{-1} D_{SLB}^{(k)}(R) \triangleq (\det B)^{1/k}(2e\pi)^{-1}2^{-2(rk-h(X))/k} \tag{12}$$

where $R = kr$ and $h(X)$ is the differential entropy of X, i.e.,

$$h(X) = -\int f_X(x) \log f_X(x)\, dx$$

where $f_X(x)$ is the pdf for X. The differential entropy $h(X)$ of the autoregressive source model is difficult to evaluate exactly for fixed k, but it is shown in the Appendix that

$$h(X) \geq kh(\sigma Z) = k(h(Z) + \log \sigma) \tag{13}$$

where

$$h(Z) = h(Z_0) = -\int f_{Z_0}(z) \log f_{Z_0}(z)\, dz$$

is the differential entropy of the unit variance random variable Z_0. Thus, we have the bound

$$\rho_k(r) \geq D^{(k)}(f) \triangleq \sigma^2 (\det B)^{1/k}(2e\pi)^{-1}2^{-2(r-h(Z))}. \tag{14}$$

We shall refer to this bound simply as the distortion-rate bound.

For the usual squared-error distortion measure, we have that $\det B = 1$, and hence

$$\rho_k(r) \geq D^{(k)}(r) = \sigma^2(2e\pi)^{-1}2^{-2(r-h(Z))} \tag{15}$$

(as in Berger [11]), and for the Mahalanobis distortion measure, we have that $\det B = \det R^{-1}$, and hence

$$\rho_k(r) \geq D^{(k)}(r) = \sigma^2(\det R)^{-1/k}(2e\pi)^{-1}2^{-2(r-h(Z))} \tag{16}$$

which is simply the squared-error bound scaled by $(\det R)^{-1/k}$. These bounds are readily evaluated for the sources considered since, for a unit variance Laplace pdf f_Z, we have by direct evaluation that

$$h(Z) = \log \sqrt{2}e \cong 1.347 \text{ nats} = 1.942 \text{ bits} \tag{17}$$

and for the unit variance K_0 pdf f_Z, we have by computer evaluation that

$$h(Z) = \log \pi - (2/\pi) \int_0^\infty K_0(x) \log K_0(x)\, dx$$

$$\cong \log \pi + 0.200 \cong 1.345 \text{ nats} = 1.940 \text{ bits}. \tag{18}$$

To the accuracy of the figures, $h(Z)$ and, hence the distortion-rate and asymptotic quantizer bounds, are the same for both Laplace and K_0 pdf's.

For the Gaussian pdf

$$h(Z) = \tfrac{1}{2} \log 2\pi e \cong 1.419 \text{ nats} \cong 2.047 \text{ bits}. \tag{19}$$

Unfortunately, bounds from distortion-rate theory can be quite poor when the dimension k is not large. If we assume that the codebook size $N = 2^{kr}$ is large, however, then the generalization of Yamada *et al.* [14] of the asymptotic quantizer performance bounds of Gish and Pierce [12] and Gersho [13] can be used to improve the distortion-rate bound. From the above and [14, eq. (55)], if 2^{kr} is large, then

$$\rho_k(r) \geqslant \Delta^{(k)}(r) \triangleq \{e\Gamma(1 + k/2)^{2/k}/(1 + k/2)\}\, D^{(k)}(r) \tag{20}$$

where Γ is the gamma function and the bracketed term is bound from below by 1. We shall refer to this bound as the asymptotic quantizer bound.

It is shown in the Appendix that

$$\lim_{k\to\infty} e\Gamma(1 + k/2)^{2/k}/(1 + k/2) = 1, \tag{21}$$

and hence the asymptotic quantizer bound approaches the distortion-rate bound in the limit of large dimension k. For finite k, however, the asymptotic quantizer always provides a strictly better bound [14].

It is well known from rate-distortion theory that for a given autocorrelation, the distortion rate function is most conservative for the "worst case" of a Gaussian source [11], that is, for a fixed rate, the optimal achievable distortion is larger for the Gaussian than for any other density. From (14), (19), and (20), it is clear that the Gaussian density plays a similar "worst case" role in the bounds developed here. We henceforth focus on the more optimistic bounds for the simulated speech cases of K_0 and Laplacian densities, as even these bounds will be seen to be overly conservative for real speech.

V. Quantizer Performance Bounds

A problem with using the average distortion $\overline{D}(q)$ of (11) as a measure of quantizer performance is its dependence on the input signal amplitude. A common and useful normalization in the case of squared-error distortion is to divide by the input signal variance (or energy, if the signal has zero mean) and use either $\overline{D}(q)/R(0)$, $R(0) = E[(X - EX)^2]$ as a measure of distortion or $R(0)/\overline{D}(q)$—the signal to (quantization) noise ratio (SNR)—as a measure of quality.

