also be decoded by the better receiver. Once both receivers have removed the digital component from the received signal, the problem becomes that of transmitting a Gaussian source over a Gaussian broadcast channel. For this problem, analog transmission is optimal.

III. Conclusion

The distortion region for joint source–channel coding over the broadcast channel is yet unknown. Here we derived inner and outer bounds for this region in one special case. These bounds are asymptotically tight. In [13], we also discuss the case of dependent B and N.

We believe that some of the ideas we used can be developed further. For example, consider a general analog source, which is quantized by a vector quantizer. We can regard the quantizer output as the "digital" part of the source, and the quantization error as the "analog" part. We can then construct an encoding scheme, similar to the one presented here, and analyze it with similar tools [15].

Acknowledgment

The authors wish to thank the anonymous reviewers for some useful comments and corrections. R. Zamir acknowledges an early discussion with E. Shlomot which inspired some of the ideas presented.

References

[1] T. M. Cover, "Broadcast channels," *IEEE Trans. Inform. Theory*, vol. IT-18, pp. 2–14, Jan. 1972.

[2] M. D. Trott, "Unequal error protection codes: Theory and practice," in *Proc. IEEE Information Theory Workshop*, Haifa, Israel, June 1996, p. 11.

[3] W. H. R. Equitz and T. M. Cover, "Successive refinement of information," *IEEE Trans. Inform. Theory*, vol. IT-37, pp. 269–274, Mar. 1991.

[4] B. Rimoldi, "Successive refinement of information: Characterization of the achievable rates," *IEEE Trans. Inform. Theory*, vol. 40, pp. 253–259, Jan. 1994.

[5] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.

[6] T. J. Goblick, "Theoretical limitations on the transmission of data from analog sources," *IEEE Trans. Inform. Theory*, vol. IT-11, pp. 558–567, Oct. 1965.

[7] B. Chen and G. Wornell, "Analog error-correcting codes based on chaotic dynamical systems," *IEEE Trans. Commun.*, vol. 46, pp. 881–890, July 1998.

[8] U. Mittal and N. Phamdo, "Joint source-channel codes for broadcasting and robust communication," *IEEE Trans. Inform. Theory*, to be published.

[9] S. Shamai (Shitz), S. Verdú, and R. Zamir, "Systematic lossy source-channel coding," *IEEE Trans. Inform. Theory*, vol. 44, pp. 564–579, Mar. 1998.

[10] P. P. Bergmans, "A simple converse for broadcast channels with additive white Gaussian noise," *IEEE Trans. Inform. Theory*, vol. IT-20, pp. 279–280, Mar. 1974.

[11] T. Berger, *Rate-Distortion Theory: A Mathematical Basis for Data-Compression*. Englewood Cliffs, NJ: Prentice-Hall, 1971.

[12] Z. Reznic, R. Zamir, and M. Feder, "Joint source-channel coding of a Gaussian mixture source over the Gaussian Broadcast channel," in *Proc. 36th Annu. Allerton Conf. Communication, Control and Computing*, 1998, pp. 800–809.

[13] ——, Joint source-channel coding of a Gaussian mixture source over the Gaussian broadcast channel. Tel-Aviv Univ., Tech. Rep. EES2001-1, Tel-Aviv, Israel. [Online]. Available: http://www.eng.tau.ac.il/Pages/Departments/systems.html

[14] C. E. Shannon, "A mathematical theory of communication," *Bell. Syst. Tech. Journal*, vol. 27, pp. 379–423, 623–656, 1948.

[15] Z. Reznic and R. Zamir, "On the transmission of analog sources over channels with unknown SNR ," paper. Submitted to 2002 IEEE International Symposium on Information Theory, Lausanne, Switzerland, June 30–July 5, 2002.

Multiple Description Vector Quantization With a Coarse Lattice

Vivek K Goyal, *Member, IEEE*, Jonathan A. Kelner, and Jelena Kovačević, *Fellow, IEEE*

***Abstract*—A multiple description (MD) lattice vector quantization technique for two descriptions was recently introduced in which fine and coarse codebooks are both lattices. The encoding begins with quantization to the nearest point in the fine lattice. This encoding is an inherent optimization for the decoder that receives both descriptions; performance can be improved with little increase in complexity by considering all decoders in the initial encoding step. The altered encoding relies only on the symmetries of the coarse lattice. This allows us to further improve performance without a significant increase in complexity by replacing the fine lattice codebook with a nonlattice codebook that respects many of the symmetries of the coarse lattice. Examples constructed with the two-dimensional (2-D) hexagonal lattice demonstrate large improvement over time sharing between previously known quantizers.**

***Index Terms*—Codebook optimization, high-rate source coding, lattice vector quantization.**

I. Introduction

By using the additional structure of a lattice codebook, lattice vector quantizers can be implemented much more efficiently than their more general counterparts. By labeling the points of a lattice with ordered pairs of points in a sublattice, Servetto, Vaishampayan, and Sloane (SVS) create the two descriptions of a *multiple description* (MD) lattice vector quantizer that achieves similar performance gains over unconstrained MD vector quantizers [1], [2]. However, these quantizers turn out to be inherently optimized for the central decoder. This correspondence describes the result of modifying the encoding and decoding used by SVS to minimize a weighted combination of central and side distortions, while keeping the index assignments generated by their elegant theory. This generalization creates a continuum of quantizers for each SVS quantizer, improving the convex hull of operating points. It does this while retaining most of the computational advantages of lattice codebooks.

A. MD Coding

In an MD coding scenario, sequences of symbols are sent separately on two or more channels. Each sequence of channel symbols is called a *description*, and decoders are designed for each nonempty subset of the descriptions. Such a system with two channels is depicted in Fig. 1. This is a generalization of usual "single description" source coding.

We shall focus on the case in which the encoder receives an independent and identically distributed (i.i.d.) sequence of source symbols $\{X_k\}_{k=1}^{K}$ to communicate to three receivers over two noiseless (or error-corrected) channels. One decoder (the *central decoder*) receives information sent over both channels while the remaining two decoders

Manuscript received September 11, 2000; revised October 1, 2001. The material in this correspondence was presented in part at the IEEE Data Compression Conference, Snowbird, UT, March 2000. This work was initiated and largely completed while V. K Goyal and J. A. Kelner were with Bell Labs.

V. K Goyal is with Digital Fountain, Inc., Fremont, CA 94538 USA (e-mail: v.goyal@ieee.org).

J. A. Kelner is with Harvard University, Cambridge, MA 02138 USA (e-mail: kelner@fas.harvard.edu).

J. Kovačević is with Bell Labs, Lucent Technologies, Murray Hill, NJ 07974 USA (e-mail: jelena@bell-labs.com).

Communicated by P. A. Chou, Associate Editor for Source Coding.

