

Generalized Multiple Description Coding With Correlating Transforms

Vivek K Goyal, *Member, IEEE*, and Jelena Kovačević, *Senior Member, IEEE*

Abstract—Multiple description (MD) coding is source coding in which several descriptions of the source are produced such that various reconstruction qualities are obtained from different subsets of the descriptions. Unlike multiresolution or layered source coding, there is no hierarchy of descriptions; thus, MD coding is suitable for packet erasure channels or networks without priority provisions. Generalizing work by Orchard, Wang, Vaishampayan, and Reibman, a transform-based approach is developed for producing M descriptions of an N -tuple source, $M \leq N$. The descriptions are sets of transform coefficients, and the transform coefficients of different descriptions are correlated so that missing coefficients can be estimated. Several transform optimization results are presented for memoryless Gaussian sources, including a complete solution of the $N = 2, M = 2$ case with arbitrary weighting of the descriptions. The technique is effective only when independent components of the source have differing variances. Numerical studies show that this method performs well at low redundancies, as compared to uniform MD scalar quantization.

Index Terms—Erasure channels, integer-to-integer transforms, packet networks, robust source coding.

I. INTRODUCTION

SOURCE coding researchers are demanding consumers of communication systems. They ask for every bit they produce to be reliably delivered. Depending on what is known about the channel, this may be possible in Shannon's sense, but at what cost? At the very least, depending on the acceptable probability of failure and on how close the rate is to the channel capacity, large block sizes and complex encoding and decoding may be needed. Furthermore, compression may greatly increase the sensitivity to any remaining uncorrected errors.

A simple example is an English language text file. If a handful of characters are deleted at random, the reader may be distracted, but the meaning is likely to be fully conveyed. On the other hand, losing a few random bytes of a Lempel–Ziv compressed version of the story could be catastrophic. If the compression is by a factor of, say, ten, the effect is much more pronounced than the loss of ten times as many bytes. This suggests that if channel coding cannot make the probability of error zero, it may be beneficial to leave the data uncompressed.

Manuscript received March 2, 1999; revised February 20, 2001. The material in this paper was presented in part at the IEEE Data Compression Conference, Snowbird, UT, Mar. 1998.

V. K Goyal was with Bell Laboratories, Murray Hill, NJ 07974 USA. He is now with Digital Fountain, Fremont, CA 94538 USA (e-mail: v.goyal@ieee.org).

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

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

Publisher Item Identifier S 0018-9448(01)05462-1.

To not compress is a rather extreme reaction to the possibility of a bit error or erasure. A more temperate approach is to account for the possibility of (uncorrected) channel impairments in the source coding design. This paper addresses the problem of multiple description (MD) source coding, which can be cast as a source coding method for a channel whose end-to-end performance (with channel coding) includes uncorrected erasures. This channel is encountered in a packet communication system that has effective error detection but does not have retransmission of incorrect or lost packets. A transform-based method for MD source coding is introduced. This is a generalization to more arbitrary transforms of a technique for two channels (or packets, or descriptions) introduced by Orchard, Wang, Vaishampayan, and Reibman [1].

The organization of the paper is as follows. First, MD coding and its connection to packet communication are described in the remainder of this section. Then, in Section II, a few theoretical and practical results in MD coding are reviewed. Prior to the general analysis and optimization of Section IV, the two-tuple case is thoroughly developed in Section III. Applications to image and audio coding are summarized in Section V.

A. The Multiple Description Problem

At the September 1979 Shannon Theory Workshop, the following question was posed by Gersho, Ozarow, Witsenhausen, Wolf, Wyner, and Ziv [2].¹ If an information source is described by two separate descriptions, what are the concurrent limitations on qualities of these descriptions taken separately and jointly? Known at that time in Bell Laboratories as the channel splitting problem, this came to be known as the *multiple description (MD) problem* in the information theory community. The primary theoretical results in this area were provided in the 1980s by the aforementioned researchers along with Ahlswede, Berger, Cover, El Gamal, and Zhang.

MD coding refers to the scenario depicted in Fig. 1. An encoder is given a sequence of source symbols $\{X_k\}_{k=1}^N$ 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 (the *side decoders*) receive information only over their respective channels. The transmission rate over channel i is denoted by R_i , $i = 1, 2$; i.e., signaling over channel i uses at most

¹We may infer from [3]–[5] and private communications that the problem originated with Gersho and Witsenhausen. Gersho proposed a scheme (unpublished [6]) that utilized the redundancy in coded speech waveforms to mitigate the effects of a channel breakdown. Witsenhausen recognized that the achievable distortions are better than one might guess, even for a memoryless source.




Fig. 1. Scenario for MD source coding with two channels and three receivers. The general case has M channels and $2^M - 1$ receivers.

$2^N R_i$ symbols. Denoting by $\{\hat{X}_k^{(i)}\}_{k=1}^N$ the reconstruction sequence produced by decoder i , there are three distortions

$$D_i = \frac{1}{N} \sum_{k=1}^N E \left[\delta_i \left(X_k, \hat{X}_k^{(i)} \right) \right], \quad \text{for } i = 0, 1, 2$$

where the $\delta_i(\cdot, \cdot)$'s are potentially distinct, nonnegative, real-valued distortion measures.

The central theoretical problem is to determine the set of achievable values (in the usual Shannon sense) for the quintuple $(R_1, R_2, D_0, D_1, D_2)$. Decoder 1 receives R_1 bits and hence cannot have distortion less than $D(R_1)$, where $D(\cdot)$ is the distortion-rate function of the source. Making similar arguments for the other two decoders gives the following bounds on the achievable region:²

$$D_0 \geq D(R_1 + R_2) \quad (1)$$

$$D_1 \geq D(R_1) \quad (2)$$

$$D_2 \geq D(R_2). \quad (3)$$

Achieving equality simultaneously in (1)–(3) would imply that an optimal rate $R_1 + R_2$ description can be partitioned into optimal rate R_1 and rate R_2 descriptions. Unfortunately, this is usually not true because optimal individual descriptions at rates R_1 and R_2 are similar to each other and hence redundant when combined. Making descriptions individually good and yet sufficiently different is the fundamental tradeoff in this problem.

The MD problem can be generalized to more than two channels and more than three receivers. The natural extension is to M channels and $2^M - 1$ receivers—one receiver for each nonempty subset of channels. This generalization was considered by Witsenhausen [3] for the restricted case where the source has finite entropy rate and lossless communication is required when any k , $k < M$, of the channels are lost. Normalizing the source rate to one and assuming equal usage of each channel, each channel must accommodate a rate of $1/(M - k)$. (The rate cannot be lowered because the sum of the rates of the received channels must be at least one.) This bound is achieved by using truncated Reed–Solomon codes. A similar result holds in the more general setting of [7]. The situation with three channels and seven decoders was studied by Zhang and Berger [8].

The MD codes constructed in this paper apply to the generalized MD problem. Unfortunately, no tight achievability bounds are known for generalized MD coding.

²Since the distortion metrics may differ, the use of a single symbol $D(\cdot)$ for the distortion-rate function of the source is a slight abuse of notation. In this paper, squared-error per component is used exclusively, so there is no ambiguity.

B. Applicability to Packet Networks

Recently, the problem of transmitting data over heterogeneous packet networks has received considerable attention. A typical scenario might require data to move from a fiber link to a wireless link, which necessitates dropping packets to accommodate the lower capacity of the latter. If the network is able to provide preferential treatment to some packets, then the use of a multiresolution or layered source coding system is the obvious solution. But what if the network will not look inside packets and discriminate? Then packets will be dropped at random, and it is not clear how the source (or source-channel) coding should be designed. If packet retransmission is not an option (e.g., due to a delay constraint or lack of a feedback channel), one has to devise a way of making the packets that are received meaningful. The situation is similar if packets are lost due to transmission errors or congestion.

Drawing an analogy between packets and channels, packet communication with M packets is equivalent to generalized MD coding with M channels. Each of the $2^M - 1$ nonempty subsets of the packets leads to a potentially distinct reconstruction with some distortion. The case where no packet is received is ignored to maintain the analogy with the classical MD problem and because the source coding in that case is irrelevant. A recent surge of interest in MD coding seems to be due primarily to this application (see [1], [9]–[12]) and yet the present work was among the first to effectively use more than two packets [13]–[17].

Retransmission of lost packets, when feasible, is an effective technique for maintaining constant quality despite losses. The use of a retransmission protocol, such as TCP [18], requires at a minimum that a feedback channel is available to indicate which packets have been successfully received. Even if feedback is available, many factors may preclude the retransmission of lost or corrupted packets. Retransmission adds delay and necessitates additional buffering for streaming media. Retransmission is generally not feasible in broadcast environments because of the so-called feedback implosion problem whereby the loss of a single packet may spark many retransmission requests.

From an information-theoretic perspective, an idealized model is to assume that packet losses are independent and identically distributed (i.i.d.) with a known probability p and that the message sequence is arbitrarily long. Then, assuming that the packets have fixed payload of one unit, the capacity of the channel is $1 - p$ per channel use. Furthermore, this capacity can be attained by choosing a sequence of good (n, k) block codes with rate $k/n < 1 - p$, with $n, k \rightarrow \infty$. Attaining error-free transmission at a rate arbitrarily close to the channel capacity is intimately tied to having an arbitrarily long message sequence. For any finite code length, the probability of failure is nonzero.

The emphasis in this paper is on situations in which long block codes cannot be used, which is often the case. For example, consider a network using Internet Protocol, Version 6 (IPv6) [19]. An IPv6 node is required to handle 576-byte packets without fragmentation, and it is recommended that larger packets be accommodated.³ Accounting for packet head-

³Without the “Jumbo Payload” option, the maximum packet size is 65 575 bytes (65 535-byte payload).

ers, a 576-byte packet may have a payload as large as 536 bytes. With packets of this size, a typical image associated with a WWW page may be communicated in a handful of packets; say, eight.⁴ One cannot use the law of large numbers to analyze channel codes with only eight output symbols. Another reason for using short channel codes is to keep buffering requirements and delay small.

Aside from these issues, we should note that traditional channel coding approaches aim to completely counteract all channel impairments; but, for communication subject to a smooth distortion measure, we may impose a much less stringent requirement with the hope of achieving satisfactory performance with less coding overhead. The philosophy of replacing a traditional source code and channel code with an MD source code is to *mitigate* the effect of erasures even if not a single erasure can be *corrected*.

II. SURVEY OF MD CODING

A. Historical Notes

At the previously mentioned September 1979 meeting, Wyner presented preliminary results on MD coding obtained with Witsenhausen, Wolf, and Ziv for a binary source and Hamming distortion. At that very meeting, Cover and El Gamal determined and reported the achievable rate region later published in [2]. Ozarow's contribution was to show the tightness of the El Gamal–Cover region for memoryless Gaussian sources and squared-error distortion [5]. Subsequently, Ahlswede [21] showed that the El Gamal–Cover region is tight in the “no excess rate sum” case (where there is equality in (1)), and Zhang and Berger [8] showed that this region is not tight when there is excess rate. The complementary situation, where (2) and (3) hold with equality, is called the “no excess marginal rate” case and was also studied by Zhang and Berger [22].

MD coding includes as a special case the better known *successive refinement* or *multiresolution* coding. The successive refinement problem can also be described by Fig. 1, but the interest is only in characterizing achievable (R_1, R_2, D_0, D_1) . In other words, no attempt is made to estimate the source from channel 2 alone; or channel 1 is always present. Successive refinement was first studied by Koshelev [23]–[25]. The conditions for perfect successive refinement—where (1) and (2) hold with equality—are described in [26]. The result follows from the tightness of the achievable region of [2] for the no excess rate sum case [21] (see also [27]).

Finally, it should also be noted that there is a substantial literature on the error sensitivity of compressed data and more generally on tradeoffs between source and channel coding. The reader is referred to [28]–[32] for a sampling of the results.

B. Theoretical Bounds for a Memoryless Gaussian Source

The achievable rate-distortion region is completely known only for a memoryless Gaussian source with mean-squared error

(MSE) distortion. This result, obtained by Ozarow [5] is summarized by the following theorem.⁵

Theorem 1: Let X_1, X_2, \dots be a sequence of i.i.d. unit variance Gaussian random variables. The achievable set of rates and MSE distortions is the union of points satisfying

$$D_i \geq 2^{-2R_i}, \quad i = 1, 2 \quad (4)$$

$$D_0 \geq 2^{-2(R_1+R_2)} \cdot \gamma(D_1, D_2, R_1, R_2) \quad (5)$$

where

$$\gamma = \frac{1}{1 - \left(\sqrt{(1-D_1)(1-D_2)} - \sqrt{D_1 D_2 - 2^{-2(R_1+R_2)}} \right)^2}$$

for $D_1 + D_2 < 1 + 2^{-2(R_1+R_2)}$ and $\gamma = 1$ otherwise.

Since Theorem 1 is the key result in coding continuous-valued sources, the region defined therein warrants a closer look. The bounds (4) are simply the side-channel rate-distortion bounds, a repeat of (2) and (3). In the final inequality (5), the central distortion must exceed the rate-distortion bound by the factor γ .

A few examples will clarify the behavior of γ and the resulting properties of the achievable region. First, suppose that the descriptions are individually good, yielding $D_1 = 2^{-2R_1}$ and $D_2 = 2^{-2R_2}$. Then we may write

$$D_0 \geq D_1 D_2 \frac{1}{1 - (1-D_1)(1-D_2)} = \frac{D_1 D_2}{D_1 + D_2 - D_1 D_2}.$$

A further chain of inequalities gives $D_0 \geq \min\{D_1, D_2\}/2$, so the joint description is only slightly better than the better of the two individual descriptions.

On the other hand, suppose the joint description is as good as possible, so $D_0 = 2^{-2(R_1+R_2)}$. Then $\gamma = 1$, and thus

$$D_1 + D_2 \geq 1 + 2^{-2(R_1+R_2)}. \quad (6)$$

Recall that a distortion value of 1 is obtained with no information, simply estimating the source by its mean. For anything but a very low rate, (6) implies a very poor reconstruction for at least one of the side decoders.

Intermediate to these two extremes, the region of Theorem 1 is easier to understand in the *balanced* case, where $R_1 = R_2$ and $D_1 = D_2$. We will take two points of view to improve our qualitative understanding of the bounds. The first relates exponential decay exponents in the three distortions and the second bounds the side distortion based on the gap between D_0 and $2^{-2(R_1+R_2)}$.

First, estimate γ under the assumptions $R_1 = R_2 \gg 1$ and $D_1 = D_2 \ll 1$. Then

$$\begin{aligned} \frac{1}{\gamma} &= 1 - \left((1-D_1) - \sqrt{D_1^2 - 2^{-4R_1}} \right)^2 \\ &\approx 1 - ((1-D_1) - D_1)^2 = 4D_1 - 4D_1^2 \approx 4D_1. \end{aligned}$$

⁵Ozarow [5] neglects the high side distortion regime where

$$D_1 + D_2 \geq 1 + 2^{-2(R_1+R_2)}.$$

In the converse, the developments in [5, p. 1917] require (in Ozarow's notation) $\Pi \geq \Delta$, but this is not specified.

⁴This is not an arbitrary figure. Based on statistics in [20], the average WWW image is only 3.9 kbytes; an image of this size would not be fragmented into more than eight packets. Of course, this depends on the sampling of the WWW and may change.




Fig. 2. Lower bound (8) on side distortions $D_1 = D_2$ when the base rate is r , i.e., the central distortion is $D_0 = 2^{-2r}$. At high redundancies, the curves merge because the bound is independent of r .

Substituting $\gamma = (4D_1)^{-1}$ in (5) gives $D_0 \geq 2^{-4R_1}(4D_1)^{-1}$, so the product of central and side distortions is approximately lower-bounded by $4^{-1}2^{-4R_1}$. If $D_1 = D_2 \approx 2^{-2(1-\alpha)R_1}$ with $0 < \alpha \leq 1$, the best decay of the central distortion is

$$D_0 \approx 4^{-1}2^{-2(1+\alpha)R_1}.$$

This shows that the penalty in the exponential rate of decay of D_1 (the difference from the optimal decay rate, indicated by α being positive) is precisely the increase in the rate of decay of D_0 .

The second way to interpret the tradeoff between central and side distortion begins with rearranging (5), under the assumption of equal side distortion $D_1 = D_2$, to obtain [33]

$$D_1 \geq \min \left\{ \frac{1}{2} \left[1 + D_0 - (1 - D_0) \sqrt{1 - 2^{-2(R_1+R_2)}/D_0} \right], \right. \\ \left. 1 - \sqrt{1 - 2^{-2(R_1+R_2)}/D_0} \right\}. \quad (7)$$

Now separate the coding rate into a *base rate* r and a *redundancy* ρ^* . The base rate corresponds to the quality of the joint description and the redundancy is the additional rate added to make the side distortions low. Mathematically, $D_0 = 2^{-2r}$ and $\rho^* = R_1 + R_2 - r$. Substituting in (7) gives

$$D_1 \geq \begin{cases} \frac{1}{2} [1 + 2^{-2r} - (1 - 2^{-2r}) \sqrt{1 - 2^{-2\rho^*}}], & \text{for } \rho^* \leq r - 1 + \log_2(1 + 2^{-2r}) \\ 1 - \sqrt{1 - 2^{-2\rho^*}}, & \text{for } \rho^* > r - 1 + \log_2(1 + 2^{-2r}). \end{cases} \quad (8)$$

The second expression comes into play when the redundancy is so large that there is no conflict in achieving the central and side distortions. This bound is plotted in Fig. 2 for several values of the base rate r .

One interesting thing to glean from (8) is the slope of the low-redundancy D_1 versus ρ^* characteristic

$$\frac{\partial D_1}{\partial \rho^*} = -\frac{1 - 2^{-2r}}{2} \frac{2^{-2\rho^*} \ln 2}{\sqrt{1 - 2^{-2\rho^*}}}.$$

At $\rho^* = 0^+$, the slope is infinite. Consider starting with a system designed only to minimize central distortion. Assuming that side distortions are also of interest, the infinite slope means that a small additional rate will have much more impact if dedicated to reducing the side distortion than if dedicated to reducing the central distortion.