As a natural generalization of this normalization, we define

$$\gamma = \min_u k^{-1} E(d(X, u)), \tag{22}$$

the smallest average distortion achievable when no information is sent, that is, the average distortion using the optimum zero rate code. We will focus on the normalized average distortion defined by

$$\overline{d}(q) = \overline{D}(q)/\gamma \tag{23}$$

as the measure of performance of a quantizer. The expectations of (11) and (22) will be probabilistic averages if a distribution is assumed and will be estimated by a time average if a training sequence is used.

Analogous to (10), the minimizing u in (22) is simply EX, and hence

$$\gamma = k^{-1}E((X - EX)^t B(X - EX)). \tag{24}$$

If B is the identity matrix, then $\gamma = R(0)$, and hence $1/\overline{d}(q)$ is the usual SNR. In the general case, $1/\overline{d}(q)$ plays a role analogous to that of an SNR.

For the Mahalanobis distortion, $B = R^{-1}$, and hence $\gamma = 1$.

VI. Experimental Results

Quantizers were designed and tested for a variety of rates, dimensions, and data sets. In all of the experiments except for the K_0 simulated speech, duplicate runs were made on separate computers using different programs as a check on the results.

The first series of experiments consisted of the design of vector quantizers of 1 bit/sample and dimensions 1–8 and 2 bits/sample and dimensions 1–4 for both real speech and simulated speech. The training sequence for real speech consisted of the first 128 000 samples or approximately 20 s of the training sequence of [23] described in Section III, and hence contained a single male speaker. The simulated speech training sequences consisted of 60 000 samples produced by a random number generator. The final average normalized distortion is plotted in decibels for the entire training sequence together with the asymptotic quantizer bound and distortion rate bounds in Figs. 1–4—Figs. 1 and 3 are for squared-error distortion and 1 bit/sample and 2 bits/sample, respectively. Figs. 2 and 4 are similar plots for the Mahalanobis distortion. The real speech tests were run on a PDP 11/34 at Stanford University and a PDP 11/70 at Time and Space Processing, Inc. The results for simulated speech were conducted on a Cyber computer at California State University at Long Beach and on a Cromemco System 3 at Stanford.

In this test series, the training sequence was longer for the real speech simply because of greater available computer time. Longer training sequences are always better, but the availability of the Cyber and the speed of the Cromemco necessitated shorter sequences.

The digitized speech was synthesized for all of these experiments. The subjective quality of the squared-error distortion results was almost identical to the corresponding Mahalanobis

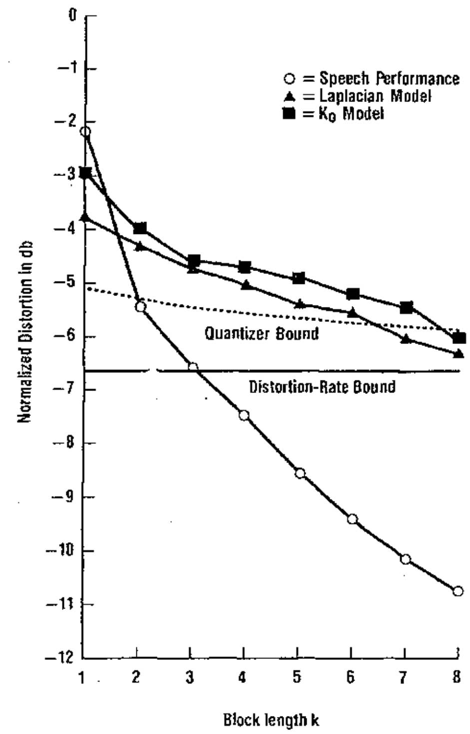

Fig. 1. Design performance of vector quantizers under MSE distortion measure at a rate of 1 bit/sample.

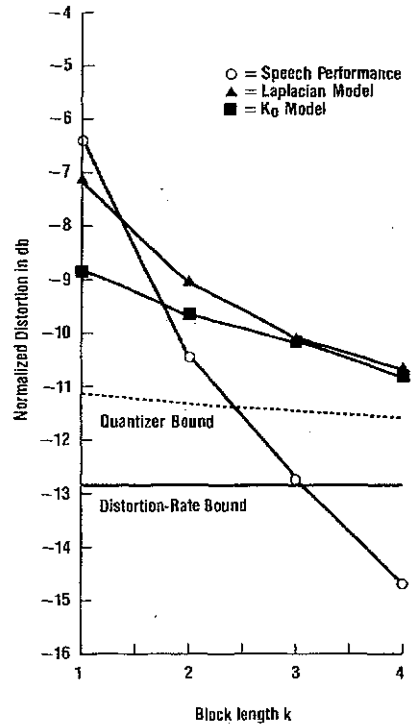

Fig. 3. Design performance of vector quantizers under MSE distortion measure at a rate of 2 bits/sample.