Publisher Item Identifier S 0018-9448(02)00857-X.

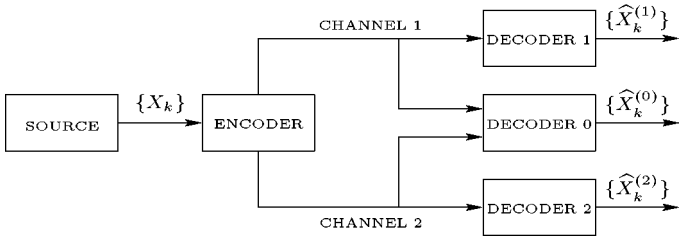

Fig. 1. Scenario for MD source coding with two channels and three receivers. The general case has M channels and $2^M - 1$ receivers, but we limit our attention to the two-channel case shown.

(the *side decoders*) receive information only over their respective channels. The transmission rate over channel i is denoted by R_i, $i = 1, 2$, in bits per source sample. The reconstruction sequence produced by decoder i is denoted by $\{\widehat{X}_k^{(i)}\}_{k=1}^K$. Using a single-letter distortion measure d, there are three expected average distortions

$$D_i = E\left[\frac{1}{K}\sum_{k=1}^{K} d\left(X_k, \widehat{X}_k^{(i)}\right)\right], \qquad \text{for } i = 0, 1, 2.$$

There are two main situations in which one might use MD coding. The first is in a broadcasting or multicasting scenario, where receivers can be categorized by the subset of the channels that they receive, and one would like to control the qualities of source approximations available to each category. In this case, the three decoders shown in Fig. 1 represent different users. The second is in a point-to-point communication system, where there are several channels that either work perfectly or fail completely to connect the sender to the receiver. If the receiver knows which channels are working but the sender does not, MD coding fits nicely. The decoders shown then represent different states of the receiver.

Two-channel MD coding was introduced as a theoretical problem in Information Theory in 1979: Given a source and distortion measure, what is the set of achievable quintuples $(R_1, R_2, D_0, D_1, D_2)$? The most important Shannon theory results were obtained by El Gamal and Cover [3], Ozarow [4], Ahlswede [5], and Zhang and Berger [6]; these are summarized in [7] and [8]. Theoretical results for MD coding with more than two channels have only recently been put forth [9]–[11]. The origins of MD coding lie in the "channel splitting" problem studied at Bell Labs in the late 1970s. This history and many practical techniques and applications are reviewed in [12].

The present work is more practical—it concerns generating MD codes for continuous-valued sources in blocks of length K (K-dimensional vector quantizers) that can be designed and implemented easily. Entropy-coding of quantizer outputs is assumed, as is a moderate to high rate and squared error distortion. Under these assumptions, the MD lattice vector quantization (MDLVQ) technique of SVS [1] is a very attractive way to produce two descriptions. The encoder does a nearest neighbor search in a lattice Λ and then applies an index assignment mapping ℓ to get an ordered pair of descriptions in a sublattice Λ'. The index assignment is constrained to make abundant use of the symmetries of the system. The nearest neighbor search is simple because of the lattice structure, and the index assignment ℓ is simple because of its many symmetries.

The techniques introduced in this correspondence are seemingly minor modifications of the SVS technique; in fact, the mapping ℓ—the key to the SVS construction—is unchanged. However, by introducing greater flexibility in the encoder (not necessarily encoding to the nearest element of Λ) and codebook (keeping Λ' a lattice but reducing restrictions on Λ), more potential operating points are created. These changes increase the encoding and design complexities, but these complexities remain much lower than if no lattices are used, as in [13]. With the slight increase in complexity, the convex hull of the operating points is improved. Attention here is limited to two channels, but the modifications could be applied to MDLVQ for any number of channels. Some examples of index assignments for three channels are given in [14].

The remainder of this correspondence is organized as follows. Section II describes how lattice codebooks simplify vector quantization (VQ) encoding and reviews the SVS technique for two-channel MDLVQ. The effects of changing the encoder without changing the codebooks are explored in Section III. The improvements afforded by altering only the encoding motivate the iterative improvement of the codebooks discussed in Section IV. In all cases, the alterations are demonstrated with two-dimensional (2-D) MDLVQs based on the hexagonal lattice. This makes it possible not only to show numerical improvements, but to visualize the changing encoder partitions. As in the work of SVS, the analyses and computations are based on a high rate assumption. Applying MDLVQ at low rates revealed advantages to the modified encoding, which are discussed in Section VI. This should not be overemphasized, however, because lattice constraints are less beneficial at low rates. Section VII summarizes the contributions.

II. MD Lattice VQ

The complexity of source coders—even single description ones—can grow very quickly with the dimension K. Assume for the moment fixed-rate source coding at rate R bits per component. Handling vectors of dimension K all at once implies having a codebook with 2^{RK} entries. An exhaustive search through such a codebook to find the nearest codeword has complexity exponential in K and R. Methods that reduce the number of operations needed to find the nearest codeword require more storage [15]. To reduce complexity, it is common to either constrain the codebook so that searching for the nearest codeword is much simpler or use a search technique that does not necessarily find the nearest codeword.

Using a codebook that is a lattice or a large subset of a lattice is called *lattice VQ* (LVQ). LVQ greatly simplifies optimal encoding; for the lattices considered in [16], the complexity of encoding is proportional to K, $K \log K$, or $K^2 \log K$—much better than exponential complexity.

For moderate- or high-rate, entropy-constrained VQ, LVQ is attractive even without invoking complexity issues. Using high-resolution analysis, it can be shown that the optimal quantizer point density is constant.[1] For scalar quantization, this gives the optimality of uniform quantization first observed by Gish and Pierce [18]. Evenly spaced (scalar) codewords give the only one-dimensional (1-D) lattice $\mathbb{Z}$. For vectors, constant point density leaves room for lattice and nonlattice solutions. Gersho [19] conjectured that at high rates, the optimal quantizer for a uniformly distributed random variable induces a partition that is a tessellation. This conjecture, which remains open, suggests there is no performance loss in limiting attention to LVQ.

[1]For details, see [17].

For our purposes, it is sufficient to understand that the distortion of entropy-constrained LVQ with lattice codebook Λ, under high-resolution assumptions, can be approximated as

$$D(R) = G(\Lambda)2^{2h(X)}2^{-2R} \tag{1}$$

where $h(X)$ is the differential entropy of the source. In this formula

$$G(\Lambda) = \frac{1}{K}\frac{\int_{V_0}\|x\|^2\,dx}{\left(\int_{V_0}dx\right)^{(K+2)/K}}$$

is called the *dimensionless second moment* of the lattice and

$$V_0 = \{x \in \mathbb{R}^K\colon \|x-\lambda\| \geq \|x\| \ \text{ for all } \lambda \in \Lambda\}$$

is the Voronoi cell of the origin. The importance of (1) is that it separates the influences of rate, source, and lattice.