The rate-distortion region is not completely known for any non-Gaussian source. Zamir [34] has found inner and outer bounds to the achievable rate region for MD coding of any continuous-valued memoryless source with squared-error distortion. This is an extension of Shannon's bounds on rate-distortion functions (see [35], [60], [36]).

C. Practical Codes

The bounds of the previous section can be approached when a long sequence of source symbols is coded. In contrast, the focus in this paper is on simple codes for finite-length source vectors.

1) *MD Scalar Quantization (MDSQ)*: The first method for MD coding of a memoryless, continuous-valued source was MD scalar quantization (MDSQ). MDSQ is the use of two separate scalar quantizers to give two descriptions of a scalar source sample, with an additional central decoder that makes use of both descriptions. This approach was expounded and popularized by Vaishampayan [37], though Reudink introduced very similar methods in earlier unpublished work [38].

The simplest example is to have scalar quantizers with nested thresholds, as shown in Fig. 3(a). Each quantizer outputs an index that can be used by itself to estimate the source sample. Using $Q_i: \mathbb{R} \rightarrow \{1, 2, \dots, 6\}$, $i = 1, 2$, to denote the encoding map of quantizer i , the reconstruction knowing $Q_i(x) = k_i$ should be the centroid of the cell $Q_i^{-1}(k_i)$. The central decoder has both $Q_1(x) = k_1$ and $Q_2(x) = k_2$ and thus reconstructs to the centroid of the intersection cell $Q_2^{-1}(k_1) \cap Q_2^{-1}(k_2)$. In the example, the intersection cells are half as big as the individual quantizer cells, so the central distortion is about a quarter of the side distortions. Asymptotically, if the side rates are $R_1 = R_2 = R$, then D_0 , D_1 , and D_2 are all $O(2^{-2R})$. This is optimal decay for D_1 and D_2 , but far from optimal for D_0 .

Recalling the discussion following Theorem 1, it should be possible to speed the decay of D_0 at the expense of slowing the decay of D_1 and/or D_2 . Let n_i , $i = 1, 2$, denote the number of cells in quantizer Q_i . Let n_0 denote the number of intersections between cells of Q_1 and Q_2 that are nonempty. Notice that in Fig. 3(a), $n_0 = n_1 + n_2 - 1$. When $n_1 = n_2 = n$, the exponential rate of decay of D_0 is changed only if n_0 grows faster than linearly with n . Accomplishing this requires something never seen in single-description scalar or vector quantization: disconnected partition cells. The maximum number of central decoder partition cells is $n_0 = n_1 n_2$. This occurs when each $Q_1^{-1}(k_1) \cap Q_2^{-1}(k_2)$ is nonempty, as shown in Fig. 3(b). Quantizer Q_2 is individually poor. The asymptotic performance




Fig. 3. Three MD scalar quantizers. (a) The simplest form of MDSQ, with nested quantization thresholds. When $R_1 = R_2 = R$, all three distortions, D_0 , D_1 , and D_2 , are $O(2^{-2R})$. (b) An MDSQ which minimizes D_0 for a given rate. Asymptotically, $D_0 = O(2^{-4R})$, but at least one of D_1 and D_2 must be $O(1)$. (c) An MDSQ based on Vaishampayan's "modified nested" index assignment [37]. This construction systematically trades off central and side distortions while maintaining optimal joint asymptotic decay of distortion with rate.

of this scheme with $R_1 = R_2 = R$ is at the opposite extreme of the previous example; $D_0 = O(2^{-4R})$, but at least one of D_1 and D_2 is $O(1)$.

Given a desired partitioning for the central encoder,⁶ the crux of MDSQ design is the assignment of indexes to the individual quantizers. Vaishampayan's main results in [37] are this observation and an idealized index assignment scheme that gives the optimal combined exponential decay rates for the central and side distortions. An MDSQ designed with Vaishampayan's "modified nested" index assignment is shown in Fig. 3(c). In contrast to the MDSQ of Fig. 3(b), the side distortions are approximately equal and the quantizers "refine each other" in a symmetric fashion. For a given number of side levels n , the central distortion is smaller—at the cost of higher side distortions—than for an MDSQ as in Fig. 3(a).

The optimal design for fixed-rate MDSQ in [37] was extended to entropy-constrained MDSQ in [39]. One of the most satisfying aspects of the theory of MDSQ is that at high rates, all the decay exponent tradeoffs discussed in Section II-B can be obtained. Furthermore, the factor by which the distortion product D_0D_1 exceeds the bounds of Theorem 1 is approximately constant [40]. For $R = R_1 = R_2$ and $D_1 = D_2$, the bound for large rates is approximately

$$D_0D_1 \geq \frac{1}{4}2^{-4R} \quad (9)$$

⁶There is no "central encoder," but Q_1 and Q_2 effectively implement a quantizer with cells given by the intersections of the cells of the individual quantizers. An MDSQ could be viewed as a single quantizer that outputs two-tuple indexes.

while the performance of optimal fixed-rate MDSQ is

$$D_0D_1 \approx \frac{1}{4} \left(\frac{\sqrt{3}\pi}{2} \right)^2 2^{-4R}$$

and the performance of optimal entropy-constrained MDSQ is

$$D_0D_1 \approx \frac{1}{4} \left(\frac{\pi e}{6} \right)^2 2^{-4R}. \quad (10)$$

The factors in parentheses are familiar from high-resolution quantization theory (see, e.g., [41] and the references therein).

These features of high-rate MDSQ, however, tell us very little about how MDSQ performs at low rates. Later, in Section III-E, the performance of MD based on correlating transforms will be compared to the performance of MDSQ at low and intermediate rates. These numerical comparisons are based on a family of uniform MDSQs (UMDSQs) described in detail in [33]. The encoder for a UMDSQ operates as follows. A source sample is first uniformly quantized by rounding off to the nearest multiple of a step size Δ ; then the index output by the scalar quantizer is mapped to a pair of indexes using an index assignment based on [40]; finally, these indexes are entropy coded. In all of the numerical results presented here, rate is measured by entropy rather than by the average code length of a particular code. Seven index assignments are used, giving seven operating points for each value of Δ ; in the notation of [40], the index assignments could be labeled $k = 0, \frac{1}{2}, \dots, 3$. At the decoder, we allow either centroid reconstruction with respect to a uniform source density (which is called midpoint reconstruction, since it is the natural extension of midpoint reconstruction to MD) or centroid reconstruction with respect to the actual Gaussian density.

The performance of UMDSQ is compared to the bounds of Theorem 1 in Fig. 4. The solid curve in Fig. 4(a) is the side distortion bound (7) evaluated for central distortion $D_0 = 2^{-12}$. The dashed line is the high-rate approximation of MDSQ performance (10) and the dotted line is the high-rate bound (9). Operating points obtained with midpoint and centroid reconstruction are marked with dots and circles, respectively.

At this low central distortion (high base rate $r = 6$ bits), all seven considered index assignments give reasonably useful operating points and centroid reconstruction does not greatly improve the performance over midpoint reconstruction. Note also that the performance is well-approximated by high-rate analysis. However, at low and very high redundancies the performance is worse than predicted by (10). The highest redundancy point is obtained with identical descriptions over the two channels, which is not very clever.

At lower base rates, the initial quantization is more coarse; hence, less of the index assignments are useful and it is more important to have centroid reconstruction. Fig. 4(b) shows performance and bounds for $D_0 = 2^{-4}$. With midpoint reconstruction, only the repetition of identical information over each channel and the simplest nested index assignment of Fig. 3(a) give side distortion less than 1 (other operating points are not shown); with centroid reconstruction there are only three useful operating points. Fig. 4(b) clearly indicates the deficiency of




Fig. 4. Performance of entropy-constrained uniform MDSQ. At low central distortion or high base rate (part (a)), the high-rate approximate bound (9) approximates the true bound at low redundancies and the high-resolution approximate performance (10) is accurate. At higher central distortion or lower base rate (part (b)), the true bound is considerably more stringent than the high-rate approximation, and the gap between actual performance and the high-resolution approximation is larger. Only useful operating points are given, i.e., those with side distortion less than 1 and no operating point with lower side distortion and lower redundancy. At low base rates, uniform MDSQ gives few useful operating points. (a) Central distortion $D_0 = 2^{-12}$. (b) Central distortion $D_0 = 2^{-4}$.

UMDSQ in providing useful operating points for low base rates and redundancies.

Applications of MDSQ are described in [9], [42], [12]. Joint optimization of an orthogonal transform and MDSQ is discussed in [43]. For the purposes of this paper, it is sufficient to note that an identity transform is optimal for a source vector with independent components. Note that Batllo and Vaishampayan use the term *MD transform coding* (MDTC) to refer to MDSQ applied to coefficients from an orthogonal transform. The same term could be used to describe the techniques developed here or in [14], [16], [17]; however, to avoid confusion the term *MD correlating transform* (MDCT) is used for the present technique.

2) *Pairwise Correlating Transforms*: A considerably different approach to MD coding was introduced by Wang, Orchard, and Reibman [10]. Instead of using MDSQ to produce two indexes that describe the same quantity, the MD character

is achieved with a linear transform that introduces correlation between a pair of random variables; quantization is treated as secondary.

Let X_1 and X_2 be independent zero-mean Gaussian random variables with variances $\sigma_1^2 > \sigma_2^2$. For conventional (single-description) source coding, there would be no advantage to using a linear transform prior to quantization. Assuming high-rate entropy-coded uniform quantization, the MSE distortion per component at R bits per sample would be given by⁷

$$D_0 = \frac{\pi e}{6} \sigma_1 \sigma_2 2^{-2R}.$$

This is the best single-description performance that can be obtained with scalar quantization.

Now suppose that the quantized versions of X_1 and X_2 are sent on channels 1 and 2, respectively, in an MD system. Since X_1 and X_2 are independent, side decoder 1 cannot estimate X_2 , aside from using its mean. Thus,

$$D_1 = \frac{\pi e}{12} \sigma_1 \sigma_2 2^{-2R} + \frac{1}{2} \sigma_2^2$$

and, similarly,

$$D_2 = \frac{\pi e}{12} \sigma_1 \sigma_2 2^{-2R} + \frac{1}{2} \sigma_1^2.$$

Assume for the moment that each channel is equally likely to fail. Then, instead of concerning ourselves with D_1 and D_2 separately, we will use the average distortion when one channel is lost

$$\bar{D}_1 = \frac{1}{2} (D_1 + D_2) = \frac{1}{4} (\sigma_1^2 + \sigma_2^2) + \frac{\pi e}{12} \sigma_1 \sigma_2 2^{-2R}. \quad (11)$$

\bar{D}_1 could be reduced if side decoder i had some information about X_j , $i \neq j$. This can be accomplished by transmitting not X_i 's, but correlated transform coefficients. The simplest possibility, as proposed in [10], is to transmit quantized versions of Y_1 and Y_2 given by

$$\begin{bmatrix} Y_1 \\ Y_2 \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}.$$

Since $[Y_1 \ Y_2]^T$ is obtained with an orthonormal transformation and we are using MSE distortion, the distortion in approximating the X_i 's equals the distortion in approximating the Y_i 's. The variances of Y_1 and Y_2 are both $(\sigma_1^2 + \sigma_2^2)/2$, so the central decoder performance is

$$D_0 = \frac{\pi e}{6} \left(\frac{\sigma_1^2 + \sigma_2^2}{2} \right) 2^{-2R}$$

which is worse than the performance without the transform by a constant factor of⁸

$$\Gamma = \frac{(\sigma_1^2 + \sigma_2^2)/2}{\sigma_1 \sigma_2}.$$

Now consider the situation at side decoder 1. The distortion is approximately equal to the quantization error plus the distortion in estimating Y_2 from Y_1 . Since Y_1 and Y_2 are jointly Gaussian,

⁷Standard high-rate operational rate-distortion estimates will be used throughout the paper. The reader is referred to [41] for details.

⁸This factor is like a coding gain, but it is the *increase* in distortion from coding correlated quantities.

$Y_2 | Y_1 = y_1$ is Gaussian and $E[Y_2 | Y_1 = y_1]$ is a linear function of y_1 . Specifically, $Y_2 | Y_1 = y_1$ has mean

$$(\sigma_1^2 + \sigma_2^2)^{-1}(\sigma_1^2 - \sigma_2^2)y_1$$

and variance

$$2(\sigma_1^2 + \sigma_2^2)^{-1}\sigma_1^2\sigma_2^2.$$

Thus,

$$\bar{D}_1 \approx \underbrace{\frac{\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2}}_{\frac{1}{2}E[|Y_2 - \hat{Y}_2|^2]} + \underbrace{\frac{\pi e}{12} \left(\frac{\sigma_1^2 + \sigma_2^2}{2} \right)}_{\frac{1}{2}E[|Y_1 - \hat{Y}_1|^2]} 2^{-2R}. \quad (12)$$

Comparing (11) and (12), the constant term has been reduced by the factor Γ^2 , but the exponential term is increased by the factor Γ . This method thus improves the side distortion while it degrades the central distortion by a related amount. By using other orthogonal transforms, intermediate tradeoffs can be obtained. As demonstrated in Section III, nonorthogonal transforms allow yet more operating points, including more extreme tradeoffs. The factor Γ approaches unity as σ_1/σ_2 approaches unity; the pairwise correlation has no effect when the joint density of X_1 and X_2 has spherical symmetry. Even with nonorthogonal transforms, the potential reduction in side distortion vanishes as $\sigma_1/\sigma_2 \rightarrow 1$.

The high-rate asymptotic behavior of the pairwise correlating transform method is not interesting because, independent of the choice of the transform, $D_0 = O(2^{-2R})$ and $D_1 = D_2 = O(1)$. However, in practice, high rates are not necessarily important and constant factors can be very important. As a case in point, this method was introduced in the context of loss-resilient image coding, where it was shown to be successful [10]. Subsequently, this method was extended to nonorthogonal transforms [1]. The contributions of this paper are to generalize this method to communicating N variables over M channels, $M \leq N$, where the channel failures may be dependent and may have unequal probabilities, and to provide new results for the two-channel case.

III. CODING OF TWO-TUPLES

This section takes an in-depth look at the MD coding of two-tuple source vectors using a discrete correlating transform. The generalization to arbitrary-length vectors is considered in Section IV.

A. Intuitive Development

The limitations of the pairwise correlating transform method of Section II-C2 led to the work reported in [1]. As in Section II-C2, let $X = [X_1, X_2]^T$ where X_1 and X_2 are independent zero-mean Gaussian random variables with variances $\sigma_1^2 > \sigma_2^2$. Let $\{e_1, e_2\}$ denote the standard basis of \mathbb{R}^2 . Any level curve of the joint probability density function (pdf) of X is an ellipse with principal axis aligned with e_1 and secondary axis aligned with e_2 . Using the standard basis corresponds to representing X by $(\langle X, e_1 \rangle, \langle X, e_2 \rangle)$ (see Fig. 5(a)).

Now imagine that uniform scalar quantized versions of $\langle X, e_1 \rangle$ and $\langle X, e_2 \rangle$ are used as descriptions. It was demonstrated in Section II-C2 that, for a given total rate, the average




Fig. 5. Basis configurations for correlating transforms. (a) The standard basis gives high side distortion when the component with high variance is lost. (b) Basis for the original correlating transform of [10]. (c) Generalization to arbitrary orthogonal bases. (d) When descriptions are lost with equal probability, the best bases are symmetric with respect to the principal axis of the source density.

of the side distortions $\bar{D}_1 = (D_1 + D_2)/2$ can be decreased in exchange for an increase in the central distortion D_0 by using the representation

$$\left(\left\langle X, \frac{1}{\sqrt{2}} [1, 1]^T \right\rangle, \left\langle X, \frac{1}{\sqrt{2}} [-1, 1]^T \right\rangle \right). \quad (13)$$

But this is only a single operating point, whereas one would like to be able to trade off D_0 and \bar{D}_1 in a continuous manner.

Recognizing (13) as $(\langle X, G_{\pi/4}e_1 \rangle, \langle X, G_{\pi/4}e_2 \rangle)$, where G_θ is a Givens rotation of angle θ (see Fig. 5(b)), a natural extension is to consider all representations of the form

$$(\langle X, G_\theta e_1 \rangle, \langle X, G_\theta e_2 \rangle), \quad \text{for } 0 \leq \theta \leq \pi/4.$$

This indeed creates a continuous tradeoff between D_0 and \bar{D}_1 . However, it has an undesirable asymmetry. For $0 \leq \theta < \pi/4$, the side distortions are not equal. It is an easy exercise to calculate D_1 and D_2 , but instead let us look geometrically at why they are unequal. D_1 is the variation of X which is *not* captured by $\langle X, G_\theta e_1 \rangle$, or the variation perpendicular to $G_\theta e_1$.⁹ Similarly, D_2 is the variation perpendicular to $G_\theta e_2$. Now since $G_\theta e_1$ and $G_\theta e_2$ are not symmetrically situated with respect to the pdf of X (except for $\theta = \pi/4$), D_1 and D_2 are unequal (see Fig. 5(c)).

This in itself does not imply that the scheme can be improved, but since we are trying to have two channels of equal importance¹⁰ we might expect equal side distortions. Based on the geometric observation above, it makes sense to represent X by

$$(\langle X, G_\theta e_1 \rangle, \langle X, G_{-\theta} e_1 \rangle), \quad \text{for } 0 < \theta < \pi/2. \quad (14)$$

⁹We are neglecting quantization error at this point because D_1 and D_2 are equally affected by quantization.

¹⁰“Equal importance” comes from the equal weighting of D_1 and D_2 in \bar{D}_1 . Later the weights will be arbitrary.

Furthermore, in order to be capturing most of the principal component of the source, the basis should be skewed toward e_1 , so θ should be between 0 and some maximum value $\theta_{\max} < \pi/2$ (see Fig. 5(d)). This yields $D_1 = D_2$, but introduces a new problem.