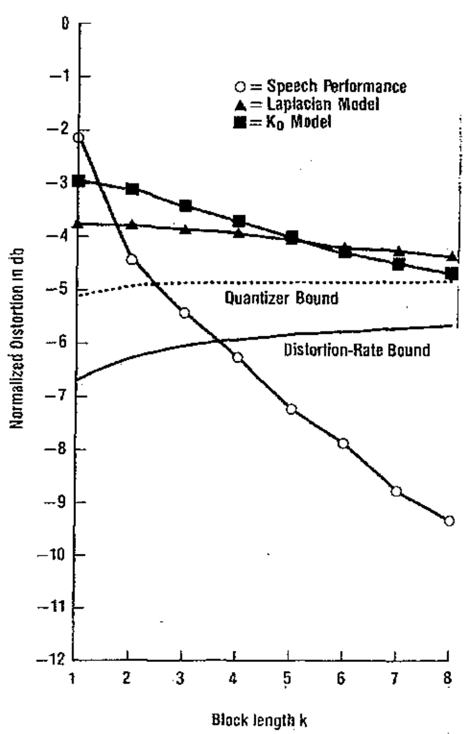

Fig. 2. Design performance of vector quantizers under Mahalanobis distortion measure at a rate of 1 bit/sample.

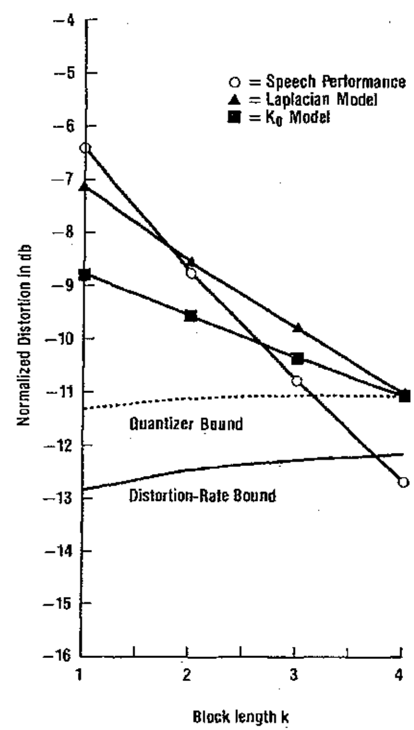

Fig. 4. Design performance of vector quantizers under Mahalanobis distortion measure of a rate of 2 bits/sample.

TABLE II
SNR INSIDE AND OUTSIDE 128K TRAINING SEQUENCE

SNR(dB) \ k	1	2	3	4	5	6	7	8
Inside	2.16	5.44	6.56	7.48	8.55	9.40	10.15	10.75
Outside	2.04	5.26	6.00	6.94	7.37	7.82	8.05	8.16

(a) 1 bit per sample

SNR(dB) \ k	1	2	3	4
Inside	6.41	10.44	12.72	14.68
Outside	6.37	9.71	11.70	12.38

(b) 2 bits per sample

distortion results. As the Mahalanobis distortion is more complicated to implement, further tests were not conducted for it.

Three properties common to all of the figures are apparent. First, the relative performance improvements achievable for the real speech are much larger than those achieved by the model. For example, an improvement of roughly 8 dB is achieved between optimum PCM ($k = 1$) and a rate 8 bits/block ($k = 8$ for 1 bit/sample, $k = 4$ for 2 bits/sample) using MSE vector quantization. The corresponding improvement is only about 3–4 dB on the simulated speech. Second, the performance of the vector quantizers on real speech exceeds that of information theoretic bounds based on the model by several decibels, indicating that model-based theoretical bounds are not appropriate for predicting performance on real speech. Third, the performance of the quantizers designed for the models are quite close to the theoretical quantizer optimal performance bound for the larger k. As the number of levels in the largest rate codes is $2^8 = 256$, the bounds should be good approximations. This suggests that the design algorithm is yielding codes that are approximately globally optimal for the simulated speech sources and not merely locally optimum. Not surprisingly, the differences between the design performance and the theoretical bounds are much smaller for these highly correlated results than in the case of Gauss-Markov sources [24] and memoryless sources [21]—the more memory a source has, the better it can be compressed.

The relative subjective quality of the speech clearly increased with dimension. In the 1 bit/sample case, it went from very noisy, partially, or barely intelligible quality at dimension 1 to an easily intelligible, but somewhat noisy quality at dimension 8. The 2 bit/sample case yielded significantly better quality at all dimensions. The dimension 4, 2 bit/sample example exhibited little audible quantization noise on the recording equipment used.