We will be concerned only with how the lattice and variations thereof affect distortion. We consider only the distortion in one representative cell assuming a uniformly distributed source. (Having an approximately uniform distribution over any single small cell is part of the high-resolution assumption.) The assignment of binary strings to codewords (entropy coding) is ignored, as it affects only the rate.

A. The SVS Technique

Single description LVQ with a lattice Λ is described by an encoder mapping $\alpha\colon \mathbb{R}^K \to \Lambda$ such that $\lambda = \alpha(x) = \operatorname{argmin}_{y\in\Lambda}\|x-y\|$. In order to produce two descriptions, SVS introduce two new elements.

1) *A sublattice Λ' of Λ that is geometrically similar to Λ.* Geometric similarity means that $\Lambda' = c\Lambda U$ for some positive scalar c and some orthogonal matrix U. In addition, no elements of Λ should lie on the boundaries of Voronoi regions of Λ'. Examples of admissible sublattices are shown in Fig. 2. Λ' is called the *coarse lattice* and Λ the *fine lattice.* An important parameter of the construction is the *index* $N = |\Lambda/\Lambda'|$ of the sublattice. The index is the number of elements of the fine lattice in each Voronoi cell of the coarse lattice. Note also that the points in the Voronoi cell are classified into orbits; any two points in an orbit are related by an automorphism of Λ'.
2) *An injective index assignment mapping* $\ell\colon \Lambda \to \Lambda' \times \Lambda'$. The generation of ℓ exploits symmetries of Λ', so ℓ can be implemented efficiently.

An ordered pair of descriptions is generated by composing the two mappings

$$(\lambda_1', \lambda_2') = \ell(\lambda) = \ell(\alpha(x)).$$

Because ℓ is injective, the central decoder can compute

$$\hat{x}^{(0)} = \ell^{-1}(\lambda_1', \lambda_2') = \lambda$$

which is the element of the fine lattice Λ closest to x. The side decoders are identity mappings, so that $\hat{x}^{(i)} = \lambda_i'$, $i = 1, 2$.

As long as the index assignment is an injection, the central distortion is fixed by the fine lattice. The side distortions, however, depend on the index assignment. Of course, since the side distortion D_i depends on the distance between x and λ_i', the goal should be to make λ_1' and λ_2' as close to λ as possible.

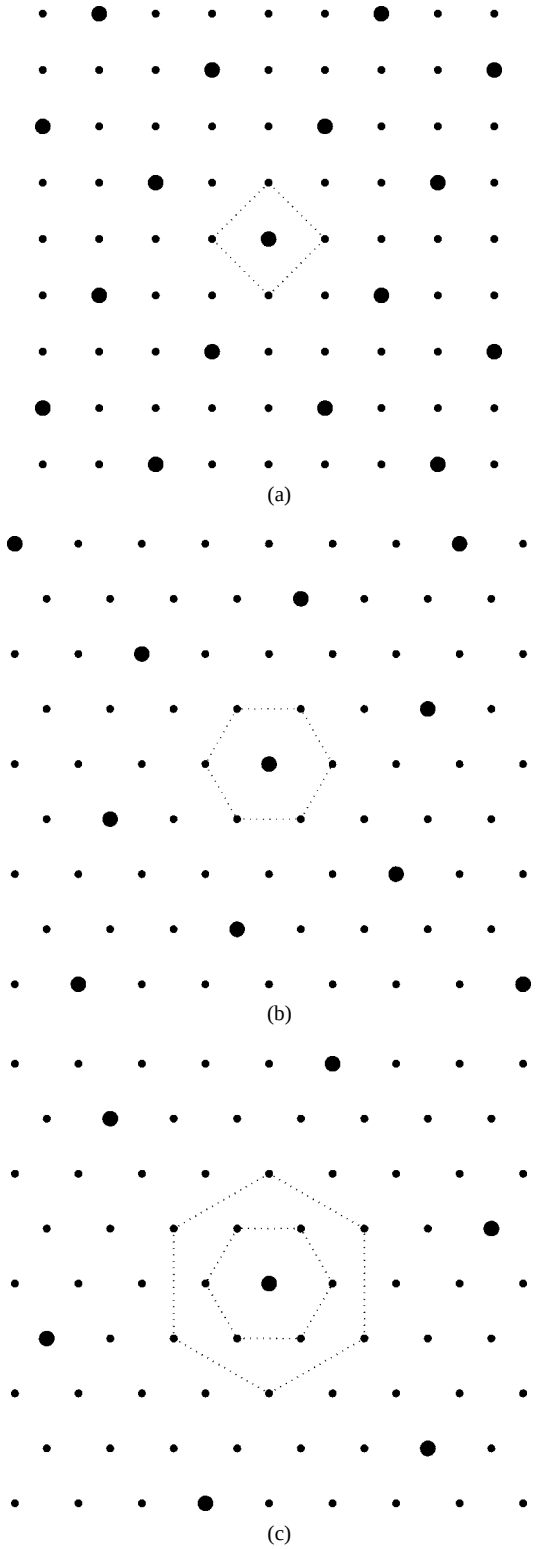

Fig. 2. Examples of geometrically similar sublattices. All the marked points are in the lattice Λ, and the heavier points are also in the sublattice Λ'. (a) $\mathbb{Z}^2$ lattice and index-5 sublattice. Each Voronoi cell of Λ' has one four-point orbit of elements of Λ. (b) A_2 (hexagonal) lattice and index-7 sublattice. There is one six-point orbit. (c) A_2 (hexagonal) lattice and index-13 sublattice. There are two orbits with six points each.

A first guess may be to make λ_1' and λ_2' the elements of Λ' closest to λ. This will not work, except with index $N = 1$ ($\Lambda' = \Lambda$), because ℓ will not be invertible. In general, the side distortion with the best assignment ℓ increases monotonically with the index N. Examples of