The representation of X by (14) is (for $\theta \neq \pi/4$) a nonorthogonal basis expansion. The uniform scalar quantization of such a representation produces nonsquare partition cells.¹¹ These partition cells have higher normalized second moments than square cells, and are thus undesirable [44]. The insight attributed to Vaishampayan is that a correlating transform can be applied *after* quantization has been performed in an orthogonal basis representation. This ensures that the partition cells are square, regardless of the transform, so the transform affects only the rate and side distortions. The advantage of this approach over the original pairwise correlating method was shown in [1]. The geometric view of Fig. 5 also helps us understand the significance of the ratio σ_1/σ_2 . When this ratio approaches unity, the level curves become circles. It then does not matter—as far as the side distortion is concerned—what the basis is: the variation perpendicular to either basis vector is invariant to the basis. This fact holds equally well for orthogonal and nonorthogonal bases.

If we do not want the two descriptions to be equally important, for example, if they are sent over links with different failure probabilities, then we expect the optimal representation to be different from (14). Recalling $\sigma_1 > \sigma_2$, we would expect the description over the more reliable link to be closer to the e_1 direction, as this captures most of the energy of the source. This intuition is vindicated by the detailed analysis that follows. The detailed analysis does not use this parameterization of the basis; it turns out to be more convenient to have determinant-1 transforms than to have normalized basis vectors.

B. Detailed Analysis

Now consider in detail sending a two-tuple source over two channels. The scenario is shown in Fig. 6, where $[\cdot]_{\Delta}$ represents quantizing to the nearest multiple of Δ .¹² Recall that the transform is after the quantization so that square partition cells are maintained. The source $x = [x_1, x_2]^T$ has independent, zero-mean Gaussian components with variances σ_1^2 and σ_2^2 , respectively. Without loss of generality, we assume $\sigma_1^2 \geq \sigma_2^2$. The transform \hat{T} produces descriptions y_1 and y_2 through $y = [y_1, y_2]^T = \hat{T}([x]_{\Delta})$, where \hat{T} is a discrete version of the continuous transform

$$T = \begin{bmatrix} a & b \\ c & d \end{bmatrix}, \quad \text{with } \det T = 1.$$

The derivation of \hat{T} from T is by first factoring T into a product of upper- and lower-triangular matrices, e.g.,

$$T = T_1 T_2 T_3 = \begin{bmatrix} 1 & (a-1)/c \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ c & 1 \end{bmatrix} \begin{bmatrix} 1 & (d-1)/c \\ 0 & 1 \end{bmatrix}$$

¹¹In higher dimensions, nonhypercubic cells.

¹²The use of equal quantization step sizes for each component is justified by high-resolution analysis. As long as the rate allocated to the lower variance component is at least about 1 bit/sample, an optimal rate allocation only negligibly reduces the distortion [45].




Fig. 6. Correlating transform structure for MD coding of a two-tuple source.

and then computing the transform with intermediate roundings

$$\hat{T}(x) = [T_1 [T_2 [T_3 x]_{\Delta}]_{\Delta}]_{\Delta}. \quad (15)$$

The properties of \hat{T} are described in Appendix I. Though not shown explicitly, it is assumed that y_1 and y_2 are entropy coded.

The coding structure presented here is a generalization of the method proposed by Orchard, Wang, Vaishampayan, and Reibman [1]. They considered coding of two-tuples with the transform

$$T = \begin{bmatrix} 1 & \beta \\ -(2\beta)^{-1} & 1/2 \end{bmatrix} \quad (16)$$

approximated by

$$\hat{T}(x) = \left[\left[\begin{bmatrix} 1 & 0 \\ -(2\beta)^{-1} & 1 \end{bmatrix} \left[\begin{bmatrix} 1 & \beta \\ 0 & 1 \end{bmatrix} x \right]_{\Delta} \right]_{\Delta} \right]_{\Delta}. \quad (17)$$

The mysterious form of (16) provided the initial motivation for this work. When both side distortions are equally weighted, (16) is in the set of optimal transforms; otherwise, it is not. In either case, (15) gives an additional degree of freedom that may be exploited.

1) *Relationship Between Continuous Transform T and Discrete Transform \hat{T}* : The discrete transform \hat{T} produces the descriptions, but the analysis and optimizations are based on the continuous transform T , so it is important to understand the relationship between the two. \hat{T} is a one-to-one and onto mapping $\Delta\mathbb{Z}^2 \rightarrow \Delta\mathbb{Z}^2$. \hat{T} relabels the quantized source vector $[x]_{\Delta}$ with another value in the lattice $\Delta\mathbb{Z}^2$ in a manner that gives $T[x]_{\Delta} \approx \hat{T}([x]_{\Delta})$. An error bound for the approximation is given in Appendix I.

Since \mathbb{Z}^2 has only a handful of reflective and rotational symmetries, there are only a few one-to-one and onto linear mappings $\Delta\mathbb{Z}^2 \rightarrow \Delta\mathbb{Z}^2$. The transforms used here are not linear, but they approximate linear transforms. As an example, consider

$$T = \begin{bmatrix} 6/5 & 5/12 \\ -6/5 & 5/12 \end{bmatrix} \quad (18)$$

and the associated discrete transform (see Appendix I)

$$\hat{T}(x) = \left[\left[\begin{bmatrix} 1 & 0 \\ -7/5 & 1 \end{bmatrix} \left[\begin{bmatrix} 1 & 5/12 \\ 0 & 1 \end{bmatrix} \left[\begin{bmatrix} 1 & 0 \\ 12/25 & 1 \end{bmatrix} x \right]_{\Delta} \right]_{\Delta} \right]_{\Delta} \right]_{\Delta}. \quad (19)$$

The behavior of \hat{T} is shown in Fig. 7. In this plot each $([x_1]_{\Delta}/\Delta, [x_2]_{\Delta}/\Delta)$ position is labeled with $(y_1/\Delta, y_2/\Delta)$,




Fig. 7. Example of the index relabeling by the discrete transform \hat{T} , given by (19). The solid curves connect the grid points in the inverse images of the y_1 and y_2 axes. Though the curves are somewhat erratic, they follow the corresponding curves for the continuous transform T , shown with dashed lines.

where $y = \hat{T}(x)$. The solid curves connect the (x_1, x_2) pairs on the grid that map to $(y_1, 0)$ or $(0, y_2)$; they are like inverse images of the axes. These curves are somewhat erratic, but they approximate the straight lines

$$x_1 = \frac{25}{72}x_2 \quad \text{and} \quad x_1 = -\frac{25}{72}x_2 \quad (20)$$

which are the points with at least one component of Tx equal to zero. The lines (20) are shown in Fig. 7 as dashed lines. This shows how \hat{T} is a relabeling of the points $\Delta\mathbb{Z}^2$, but at the same time approximates a linear transform.

The reason to bother with the discrete transform is that it allows the use of a nonorthogonal transform without enhancing the quantization error. Fig. 8 shows the partitioning induced by using the continuous transform directly as $y = [Tx]_\Delta$. When the source probability density is approximately uniform over each partition cell—as is the case when the density is smooth and Δ is small—the quantization error is proportional to the normalized second moment of the lattice cell. Thus, nonsquare partition cells are suboptimal. In this particular case, the enhancement of quantization error due to nonsquare cells is by a factor of $33744481/12960000 \approx 2.6$. This means that the representation $[Tx]_\Delta$ has about 2.6 times as much distortion as $\hat{T}([x]_\Delta)$; the latter representation is used in this work.

2) *Computation of Side Rates and Distortions:* If Δ is small, the difference between $\hat{T}(x)$ and Tx is also small. For this reason, a valid high-rate analysis of the actual system, which uses \hat{T} , can be based on the properties of T . The choice of T determines both the rate, through the entropies of y_1 and y_2 , and the distortion when reconstructing from y_1 or y_2 alone. To reiterate, using $\hat{T}([x]_\Delta)$ in place of $[Tx]_\Delta$ is to reduce the




Fig. 8. Example of partitioning induced by a continuous, nonorthogonal transform followed by scalar quantization. The transform T is given by (18). Each cell is labeled by $(y_1/\Delta, y_2/\Delta)$, where $y = [Tx]_\Delta$. The nonsquare partition cells increase the quantization error by a factor of about 2.604.

central distortion by a constant factor; but, for the sake of estimating rates and correlations, the difference between $\hat{T}([x]_\Delta)$ and $[Tx]_\Delta$ is vanishingly small.

The distortion when both y_1 and y_2 are received is independent of T because, in this case, the initial quantization is the only source of distortion. Under the fine-quantization assumption used throughout

$$D_0 = \frac{\Delta^2}{12}. \quad (21)$$

Let us first look at the rates associated with y_1 and y_2 . The continuous transform gives

$$z_1 = (Tx)_1 = ax_1 + bx_2 \quad (22)$$

$$z_2 = (Tx)_2 = cx_1 + dx_2. \quad (23)$$

These are Gaussian random variables with variances $a^2\sigma_1^2 + b^2\sigma_2^2$ and $c^2\sigma_1^2 + d^2\sigma_2^2$, respectively. Since $y_i \approx z_i$, $i = 1, 2$, we may estimate the entropy of y_i as the entropy of a quantized Gaussian random variable. A Gaussian random variable with variance σ^2 quantized with a bin width Δ has entropy given approximately by [46, Ch. 9]

$$\frac{1}{2} \log \sigma^2 + k_\Delta, \quad \text{where } k_\Delta = \frac{1}{2} \log \frac{2\pi e}{\Delta}. \quad (24)$$

Thus, we have rate estimates

$$R_1 = H(y_1) \approx \frac{1}{2} \log (a^2\sigma_1^2 + b^2\sigma_2^2) + k_\Delta, \quad (25)$$

$$R_2 = H(y_2) \approx \frac{1}{2} \log (c^2\sigma_1^2 + d^2\sigma_2^2) + k_\Delta. \quad (26)$$

All logarithms are base-two and all rates are in bits.

Now we turn to a computation of the side distortions. We have assumed that the effect of a lost component makes the quantization error negligible; therefore, D_i is approximately the distortion in estimating x from z_i .

Suppose only z_1 is received. To minimize the MSE distortion, the reconstruction at the first side decoder is $\hat{x}^{(1)} = E[x|z_1]$. Since x and z_1 are jointly Gaussian, the conditional variable $x|z_1$ is also Gaussian and computations of $\hat{x}^{(1)}$ and D_1 are simple. Noting that $[z_1, z_2]^T$ is a linear function of x ,

$$\hat{x}^{(1)} = T^{-1} \begin{bmatrix} z_1 \\ \hat{z}_2 \end{bmatrix} \quad (27)$$

where

$$\hat{z}_2 = E[z_2|z_1] = \frac{E[z_1 z_2]}{E[z_1^2]} z_1 = \frac{ac\sigma_1^2 + bd\sigma_2^2}{a^2\sigma_1^2 + b^2\sigma_2^2} z_1.$$

The error in z_2 , $\eta = z_2 - \hat{z}_2$, is a zero-mean Gaussian random variable with variance

$$\sigma_\eta^2 = E[z_2^2] - \frac{E[z_1 z_2]^2}{E[z_1^2]} = \frac{\sigma_1^2 \sigma_2^2}{a^2\sigma_1^2 + b^2\sigma_2^2}. \quad (28)$$

The error in the estimate of x can now be written in terms of T^{-1} and η as

$$x - \hat{x}^{(1)} = T^{-1} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} - T^{-1} \begin{bmatrix} z_1 \\ \hat{z}_2 \end{bmatrix} = T^{-1} \begin{bmatrix} 0 \\ \eta \end{bmatrix} = \begin{bmatrix} -b \\ a \end{bmatrix} \eta. \quad (29)$$

By combining (28) and (29), we can complete the computation of D_1 , the distortion per component in reconstructing from y_1

$$\begin{aligned} D_1 &= \frac{1}{2} E \left[\|x - \hat{x}^{(1)}\|^2 \right] = \left\| \begin{bmatrix} -b \\ a \end{bmatrix} \right\|^2 E[\eta^2] \\ &= \frac{(a^2 + b^2) \sigma_1^2 \sigma_2^2}{2(a^2\sigma_1^2 + b^2\sigma_2^2)}. \end{aligned}$$

A similar computation yields

$$D_2 = \frac{1}{2} E \left[\|x - \hat{x}^{(2)}\|^2 \right] = \frac{(c^2 + d^2) \sigma_1^2 \sigma_2^2}{2(c^2\sigma_1^2 + d^2\sigma_2^2)}.$$

C. Transform Optimization

Using the rate and distortion expressions from Section III-B2), we can now identify the best transforms for MDCT coding. Define four system states by whether or not each description is received and denote the probabilities of these states as in Table I. A transform is considered optimal if for fixed average rate $R = (R_1 + R_2)/2$ the average distortion is minimized.

From (25) and (26), it appears that the rates can be made arbitrarily small by making $|a|$, $|b|$, $|c|$, and $|d|$ small. However, the constraint $\det T = 1$ prevents this; it is straightforward to show that the minimum average rate is $R^* = \frac{1}{4} \log \sigma_1^2 \sigma_2^2 + k_\Delta$. As in [1], we refer to $\rho = R - R^*$ as the *redundancy*. Using (25) and (26)

$$\rho = \frac{1}{4} \log \frac{(a^2\sigma_1^2 + b^2\sigma_2^2)(c^2\sigma_1^2 + d^2\sigma_2^2)}{\sigma_1^2 \sigma_2^2}. \quad (30)$$

Our optimization will be to minimize the average distortion for fixed, nonnegative redundancy ρ .

In a single-description situation, there would be no reason to have a rate higher than R^* . Thus, ρ is the rate added to improve

TABLE I
PROBABILITIES OF SYSTEM STATES IN TRANSFORM OPTIMIZATION FOR MDCT CODING OF TWO VARIABLES OVER TWO CHANNELS

		Description 1	
		not received	received
Description 2	not received	$p_{0,0}$	$p_{1,0}$
	received	$p_{0,1}$	$p_{1,1}$

the reconstruction from one description. This ρ is redundancy relative to optimal transform coding with scalar quantization; it differs from the absolute excess rate $\rho^* = R - r = R - R(D_0)$ of Section II by the redundancy of uniform scalar quantization, approximately $\frac{1}{2} \log_2 \left(\frac{\pi e}{6} \right) \approx 0.255$ bits. Though ρ is analogous to an increase in rate from channel coding, it is a fixed increase, not proportional to the base source-coding rate.

1) *General Solution:* In the most general situation, channel outages may have unequal probabilities and may be dependent. Using the probabilities of the system states, the overall average distortion is

$$\begin{aligned} \bar{D} &= \left[\underbrace{p_{1,1} D_0}_{\text{both received}} + \underbrace{p_{0,0} \frac{1}{2} (\sigma_1^2 + \sigma_2^2)}_{\text{neither received}} \right] \\ &+ \left[\underbrace{p_{1,0} D_1}_{\text{Description 1 received}} + \underbrace{p_{0,1} D_2}_{\text{Description 2 received}} \right]. \quad (31) \end{aligned}$$

The first pair of terms in (31) does not depend on the transform; we can consider only the second bracketed term in the optimization

$$\bar{D}' = p_{1,0} \frac{(a^2 + b^2) \sigma_1^2 \sigma_2^2}{2(a^2\sigma_1^2 + b^2\sigma_2^2)} + p_{0,1} \frac{(c^2 + d^2) \sigma_1^2 \sigma_2^2}{2(c^2\sigma_1^2 + d^2\sigma_2^2)}.$$

First note that if the source pdf is circularly symmetric, i.e., $\sigma_1 = \sigma_2$, then $D_1 = D_2 = \sigma_1^2/2$, independent of T . In this case, the side distortions cannot be reduced with the addition of redundancy, so the MDCT technique is useless. Henceforth we assume $\sigma_1 > \sigma_2$.

For a given value of ρ , the admissible transforms are simultaneous solutions of

$$(a^2\sigma_1^2 + b^2\sigma_2^2)(c^2\sigma_1^2 + d^2\sigma_2^2) = \sigma_1^2 \sigma_2^2 2^{4\rho} \quad (32)$$

$$ad - bc = 1. \quad (33)$$

Here (32) is a rearrangement of (30) and (33) is $\det T = 1$. There are several branches of solutions.¹³ First suppose $a = 0$. In this case, the transforms are of the form $\begin{bmatrix} 0 & b \\ -1/b & a \end{bmatrix}$ and substituting in (22) and (23) gives

$$\begin{aligned} z_1 &= bx_2 \\ z_2 &= -\frac{1}{b} x_1 + dx_2. \end{aligned}$$

Since channel 1 carries only x_2 , the side distortion at decoder 1 is $D_1 = \sigma_1^2/2$. This is equal to the larger of the two side distortions obtained without any transform; hence, it is poor per-

¹³The multiplicity of solutions is not present in conventional transform coding with orthogonal linear transforms, where for minimum rate ($\rho = 0$) the optimal transform is unique up to reflections and permutations of coordinates. Uses for the extra design freedom in using discrete transforms are discussed in [47], [48].

formance. It can be shown that transforms of this form are in the optimal set only when $p_{1,0} = 0$ or $d = 0$. Another special case is where $c = 0$. This gives transforms of the form $\begin{bmatrix} a & b \\ 0 & 1/a \end{bmatrix}$, which, by analogous reasoning, are not useful.¹⁴

Now, assuming $a \neq 0$ and $c \neq 0$, we may substitute $d = (1 + bc)/a$ into (32) and rearrange to get

$$a^4 \left(\frac{\sigma_1}{\sigma_2} \right)^4 + a^2 \frac{b^2 c^2 + (1 + bc)^2 - 2^{4\rho}}{c^2} \left(\frac{\sigma_1}{\sigma_2} \right)^2 + \frac{b^2(1 + bc)^2}{c^2} = 0.$$

Solving this quadratic in a^2 and choosing signs appropriately gives¹⁵

$$a = \frac{1}{2c} \frac{\sigma_2}{\sigma_1} \left(\sqrt{2^{4\rho} - 1} + \sqrt{2^{4\rho} - 1 - 4bc(bc + 1)} \right). \quad (34)$$

When this value of a is used, \bar{D}' depends only on the product $b \cdot c$, not on the individual values of b and c . The value of bc that minimizes \bar{D}' is

$$(bc)_{\text{opt}} = -\frac{1}{2} + \frac{1}{2} \left(\left(\frac{p_{1,0}}{p_{0,1}} + 1 \right)^2 - 4 \left(\frac{p_{1,0}}{p_{0,1}} \right) 2^{-4\rho} \right)^{-1/2} \cdot \left(\frac{p_{1,0}}{p_{0,1}} - 1 \right) \quad (35)$$

independent of σ_1/σ_2 .