The next experiments comprised testing the codebooks designed above on the 78 600 sample open test sequence of [23] containing a speaker not in the training sequence. Subjectively there was a clear increase in quantization noise and buzziness of the speech, especially at the smaller dimensions and at 1 bit/sample. The quantitative difference between the design SNR and the test sequence SNR are summarized in Table II for the squared-error distortion (the design values are for the source experiments depicted in Figs. 1 and 3). The quantitative results reinforce the quality loss of the subjective tests. Note that over 2 dB is lost at the larger dimensions by using a quantizer designed for one speaker on another.

The theoretical results of [22], the experimental results of [21], and common sense all suggest that a longer training sequence with more speakers should yield codes with better performance (on the average) outside of the training sequence. Hence the design was repeated on the full 640 000 sample, five speaker training sequence (about 100 s) of [23] on DEC VAX's at Stanford University and at National Semiconductor. The designed codes were then tested on the 78 600 test sequence. In addition, the National Semiconductor experiment tested the designed codes on the first 128 000 samples of the training sequence, that is, the first speaker only.

The two experiments differed in the means of generating initial guesses and in convergence thresholds. The National Semiconductor experiments used the product code initialization introduced here. The Stanford experiments used the splitting technique of [21], [24]. The results of the National Semiconductor tests and the Stanford University tests are given in Tables III and IV, respectively.

Observe that the design SNR is uniformly smaller for the larger training sequence—reflecting the fact that one codebook must now work for five speakers instead of one. On the other hand, the performance of the resulting code on the test sequence is everywhere better—both quantitatively and subjectively—with the codes based on the larger training sequence. Furthermore, the differences between the design and test values are smaller. Somewhat surprisingly, the code based on

TABLE III
SNR INSIDE AND OUTSIDE 640K TRAINING SEQUENCE: PRODUCT CODE INITIALIZATION

SNR(dB) \ k	1	2	3	4	5	6	7	8
Inside (full)	2.01	5.25	6.04	7.11	7.82	8.51	9.13	9.70
Inside (1st 128K samples)	2.16	5.42	6.33	7.48	8.28	8.95	9.54	10.04
Outside	2.05	5.28	6.03	6.99	7.47	8.07	8.47	8.76

(a) 1 bit per sample

SNR(dB) \ k	1	2	3	4
Inside (full)	6.42	9.93	12.18	12.37
Inside (1st 128K samples)	6.33	16.31	14.31	14.57
Outside	6.45	9.68	11.83	11.58

(b) 2 bits per sample

the full sequence is nearly optimal for the first speaker alone (compare the "Inside" results of Table II with the "Inside" (1st 128K samples) of Table III).

Observe also the variation among the SNR's of one code on different speakers. This simply reflects the well-known fact that short-term SNR's can vary widely for a given coding scheme.

In the original experiments, a convergence threshold of $\epsilon = 0.005$ was used for both designs. The 1 bit/sample results were quite close, as indicated in Tables III(a) and IV(a). For 2 bits/sample, however, the product code initial guess resulted in noticeably inferior codes, and hence the threshold for the National Semiconductor 2 bit/sample experiment was lowered to $\epsilon = 0.0025$ to force further iterations. The resulting performance is given in Table III(b) and compares well with Table IV(b) except for the $k = 4$ case. Attempts to improve the $k = 4$ performance of Table III(b) by further iterations yielded no reduction in performance, indicating that the product code had led to a local minimum distinct from and inferior (by over 1 dB) to that of the splitting technique. Note further that the $k = 4$ code of Table III performed poorly on the test sequence, worse than that of the Table III $k = 3$ code! This emphasizes the fact that the choice of initial codes for the algorithm can be crucial to the resulting performance. Observe, however, that the results of Table IV are not all better than those of Table III.

VII. Comparisons with Other Systems

The vector quantization systems developed in the previous sections were designed to minimize the long run mean squared-error, and hence maximize the long run SNR. As discussed in the Introduction, we therefore use the long run SNR for quantitative comparisons with other waveform coding systems of comparable rates. For these comparisons, we use the results of the 640 000 sample training sequence and the 76 800 test sequence. All of the comparisons are for the 2 bit/sample case as that was the minimum rate considered in the references.

Paez and Glisson [5] developed PCM systems for 2 bits/sample and more using the Lloyd–Max scalar quantizer design procedure on several model densities. Their best system resulted from a gamma model of the marginal density, and yielded an SNR of 6.04 dB when applied to real speech. Lloyd's Method I on the original speech yielded a slightly better 6.4–6.5 dB in both design and open tests. Thus, the model-based design does fairly well for $k = 1$ (as is also suggested by Fig. 3). It is worth observing, however, that the actual codes can be quite different. The code based on the real speech is highly asymmetric, with a value near zero occurring over 80 percent of the time, while the model-based codes have levels symmetric about the origin.