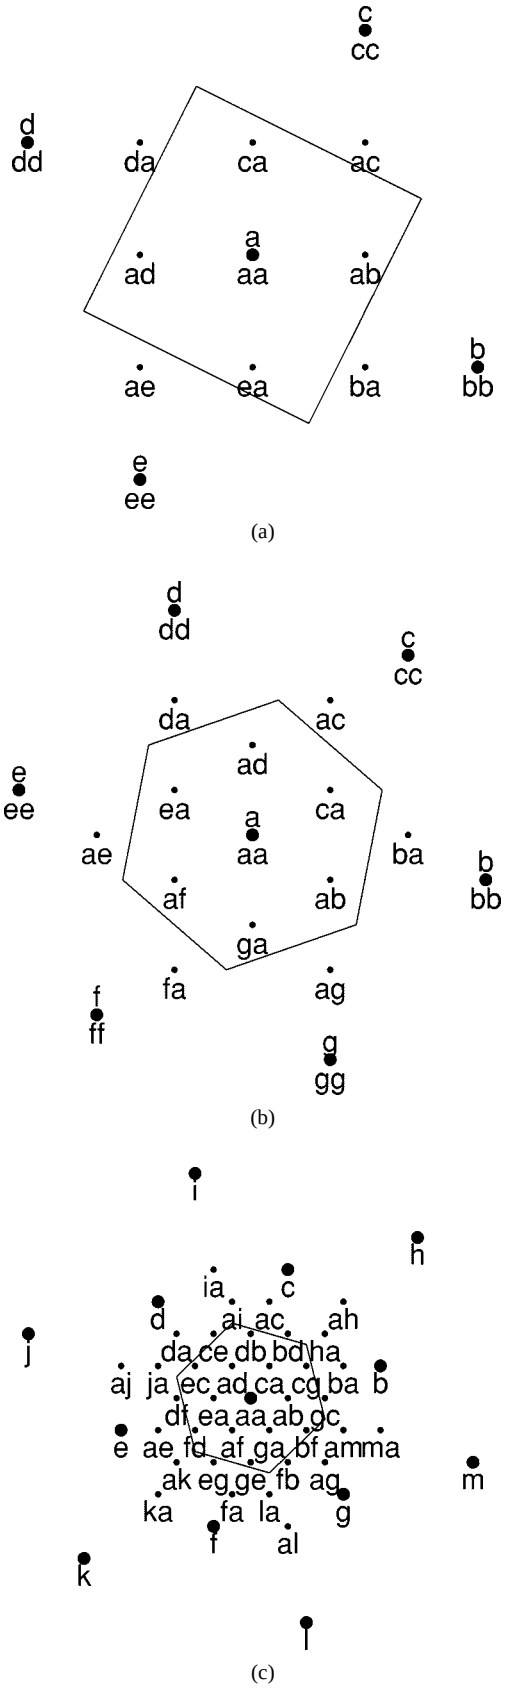

Fig. 3. Examples of optimal index assignments. The neighborhood of the origin is shown, and only points relevant to the encoding of points in the Voronoi cell (with respect to Λ') of the origin are labeled. Singlet and doublet labels give names to points of Λ' and Λ, respectively. (a) $\mathbb{Z}^2$ lattice and index-5 sublattice. (b) A_2 (hexagonal) lattice and index-7 sublattice [1]. (c) A_2 (hexagonal) lattice and index-13 sublattice.

optimal assignments are shown in Fig. 3. All further examples are based on the A_2 lattice and index-7 or index-13 sublattice with the assignment shown in Fig. 3(b) or (c). For details on the optimization of ℓ, the reader is referred to [1], [2].

The original SVS work considers only the design of index assignments that produce equal side distortions from equal-rate channels (the *balanced* case). The unbalanced case was explored by Diggavi, Sloane, and Vaishampayan [20]. Only the balanced case is considered explicitly here, though the principles apply equally well to the unbalanced case.

Let us now use the index assignment in Fig. 3(b) to reiterate the steps in MDLVQ. The first step is to quantize a source vector x to the nearest element of Λ, which we call λ. For example, the nearest fine lattice point may be the one labeled ca. This label is an ordered pair of names of points of Λ'. The label c is sent on Channel 1 and a is sent on Channel 2. The central decoder uses c and a, known in their proper order, to determine λ. The first and second side decoders reconstruct to c and a, respectively, which are nearby points in the coarse lattice.

The sublattice index fixes the resolution of Λ' relative to the resolution of Λ. The index assignment then determines the maximum distance between λ_1' and λ_2' when $(\lambda_1', \lambda_2') = \ell(\lambda)$. It is desirable for this maximum distance to be small because it reflects the quality of the side reconstructions. For example, with the index-7 sublattice labeling shown in Fig. 3(b), each side decoder finds, at worst, the sublattice point second closest to λ (at worst third closest to x); in the index-13 case shown in Fig. 3(c), a side reconstruction can be as bad as the fourth closest to λ (fifth closest to x).

B. Choosing the Scaling and Sublattice Index

As initially presented, the fine lattice Λ and, hence, the central distortion D_0 are fixed. The choice of Λ' and ℓ determine the rates and side distortions. We could equally well require the rates to remain constant and study the tradeoff between central and side distortions. Since the sublattice indexes are sent independently over the two channels, the rates are given by $R_i = H(\lambda_i')$. To keep the rates constant as N is varied, we may assume that Λ and Λ' are scaled to maintain a fixed distance between neighbors in Λ'.

Starting with the A_2 lattice, SVS have computed optimal index assignments for all admissible sublattice indexes up to 127. Keeping the rates fixed and normalizing the distortions in a manner discussed below gives the operating points shown in Fig. 4. Added to the SVS data is a point labeled "$N = 1$" which represents repeating the same information over both channels. The distortions are normalized to 0 dB for $N = 1$. In this log–log plot, the SVS operating points lie approximately on a straight line, confirming that the product $D_0 D_1$ is approximately constant. Repeating information on both channels is several decibels worse.

The relative importance of central and side distortions depends on the application. Yet we assert that *generally it is the low indexes that are important.* As demonstrated in Fig. 4, the gap between central andside distortions grows quickly as the index is increased. At index 13, the central and side distortions differ by 16.8 dB. With a larger gap, it seems unlikely that central and side distortions can both be low enough for a useful representation yet neither so low that rate has been wasted. When the sublattice index is low, the modifications proposed here do not greatly increase encoding complexity.

One reasonable way to choose the sublattice index is to minimize a weighted average of central and side distortions. Let

$$\overline{D}_p = \frac{1-p}{1+p} D_0 + \frac{p}{1+p} (D_1 + D_2) \tag{2}$$

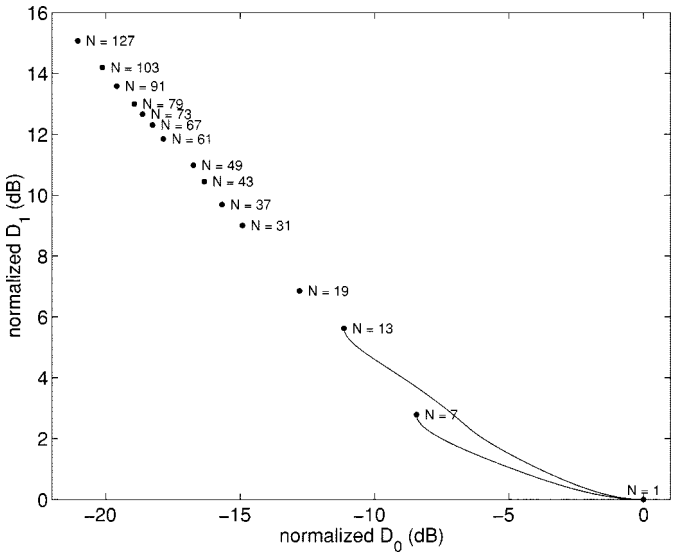

Fig. 4. Operating points of MDLVQ for the A_2 (hexagonal) lattice. Distortions are normalized to 0 dB when a description is repeated over both channels. Marked points, taken from [1, Fig. 4], are optimal for the given index N. ($N = 1$ is not allowed in [1], but simply indicates the repetition of information on both channels.) The modified encoding of Section III, applied with $N = 7$ and $N = 13$, gives the two indicated continuums of operating points; these are improved by the optimizations of Section IV.