To summarize, an optimal transform for a given redundancy ρ can be found through the following steps. First, choose any nonzero b . The corresponding value of c is determined from (35). Then, substituting in (34) gives a . Finally, d is chosen to make $\det T = 1$: $d = (1 + bc)/a$. The performance, measured by the tradeoff between ρ and \bar{D} , does not depend on the choice of b . This degree of freedom can be used to control the split of the total rate $R^* + \rho$ between R_1 and R_2 or to simplify \hat{T} by making $a = 1$ or $d = 1$. Other uses are described in [47], [48].

It is easy to check that $(bc)_{\text{opt}}$ ranges from -1 to 0 as $p_{1,0}/p_{0,1}$ ranges from 0 to ∞ . The limiting behavior can be explained as follows. Suppose $p_{1,0} \gg p_{0,1}$, i.e., channel 1 is much more reliable than channel 2. Since $(bc)_{\text{opt}}$ approaches 0 , ad must approach 1 , and hence one optimally sends x_1 (the larger variance component) over channel 1 (the more reliable channel), and *vice versa*. This is the intuitive, layered solution. The MD approach is most useful when the channel failure probabilities are comparable, but this demonstrates that the MD framework subsumes layered coding.

2) *Equal Channel Failure Probabilities*: Suppose the channels are equally likely to fail, so $p_{1,0} = p_{0,1}$. Then $(bc)_{\text{opt}} =$


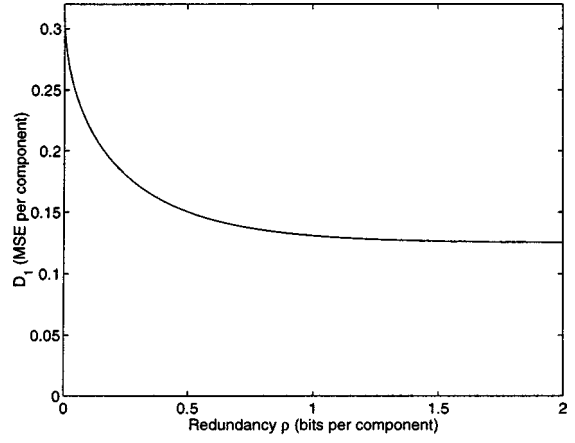


Fig. 9. Optimal R - D_0 - D_1 tradeoffs for $\sigma_1 = 1$, $\sigma_2 = 0.5$. (a) Relationship between redundancy ρ and D_1 . (b) Relationship between D_0 and D_1 for various rates.

$-1/2$, independent of ρ . The optimal set of transforms is described by

$$\begin{aligned} a &= \pm b \sigma_1^{-1} \sigma_2 \left(2^{2\rho} + \sqrt{2^{4\rho} - 1} \right), & b &\neq 0 \\ c &= -(2b)^{-1}, & d &= (2a)^{-1} \end{aligned} \quad (36)$$

and using a transform from this set gives

$$D_1 = D_2 = \frac{1}{2} \sigma_2^2 + \frac{\sigma_1^2 - \sigma_2^2}{4 \cdot 2^{2\rho} (2^{2\rho} + \sqrt{2^{4\rho} - 1})}. \quad (37)$$

This relationship is plotted in Fig. 9(a). Notice that, as expected, D_1 starts at a maximum value of $(\sigma_1^2 + \sigma_2^2)/4$ and asymptotically approaches a minimum value of $\sigma_2^2/2$. By combining the rate expressions (25) and (26) and the D_0 expression (21) with (37), one can find the relationship between R , D_0 , and D_1 . For various values of R , the tradeoff between D_0 and D_1 is plotted in Fig. 9(b).

The solution for the optimal set of transforms (36) has an extra degree of freedom which does not affect the ρ versus D_1 performance. Fixing $a = 1$ gives the transforms suggested in [1] and allows $\hat{T}(\cdot)$ to be implemented with two lifting steps instead

¹⁴The only useful transforms in these two branches of solutions of (32) and (33) are those that add no redundancy, e.g., $T = I$.

¹⁵In the solution for a^2 , one of the two branches of the quadratic is valid. Then, in the square root of this expression we arbitrarily choose $a \geq 0$ since the sign of a does not affect ρ or \bar{D}_1 .




Fig. 10. Geometric interpretations. (a) When $\sigma_1 > \sigma_2$, the optimality condition ($ad = 1/2$, $bc = -1/2$) is equivalent to $\theta_1 = \theta_2 < \theta_{\max} = \arctan(\sigma_1/\sigma_2)$. (b) If in addition to the optimality condition we require the output streams to have equal rate, the analysis vectors are symmetrically situated to capture the dimension with greatest variation. At $\rho = 0$, $\theta_1 = \theta_2 = \theta_{\max}$; as $\rho \rightarrow \infty$, φ_1 and φ_2 close on the x_1 -axis.

of three. This degree of freedom can also be used to control the partitioning of the rate between channels.

3) *Geometric Interpretation:* When the channels are equally likely to fail, minimizing the average distortion automatically makes the two side distortions, D_1 and D_2 , equal. Since the channel rates are additive in the constraint, this is an expected result; with a Lagrangian approach, matching the slopes also matches the distortions. This has a clear geometric interpretation also.

Neglecting quantization, the transmitted representation of x is given by $y_1 = \langle x, \varphi_1 \rangle$ and $y_2 = \langle x, \varphi_2 \rangle$, where $\varphi_1 = [a, b]^T$ and $\varphi_2 = [c, d]^T$. For optimality with equal channel failure probabilities, $ad = 1/2$ and $bc = -1/2$. One can then show that φ_1 and φ_2 form the same (absolute) angles with the positive x_1 -axis. For convenience, suppose $a > 0$ and $b < 0$; then $c > 0$ and $d > 0$, as shown in Fig. 10(a). Let θ_1 and θ_2 be the angles by which φ_1 and φ_2 are below and above the positive x_1 -axis, respectively. Then $\tan \theta_1 = -b/a = d/c = \tan \theta_2$. If we assume $\sigma_1 > \sigma_2$, then the maximum angle (for $\rho = 0$) is $\arctan(\sigma_1/\sigma_2)$ and the minimum angle (for $\rho \rightarrow \infty$) is zero. This means that for $\rho = 0$ the analysis basis has vectors at angles $\arctan(\sigma_1/\sigma_2)$ from the x_1 -axis; as ρ is increased, the analysis basis closes in on the x_1 -axis. Thus, as the redundancy is increased, x_1 is emphasized over x_2 because it has higher variance (see Fig. 10(b)).

4) *Optimal Transforms that Give Balanced Rates:* The transforms of [1] give descriptions with unequal rates. In practice, this can be remedied through time-multiplexing. An alternative is to use the “extra” degree of freedom to make $R_1 = R_2$. Doing this is equivalent to requiring $|a| = |c|$ and $|b| = |d|$, and yields

$$a = \pm \sqrt{\frac{1}{2} \frac{\sigma_2}{\sigma_1} (2^{2\rho} + \sqrt{2^{4\rho} - 1})}$$

and

$$b = \pm \frac{1}{2a} = \pm \sqrt{\frac{1}{2} \frac{\sigma_1}{\sigma_2} (2^{2\rho} - \sqrt{2^{4\rho} - 1})}.$$

These balanced-rate transforms will be used frequently in the sequel, so we introduce a notation for them. For any nonzero α , let

$$T_\alpha = \begin{bmatrix} \alpha & (2\alpha)^{-1} \\ -\alpha & (2\alpha)^{-1} \end{bmatrix}. \quad (38)$$

When there are no erasures, the reconstruction uses

$$T_\alpha^{-1} = \begin{bmatrix} (2\alpha)^{-1} & -(2\alpha)^{-1} \\ \alpha & \alpha \end{bmatrix}.$$

Evaluating (27) shows that the optimal linear estimates from y_1 and y_2 , neglecting quantization noise, are

$$\hat{x}^{(1)} = \frac{2\alpha}{4\alpha^4\sigma_1^2 + \sigma_2^2} \begin{bmatrix} 2\alpha^2\sigma_1^2 \\ \sigma_2^2 \end{bmatrix} y_1$$

and

$$\hat{x}^{(2)} = \frac{2\alpha}{4\alpha^4\sigma_1^2 + \sigma_2^2} \begin{bmatrix} -2\alpha^2\sigma_1^2 \\ \sigma_2^2 \end{bmatrix} y_2,$$

respectively.

In the optimizations leading to (36), we have used a *mean* squared error measure for the side distortion; no distinction is made between transforms that yield the same MSE but different *distributions* of squared error. Additional robustness is implied by a low variance of squared error or, equivalently, low expected fourth power of the Euclidean norm of the error. The balanced-rate transforms (38) give this desirable property [49]. Another benefit of the balanced-rate transforms is that the descriptions are identically distributed; the same entropy code can be applied to each, reducing memory requirements [48].

D. Modifications for Centroid Reconstruction

The analyses in this paper are focused on the case in which distortion due to quantization is small and linear estimation is used in the decoder. However, because the quantization is in coordinates in which the components are independent, it is not too difficult to compute true (nonlinear) minimum MSE reconstructions.

Without loss of generality, consider the case that decoder 1 has received description y_1 . The initial quantization of the source vector x followed by an invertible transform indicates that x lies in a particular union of square cells (see, e.g., Fig. 7)

$$x \in \mathcal{C}_{y_1}^{(1)} = \bigcup_{k_2=-\infty}^{\infty} \mathcal{C}_{y_1/\Delta, k_2}$$

where

$$\mathcal{C}_{k_1, k_2} = \left\{ x \mid \hat{T}([x]_\Delta) = (k_1, k_2)\Delta \right\}.$$

(Read \mathcal{C}_{k_1, k_2} as the set that maps to (k_1, k_2) and $\mathcal{C}_k^{(i)}$ as the set that maps to i th component equal to k .) The estimate that minimizes MSE is the conditional expectation of x given $x \in \mathcal{C}_{y_1}^{(1)}$. This is not too difficult because each \mathcal{C}_{k_1, k_2} in $\mathcal{C}_{y_1}^{(1)}$ makes independent contributions to the estimates of each coordinate.

Written in a general form, the optimal estimate is given by

$$\hat{x}^{(1)} = E \left[x \mid x \in \mathcal{C}_{y_1}^{(1)} \right] = \frac{\int_{\mathcal{C}_{y_1}^{(1)}} x f_{x_1, x_2}(x_1, x_2) dx_1 dx_2}{\int_{\mathcal{C}_{y_1}^{(1)}} f_{x_1, x_2}(x_1, x_2) dx_1 dx_2}$$

where f is the joint pdf of the source vector. This computation is simplified first by decomposing $\mathcal{C}_{y_1}^{(1)}$ and then by the fact that

the integral of f over each \mathcal{C}_{k_1, k_2} is separable. Let p_{k_1, k_2} be the probability of $x \in \mathcal{C}_{k_1, k_2}$. Then

$$\begin{aligned} p_{k_1, k_2} &= \int_{\mathcal{C}_{k_1, k_2}} f_{x_1, x_2}(x_1, x_2) dx_1 dx_2 \\ &= \left(\operatorname{erf} \left(2^{-1/2} \sigma_1^{-1} \left(j_1 + \frac{1}{2} \right) \Delta \right) \right. \\ &\quad \left. - \operatorname{erf} \left(2^{-1/2} \sigma_1^{-1} \left(j_1 - \frac{1}{2} \right) \Delta \right) \right) \\ &\quad \cdot \left(\operatorname{erf} \left(2^{-1/2} \sigma_2^{-1} \left(j_2 + \frac{1}{2} \right) \Delta \right) \right. \\ &\quad \left. - \operatorname{erf} \left(2^{-1/2} \sigma_2^{-1} \left(j_2 - \frac{1}{2} \right) \Delta \right) \right) \end{aligned}$$

where $\hat{T}((j_1, j_2)\Delta) = (k_1, k_2)\Delta$. Similarly simple integrals give the contribution of a cell to each first moment

$$\begin{aligned} m_{k_1, k_2}^{(1)} &= \int_{\mathcal{C}_{k_1, k_2}} x_1 f_{x_1, x_2}(x_1, x_2) dx_1 dx_2 \\ &= \sigma_1^2 (2\pi)^{-1/2} \left(e^{-2^{-1} \sigma_1^2 (j_1 - \frac{1}{2})^2 \Delta^2} - e^{-2^{-1} \sigma_1^2 (j_1 + \frac{1}{2})^2 \Delta^2} \right) \\ &\quad \cdot \left(\operatorname{erf} \left(2^{-1/2} \sigma_2^{-1} \left(j_2 + \frac{1}{2} \right) \Delta \right) \right. \\ &\quad \left. - \operatorname{erf} \left(2^{-1/2} \sigma_2^{-1} \left(j_2 - \frac{1}{2} \right) \Delta \right) \right) \\ m_{k_1, k_2}^{(2)} &= \int_{\mathcal{C}_{k_1, k_2}} x_2 f_{x_1, x_2}(x_1, x_2) dx_1 dx_2 \\ &= \sigma_2^2 (2\pi)^{-1/2} \left(e^{-2^{-1} \sigma_2^2 (j_2 - \frac{1}{2})^2 \Delta^2} - e^{-2^{-1} \sigma_2^2 (j_2 + \frac{1}{2})^2 \Delta^2} \right) \\ &\quad \cdot \left(\operatorname{erf} \left(2^{-1/2} \sigma_1^{-1} \left(j_1 + \frac{1}{2} \right) \Delta \right) \right. \\ &\quad \left. - \operatorname{erf} \left(2^{-1/2} \sigma_1^{-1} \left(j_1 - \frac{1}{2} \right) \Delta \right) \right). \end{aligned}$$

The optimal estimate is then given elementwise by

$$\hat{x}_i^{(1)} = \frac{\sum_{k_2=-\infty}^{\infty} m_{y_1/\Delta, k_2}^{(i)}}{\sum_{k_2=-\infty}^{\infty} p_{y_1/\Delta, k_2}}.$$

Since the Gaussian pdf has light tails, the infinite sums can be truncated to a small number of terms.

In an application in which \hat{T} is fixed, a table of optimal reconstructions could be precomputed for each side decoder. (A similar precomputation would be needed for MDSQ with centroid reconstruction.) To allow \hat{T} to be varied easily, one could precompute each appreciable cell probability and cell moment. These can then be easily combined to form the optimal estimate. In this way, one can take advantage of the smoothly varying redundancy–distortion tradeoff provided by correlating transforms.

E. Numerical Results and Comparisons

To conclude this section on coding two-tuples, we create a set of numerical calculations to compare systems with correlating transforms against those with MDSQs. These computed performances are also compared to estimated performance (37) and a bound derived from Theorem 1. MATLAB code to generate these results is available on-line in conjunction with [33].

The numerical results given in Fig. 11 consider two sources and two base coding rates, for a total of four comparisons. The sources have independent components with $\sigma_1 = 1$ and $\sigma_2 \in \{1/2, 1/4\}$. The base rates are 2 and 3 bits/component. Robustness is measured by an equally weighted average of side distortions $(D_1 + D_2)/2$, so the transforms (38) are among the optimal transforms; these are used in the computations. Operating points for MDCT are obtained by a fine sampling of α .

The operating points for MDSQ are obtained by trying each of seven possible UMDSQ index assignments (see Section II-C1)) on each of the two components, giving 49 operating points. Only the “useful” ones are shown; those with average side distortion per component more than $(\sigma_1^2 + \sigma_2^2)/2$ are discarded, as are those strictly worse than any other operating point. This strategy is essentially prescribed by the theory since using no transform is optimal according to [43].

Theorem 1 does not directly give us performance bounds for the sources that we consider. However, it does give bounds to apply separately to each component. The approximate bounds shown in Fig. 11 are obtained by fixing the central distortion for each component to match that of the other systems and then optimally allocating redundancy across the pair. This redundancy allocation is straightforward because of the monotonicity of the bound (8).

An initial observation from the graphs is that the analytical performance curves for MDCT follow the actual performance with midpoint reconstruction. The gap is due to neglecting quantization error, and diminishes as the base rate is increased. We may also observe a general similarity in shape between the approximate bound and the performance curves.

In the examples shown, the comparison between MDCT and MDSQ is summarized as follows: *The correlating transform method is better at low redundancies and MDSQ is better at high redundancies.* The near-continuum of operating points is another potential advantage of MDCT. A more nuanced view recognizes that the method of choice depends also on the base rate and the ratio σ_1/σ_2 .

First, the base rate. Increasing the base rate improves the performance of MDSQ, but has little effect on the performance of MDCT; thus, MDSQ becomes more suitable as the base rate is increased. Considering MDSQ, as the base rate is increased the number of useful index assignments increases (see Fig. 4). Also, the central distortion decreases and, since the side distortion approaches the central distortion, the side distortion drops more quickly as a function of ρ . On the other hand, as shown in (37), the performance of MDCT is largely independent of the base rate.

Second, the ratio σ_1/σ_2 . It is clear from (37) and the performance plots that the usefulness of MDCT increases as the ratio σ_1/σ_2 increases. However, at the same time the performance of MDSQ improves and the approximate bound decreases. In both cases, this is because redundancy can be allocated to the component with larger variance. When $\sigma_1/\sigma_2 = 1$, MDCT is certainly useless, but large σ_1/σ_2 does not imply that MDCT is better than MDSQ.