Paez and Glisson also use the Lloyd–Max technique to design quantizers inside the feedback loop of a DPCM system using a model for the error sequence. Their best performance is obtained using a Laplace model which yields an SNR of 11.67 dB at 2 bits/sample. A dimension three 2 bit/sample vector quantizer yields a slightly better performance of 11.85 dB outside the training sequence and a design distortion of 12.13 dB. A dimension 4 code improves these values to 12.70 and 13.48, respectively.

The best low rate transform coder of Campanella and Robinson [4] (using a Karhunen–Loeve expansion as in [15], [16]) yields an SNR of 12.6 dB at 14 kbits/s. Their Fourier transform coder is inferior by about 4 dB. A vector

TABLE IV
SNR INSIDE AND OUTSIDE 640K TRAINING SEQUENCE: SPLITTING INITIALIZATION

SNR(dB) \ k	1	2	3	4	5	6	7	8
Inside	2.01	5.24	6.09	7.09	7.88	8.53	9.08	9.74
Outside	2.05	5.25	6.01	6.99	7.57	8.05	8.41	8.82

(a) 1 bit per sample

SNR(dB) \ k	1	2	3	4
Inside	6.50	9.98	12.13	13.48
Outside	6.41	9.72	11.85	12.70

(b) 2 bits per sample

dimension of $k = 16$ was used. This performance is close to that of our vector quantizer of dimension $k = 4$ and 2 bits/sample (Table IV). This shows that for speech, higher dimension transform coders with an optimum transform perform quite close to optimized vector quantizers of smaller dimensions. Similar comparisons were found for Gauss–Markov sources [24].

Crochiere *et al.* [7] reported performance at 2 bits/sample of 10.9 dB for an adaptive predictive PCM (ADPCM) system and 11.1 dB for their subband coder of the same rate. Our vector quantizers provide at $k = 3$ about 0.7 dB improvement in open tests and about 1 dB in design performance. At $k = 4$, there is roughly 1.5 dB improvement in the open test, and about 2.4 dB in the design performance over the reported subband coder performance.

Of course, the above comparisons must be taken with appropriate caution, especially when the SNR is computed and compared for short segments. In all cases considered, however, a block length of 3 or 4 sufficed for a 2 bit/sample vector quantizer to provide slightly better long term MSE performance than that of the other systems considered.

VIII. Comments

Figs. 1–4 demonstrate the relative performance of vector quantizers designed for real speech and simulated speech. Tables II–IV provide similar results for real speech for longer training sequences and indicate the performance of using the codes on open tests, that is, on speakers not in the training sequence. Comparison of these results with each other and theoretical bounds based on information theory yield several conclusions. Roughly 8 dB in performance improvement is achieved in design performance and 6 dB in open test performance by the 8 bits/block codes over the optimum PCM. This improvement is greater by at least 4 dB than that predicted by design experiments on the simulated speech. In addition, the performance achieved was several decibels better than that predicted by model based information theory bounds. This suggests that except for PCM, theoretical performance bounds based on the model are of little use in predicting performance for real speech—even though the model is of order 10 and the largest block length is 8.

It was also found that longer training sequences with more speakers yield lower design performance, but higher open test performance for the cases considered. Furthermore, the code obtained for the larger sequence gave performance for the first speaker close to that of the code designed only for the first speaker. These results reinforce the theoretical convergence results of [23].

The performance on the simulated speech was close to that of approximate bounds to the optimal achievable distortion at a fixed dimension for the larger dimensions, suggesting that on these sources the vector quantizers were nearly optimal. In the 1 bit/sample MSE case, the results were within 1 dB of the distortion rate bound, the optimal achievable performance for any dimension.

The performances of the vector quantizers were compared with several popular waveform coding systems of 2 bits/sample. The higher rate codes were shown to provide slight improvements in long term SNR over the best of the popular systems. This shows that existing systems are performing fairly near the optimal performance achievable using short vector quantizers. It also shows that short vector quantizers in principle provide an alternative technique for achieving this performance.

The vector quantizers considered here have one clear disadvantage with respect to the other systems in that they have not been implemented in real-time hardware. The quantitative comparisons of this work are based on average distortion and not implementation complexity. The full search vector quantizers are not well suited to software or firmware use because of the large searches required. For example, 1 bit/sample full search quantizer requires the computation of 256 energies of eight-dimensional error vectors, and the selection of a codeword with a minimum energy for each vector consisting of eight consecutive samples. Fast search algorithms using tree-searched vector quantizers have been shown to provide much reduced search effort (16 energies in the above code example) with only small losses of performance for voice coding [23] and waveform coding of Gauss-Markov sources [24]. Other vector quantizer encoder structures are currently being studied that provide other complexity/performance tradeoffs. The full search vector quantizers remain of interest, however, as a benchmark for comparison for other schemes of comparable rates and dimensions. In addition, the general vector quantizer structure has several possible advantages over the popular schemes.