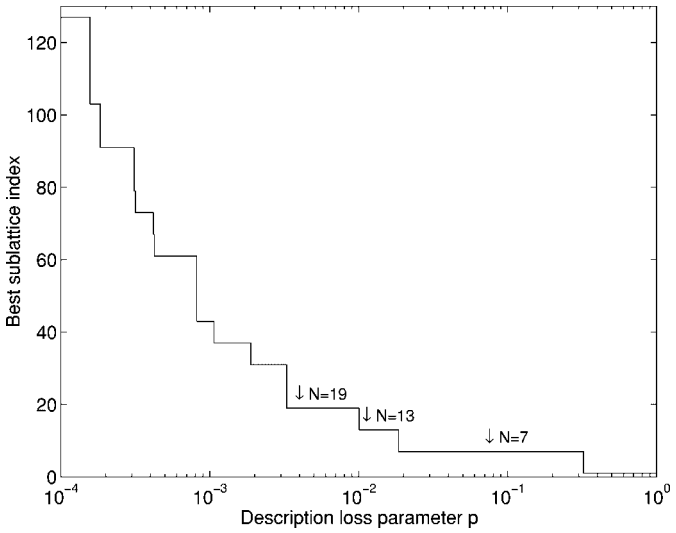

Fig. 5. The sublattice index that minimizes $\overline{D}_p$ as a function of the parameter p for the 2-D hexagonal lattice.

with $p \in [0, 1]$. The form of (2) is inspired by the interpretation that if descriptions are lost independently with probability p then $\overline{D}_p$ is the expected distortion conditioned on receiving at least one description. Fig. 5 shows the best index (for a sublattice of A_2) among the range of indexes considered in [1] as a function of p. Note the transition from $N = 13$ to $N = 7$ at $p \approx 0.0185$. Index 7 and index 1 (repetition) are optimal for all $p > 0.0185$, confirming that the low indexes are the interesting ones.

III. Altering the Encoding: More and Better Operating Points

By optimizing the sublattice index, we have done all we can to minimize $\overline{D}_p$ within the constraints of MDLVQ. However, allowing greater flexibility in the encoder and decoder makes it possible to reduce $\overline{D}_p$. First, we keep Λ, Λ', ℓ, and the decoder unchanged and change only the encoding rule.

Once the possibility of change is raised, the optimal encoder is obvious. The SVS encoder first finds the nearest fine lattice point $\lambda \in \Lambda$; i.e., it minimizes $\|x - \lambda\|^2 = \|x - \hat{x}^{(0)}\|^2$ for every sample x. The optimal encoder minimizes

$$\frac{1-p}{1+p}\left\|x-\hat{x}^{(0)}\right\|^2 + \frac{p}{1+p}\left(\left\|x-\hat{x}^{(1)}\right\|^2 + \left\|x-\hat{x}^{(2)}\right\|^2\right)$$

for every sample x. Introducing the notation $\ell(\lambda) = (\ell_1(\lambda), \ell_2(\lambda))$, the optimal encoder is

$$\alpha_p(x) = \operatorname*{argmin}_{\lambda\in\Lambda}\left\{\frac{1-p}{1+p}\|x-\lambda\|^2 + \frac{p}{1+p}\left(\|x-\ell_1(\lambda)\|^2 + \|x-\ell_2(\lambda)\|^2\right)\right\}. \quad (3)$$

Similar optimal encoders can be found for objective functions other than $\overline{D}_p$.

Except in the trivial case, where $N = 1$, $\Lambda = \Lambda'$, and ℓ is the diagonal map, encoder (3) is more complex than the SVS encoder. However, the increase in complexity depends on the index N and, as we have argued, small indexes are very often best. In the A_2 lattice, index-7 example, α_p can be implemented by quantizing to the nearest point in Λ' and searching among 13 possibilities. The search itself is very simple. Furthermore, this search procedure can be made even more efficient for large N by using the symmetries of the system [21].

The first and fourth columns in Fig. 6 are the partitions induced by α_p for a few values of p. With $p = 0$, the encoder is precisely the SVS encoder and the partition is into Voronoi cells of Λ. As p increases, certain "central" cells emerge that are larger than the others. These are the cells of fine lattice points that also lie on the sublattice. The change in shape is intuitive: as more weight is placed on the side reconstructions, it becomes advantageous to quantize to the points with $\ell_1(\lambda) = \ell_2(\lambda) = \lambda$; the remaining points (those in $\Lambda \setminus \Lambda'$) have either large $\|x - \ell_1(\lambda)\|$ or large $\|x - \ell_2(\lambda)\|$. For $p = 1$, the encoder uses the Voronoi partition of Λ'.

Since p can be chosen independently of all the other components of the system, encoding with (3) gives a continuum of operating points for each SVS quantizer. For A_2 in the index-7 and -13 cases, these operating points are given by the curves in Fig. 4. They are repeated on a linear scale in Fig. 7. (They are the top curves in each graph.) Note that the performance is much better than time sharing between the SVS points and repetition. In addition to minimizing $\overline{D}_p$ over a wider class of quantizers, the convex hull of (D_0, D_1) operating points is improved; thus, almost any objective function of interest should be reduced including the product $D_0 D_1$ by which MD systems are often measured.

IV. Altering the Decoding: Further Improved Performance

Once we have replaced the SVS encoder α with slightly more complicated encoding, performance can be improved further with almost no change in encoding and decoding complexity. The encoder is optimal given the decoders, but all the decoders are not optimal. As demonstrated by the first and fourth columns of partition diagrams in Fig. 6, the side decoders are optimal but the central decoder is not.

The points in $\Lambda \setminus \Lambda'$ are not at the centroids of their corresponding cells. The central decoder should be adjusted to reconstruct these cells to their centroids. This requires only a single displacement to be stored for each orbit. For A_2 in the index-7 and -13 cases, the improvement is shown by the middle curves of Fig. 7.