Another observation from Fig. 11 is that there is a limit to how much redundancy can be added. Though obscured by the analysis that neglects quantization, this is quite clear: There is an

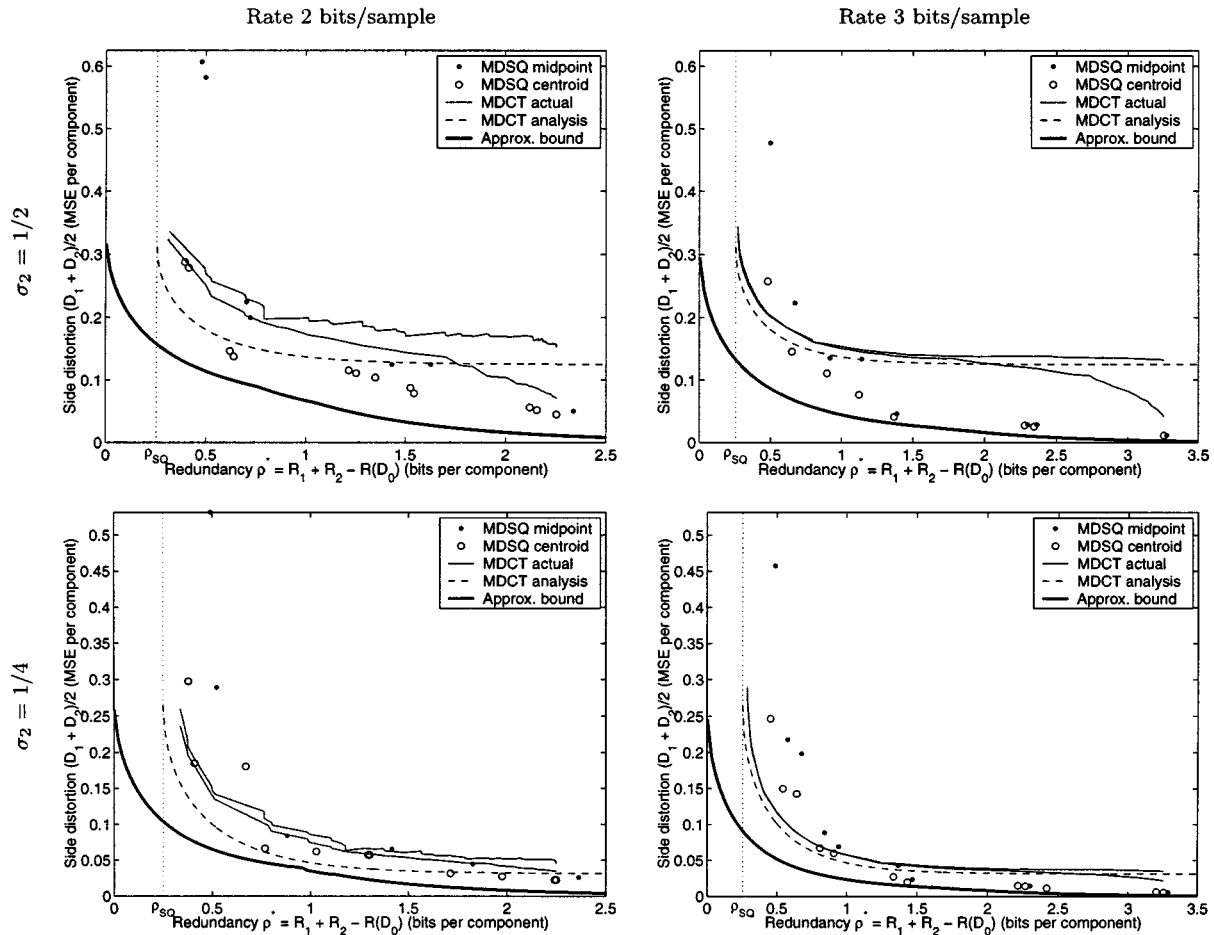


Fig. 11. Comparison between MDCT and MDSQ with estimated performance (37) (dashed) and an approximate bound (bold) also shown. In each graph, the upper and lower MDCT curves are for midpoint and centroid reconstruction, respectively. In all cases, the source has independent components with $\sigma_1 = 1$. The two rows have $\sigma_2 = 1/2$ and $1/4$. The two columns have base rates 2 and 3 bits/sample. The vertical dotted lines indicate the redundancy, labeled ρ_{SQ} , due simply to the use of scalar quantization. This portion of the redundancy is approximately 0.255 bit. The comparisons show that correlating transforms are good at low redundancies.

initial representation $[x]_{\Delta}$ and the redundancy can be no greater than to repeat this representation over each channel. When α is large enough, each of the cells with appreciable probability is uniquely identified separately by y_1 and y_2 . The MDCT performance, with centroid reconstruction, becomes the same as the repetition code incarnation of MDSQ.

Finally, the reconstruction method—midpoint or centroid—makes a big difference for low-rate MDSQ. This suggests high sensitivity to the source density; care should be taken with non-Gaussian sources.

IV. CODING LARGER VECTORS

We now turn to the communication of vectors with more than two components. This section describes a method for MD coding based on using transform coefficients or sets of transform coefficients as descriptions. The reconstruction from a proper subset of the descriptions exploits a statistical correlation between transform coefficients. Thus, this technique may be dubbed statistical channel coding for an erasure channel. A square transform is used, so for coding an N -dimensional source at most N descriptions are produced. The method is a generalization of the pairwise correlating transforms of

Orchard, Wang, Vaishampayan, and Reibman [1] and of our work of the previous section to $N > 2$.

It is not difficult to extend the calculations in Section III to find optimization criteria for the general case; this is done in Section IV-B. It is considerably more difficult to find the optimal transforms, so we are only able to do this in certain cases.

A. Design

Let $\{x_k\}$ be an i.i.d. sequence of zero-mean jointly Gaussian vectors in \mathbb{R}^N with a known distribution.¹⁶ Because we could use a Karhunen–Loève transform (KLT) at the encoder, we may assume without loss of generality that the components of x_k are independent with variances $\sigma_1^2 \geq \sigma_2^2 \geq \dots \geq \sigma_N^2$.

In correlating transform-based MD coding, each source vector x is processed as follows.

- 1) x is quantized with an unbounded uniform scalar quantizer with step size Δ ; i.e., $x_i^q = [x_i]_{\Delta}$, where $[\cdot]_{\Delta}$ denotes rounding to the nearest multiple of Δ .
- 2) The vector $x^q = [x_1^q, x_2^q, \dots, x_N^q]^T$ is transformed with an invertible, discrete transform $T: \Delta\mathbb{Z}^N \rightarrow \Delta\mathbb{Z}^N$, $y =$

¹⁶Note that it is the vectors, not the scalar components of the vectors, that are i.i.d.

$\hat{T}(x^q)$. \hat{T} is within a certain quasilinear class described below.

- 3) The components of y are placed into M sets (in an *a priori* fixed manner). These sets form the M descriptions.
- 4) The M sets of coefficients, i.e., the descriptions, are entropy coded independently from each other. This stage may use block coding *within one description* to improve efficiency.

The transform \hat{T} is a discrete transform derived from a linear transform T , with $\det T = 1$. In extending (15) to more than two components, first T is factored into matrices with unit diagonals and nonzero off-diagonal elements only in one row or column: $T = T_1 T_2 \cdots T_k$. The discrete transform is then given by

$$\hat{T}(x^q) = [T_1 [T_2 \cdots [T_k x^q]_{\Delta}]_{\Delta}]_{\Delta}.$$

This construction of the transform ensures that \hat{T} is invertible on $\Delta\mathbb{Z}^N$. See Appendix I for details.

The analysis and optimization are again based on high-rate (or fine-quantization, small Δ) approximations. These approximations facilitate the analytical treatment of the design of T , or at least the generation of an optimization criterion.

When all the components of y are received, the reconstruction process is to (exactly) invert the transform \hat{T} to get $\hat{x} = x^q$. If some components of y are lost, they are estimated from the received components using the correlation introduced by the transform \hat{T} .

Recall that the variances of the components of x are $\sigma_1^2, \sigma_2^2, \dots, \sigma_N^2$ and denote the correlation matrix of x by

$$R_x = \text{diag}(\sigma_1^2, \sigma_2^2, \dots, \sigma_N^2).$$

With fine quantization, the correlation matrix of y is $R_y = TR_x T^T$. By renumbering the variables if necessary, assume that $y_1, y_2, \dots, y_{N-\ell}$ are received and $y_{N-\ell+1}, \dots, y_N$ are lost. Partition y into “received” and “not received” portions as $y = [\tilde{y}_r, \tilde{y}_{nr}]^T$ where

$$\tilde{y}_r = [y_1, y_2, \dots, y_{N-\ell}]^T$$

and

$$\tilde{y}_{nr} = [y_{N-\ell+1}, \dots, y_{N-1}, y_N]^T.$$

The minimum MSE estimate of x given \tilde{y}_r is $E[x | \tilde{y}_r]$, which has a simple closed form because x is a jointly Gaussian vector. Using the linearity of the expectation operator gives the following sequence of calculations:

$$\begin{aligned} \hat{x} &= E[x | \tilde{y}_r] = E[T^{-1}Tx | \tilde{y}_r] = T^{-1}E[Tx | \tilde{y}_r] \\ &= T^{-1}E\left[\begin{bmatrix} \tilde{y}_r \\ \tilde{y}_{nr} \end{bmatrix} \middle| \tilde{y}_r\right] = T^{-1}\begin{bmatrix} \tilde{y}_r \\ E[\tilde{y}_{nr} | \tilde{y}_r] \end{bmatrix}. \end{aligned} \quad (39)$$

If the correlation matrix of y is partitioned compatibly with the partitioning of y as

$$R_y = \begin{bmatrix} R_1 & B \\ B^T & R_2 \end{bmatrix}$$

then $\tilde{y}_{nr} | \tilde{y}_r$ is a Gaussian random variable with mean $B^T R_1^{-1} \tilde{y}_r$ and correlation matrix $R_2 - B^T R_1^{-1} B$. Thus

$$E[\tilde{y}_{nr} | \tilde{y}_r] = B^T R_1^{-1} \tilde{y}_r$$

and the reconstruction is

$$\hat{x} = T^{-1} \begin{bmatrix} I \\ B^T R_1^{-1} \end{bmatrix} \tilde{y}_r. \quad (40)$$

B. General Optimization Criteria

As before, the choice of the transform T determines the performance of the system. This section develops the relationships between the transform, rates, and distortions necessary to design T .

Estimating the rate is straightforward. Since the quantization is fine, y_i is approximately the same as $[(Tx)_i]_{\Delta}$, i.e., a uniformly quantized Gaussian random variable. With $R_y = TR_x T^T$, the variance of y_i is $\sigma_{y_i}^2 = (R_y)_{ii}$. Thus, using (24), the rate per component is

$$R = \frac{1}{N} \sum_{i=1}^N H(y_i) \approx k_{\Delta} + \frac{1}{2N} \sum_{i=1}^N \log \sigma_{y_i}^2 = k_{\Delta} + \frac{1}{2N} \log \prod_{i=1}^N \sigma_{y_i}^2.$$

The minimum rate occurs when

$$\prod_{i=1}^N \sigma_{y_i}^2 = \prod_{i=1}^N \sigma_i^2$$

and at this rate the components of y are uncorrelated. As in the two-tuple case, $T = I$ is not the only transform which achieves the minimum rate. In fact, an arbitrary split of the total rate among the different components of y is possible [48]. This is a justification for using a total rate constraint in our following analyses. The excess over the minimum rate is the redundancy

$$\rho = \frac{1}{2N} \log \prod_{i=1}^N \sigma_i^{-2} \sigma_{y_i}^2. \quad (41)$$

When all the transform coefficients are received, the distortion is due only to the initial quantization; as in the two-tuple case, (21) holds. The distortion when reconstructing from a proper subset of the descriptions is more complicated.

With more than two descriptions, our earlier notation becomes inadequate. With M descriptions, there are $2^M - 1$ nontrivial reconstructions, each with a potentially distinct distortion. Assign to each channel a state $s_i \in \{0, 1\}$ to denote whether the channel is received (1) or not received (0). For any system state $S = s_1 \times s_2 \times \cdots \times s_M$, the distortion will be denoted $D_{(s_1, s_2, \dots, s_M)}$.

The optimization cost function is a weighted average distortion

$$\bar{D} = \sum_{s_i \in \{0, 1\}, 1 \leq i \leq M} \alpha_{(s_1, s_2, \dots, s_M)} D_{(s_1, s_2, \dots, s_M)} \quad (42)$$

where

$$\sum_{s_i \in \{0, 1\}, 1 \leq i \leq M} \alpha_{(s_1, s_2, \dots, s_M)} = 1.$$

In the simplest case, $\alpha_{(s_1, s_2, \dots, s_M)}$ could be the probability of state (s_1, s_2, \dots, s_M) . In this case, \bar{D} is the overall average MSE. Other meaningful choices are available. For example, if a certain minimum quality is required when k of M channels are received, then the $\binom{M}{k}$ states with $\sum s_i = k$ can be assigned equal weights of $\binom{M}{k}^{-1}$, with the remaining states having no weight. In this case, the optimization will make the distortion equal in each of the states with k channels received and this distortion will be an upper bound for the distortion when more than k channels are received.

We now compute a generic distortion factor $D_{(s_1, s_2, \dots, s_M)}$. This computation was almost completed in the development of (40). Renumbering the variables if necessary, assume the last ℓ components of y are lost and partition the received vector as before. Let η be the error in predicting \tilde{y}_{nr} from \tilde{y}_{r} , $\eta = \tilde{y}_{\text{nr}} - E[\tilde{y}_{\text{nr}} | \tilde{y}_{\text{r}}]$, which is a Gaussian random variable with zero mean and correlation matrix $A = R_2 - B^T R_1^{-1} B$. Substituting $\tilde{y}_{\text{nr}} - \eta$ for $E[\tilde{y}_{\text{nr}} | \tilde{y}_{\text{r}}]$ in (39) gives

$$\hat{x} = T^{-1} \begin{bmatrix} \tilde{y}_{\text{r}} \\ \tilde{y}_{\text{nr}} - \eta \end{bmatrix} = x + T^{-1} \begin{bmatrix} 0 \\ -\eta \end{bmatrix}$$

so

$$\|x - \hat{x}\|^2 = \left\| T^{-1} \begin{bmatrix} 0 \\ -\eta \end{bmatrix} \right\|^2 = \eta^T U^T U \eta$$

where U is comprised of the last ℓ columns of T^{-1} . Finally,

$$\begin{aligned} E\|x - \hat{x}\|^2 &= E[\text{tr}(\eta^T U^T U \eta)] = E[\text{tr}(\eta \eta^T U^T U)] \\ &= \text{tr}(A U^T U). \end{aligned} \quad (43)$$

The problem is to minimize \bar{D} subject to a constraint on R . The expressions given in this section can be used to numerically determine transforms to realize this goal. Analytical solutions are possible in certain special cases. Some of these are outlined in the following sections.

C. Techniques for Two Channels

The simplest generalization of sending two variables over two channels is to keep the number of channels the same, but to increase the number of variables N . The significance of having two channels is that the transform coefficients must be placed in two sets. The distortion expression (42) has just $2^2 = 4$ terms—not 2^N terms—because each set is either received in full or lost in full.

The general solution for sending two variables over two channels can be used to derive methods for sending more variables over two channels. These methods use at most $\lfloor N/2 \rfloor$ transforms of size 2×2 in parallel and thus have complexity that is only linear in the number of variables. For simplicity, it is assumed that the channels are equally likely to fail.

1) *Three Variables:* The natural first step is to consider the transmission of three variables. Suppose y_1 is transmitted on channel 1 and (y_2, y_3) is transmitted on channel 2. We could start as before, designing a 3×3 transform with determinant 1 to minimize the distortion given by (42) and (43). The eight free parameters make this a difficult optimization. A much easier way to determine the optimal performance is to first reduce the number of parameters without risking a loss in performance. It turns out that it is sufficient to send one of the original variables

as y_3 and to use an optimal 2×2 transform to produce y_1 and y_2 . This assertion is formalized by the following theorem.

Theorem 2: Consider MD transform coding where y_1 is sent on channel 1 and (y_2, y_3) is sent on channel 2. To minimize the average side distortion D_1 with an upper bound on redundancy ρ , it is sufficient to optimize over transforms of the form

$$\begin{bmatrix} \tilde{T} & 0_{2 \times 1} \\ 0_{1 \times 2} & 1 \end{bmatrix} P \quad (44)$$

with $\tilde{T} \in \mathbb{R}^{2 \times 2}$, $\det \tilde{T} = 1$, and P a permutation matrix.

Proof: See Appendix II-A. \square

Theorem 2 reduces the number of design parameters from eight to three and makes the design of an optimal transform a simple application of the results of Section III. A transform of the form (44) shuffles the input variables and then correlates the first two. Since the order of the elements in a correlated pair does not matter, the permutation can be limited to one of the following three:

$$P_1 = I, \quad P_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \quad \text{and} \quad P_3 = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}.$$

Let us consider a generic choice among the three permutations by assuming that the outputs of the permutation have variances ς_1^2 , ς_2^2 , and ς_3^2 ; i.e., applying the permutation represented by P to $(\sigma_1^2, \sigma_2^2, \sigma_3^2)$ gives $(\varsigma_1^2, \varsigma_2^2, \varsigma_3^2)$. Recall the original ordering of the variances ($\sigma_1^2 \geq \sigma_2^2 \geq \sigma_3^2$) and notice that the three permutations under consideration preserve the ordering of the first two components: $\varsigma_1^2 \geq \varsigma_2^2$.

The component with variance ς_3^2 is sent over channel 2 without any channel protection. (It is uncorrelated with the other components.) Since channel 2 is lost half of the time and the distortion is measured per component, this component contributes $\varsigma_3^2/6$ to D_1 , independent of \tilde{T} . Now the optimization of \tilde{T} in (44) is precisely the problem of the previous section. Thus, \tilde{T} can be chosen in the form of (38) and

$$D_1 = \frac{1}{6} \varsigma_3^2 + \frac{1}{3} \varsigma_2^2 + \frac{\varsigma_1^2 - \varsigma_2^2}{6 \cdot 2^{3\rho} (2^{3\rho} + \sqrt{2^{6\rho} - 1})}. \quad (45)$$

The second and third terms of (45) come from evaluating (37) and rescaling the redundancy and distortion to account for the change from two to three components.