1) Given a code structure, the code is optimized by a straightforward iterative computer procedure without the need for human injected "bells and whistles" in the design.

2) Unlike DPCM and similar systems, there is no feedback loop and no predictor is required. Hence, channel errors cannot propagate past the block length (here of eight or fewer samples).

3) There are no explicit models required for the optimization procedure.

4) The algorithm of [21] can be run for quite general distortion measures. For example, minimizing a distortion measure of the form $d(x, \hat{x}) = \|x - \hat{x}\|^2/\|x\|^2$ would maximize the average short term or segmented SNR, a currently fashionable fidelity criterion.

5) Shorter block lengths are required for a given performance than for other vector quantizer schemes such as transform coders.

Appendix

Centroids

To find the vector u that minimizes

$$\Delta(u) = \sum_{i=1}^{K} d(x_i, u) = \sum_{i=1}^{K} (x_i - u)^t B (x_i - u)$$

for a symmetric positive definite B, first form the triangular decomposition $B = G^t G$ where G is lower triangular and positive definite. We then have that

$$\Delta(u) = \sum_{i=1}^{K} \|Gx_i - Gu\|^2.$$

From Linde *et al.* [21], this is minimized if

$$Gu = K^{-1} \sum_{i=1}^{K} Gx_i,$$

and hence since G is invertible, $\Delta(u)$ is minimized by

$$u = K^{-1} \sum_{i=1}^{K} x_i,$$

the Euclidean center of gravity of the $\{x_i\}$.

Differential Entropy of Autoregressive Processes

Given an autoregressive source $\{X_n\}$ as in Section III, we wish to show that for any k and vector $X = (X_0, \cdots, X_{k-1})$, $h(X) \geq k(h(Z) + \log \sigma^2)$. First fix initial conditions $X_i = x_i$, $i = -M, \cdots, -1$, and observe that then $X = \sigma AZ + b$ where A is a lower triangular matrix with elements $(A)_{i,j} = a_{i-j}$ where $a_{i-j} = 0$ for $i < j$ or $|i - j| > M$ and b depends on the initial conditions and the a_i. From Problem 1.33 of McEliece [35], the conditional differential entropy

$$h(X|X_{-m} = x_{-m}, \cdots, X_{-1} = x_{-1}) = h(Z) + \log |J|$$

where J is the Jacobian of the transformation $f(Z) = \sigma AZ + b$. The Jacobian is given by $\det(\sigma A) = \sigma^k$. Since the Z_i are independent, $h(Z) = kh(Z_0)$, and hence the conditional differential entropy above is given by $k\{h(Z) + \log \sigma\}$ regardless of the initial conditions, and hence

$$\begin{aligned} h(X|X_{-m}, \cdots, X_{-1}) &= - \int f_{X_{-m}, \cdots, X_{-1}}(x_{-m}, \cdots, x_{-1}) \\ &\quad h(X|X_{-m} = x_{-m}, \cdots, X_{-1} = x_{-1})\, dx_{-m} \cdots dx_{-1} \\ &= k\{h(Z) + \log \sigma\}. \end{aligned} \tag{A1}$$

From McEliece [35, Sect. 1.3], $h(X) \geq h(X|X_{-m}, \cdots, X_{-1})$ and the desired result follows.

Proof of (21)

From Gradshteyn and Ryzhik [36, p. 937, eq. 8.327], as $k \to \infty$,

$$\Gamma(1 + k/2) \cong (1 + k/2)^{(1+k)/2} e^{-(1+k/2)} \sqrt{2\pi}$$

whence

$$e\Gamma(1 + k/2)^{2/k}/(1 + k/2) \cong (1 + k/2)^{1/k} e^{-2/k} (2\pi)^{1/k}$$

$$e\Gamma(1 + k/2)^{2/k}/(1 + k/2) \cong (1 + k/2)^{1/k} \cdot e^{-2/k}(2\pi)^{1/k} \xrightarrow[k \to \infty]{} 1.$$

References

[1] N. S. Jayant, "Digital coding of speech waveforms: PCM, DPCM, and DM quantizers," *Proc. IEEE*, vol. 62, no. 5, 1974.

[2] L. D. Davisson and R. M. Gray, Eds., *Benchmark Papers in Data Compression*. Stroudsburg, PA: Hutchinson and Ross, 1976.

[3] J. L. Flanagan, M. R. Schroeder, B. S. Atal, R. E. Crochiere, N. S. Jayant, and J. M. Tribolet, "Speech coding," *IEEE Trans. Commun.*, vol. COM-27, pp. 710-736, Apr. 1979.

[4] S. J. Campanella and G. S. Robinson, "A comparison of orthogonal transformations for digital speech processing," *IEEE Trans. Commun. Technol.*, vol. COM-19, pp. 1045-1049, 1971.