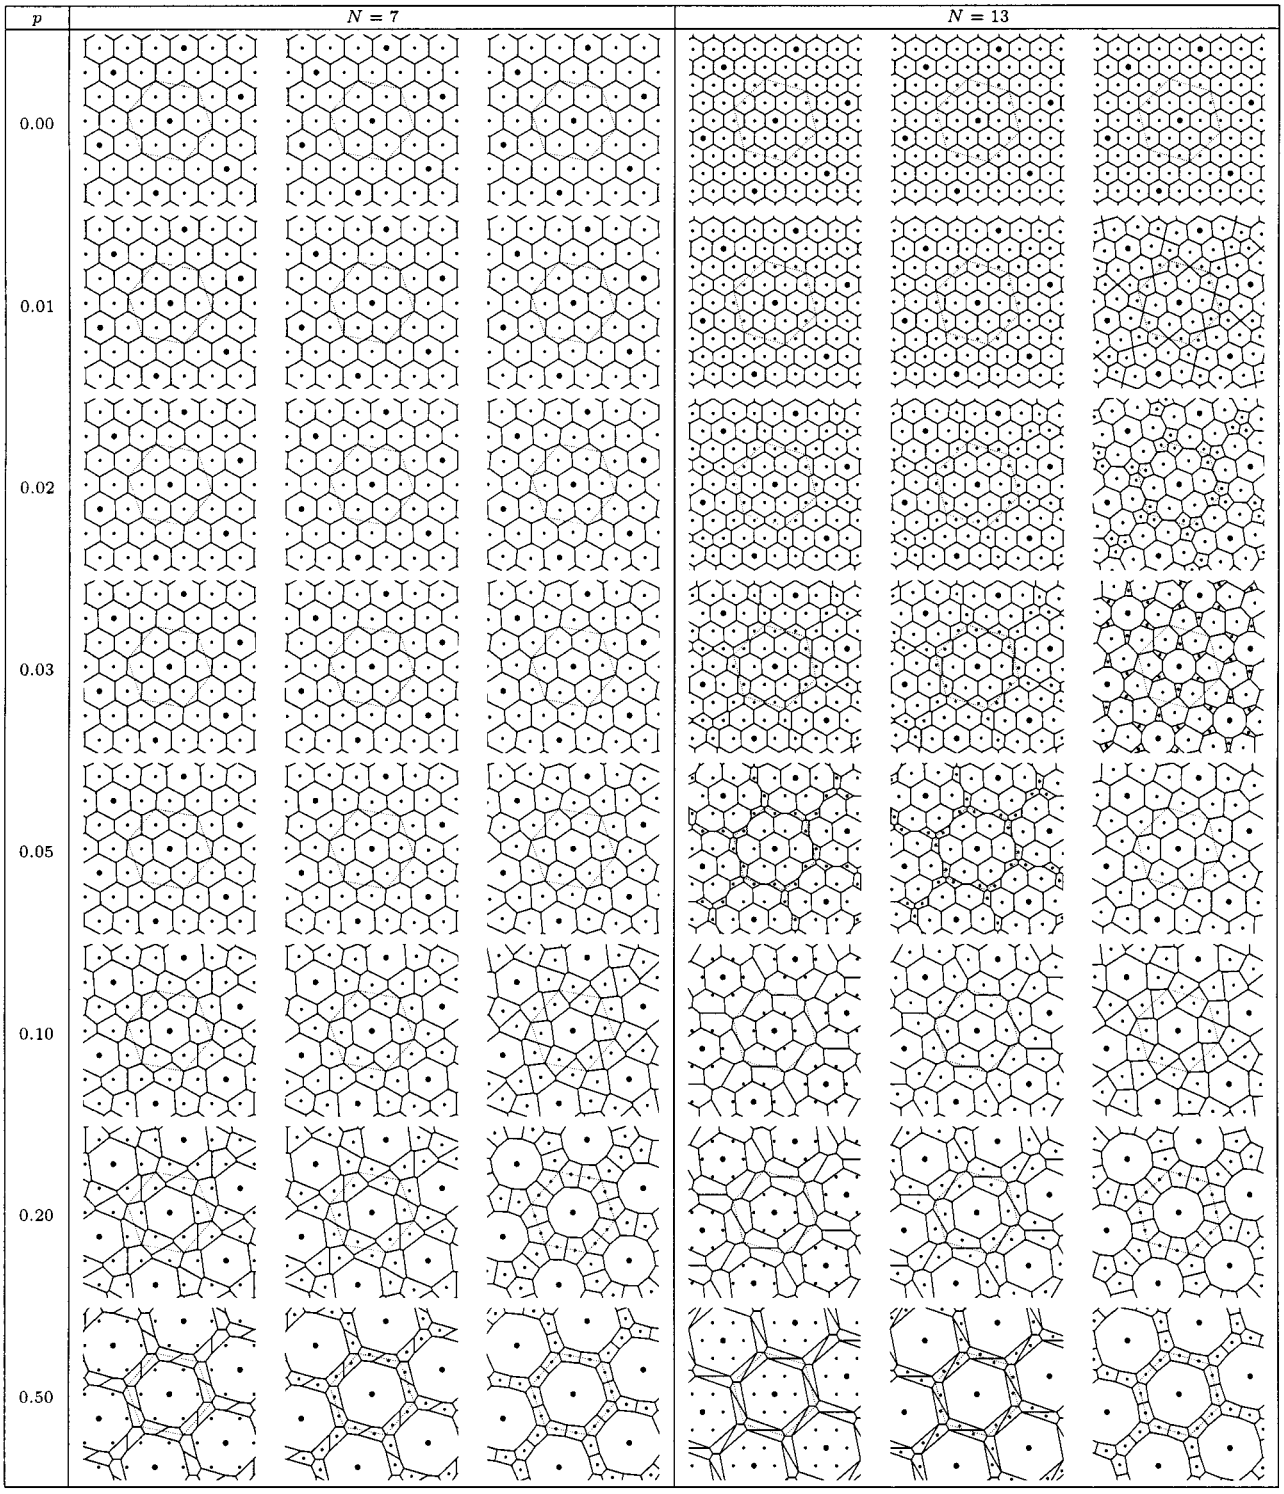

Fig. 6. Partitions created by MDLVQ and the three modified versions proposed in this correspondence for various values of the parameter p. Underlying all the partitions is the hexagonal lattice and the optimal index assignments of SVS. For each value of the sublattice index N, the left column shows the partitioning with modified encoding, the middle column gives the same partition with centroid reconstructions marked, and the right column shows the result of an iterative optimization of the points in $\Lambda \setminus \Lambda'$. In each diagram, a dotted hexagon demarkates a fundamental cell that tiles $\mathbb{R}^2$. An animated version of this figure, with many more values of p, can be found at http://lcavwww.epfl.ch/~goyal/MDVQ/.