Now we can choose the best permutation; i.e., the permutation yielding the lowest side distortion. The three permutations give the following average side distortions, respectively:

$$(D_1)_1 = \frac{1}{6} \sigma_3^2 + \frac{1}{3} \sigma_2^2 + \frac{\sigma_1^2 - \sigma_2^2}{6 \cdot 2^{3\rho} (2^{3\rho} + \sqrt{2^{6\rho} - 1})}$$

$$(D_1)_2 = \frac{1}{6} \sigma_2^2 + \frac{1}{3} \sigma_3^2 + \frac{\sigma_1^2 - \sigma_3^2}{6 \cdot 2^{3\rho} (2^{3\rho} + \sqrt{2^{6\rho} - 1})}$$

$$(D_1)_3 = \frac{1}{6} \sigma_1^2 + \frac{1}{3} \sigma_3^2 + \frac{\sigma_2^2 - \sigma_3^2}{6 \cdot 2^{3\rho} (2^{3\rho} + \sqrt{2^{6\rho} - 1})}.$$

The best permutation is P_2 because

$$\begin{aligned} (D_1)_1 - (D_1)_2 &= \frac{1}{6} (\sigma_2^2 - \sigma_3^2) \left(1 - \frac{1}{2^{3\rho} (2^{3\rho} + \sqrt{2^{6\rho} - 1})} \right) \geq 0 \end{aligned}$$

and

$$(D_1)_3 - (D_1)_2 = \frac{1}{6}(\sigma_1^2 - \sigma_2^2) \left(1 - \frac{1}{2^{3\rho}(2^{3\rho} + \sqrt{2^{6\rho} - 1})} \right) \geq 0.$$

To summarize, transforms of the form

$$\begin{bmatrix} \alpha & 0 & (2\alpha)^{-1} \\ -\alpha & 0 & (2\alpha)^{-1} \\ 0 & 1 & 0 \end{bmatrix}$$

attain the optimal performance

$$D_1 = \frac{1}{6}\sigma_2^2 + \frac{1}{3}\sigma_3^2 + \frac{\sigma_1^2 - \sigma_3^2}{6 \cdot 2^{3\rho}(2^{3\rho} + \sqrt{2^{6\rho} - 1})}.$$

This performance is matched by many other transforms, but not surpassed.

2) *Four Variables*: We now move on to communicating four variables over two channels. The problem is similar to the one we just solved if one channel carries y_1 and the other carries (y_2, y_3, y_4) . In this case, a result analogous to Theorem 2 holds, revealing that it is sufficient to consider transforms of the form

$$T = \begin{bmatrix} \tilde{T} & 0_{2 \times 2} \\ 0_{2 \times 2} & I_{2 \times 2} \end{bmatrix} P$$

where \tilde{T} is a 2×2 correlating transform and P is one of six permutations.

The best choice of permutation causes the correlating transform to be applied to the components with the largest and smallest variances. The result is a transform of the form

$$\begin{bmatrix} \alpha & 0 & 0 & (2\alpha)^{-1} \\ -\alpha & 0 & 0 & (2\alpha)^{-1} \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

and optimal performance is given by

$$D_1 = \frac{1}{8}\sigma_2^2 + \frac{1}{8}\sigma_3^2 + \frac{1}{4}\sigma_4^2 + \frac{\sigma_1^2 - \sigma_4^2}{8 \cdot 2^{4\rho}(2^{4\rho} + \sqrt{2^{8\rho} - 1})}.$$

Let us now consider the case where each channel carries two coefficients; for concreteness, one carries the odd-numbered coefficients and the other carries the even-numbered coefficients.

The transmission over two channels and the allocation of the coefficients to channels does not place any limitation on the transform. However, we can again place a simplifying limitation on the transform without loss of optimality. It is sufficient to consider pairing the input coefficients, applying a 2×2 correlating transform to each pair, and sending one output of each 2×2 subtransform over each channel. This is justified by the following theorem.

Theorem 3: Consider MD transform coding where (y_1, y_3) is sent on channel 1 and (y_2, y_4) is sent on channel 2. To minimize the average side distortion D_1 with an upper bound on ρ , it is sufficient to optimize over transforms of the form

$$\begin{bmatrix} \tilde{T}_1 & 0_{2 \times 2} \\ 0_{2 \times 2} & \tilde{T}_2 \end{bmatrix} P \quad (46)$$

with $\tilde{T}_i \in \mathbb{R}^{2 \times 2}$, $\det \tilde{T}_i = 1$, $i = 1, 2$, and P a permutation matrix.

Proof: See Appendix II-B. \square

Since the canonical building blocks defined in (38) solve the problem of designing \tilde{T}_i , $i = 1, 2$, we may write the transform as

$$T = \begin{bmatrix} T_{\alpha_1} & 0 \\ 0 & T_{\alpha_2} \end{bmatrix} P.$$

We only have to select the pairing and two transform parameters.

The ordering of a pair does not affect the possible performance, so there are three permutations of interest

$$P_1 = I, \quad P_2 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \text{and} \quad P_3 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

Let us consider a generic choice among the three by assuming that the inputs to transform T_{α_1} have variances ζ_1^2 and ζ_2^2 , $\zeta_1^2 \geq \zeta_2^2$; and the inputs to transform T_{α_2} have variances ζ_3^2 and ζ_4^2 , $\zeta_3^2 \geq \zeta_4^2$. Denote the redundancies associated with the pairs ρ_1 and ρ_2 , respectively. The redundancy and distortion are both additive between the pairs, so the problem is to minimize

$$D_1 = \frac{1}{2}(d_1 + d_2) \quad (47)$$

subject to the constraint¹⁷

$$\frac{1}{2}(\rho_1 + \rho_2) \leq \rho \quad (48)$$

where d_1 and d_2 are the side distortions contributed by the first and second pair, respectively. According to (37)

$$d_1 = \frac{1}{2}\zeta_2^2 + \frac{\zeta_1^2 - \zeta_2^2}{4 \cdot 2^{2\rho_1}(2^{2\rho_1} + \sqrt{2^{4\rho_1} - 1})}$$

$$d_2 = \frac{1}{2}\zeta_4^2 + \frac{\zeta_3^2 - \zeta_4^2}{4 \cdot 2^{2\rho_2}(2^{2\rho_2} + \sqrt{2^{4\rho_2} - 1})}.$$

Since d_1 and d_2 are strictly decreasing functions of ρ_1 and ρ_2 , respectively, (48) can be replaced by an equality. The optimal split occurs when both pairs operate at the same distortion-redundancy slope. However, since $\partial d_i / \partial \rho_i$ is complicated, there is no simple closed form for operating at the same slope.¹⁸

Let

$$c(\rho) = \frac{1}{2^{2\rho}(2^{2\rho} + \sqrt{2^{4\rho} - 1})}.$$

For large ρ , the 1 in the denominator becomes negligible, so

$$c(\rho) \approx \frac{1}{2^{2\rho}(2^{2\rho} + \sqrt{2^{4\rho}})} = \frac{1}{2 \cdot 2^{4\rho}}. \quad (49)$$

The error made in approximation (48) is shown in Fig. 12.

¹⁷The factor of 1/2 is present as normalization because redundancy is measured per component. This applies also to the 1/2 in (47). Note also that d_i 's are *partial* side distortions; hence, the use of a new symbol.

¹⁸Matching the slopes gives

$$\alpha^4 = \frac{\gamma(16\alpha_1^8\zeta_1^4 - \zeta_2^4) + \sqrt{\gamma^2(16\alpha_1^8\zeta_1^4 - \zeta_2^4)^2 + 64\alpha_1^8\zeta_3^4\zeta_4^4}}{32\alpha_1^4\zeta_3^4}$$

where

$$\gamma = \frac{\zeta_3^2\zeta_4^2(\zeta_3^2 - \zeta_4^2)}{\zeta_1^2\zeta_2^2(\zeta_1^2 - \zeta_2^2)}.$$

This exact relationship is used in generating Fig. 13, but is too complicated to use in drawing general conclusions about best pairings.




Fig. 12. Illustration of the accuracy of the approximation of $c(\rho)$ given by (49).

Using (49) it becomes feasible to optimally allocate redundancies analytically. Operating at equal slopes means that $\partial d_1/\partial \rho_1 = \partial d_2/\partial \rho_2$, yielding, together with (48), two linear equations with two unknowns:

$$\begin{aligned} \rho_1 - \rho_2 &= \frac{1}{4} \log \frac{\varsigma_1^2 - \varsigma_2^2}{\varsigma_3^2 - \varsigma_4^2} \\ \rho_1 + \rho_2 &= 2\rho. \end{aligned}$$

Solving for the optimal split of redundancies gives

$$\rho_1 = \rho + \frac{1}{8} \log \frac{\varsigma_1^2 - \varsigma_2^2}{\varsigma_3^2 - \varsigma_4^2} \quad (50)$$

$$\rho_2 = \rho - \frac{1}{8} \log \frac{\varsigma_1^2 - \varsigma_2^2}{\varsigma_3^2 - \varsigma_4^2}. \quad (51)$$

Substituting (50) and (51) into (47), using the approximation (49), gives

$$D_1 = \frac{1}{4} (\varsigma_2^2 + \varsigma_4^2) + \frac{1}{8} ((\varsigma_1^2 - \varsigma_2^2)(\varsigma_3^2 - \varsigma_4^2))^{1/2} 2^{-4\rho}.$$

To minimize D_1 for large ρ , we can immediately conclude that ς_2 and ς_4 should be the smallest of the σ 's. This eliminates permutation P_1 . The following manipulation shows that $\varsigma_1 > \varsigma_3 > \varsigma_4 > \varsigma_2$ is the ideal sorting, i.e., the largest variance and smallest variance components should be paired;

$$\begin{aligned} (\varsigma_1^2 - \varsigma_4^2)(\varsigma_2^2 - \varsigma_3^2) - (\varsigma_1^2 - \varsigma_3^2)(\varsigma_2^2 - \varsigma_4^2) \\ = -(\varsigma_1^2 - \varsigma_2^2)(\varsigma_3^2 - \varsigma_4^2) \leq 0. \end{aligned}$$

In other words, P_3 is the best permutation. We will see shortly that this "nested" pairing method generalizes to K pairs as well.

3) *2K Paired Variables*: Let us now consider transmission of $2K$ variables over the two channels with the odd and even indexed coefficients sent over channels 1 and 2, respectively. The extension of Theorem 3 to K pairs of variables would seem to naturally hold, but no proof of this has been found. Consider transforms of the following form, though we have not proven that this restriction is innocuous:

$$T = \begin{bmatrix} T_{\alpha_1} & & & \\ & T_{\alpha_2} & & \\ & & \ddots & \\ & & & T_{\alpha_K} \end{bmatrix} P$$





Fig. 13. Numerical calculations using the exact redundancy allocation solution (without approximation (49)) confirm that the nested pairing is optimal for all rates. (a) A random instance with two pairs. (b) A random instance with three pairs.

where P is a permutation matrix. Again let $\varsigma_1^2, \varsigma_2^2, \dots, \varsigma_{2K}^2$ denote the variances of the components after the permutation, with $\varsigma_{2i-1}^2 \geq \varsigma_{2i}^2$ for $i = 1, 2, \dots, K$. Denote the redundancy associated with T_{α_i} by ρ_i . We then have a similar problem as before. Minimize

$$D_1 = \frac{1}{K} \sum_{i=1}^K d_i \quad (52)$$

with

$$d_i = \frac{1}{2} \varsigma_{2i}^2 + \frac{\varsigma_{2i-1}^2 - \varsigma_{2i}^2}{4 \cdot 2^{2\rho_i} (2^{2\rho_i} + \sqrt{2^{4\rho_i} - 1})}, \quad \text{for } i = 1, 2, \dots, K$$

subject to the constraint

$$\rho = \frac{1}{K} \sum_{i=1}^K \rho_i. \quad (53)$$

Using the approximation (49) for large ρ and imposing the equal-slope conditions gives a system of K linear equations with K unknowns ($K - 1$ equations coming from equal-slope

more complicated than what has been presented thus far because there are eight degrees of freedom remaining after fixing $\det T = 1$. Even in the case of equally weighted channel failures, a closed-form solution would be more complicated than (36).

Assume $\sigma_1 > \sigma_2 > \sigma_3$ and, furthermore, that the loss probabilities are small. Then one loss occurs much more frequently than two losses. Motivated by maintaining symmetry with the coordinate axes and skewing the analysis basis toward the two components with largest variance, while keeping the rates and importances of the descriptions equal, leads to transforms described by

$$\begin{bmatrix} a & -\frac{\sqrt{3}\sigma_1 a}{\sigma_2} & -\frac{\sigma_2}{6\sqrt{3}\sigma_1^2 a^2} \\ 2a & 0 & \frac{\sigma_2}{6\sqrt{3}\sigma_1^2 a^2} \\ a & \frac{\sqrt{3}\sigma_1 a}{\sigma_2} & -\frac{\sigma_2}{6\sqrt{3}\sigma_1^2 a^2} \end{bmatrix} \quad (55)$$

for arbitrary a . Fig. 14 compares the performance of transforms chosen from this set to arbitrary numerically optimized transforms for a source with $(\sigma_1, \sigma_2, \sigma_3) = (1, 0.7, 0.4)$ when each description is equally likely to be lost. At any particular redundancy, the distortion with one loss and the distortion with two losses can be traded off. At the computed redundancies, transforms given by (55) are nearly optimal when the distortion with one lost description is more heavily weighted than the distortion with two lost descriptions.

A simple heuristic for designing systems with more than two channels is to cascade small transforms. Just as the parallel use of two-by-two transforms gave a method for sending $2K$ variables over two channels, a cascade combination of these transforms gives a method for sending 2^K variables over 2^K channels. The cascade structure simplifies the encoding, decoding, and design when compared to using a general $2^K \times 2^K$ transform.

The simplest instance of the cascade structure is shown in Fig. 15, where four variables are sent over four channels. This is equivalent to the use of a transform of the form

$$T = \begin{bmatrix} T_\gamma & 0 \\ 0 & T_\gamma \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} T_\alpha & 0 \\ 0 & T_\beta \end{bmatrix}. \quad (56)$$

Though T in (56) has only three degrees of freedom—in place of 15 in a general determinant-1 transform of this size—empirical evidence suggests that this class of transforms is sufficiently rich to give optimal performance for one lost description and nearly optimal performance for two lost descriptions.

For a numerical example, consider redundancy $\rho = 0.125$ bits/component added to a source is described by

$$(\sigma_1, \sigma_2, \sigma_3, \sigma_4) = (1, 0.8, 0.6, 0.4)$$

to mitigate erasure effects. Assume each description is equally likely to be lost. Limiting attention to two parallel 2×2 transforms as in (46) gives a relatively simple optimization over two variables, but gives the worst performance (see Fig. 16). With a cascade transform (56), the optimization is somewhat more dif-




Fig. 15. Cascade structure for MDCT coding of four variables to be transmitted over four channels. The cascade structure simplifies the design procedure by reducing 15 free parameters to three. The use of γ in both second-stage blocks is because each first-stage block produces streams of equal rate and equal importance.

ficult, but distortion can be reduced simultaneously for one and two erasures (three-erasure performance is also improved, but not shown). Allowing a general 4×4 transform makes optimization considerably more difficult, and does not further reduce the distortion with one erasure.

V. APPLICATIONS

In practice, it is unusual to see a Gaussian source with a known distribution. Nevertheless, transform coding techniques are prevalent in audio, image, and video coding. This section briefly describes applications of correlating transform-based MDs to image and audio coding. Details appear elsewhere in the literature.

A. Application to Image Coding

On the Internet, images are usually communicated with progressive source coding and retransmission of lost packets. When there are no packet losses, the receiver can reconstruct the image as the packets arrive; but when there is a packet loss, there is a large period of latency while the transmitter determines that the packet must be retransmitted and then retransmits the packet. The latency is due to the fact that the application at the receiving end uses the packets only after they have been put in the proper sequence. To combat this latency problem, it is desirable to have a communication system that is robust to arbitrarily placed packet erasures and that can reconstruct an image progressively from packets received in any order. The MDCT method of this paper seems suitable for this task.

It is precisely in the context of image communication that Wang, Orchard, and Reibman introduced correlating transforms for MD coding; their image coding results appear in [10]. (No image coding results appear in their later papers [1], [11].) Image coding results using the generalized framework presented here appear in [15]. These experiments indicate the expected behavior: at the expense of a slight increase in bit rate, image representations can be made more robust to packet losses. Conventional channel codes adding such a small amount of redundancy are not as effective.

B. Application to Audio Coding

To provide the MD feature for packetized compressed audio, correlating transforms were introduced in a well-known existing

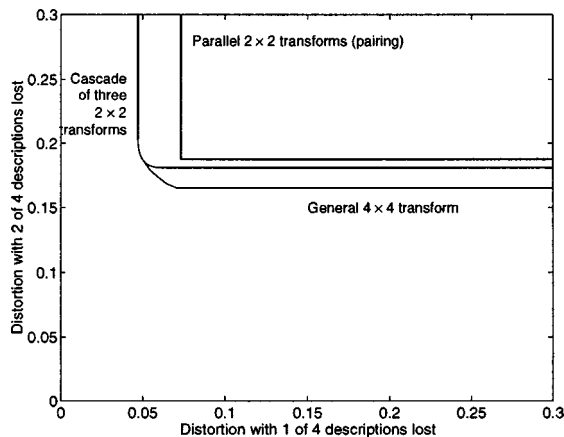


Fig. 16. Comparison between the cascade transform (56), pairing, and a general 4×4 transform for sending four variables over four channels. The source has $(\sigma_1, \sigma_2, \sigma_3, \sigma_4) = (1, 0.8, 0.6, 0.4)$, and in all three cases redundancy $\rho = 0.125$ bit/component. With the zero-erasure performance held equal, the cascade transform can simultaneously give better performance than pairing with one or two components lost. Removing restrictions further by allowing a general 4×4 transform allows a further decrease in two-erasure distortion.

coder: Bell Labs' *Perceptual Audio Coder (PAC)* [50]. Experiments confirm that the new MD audio coder is robust and gracefully degrades as the number of lost packets increases. The implementation and results are described in [51]. Since the results cannot be adequately conveyed numerically or visually, the reader is invited to hear them. Audio files are provided on-line in *aiff*, *wave*, and *next* formats at <http://cm.bell-labs.com/who/jelena/Interests/MD/AudioDemo/DemoList.html>

VI. CONCLUSION

This paper has introduced a method for generalized MD coding using correlating transforms. The framework is a generalization of the method proposed by Orchard, Wang, Vaishampayan, and Reibman [1]. In addition to extending the technique to more than two descriptions, we have provided a complete analysis and optimization of the two-description case, allowing an arbitrary pdf on the system state.