[5] M. D. Paez and T. H. Glisson, "Minimum mean squared error quantization in speech PCM and DPCM systems," *IEEE Trans. Commun. Technol.*, vol. COM-20, pp. 225-230, Apr. 1972.

[6] J. B. A. Anderson and J. B. Bodie, "Tree encoding of speech," *IEEE Trans. Inform. Theory*, vol. IT-20, pp. 379-387, 1975.

[7] R. E. Crochiere, S. M. Webber, and J. K. L. Flanagan, "Digital coding of speech in sub-bands," *Bell Syst. Tech. J.*, vol. 55, pp. 1069-1086, Oct. 1976.

[8] R. Zelinski and P. Noll, "Adaptive transform coding of speech signals," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-25, pp. 299-309, Aug. 1977.

[9] S. G. Wilson and S. Husain, "Adaptive tree encoding of speech at 8000-bits/s with a frequency weighted error criterion," *IEEE Trans. Commun.*, vol. COM-27, pp. 165-170, Jan. 1979.

[10] Y. Matsuyama and R. M. Gray, "Universal tree encoding of speech waveforms," *IEEE Trans. Inform. Theory*, vol. IT-27, pp. 31–40, Jan. 1981.
[11] T. Berger, *Rate Distortion Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1971.
[12] H. Gish and J. N. Pierce, "Asymptotically efficient quantizing," *IEEE Trans. Inform. Theory*, vol. IT-14, pp. 676–683, 1968.
[13] A. Gersho, "Asymptotically optimum block quantization," *IEEE Trans. Inform. Theory*, vol. IT-25, pp. 373–380, July 1979.
[14] Y. Yamada, S. Tazaki, and R. M. Gray, "Asymptotic performance of block quantizers with difference distortion measures," *IEEE Trans. Inform. Theory*, vol. IT-26, pp. 6–14, Jan. 1980.
[15] H. P. Kramer and M. V. Mathews, "A linear coding for transmitting a set of correlated signals," *IRE Trans. Inform. Theory*, vol. IT-2, pp. 41–46, Sept. 1956.
[16] J. Y. Huang and P. M. Schultheiss, "Block quantization of correlated Gaussian random variables," *IEEE Trans. Commun. Syst.*, vol. CS-11, pp. 289–296, Sept. 1963.
[17] T. J. Goblick, Jr. and J. L. Halsinger, "Analog source digitization: A comparison of theory and practice," *IEEE Trans. Inform. Theory*, vol. IT-13, pp. 323–326, 1967.
[18] J. B. O'Neal, Jr., "Differential pulse-code modulation (PCM) with entropy coding," *IEEE Trans. Inform. Theory*, vol. IT-22, pp. 169–174, Mar. 1970.
[19] S. P. Lloyd, "Least squares quantization in PCM," Bell Lab. Tech. Paper, Murray Hill, NJ, 1957; also, *IEEE Trans. Inform. Theory*, Mar. 1982.
[20] J. Max, "Quantizing for minimum distortion," *IRE Trans. Inform. Theory*, vol. IT-6, pp. 7–12, 1960.
[21] Y. Linde, A. Buzo, and R. M. Gray, "An algorithm for vector quantizer design," *IEEE Trans. Commun.*, vol. COM-28, pp. 84–95, Jan. 1980.
[22] R. M. Gray, J. C. Kieffer, and Y. Linde, "Locally optimal block quantizer design," *Inform. Contr.*, vol. 45, pp. 178–198, May 1980.
[23] A. Buzo, A. H. Gray, Jr., R. M. Gray, and J. D. Markel, "Speech coding based upon vector quantization," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-28, pp. 562–574, Oct. 1980.
[24] R. M. Gray and Y. Linde, "Vector quantizers and predictive quantizers for Gauss Markov sources," *IEEE Trans. Commun.*, vol. COM-30, pp. 381–389, Feb. 1982.
[25] P. C. Mahalanobis, "On the generalized distance in statistics," in *Proc. Indian Nat. Inst. Sci. (Calcutta)*, vol. 2, 1936, pp. 49–55.
[26] S. Kullback, *Information Theory and Statistics*. New York: Dover, 1968.
[27] T. W. Anderson, *An Introduction to Multivariate Statistical Analysis*. New York: Wiley, 1958.
[28] G. J. A. Hartigan, *Clustering Algorithms*. New York: Wiley, 1975.
[29] E. Diday and J. C. Simon, "Clustering analysis," in *Digital Pattern Recognition*, K. S. Fu, Ed. New York: Springer-Verlag, 1976, ch. 3.
[30] S. Tazaki, H. Ogawa, T. Gotoda, and Y. Yamada, "A quantitative estimate of the possibility of confusion of Japanese phonemes," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-27, pp. 296–298, June 1979.
[31] D. Wolf, "Statistical models of speech signals," *Inform. Syst. Theory in Digital Commun.*, NTG-Facherichte, vol. 65, VDE-Verlag, Berlin, 1978.
[32] R. J. Fontana, R. M. Gray, and J. C. Kieffer, "Asymptotically mean stationary channels," *IEEE Trans. Inform. Theory*, vol. IT-27, pp. 308–316, May 1981.
[33] L. Ya. Geronimus, *Orthogonal Polynomials*. New York: Consultants Bureau, Noble Offset Printers, 1961.
[34] J. D. Markel and A. H. Gray, Jr., "On autocorrelation with application to speech analysis," *IEEE Trans. Audio Electroacoust.*, vol. AU-21, pp. 69–79, Jan. 1973.
[35] R. J. McEliece, *The Theory of Information and Coding*. Reading, MA: Addison-Wesley, 1977.
[36] I. S. Gradshteyn and I. M. Ryzhik, *Table of Integral Series and Products*. New York: Academic, 1965.
[37] R. M. Gray, A. H. Gray, Jr., G. Rebolledo, and J. E. Shore, "Rate-distortion speech coding with a minimum discrimination information distortion measure," *IEEE Trans. Inform. Theory*, vol. IT-27, pp. 708–721, Nov. 1981.
[38] L. C. Stewart, R. M. Gray, and Y. Linde, "The design of trellis waveform coders," *IEEE Trans. Commun.*, Apr. 1982.