Of course, altering the central decoder necessitates a corresponding alteration of the encoder. We are lead to an iterative design as is typical in VQ. Note that the steps of the iteration are very simple; because of symmetries, there is just a single displacement to be tracked for each orbit. In the A_2 index-7 and -13 cases, the convergence is very fast. The performance is shown by the bottom curves in Fig. 7 and a few sample encoder partitions are shown in the third and sixth columns of Fig. 6. Note in particular that the index-13 quantizers are identical to

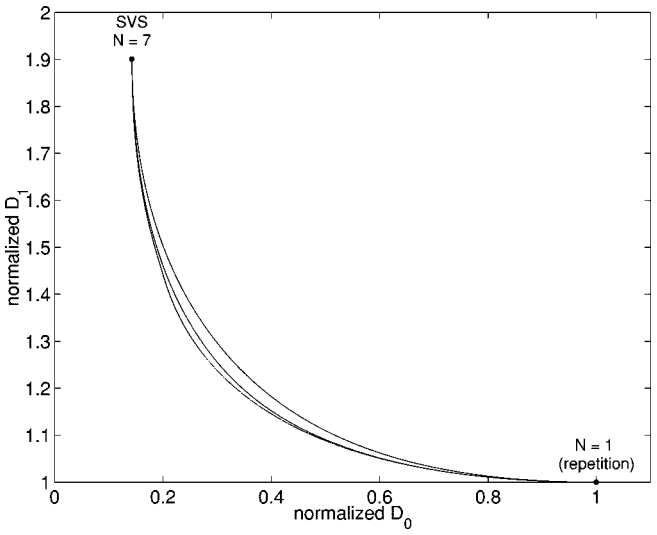

(a)

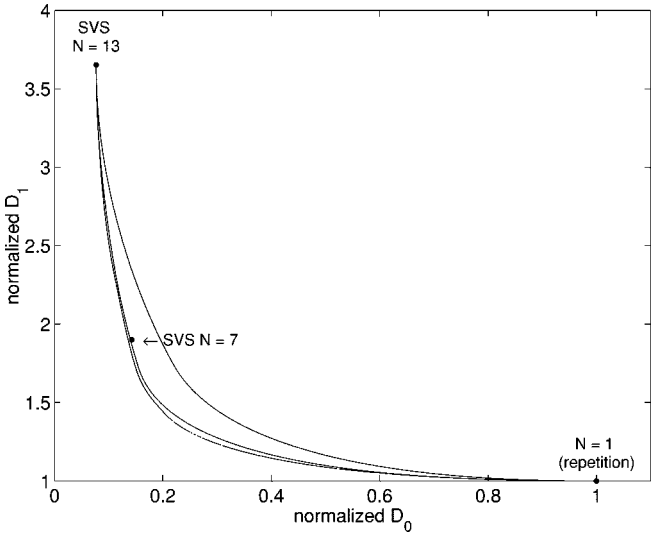

(b)

Fig. 7. Demonstrations of improved performance with the A_2 lattice and (a) $N = 7$ or (b) $N = 13$. As in Fig. 4, distortions are normalized to 0 dB when a description is repeated over both channels. The top (worst) curves arise from using only the improved encoding with the original lattice (minimization of $\overline{D}_p$). The middle curves are obtained by also improving the central decoder. Iteratively optimizing the points in $\Lambda \setminus \Lambda'$ and the central decoder gives the bottom (best) curves.

the index-7 quantizers for $p > 0.047$; outer orbit cells, labeled `ha`, `ai`, `ja`, etc., in Fig. 3(c), become empty because the corresponding coarse lattice points are far.

Because the sublattice points are not altered, the encoding complexity is hardly changed. For example, in the A_2, index-7 case, the encoder can again quantize to the nearest element of Λ' and then search among 13 possibilities.

V. Assessing the Improvement

Thus far, we have explored one method for obtaining a variety of MDLVQs: Start with a lattice Λ and vary the index of the sublattice Λ'. (Given Λ', there is nothing to be gained from looking at suboptimal index assignments mappings.) From each of these discrete MDLVQs, the methods of Sections III and IV generate a continuum of MD VQs

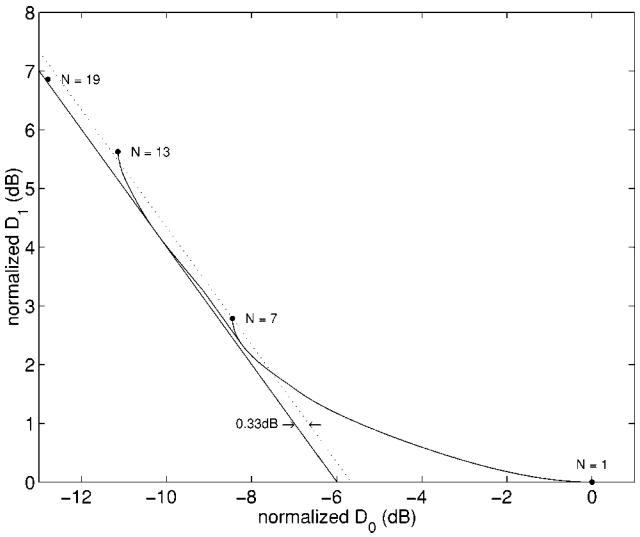

Fig. 8. Estimating the improvement in distortion product D_0D_1 from the optimization of $\Lambda \setminus \Lambda'$ in Section IV. Since distortions are plotted in decibels, straight lines with slope -1 have constant D_0D_1. The dotted line passes through the SVS index-7 point and the solid line is tangent to the operating points obtained in Section IV. The improvement is approximately 0.33 dB.

that are superior to time sharing between the given MDLVQ and repetition of one description over both channels. (Repetition can be put in the notation of MDLVQ with $\Lambda' = \Lambda$ and $\ell(\lambda) = (\lambda, \lambda)$.)

An alternative to our modifications is to choose a different starting lattice Λ_2. This will give another set of MDLVQs. If the dimensionless second moment of the new lattice is lower than that of the old lattice, most of the new operating points will be better than corresponding points in the convex hull of the original operating points. In fact, the decrease in the distortion product D_0D_1 is approximately the ratio of the dimensionless second moments $G(\Lambda_2)/G(\Lambda)$ [2]. The weakness of this approach is that reducing the dimensionless second moment will generally require an increase in vector dimension, which will increase the complexities of designing ℓ and encoding. Furthermore, the reduction of the dimensionless second moment with increasing vector dimension is slow [17], [22].

Between increasing the lattice dimension and straying from lattice encoding, which will improve performance with a lesser increase in complexity? Since we have no way to relate the improvements obtained in Sections III and IV to $G(\Lambda)$ or the dimension K, it is hard to come to any conclusions without designing and implementing both.