This method is very effective in increasing robustness with a small amount of redundancy. In contrast to transform coding systems designed purely for compression, not robustness, it produces correlated transform coefficients. As “proofs of concept,” applications of this method to image and audio coding were briefly described.

The MD scenario provides a good analogy to communication over a lossy packet network. For this reason, “description,” “channel,” and “packet” have been used interchangeably. However, this is not the only communication environment in which MD coding may be useful. Effros and Goldsmith [52] have studied the various notions of capacity for general time-varying channels. One of their results is that more information can be reliably received than can be reliably transmitted. With some thought, this is an intuitive result: it is less demanding to ask for every bit that gets across the channel to be correct than to ask for every bit that is transmitted to correctly get across the channel. For such a general channel it may be useful to use an MD source

code since all the received information will be useful, but the loss of some of the transmitted information is not catastrophic.

APPENDIX I PSEUDOLINEAR DISCRETE TRANSFORMS

Invertible discrete-domain to discrete-domain transforms are alternatively described as reversible, integer-to-integer, or invertible in finite precision [53]–[55]. These transforms are based on factorizations of matrices which make information flow in a simple, regular way. Inversion can then be achieved by reversing the information flow.

For example, one can factor any 2×2 matrix with determinant 1 into three lower- and upper-triangular matrices with unit diagonals as

$$\begin{aligned} T &= \begin{bmatrix} a & b \\ c & d \end{bmatrix} \\ &= \underbrace{\begin{bmatrix} 1 & 0 \\ (d-1)/b & 1 \end{bmatrix}}_{T_1} \underbrace{\begin{bmatrix} 1 & b \\ 0 & 1 \end{bmatrix}}_{T_2} \underbrace{\begin{bmatrix} 1 & 0 \\ (a-1)/b & 1 \end{bmatrix}}_{T_3} \\ \text{or} \\ &= \underbrace{\begin{bmatrix} 1 & (a-1)/c \\ 0 & 1 \end{bmatrix}}_{T_1} \underbrace{\begin{bmatrix} 1 & 0 \\ c & 1 \end{bmatrix}}_{T_2} \underbrace{\begin{bmatrix} 1 & (d-1)/c \\ 0 & 1 \end{bmatrix}}_{T_3}. \end{aligned}$$

Since the inverse of a block

$$\begin{bmatrix} 1 & 0 \\ x & 1 \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} 1 & y \\ 0 & 1 \end{bmatrix}$$

is simply

$$\begin{bmatrix} 1 & 0 \\ -x & 1 \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} 1 & -y \\ 0 & 1 \end{bmatrix}$$

respectively, the inverse of T can be found by reversing the order of the factors and changing the signs of the off-diagonal elements.

The more profound fact is that the simplicity of inversion remains if the off-diagonal elements represent nonlinear functions. Let $[\cdot]_{\Delta}$ represent rounding to the nearest multiple of Δ and let

$$T_1 = \begin{bmatrix} 1 & a \\ 0 & 1 \end{bmatrix}.$$

If $x \in \Delta\mathbb{Z}^2$, then

$$\begin{aligned} [T_1 x]_{\Delta} &= \left[\begin{bmatrix} 1 & a \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \right]_{\Delta} = \left[\begin{bmatrix} x_1 + ax_2 \\ x_2 \end{bmatrix} \right]_{\Delta} \\ &= \begin{bmatrix} x_1 + [ax_2]_{\Delta} \\ x_2 \end{bmatrix}. \end{aligned}$$

Thus, $[T_1 \cdot]_{\Delta}$ is an identity operator except for a nonlinear function of x_2 being added to x_1 . Direct computation shows that on the domain $\Delta\mathbb{Z}^2$, $[T_1^{-1} \cdot]_{\Delta}$ is the inverse operator. A cascade of such operations is invertible in the same manner, so a factorization $T = T_1 T_2 T_3$ yields an invertible discrete transform $\hat{T}: \Delta\mathbb{Z}^2 \rightarrow \Delta\mathbb{Z}^2$ “derived from T ” through

$$\hat{T}(x) = [T_1 [T_2 [T_3 x]_{\Delta}]_{\Delta}]_{\Delta}. \quad (57)$$

The discrete transform \hat{T} depends not only T , but the factorization of T . Among the possible factorizations, one can minimize a bound on $\|\hat{T}(x) - Tx\|$. Let

$$T_1 = \begin{bmatrix} 1 & 0 \\ a & 1 \end{bmatrix}, \quad T_2 = \begin{bmatrix} 1 & b \\ 0 & 1 \end{bmatrix}, \quad \text{and} \quad T_3 = \begin{bmatrix} 1 & 0 \\ c & 1 \end{bmatrix}.$$

For $x \in \Delta\mathbb{Z}^2$, the computation (57) involves three rounding operations. Using δ_i 's to denote the roundoff errors gives

$$\hat{T}(x) = T_1 \left(T_2 \left(T_3 x + \begin{bmatrix} 0 \\ \delta_1 \end{bmatrix} \right) + \begin{bmatrix} \delta_2 \\ 0 \end{bmatrix} \right) + \begin{bmatrix} 0 \\ \delta_3 \end{bmatrix}.$$

Expanding and using $T_1 T_2 T_3 = T$, one can compute

$$\begin{aligned} \|\hat{T}(x) - Tx\|_\infty &= \left\| T_1 T_2 \begin{bmatrix} 0 \\ \delta_1 \end{bmatrix} + T_1 \begin{bmatrix} \delta_2 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ \delta_3 \end{bmatrix} \right\|_\infty \\ &= \left\| \begin{bmatrix} b\delta_1 \\ (1+ab)\delta_1 \end{bmatrix} + \begin{bmatrix} \delta_2 \\ a\delta_2 \end{bmatrix} + \begin{bmatrix} 0 \\ \delta_3 \end{bmatrix} \right\|_\infty \\ &\leq (1 + \max\{|b|, |a| + |1+ab|\}) \frac{\Delta}{2}. \end{aligned}$$

This shows that \hat{T} approximates T in a precise sense; in particular, $\hat{T}(x) \approx Tx$ when Δ is small.

For $N \times N$ matrices, the process is similar. T is factored into a product of matrices with unit diagonals and nonzero off-diagonal elements only in one row or column: $T = T_1 T_2 \cdots T_k$. The discrete version of the transform is then given by

$$\hat{T}(x) = [T_1 [T_2 \cdots [T_k x]_\Delta]_\Delta]_\Delta. \quad (58)$$

The lifting structure ensures that the inverse of \hat{T} can be implemented by reversing the calculations in (58)

$$\hat{T}^{-1}(y) = [T_k^{-1} \cdots [T_2^{-1} [T_1^{-1} y]_\Delta]_\Delta]_\Delta.$$

The existence of such a factorization follows from the fact that any nonsingular matrix can be reduced to an identity matrix by multiplication with elementary matrices [56]. Since our original matrix has determinant 1, it can be represented as the product of the following three types of elementary matrices:

- $E_{ij}^{(\lambda)}$, to subtract a multiple λ of row j from row i ;
- P_{ij} , to exchange rows i and j ;
- $D_{ij}^{(\lambda)}$, to multiply row i by λ and row j by $1/\lambda$.

$E_{ij}^{(\lambda)}$ is already in the desired form. The remaining two can be factored as desired using the factorization of 2×2 matrices above. Thus, any matrix with determinant 1 can be factored as required.

APPENDIX II PROOFS

A. Proof of Theorem 2

The proofs of Theorems 2 and 3 utilize the following lemma.

Lemma 6: Let A , X , and Y be symmetric, real, positive-definite matrices and let $Q = Y^{1/2} A Y^{1/2}$. Suppose both X and Q each have distinct eigenvalues. Denote orthogonal eigen-decompositions of X and Q by $X = V_X \Lambda_X V_X^T$ and $Q =$

$V_Q \Lambda_Q V_Q^T$, respectively, where Λ_X and Λ_Q have decreasing diagonals. Then the optimization problem

$$\text{minimize } J(U) = \text{tr } AU^T U \quad \text{subject to } U^T X U = Y$$

is solved by

$$U_0 = V_X \Lambda_X^{-1/2} V_Q^T Y^{1/2} \quad (59)$$

yielding

$$J(U_0) = \text{tr } \Lambda_Q \Lambda_X^{-1}. \quad (60)$$

This solution for U_0 is unique up to the sign choices in defining V_X and V_Q .

In the proof of Lemma 6, the following elementary fact is used.

Lemma 7: Suppose $\text{tr } BS = 0$ for all skew-symmetric matrices S . Then B is symmetric.

Proof (Lemma 7): For any $i \neq j$, let S be the matrix with $+1$ in the (i, j) position, -1 in the (j, i) position, and the remaining elements equal to zero. Since $0 = \text{tr } BS = B_{ji} - B_{ij}$, $B_{ij} = B_{ji}$. \square

Proof (Lemma 6): First convert the constraint $U^T X U = Y$ to a simpler form. Left and right multiplying by $Y^{-1/2}$ and splitting $X = X^{1/2} \cdot X^{1/2}$ gives

$$Y^{-1/2} U^T X^{1/2} \cdot X^{1/2} U Y^{-1/2} = I.$$

With the definition $\tilde{U} = X^{1/2} U Y^{-1/2}$, we have the constraint $\tilde{U}^T \tilde{U} = I$.

The objective function is now

$$J(U) = J \left(X^{-1/2} \tilde{U} Y^{1/2} \right) \quad (61)$$

$$\begin{aligned} &= \text{tr } A Y^{1/2} \tilde{U}^T X^{-1} \tilde{U} Y^{1/2} \\ &= \text{tr } Y^{1/2} A Y^{1/2} \tilde{U}^T X^{-1} \tilde{U} \end{aligned} \quad (62)$$

$$= \text{tr } Q \tilde{U}^T X^{-1} \tilde{U} \quad (63)$$

where (62) uses the fact that cyclic permutation of factors does not affect the trace. We are left with minimizing (63) over orthogonal transforms \tilde{U} .

A differential analysis will reveal a single critical point, up to sign choices. Since A is positive definite, this critical point must be a minimum. Consider a small change to \tilde{U} , $\tilde{U}_\delta = \tilde{U} + \delta$. To obey the orthogonality constraint, we must have $\tilde{U}_\delta^T \tilde{U}_\delta = I$. Expanding $\tilde{U}_\delta^T \tilde{U}_\delta$ and neglecting the δ^2 term gives the constraint

$$\tilde{U}^T \delta + \delta^T \tilde{U} = 0. \quad (64)$$

The perturbation has the following affect on the objective function:

$$\begin{aligned} &\text{tr } Q \tilde{U}_\delta^T X^{-1} \tilde{U}_\delta \\ &= \text{tr } Q \left(\tilde{U} + \delta \right)^T X^{-1} \left(\tilde{U} + \delta \right) \\ &= \text{tr } Q \tilde{U}^T X^{-1} \tilde{U} + \text{tr } Q \tilde{U}^T X^{-1} \delta \\ &\quad + \text{tr } Q \delta^T X^{-1} \tilde{U} + \text{tr } Q \delta^T X^{-1} \delta \\ &\approx \text{tr } Q \tilde{U}^T X^{-1} \tilde{U} + 2 \text{tr } Q \tilde{U}^T X^{-1} \delta \\ &= J(U) + 2 \text{tr } Q \tilde{U}^T X^{-1} \delta \end{aligned} \quad (65)$$

where the approximation (65) results from discarding the $O(\|\delta\|^2)$ term and using $\text{tr } M = \text{tr } M^T$. Thus, a critical point of $J(\cdot)$ is a transform \tilde{U} that satisfies

$$\text{tr } Q\tilde{U}^T X^{-1}\delta = 0 \quad (66)$$

for all small δ satisfying (64).

The solutions of (64) are simple; they are $\delta = \tilde{U}S$, where S is an arbitrary skew-symmetric matrix. Substituting $\delta = \tilde{U}S$, the solutions of (66) are \tilde{U} such that $\text{tr } Q\tilde{U}^T X^{-1}\tilde{U}S = 0$ for all skew-symmetric S . Now, by Lemma 7, $Q\tilde{U}^T X^{-1}\tilde{U}$ must be symmetric. Notice the effect of transposing this matrix. Since Q and X are symmetric, the transpose is $\tilde{U}^T X^{-1}\tilde{U}Q$; thus, Q and $\tilde{U}^T X^{-1}\tilde{U}$ commute.

Diagonalizable matrices commute if and only if they are simultaneously diagonalizable [57]. Furthermore, for a matrix with distinct eigenvalues, the orthogonal transform that diagonalizes and leaves the diagonal in decreasing order is unique up to sign choices. On one hand, Q is diagonalized as

$$\underbrace{V_Q^T}_{\text{diagonalizing transform}} \cdot Q \cdot \underbrace{V_Q}_{\text{transposed transform}} = \Lambda_Q;$$

on the other hand, $\tilde{U}^T X^{-1}\tilde{U}$ is diagonalized as

$$\underbrace{V_X^T \tilde{U}}_{\text{diagonalizing transform}} \cdot \underbrace{\tilde{U}^T V_X \Lambda_X^{-1} V_X^T \tilde{U}}_{X^{-1}} \cdot \underbrace{\tilde{U}^T V_X}_{\text{transposed transform}} = \Lambda_X^{-1}.$$

Ignoring sign choices, we may equate the diagonalizing transforms using a permutation matrix P , yielding $PV_Q^T = V_X^T \tilde{U}$. The permutation will be chosen after its effect is made clear. A simple sequence of substitutions yields the optimal transform U_0

$$\begin{aligned} U_0 &= X^{-1/2} \tilde{U} Y^{1/2} = V_X \Lambda_X^{-1/2} V_X^T \cdot V_X P V_Q^T \cdot Y^{1/2} \\ &= V_X \Lambda_X^{-1/2} P V_Q^T Y^{1/2}. \end{aligned} \quad (67)$$

Evaluating $J(U_0)$ gives

$$\begin{aligned} J(U_0) &= \text{tr } A U^T U \\ &= \text{tr } \underbrace{Y^{-1/2} Q Y^{-1/2} \cdot Y^{1/2} V_Q P^T \Lambda_X^{-1/2} V_X^T \cdot V_X \Lambda_X^{-1/2} P}_{\cdot V_Q^T Y^{1/2}} \end{aligned} \quad (68)$$

$$\begin{aligned} &= \text{tr } V_Q^T Q V_Q \cdot P^T \Lambda_X^{-1} P \\ &= \text{tr } \Lambda_Q P^T \Lambda_X^{-1} P \end{aligned} \quad (69)$$

where terms are separated to emphasize substitutions, and underbraces in (68) mark terms that are subsequently commuted. Since Λ_Q and Λ_X are already sorted in the same order, (69) is minimized by choosing the identity permutation $P = I$. This finally yields (59) and (60). \square

We are now prepared to prove the theorem. The overall strategy is as follows: Starting with any transform T , we can find a transform V such that VT results in identical side distortion and at most the same redundancy as T . At the same time, VT yields a correlation R_y with a particular, simple form. This simple form in turn leads to a simple expression for the side

distortion D_1 . Lemma 6 is then used to show that the transform that yields minimum D_1 among transforms with correlation R_y has the desired form (44). Since the performance with any transform can at least be matched with a transform of the form (44), the proof is complete. Each step is now detailed.

Recall that T is an arbitrary transform. Let $R_y = T R_x T^T$, the correlation matrix of the transform coefficients when T is used. Since y_2 and y_3 are sent on the same channel, an invertible transform applied to the two will not change D_1 . However, if y_2 and y_3 are correlated, the rate can be reduced by applying a decorrelating transform. Denote a KLT for

$$\begin{bmatrix} (R_y)_{22} & (R_y)_{23} \\ (R_y)_{32} & (R_y)_{33} \end{bmatrix}$$

by \tilde{V}_1 , and let

$$V_1 = \begin{bmatrix} 1 & 0_{1 \times 2} \\ 0_{2 \times 1} & \tilde{V}_1 \end{bmatrix}.$$

Then, using $V_1 T$ in place of T does not change the side distortion D_1 , and does not increase the redundancy ρ .

After the application of V_1 , the correlation can be written in the following form:

$$V_1 T R_x T^T V_1^T = \begin{bmatrix} \varsigma_1^2 & a_1 & a_2 \\ a_1 & \varsigma_2^2 & 0 \\ a_2 & 0 & \varsigma_3^2 \end{bmatrix}.$$

This type of correlation structure cannot be produced by a transform of the desired form (44) unless $a_2 = 0$, so we simplify the correlation structure further. Let

$$\tilde{V}_2 = \frac{1}{\sqrt{a_1^2 \varsigma_3^2 + a_2^2 \varsigma_2^2}} \begin{bmatrix} \sigma_2^{-1} a_1 \varsigma_3^2 & \sigma_2^{-1} a_2 \varsigma_2^2 \\ -a_2 \sigma_2 & a_1 \sigma_2 \end{bmatrix}$$

and

$$V_2 = \begin{bmatrix} 1 & 0_{1 \times 2} \\ 0_{2 \times 1} & \tilde{V}_2 \end{bmatrix}.$$

Then

$$\begin{aligned} &V_2 V_1 T R_x T^T V_1^T V_2^T \\ &= \begin{bmatrix} \varsigma_1^2 & \sigma_2^{-1} \sqrt{a_1^2 \varsigma_3^2 + a_2^2 \varsigma_2^2} & 0 \\ \sigma_2^{-1} \sqrt{a_1^2 \varsigma_3^2 + a_2^2 \varsigma_2^2} & \sigma_2^{-2} \varsigma_2^2 \varsigma_3^2 & 0 \\ 0 & 0 & \sigma_2^2 \end{bmatrix}. \end{aligned} \quad (70)$$

(The reason for selecting V_2 with deference to σ_2 is revealed later.) With reference to (41), recall that the redundancy depends only on the product variances of the transform coefficients. Since the product of the diagonal elements of $V_2 V_1 T R_x T^T V_1^T V_2^T$ and $V_1 T R_x T^T V_1^T$ are equal, using $V_2 V_1 T$ in place of $V_1 T$ does not change the redundancy. Furthermore, since V_2 merely alters the second and third components in an invertible manner, D_1 is also unchanged. With $V = V_2 V_1$, we have found a transform such that VT is at least as good for MD coding as T , and $V T R_x T^T V^T$ has a simple form; the first step of the proof is complete.