Hüseyin Abut (S'70–M'73–SM'79) was born in Akhisar, Turkey, on November 30, 1945. He received the B.S. degree in electrical engineering from Robert College, Istanbul, Turkey, in 1968, and the M.S. and Ph.D. degrees from North Carolina State University, Raleigh, in 1970 and 1972, respectively.

In 1973 he joined the Electronics Division, Marmara Research Institute, Gebze, Turkey. During 1973–1980 he was with the Departments of Electrical Engineering and Mathematics, Boğaziçi University, Istanbul, Turkey. He was a Visiting Scholar at the Information Systems Laboratories, Stanford University, Stanford, CA, and a Visiting Associate Professor in the Department of Electrical and Computer Engineering, California State University, Long Beach, in 1979–1980. He was with the National Semiconductor Corporation Speech Laboratory during 1980–1981. Since August 1981 he has been a Professor with the Department of Electrical and Computer Engineering, San Diego State University, San Diego, CA. His current technical interests include digital speech encoding, speech analysis, synthesis and recognition, and fast computational algorithms.

Robert M. Gray (S'68–M'69–SM'77–F'80) was born in San Diego, CA, on November 1, 1943. He received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1966, and the Ph.D. degree in electrical engineering from the University of Southern California, Los Angeles, in 1969.

Since 1969 he has been with the Information Systems Laboratory and the Department of Electrical Engineering, Stanford University, Stanford, CA, where he is a Professor engaged in teaching and research in communication and information theory, and in speech compression.

From 1974 to 1980 he was a member of the Board of Governors of the IEEE Professional Group on Information Theory, and from 1977 to 1980 he was Associate Editor of the IEEE TRANSACTIONS ON INFORMATION THEORY, and he assumed the Editorship of that TRANSACTIONS with the March 1981 issue. He was the recipient of a 1981 Fellowship from the Japan Society for the Promotion of Science, and is the recipient of a 1982 Fellowship from the John Simon Gussenheim Memorial Foundation. Dr. Gray is a member of Sigma Xi, Eta Kappa Nu, the Society for Industrial and Applied Mathematics, the Institute of Mathematical Statistics, and the Société des Ingénieurs et Scientifiques de France. He holds an Advanced Class Amateur Radio License (KB6XQ).

Guillermo Rebolledo (S'76–M'82) was born in Mexico City, Mexico, on April 5, 1950. He received the degree of Ingeniero Mecanico Electricista from the University of Mexico, Mexico City, in 1974, the degree of Maestro en Ciencias from INAOE, Puebla, Mexico, in 1976, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1980 and 1982, respectively.

From 1971 to 1973 he was a Teaching Assistant in the School of Engineering and during 1973 he was a Research Assistant in the Instituto de Ingenieria, both at the University of Mexico. In 1974 he joined the Ministry of Communications of Mexico and worked as a Communications Systems Analyst for one year. From June 1978 to May 1981 he was a Research Assistant in the Information Systems Laboratory at Stanford University, working on speech compression algorithms. From March 1980 to January 1982 he was a Research Engineer at Time and Space Processing Inc., Santa Clara, CA, where he worked on problems involving signal processing and speech compression. Since February 1982 he has been with the Facultad de Ingenieria, University of Mexico. His research interests are in speech and signal processing, in data transmission, and in microprocessor applications.

Dr. Rebolledo is a member of the IEEE ASSP, Communications, and Computer Societies.