Since a complete understanding of the complexity–performance tradeoff continues to elude us, we conclude this section with an observation based on the A_2 lattice examples. Fig. 8 reveals the reduction in distortion product D_0D_1 obtained in Section IV.[2] The peak reduction is about 0.33 dB, and it appears that it would increase and hold valid for a larger range of operating points if index-19 quantizers were generated. To crudely relate this improvement to an equivalent increase in dimension, note that $G(E_6)$ is approximately 0.33 dB lower than $G(A_2)$ [22]; i.e., the improvement is roughly the same as that of changing Λ to the best known six-dimensional lattice.

VI. Low-Rate Performance

At low rates, (1) does not accurately estimate the performance of LVQ and, likewise, MD quantization systems cannot be analyzed on the basis of second moments alone. Rate and distortion computations

[2]This figure puts the data from Fig. 7 on a log–log scale so that lines with slope -1 have constant D_0D_1. Compare to Fig. 4 to see the improvement from Section III to IV.

are difficult because they require integrations involving the source density over each partition cell. Thus, one must usually be satisfied with simulations.

In work reported in [21] and [23], MDLVQ and the modified version proposed in this correspondence were applied to psychoacoustically prefiltered audio signals. Although the lattice scale relative to the variance of the source was not very fine, the central and side distortions varied with the parameter p as predicted by the high-resolution theory. In addition, the actual rates decreased relative to the predicted rates as p was increased. Thus, the advantage of the proposed method over the original SVS method was greater than predicted by high-resolution analysis.

This phenomenon is explained by Fig. 6, assuming the source density is peaked at the origin and the partition diagrams are centered at the origin. As p is increased, the probability of the central cell increases. Thus, the densities of the descriptions (the discrete variables $\ell_i(\alpha_p(x))$, $i = 1, 2$) become more peaked and the rates decrease.

VII. Conclusion

This correspondence introduces a method for two-channel MD coding that generalizes the MDLVQs developed by Servetto, Vaishampayan, and Sloane (SVS) [1]. SVS use a fine lattice Λ and a coarse sublattice Λ'. The new method uses the index assignments of SVS and a coarse lattice Λ', but allows the finer codebook Λ to deviate from a lattice. The overall result is a continuum of operating points to replace each discrete operating point of SVS and an improved convex hull of operating points.

The increase in complexity is not at all comparable to the increase that would be incurred by doing away with lattices altogether. Unconstrained techniques like those in [13] have complexity exponential in dimension and rate. As in the cases of LVQ and MDLVQ, the complexity of the proposed techniques does not depend at all on the rate. Using fast encoding techniques of LVQ, the dependence on dimension becomes much milder as well.

The principles applied in generalizing the SVS technique apply to unbalanced two-channel MD quantizers, as in [20], and to any number of channels.

Acknowledgment

The authors wish to thank S. D. Servetto for graciously providing numerical data. They also wish to thank F. Masson and G. Schuller for their collaboration in applying MDLVQ to audio, which inspired the comments on low-rate performance.

References

[1] S. D. Servetto, V. A. Vaishampayan, and N. J. A. Sloane, "Multiple description lattice vector quantization," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 1999, pp. 13–22.

[2] V. A. Vaishampayan, N. J. A. Sloane, and S. D. Servetto, "Multiple-description vector quantization with lattice codebooks: Design and analysis," *IEEE Trans. Inform. Theory*, vol. 47, pp. 1718–1734, July 2001.

[3] A. A. El Gamal and T. M. Cover, "Achievable rates for multiple descriptions," *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 851–857, Nov. 1982.

[4] L. Ozarow, "On a source-coding problem with two channels and three receivers," *Bell Syst. Tech. J.*, vol. 59, no. 10, pp. 1909–1921, Dec. 1980.

[5] R. Ahlswede, "The rate distortion region for multiple descriptions without excess rate," *IEEE Trans. Inform. Theory*, vol. IT-31, pp. 721–726, Nov. 1985.

[6] Z. Zhang and T. Berger, "Multiple description source coding with no excess marginal rate," *IEEE Trans. Inform. Theory*, vol. 41, pp. 349–357, Mar. 1995.

[7] V. A. Vaishampayan, "Design of multiple description scalar quantizers," *IEEE Trans. Inform. Theory*, vol. 39, pp. 821–834, May 1993.

[8] V. K Goyal and J. Kovačević, "Generalized multiple description coding with correlating transforms," *IEEE Trans. Inform. Theory*, vol. 47, pp. 2199–2224, Sept. 2001.

[9] R. Venkataramani, G. Kramer, and V. K Goyal, "Multiple description coding with many channels," *IEEE Trans. Inform. Theory*. In preparation; see [10], [11].

[10] ——, "Bounds on the achievable region for certain multiple description coding problems," in *Proc. IEEE Int. Symp. Information Theory*, Washington, DC, June 2001, p. 148.

[11] ——, "Successive refinement on trees: A special case of a new MD coding region," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 2001, pp. 293–301.

[12] V. K Goyal, "Multiple description coding: Compression meets the network," *IEEE Signal Processing Mag.*, vol. 18, pp. 74–93, Sept. 2001.

[13] M. Fleming and M. Effros, "Generalized multiple description vector quantization," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 1999, pp. 3–12.

[14] J. A. Kelner, V. K Goyal, and J. Kovačević, "Multiple description lattice vector quantization: Variations and extensions," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 2000, pp. 480–489.

[15] N. Moayeri and D. L. Neuhoff, "Time-memory tradeoffs in vector quantizer codebook searching based on decision trees," *IEEE Trans. Speech Audio Processing*, vol. 2, pp. 490–506, Oct. 1994.

[16] J. H. Conway and N. J. A. Sloane, "Fast quantizing and decoding algorithms for lattice quantizers and codes," *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 227–232, Mar. 1982.

[17] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Trans. Inform. Theory*, vol. 44, pp. 2325–2383, Oct. 1998.

[18] H. Gish and J. P. Pierce, "Asymptotically efficient quantizing," *IEEE Trans. Inform. Theory*, vol. IT-14, pp. 676–683, Sept. 1968.

[19] A. Gersho, "Asymptotically optimal block quantization," *IEEE Trans. Inform. Theory*, vol. IT-25, pp. 373–380, July 1979.

[20] S. N. Diggavi, N. J. A. Sloane, and V. A. Vaishampayan, "Design of asymmetric multiple description lattice vector quantizers," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 2000, pp. 490–499.

[21] G. Schuller, J. Kovačević, F. Masson, V. K Goyal, and J. A. Kelner, "Robust low-delay audio coding using multiple descriptions," paper, in preparation.

[22] J. H. Conway and N. J. A. Sloane, "Voronoi regions of lattices, second moments of polytopes, and quantization," *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 211–226, Mar. 1982.

[23] F. Masson, V. K Goyal, and J. Kovačević, "Multiple description robust low-delay speech and audio coding," Bell Labs, Lucent Technologies, Tech. Rep., 2000.