We now wish to show that for R_y of the form (70), the optimal transform has the desired form (44). In light of the uniqueness of the solution in Lemma 6, this lemma could be used to directly compute the best transform for the given R_y . This is not needed

to satisfy the statement of the theorem. More importantly, our ultimate goal is to minimize the side distortion for a given redundancy, not a given R_y . Even if two R_y 's yield the same redundancy, their corresponding minimum distortions may not be the same; after finding the minimum distortion as a function of R_y we would have to minimize over all R_y with the particular redundancy. Instead, finding only the form of the optimal transform simplifies the following computations significantly.

To simplify notation, let

$$R_y = \begin{bmatrix} \gamma_1 & a & 0 \\ a & \gamma_2 & 0 \\ 0 & 0 & \sigma_2^2 \end{bmatrix}.$$

Inverting $TR_xT^T = R_y$ gives $U^TR_x^{-1}U = R_y^{-1}$ where, as in Section IV-B, $U = T^{-1}$. When y_1 (channel 1) is lost, A in (43) is given by

$$A_1 = \gamma_1 - [a \ 0] \begin{bmatrix} \gamma_2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix}^{-1} \begin{bmatrix} a \\ 0 \end{bmatrix} = \gamma_1 - \frac{a^2}{\gamma_2}.$$

When (y_2, y_3) (channel 2) is lost, the corresponding quantity is

$$A_2 = \begin{bmatrix} \gamma_2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix} - \begin{bmatrix} a \\ 0 \end{bmatrix} \gamma_1^{-1} [a \ 0] = \begin{bmatrix} \gamma_2 - \gamma_1^{-1}a^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix}.$$

The average side distortion per component when the channels are equally likely to be lost is given by

$$D_1 = \frac{1}{6} \text{tr} AU^TU, \quad \text{where } A = \begin{bmatrix} A_1 & 0_{1 \times 2} \\ 0_{2 \times 1} & A_2 \end{bmatrix}.$$

It is now clear that we have an optimization that can be solved with Lemma 6. Identify

$$X = R_x^{-1} = \text{diag}(\sigma_1^{-2}, \sigma_2^{-2}, \sigma_3^{-2})$$

and

$$Y = R_y^{-1} = \begin{bmatrix} \tilde{Y} & 0_{2 \times 1} \\ 0_{1 \times 2} & \sigma_2^{-2} \end{bmatrix}$$

where we need not specify \tilde{Y} because we are primarily interested in sparsity. Other quantities that appear in Lemma 6 can easily be computed.²⁰ Since X is already diagonal, $V_X = I$ and $\Lambda_X^{-1/2} = \text{diag}(\sigma_1, \sigma_2, \sigma_3)$. The sparsity of Y gives

$$\begin{aligned} Q &= Y^{1/2}AY^{1/2} = \begin{bmatrix} \tilde{Y}^{1/2} & 0_{2 \times 1} \\ 0_{1 \times 2} & \sigma_2^{-1} \end{bmatrix} A \begin{bmatrix} \tilde{Y}^{1/2} & 0_{2 \times 1} \\ 0_{1 \times 2} & \sigma_2^{-1} \end{bmatrix} \\ &= \begin{bmatrix} \tilde{Y}^{1/2} \tilde{A} \tilde{Y}^{1/2} & 0_{2 \times 1} \\ 0_{1 \times 2} & 1 \end{bmatrix} \end{aligned} \quad (71)$$

so a diagonalizing transform of Q will have the form

$$V_Q = \begin{bmatrix} V_Q & 0_{2 \times 1} \\ 0_{1 \times 2} & 1 \end{bmatrix}.$$

Now substituting in (67) gives

$$U_0 = \text{diag}(\sigma_1, \sigma_2, \sigma_3)P \begin{bmatrix} V_Q^T & 0_{2 \times 1} \\ 0_{1 \times 2} & 1 \end{bmatrix} \begin{bmatrix} \tilde{Y}^{1/2} & 0_{2 \times 1} \\ 0_{1 \times 2} & \sigma_2^{-1} \end{bmatrix}.$$

The optimal transform is the inverse of U_0

$$T = \begin{bmatrix} \tilde{Y}^{-1/2} V_Q & 0_{2 \times 1} \\ 0_{1 \times 2} & \sigma_2 \end{bmatrix} P^T \text{diag}(\sigma_1^{-1}, \sigma_2^{-1}, \sigma_3^{-1}). \quad (72)$$

²⁰Arbitrary sign and permutation choices are made in diagonalizing transforms; sorting of diagonal elements is handled later.

It remains now to determine the permutation P in (72). This depends on how Λ_Q must be permuted to match the ordering of Λ_X . First note that $\Lambda_X = X$ is sorted in decreasing order. It is not necessary to precisely determine the eigenvalues of Q to find the required permutation. The sum of the eigenvalues of Q is given by

$$\begin{aligned} \text{tr} Q &= \text{tr} Y^{1/2} A Y^{1/2} = \text{tr} A Y \\ &= \text{tr} \begin{bmatrix} 1 & -a/\gamma_2 & 0 \\ -a/\gamma_1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = 3. \end{aligned}$$

From the form of (71) one of the eigenvalues is 1. In the generic case, the eigenvalues are distinct and so the remaining eigenvalues sandwich the eigenvalue 1. To counteract the different sorting of Λ_X and Λ_Q , the permutation must move the third element to the middle. With such a permutation, (72) simplifies to

$$T = \begin{bmatrix} \tilde{Y}^{-1/2} V_Q & \text{diag}(\sigma_1^{-1}, \sigma_3^{-1}) \tilde{P} & 0_{2 \times 1} \\ 0_{1 \times 2} & 1 & \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

where \tilde{P} is a 2×2 permutation that depends on whether V_Q is a clockwise or counterclockwise rotation. (The cancellation of σ_2 now retrospectively explains why this standard deviation was singled out.) The final transform is in the form (44) and thus the proof is complete.

This analysis could be pushed further to determine the optimal transform. However, this would merely be the optimal transform given R_y , not the optimal transform for given redundancy ρ . Applying this theorem, an optimal transform with an upper bound on ρ is easy to compute using the results from Section III (see Section IV-C).

B. Proof of Theorem 3

This theorem is similar to Theorem 2 so an abbreviated proof is given. The strategy of the proof is to start with an arbitrary transform T . The corresponding correlation matrix TR_xT^T may be fully dense, but there is a transform V with determinant 1 such that $VTR_xT^TV^T$ has a simple desired form and VT is no worse than T for use in the system. An application of Lemma 6 then shows that the optimal transform is of the desired form (46). The simplification of the correlation matrix is detailed below, but the application of Lemma 6 is omitted.

As in Appendix II-A, components sent over the same channel can be made uncorrelated without affecting the side distortion, while not increasing the redundancy. Using KLTs to decorrelate in such a manner gives

$$V_1 TR_x T^T V_1^T = \begin{bmatrix} \zeta_1^2 & a_{11} & 0 & a_{12} \\ a_{11} & \zeta_2^2 & a_{21} & 0 \\ 0 & a_{21} & \zeta_3^2 & a_{22} \\ a_{12} & 0 & a_{22} & \zeta_4^2 \end{bmatrix}.$$

Now we would like to find \tilde{W}_1 and \tilde{W}_2 with $\det \tilde{W}_i = 1$, $i = 1, 2$, such that

$$V_2 = \begin{bmatrix} \tilde{W}_1 & 0_{2 \times 2} \\ 0_{2 \times 2} & \tilde{W}_2 \end{bmatrix} P, \quad \text{with } P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

gives the desired correlation structure. A solution is obtained with

$$\tilde{W}_1 = \begin{bmatrix} \alpha & \beta \\ -(\alpha^2 \varsigma_1^2 + \beta^2 \varsigma_3)^{-1} \beta \varsigma_3^2 & (\alpha^2 \varsigma_1^2 + \beta^2 \varsigma_3)^{-1} \alpha \varsigma_1^2 \end{bmatrix}$$

$$\tilde{W}_2 = \begin{bmatrix} \gamma & \delta \\ -(\gamma^2 \varsigma_2^2 + \delta^2 \varsigma_4)^{-1} \delta \varsigma_4^2 & (\gamma^2 \varsigma_2^2 + \delta^2 \varsigma_4)^{-1} \gamma \varsigma_3^2 \end{bmatrix}$$

where α is a root of

$$\alpha^2 - \frac{\beta(\varsigma_3^2 \varsigma_4^2 a_{11}^2 + \varsigma_2^2 \varsigma_3^2 a_{12}^2 - \varsigma_1^2 \varsigma_4^2 a_{21}^2 + \varsigma_1^2 \varsigma_2^2 a_{22}^2)}{\varsigma_1^2 (\varsigma_4^2 a_{11} a_{21} + \varsigma_2^2 a_{12} a_{22})} \alpha - \frac{\beta^2 \varsigma_1^2}{\varsigma_3^2} = 0 \quad (73)$$

and

$$\gamma = \frac{\delta \varsigma_4^2 (\alpha a_{11} + \beta a_{21})}{\varsigma_2^2 (\alpha a_{12} + \beta a_{22})}. \quad (74)$$

The validity of this solution places no constraint on δ . The choice of β should ensure that (73) has real roots and (74) does not involve a division by zero. These requirements are easily satisfied by choosing β to have the same sign as $\varsigma_3^2 \varsigma_4^2 a_{11}^2 + \varsigma_2^2 \varsigma_3^2 a_{12}^2 - \varsigma_1^2 \varsigma_4^2 a_{21}^2 + \varsigma_1^2 \varsigma_2^2 a_{22}^2$ and eliminating a few isolated points.²¹

With $V = V_2 V_1$, we have a transform such that VT is at least as good as T ; i.e., VT gives the same average side distortion and at most the same redundancy as T . Also, $VTR_x T^T V^T$ has a simple block diagonal form. This block diagonal form permits an application of Lemma 6 which shows that the optimal transform is of the desired form (46). The application of Lemma 6 parallels its use in Appendix II-A and is thus omitted. This completes the proof.

Note that the complicated dependence of V_2 on the various parameters hinders extending this method of proof to Conjecture 5.

C. Proof of Theorem 4

First note that for high ρ , the distortion given by (54) is dominated by the first term, so the ς_{2i} 's must be the K smallest variances. Since we are interested in the pairing but not the order of the pairs, we may assign

$$\varsigma_{2i} = \sigma_{2K+1-i}, \quad \text{for } i = 1, 2, \dots, K$$

without loss of generality. Now it remains to show that

$$\varsigma_{2i-1} = \sigma_i, \quad \text{for } i = 1, 2, \dots, K \quad (75)$$

minimizes (54). The proof is completed by showing that any permutation other than (75) can be improved, or is already equivalent to (75) because of nondistinct σ_i 's.

Suppose some permutation other than (75) is used and let i^* be the smallest i for which (75) is violated. Say $\varsigma_{2i^*-1} = \sigma_j$ (instead of σ_{i^*}). Then $j > i^*$ because $j \leq i^*$ would contradict the definition of i^* . Similarly, if σ_{i^*} is paired with σ_k , then $k < 2K - i^* + 1$, for if not (75) would be violated at $i = 2K - k + 1 < i^*$. ($k = 2K - i^* + 1$ has been eliminated because this would imply that (75) is *not* violated at i^* .)

We assert that the distortion is reduced by swapping σ_{i^*} and σ_j (unless $\sigma_{i^*} = \sigma_j$ or $\sigma_k = \sigma_{2K-i^*+1}$, in which case, the distortion is unchanged and we may proceed by looking for larger i^*). The first term in (54) is unaffected by the swap, but the second term is multiplied by

$$\left[\frac{(\sigma_{i^*}^2 - \sigma_{2K-i^*+1}^2)(\sigma_j^2 - \sigma_k^2)}{(\sigma_{i^*}^2 - \sigma_k^2)(\sigma_j^2 - \sigma_{2K-i^*+1}^2)} \right]^{1/K}.$$

When the σ 's in question are distinct, this factor is less than one because of calculations for the two pair case. Since only (75) and equivalent permutations are not improved by this process, the theorem is proven.

ACKNOWLEDGMENT

The authors are grateful to M. T. Orchard, A. R. Reibman, V. A. Vaishampayan, and Y. Wang for exchanging ideas pertaining to MD coding. They wish to thank R. Han for providing image file size statistics and S. Rangan for his help in stating and proving Lemma 6. Many useful comments were provided by A. Lozano, J. E. Mazo, and the anonymous reviewers. Finally, the authors wish to thank T. M. Cover, A. Gersho, L. Ozarow, and H. S. Witsenhausen for clarifying the early history of the MD problem.

REFERENCES

- [1] M. T. Orchard, Y. Wang, V. Vaishampayan, and A. R. Reibman, "Redundancy rate-distortion analysis of multiple description coding using pairwise correlating transforms," in *Proc. IEEE Int. Conf. Image Processing*, vol. I, Santa Barbara, CA, Oct. 1997, pp. 608–611.
- [2] 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.
- [3] H. S. Witsenhausen, "On source networks with minimal breakdown degradation," *Bell Syst. Tech. J.*, vol. 59, no. 6, pp. 1083–1087, July–Aug. 1980.
- [4] J. K. Wolf, A. D. Wyner, and J. Ziv, "Source coding for multiple descriptions," *Bell Syst. Tech. J.*, vol. 59, no. 8, pp. 1417–1426, Oct. 1980.
- [5] 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.
- [6] Z. Zhang and T. Berger, "The channel splitting problem and modulo-PCM coding," Bell Labs Memo for Record (not archived), Oct. 1979.
- [7] H. S. Witsenhausen, "Minimizing the worst-case distortion in channel splitting," *Bell Syst. Tech. J.*, vol. 60, no. 8, pp. 1979–1983, Oct. 1981.
- [8] Z. Zhang and T. Berger, "New results in binary multiple descriptions," *IEEE Trans. Inform. Theory*, vol. IT-33, pp. 502–521, July 1987.
- [9] A. Ingle and V. A. Vaishampayan, "DPCM system design for diversity systems with applications to packetized speech," *IEEE Trans. Speech Audio Processing*, vol. 3, pp. 48–57, Jan. 1995.
- [10] Y. Wang, M. T. Orchard, and A. R. Reibman, "Multiple description image coding for noisy channels by pairing transform coefficients," in *Proc. IEEE Workshop Multimedia Signal Processing*, Princeton, NJ, June 1997, pp. 419–424.
- [11] —, "Optimal pairwise correlating transforms for multiple description coding," in *Proc. IEEE Int. Conf. Image Processing*, Chicago, IL, Oct. 1998.
- [12] S. D. Servetto, K. Ramchandran, V. Vaishampayan, and K. Nahrstedt, "Multiple-description wavelet based image coding," in *Proc. IEEE Int. Conf. Image Processing*, Chicago, IL, Oct. 1998.
- [13] V. K. Goyal and J. Kovačević, "Optimal multiple description transform coding of Gaussian vectors," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar.–Apr. 1998, pp. 388–397.
- [14] V. K. Goyal, J. Kovačević, and M. Vetterli, "Multiple description transform coding: Robustness to erasures using tight frame expansions," in *Proc. IEEE Int. Symp. Information Theory*, Cambridge, MA, Aug. 1998, p. 408.
- [15] V. K. Goyal, J. Kovačević, R. Aream, and M. Vetterli, "Multiple description transform coding of images," in *Proc. IEEE Int. Conf. Image Processing*, vol. 1, Chicago, IL, Oct. 1998, pp. 674–678.

²¹The case $\varsigma_4^2 a_{11} a_{21} + \varsigma_2^2 a_{12} a_{22} = 0$ should be handled separately. In this situation, one can achieve the desired correlation structure with $\tilde{W}_1 = I$.

- [16] V. K. Goyal, J. Kovačević, and M. Vetterli, "Quantized frame expansions as source-channel codes for erasure channels," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 1999, pp. 326–335.
- [17] V. K. Goyal, J. Kovačević, and J. A. Kelner, "Quantized frame expansions with erasures," *Appl. Comput. Harm. Anal.*, vol. 10, no. 3, pp. 203–233, May 2001.
- [18] J. Walrand, *Communication Networks: A First Course*, 2nd ed. Boston, MA: McGraw-Hill, 1998.
- [19] S. Deering and R. Hinden. (1995, Dec.) Internet Protocol, version 6 (IPv6) specification. Network Working Group Request for Comments 1883. [Online]. Available: <ftp://ftp.isi.edu/in-notes/rfc1883.txt>
- [20] R. Han, P. Bhagwat, R. LaMaire, T. Mummert, V. Perret, and J. Rubas, "Dynamic adaptation in an image transcoding proxy for mobile web browsing," *IEEE Pers. Commun. Mag.*, vol. 5, pp. 8–17, Dec. 1998.
- [21] R. Ahlswede, "The rate distortion region for multiple descriptions without excess rate," *IEEE Trans. Inform. Theory*, vol. IT-31, pp. 721–726, Nov. 1985.
- [22] 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.
- [23] V. Koshlev, "Multilevel source coding and data-transmission theorem," in *Proc. VII All-Union Conf. Theory of Coding and Data Transm.* Vilnius, USSR, 1978, pt. 1, pp. 85–92.
- [24] —, "Hierarchical coding of discrete sources," *Probl. Pered. Inform.*, vol. 16, no. 3, pp. 31–49, 1980.
- [25] —, "An evaluation of the average distortion for discrete scheme of sequential approximation," *Probl. Pered. Inform.*, vol. 17, no. 3, pp. 20–33, 1981.
- [26] W. H. R. Equitz and T. M. Cover, "Successive refinement of information," *IEEE Trans. Inform. Theory*, vol. 37, pp. 269–275, Mar. 1991. See also [58].
- [27] B. Rimoldi, "Successive refinement of information: Characterization of the achievable rates," *IEEE Trans. Inform. Theory*, vol. 40, pp. 253–259, Jan. 1994.
- [28] P. G. Neumann, "Efficient error-limiting variable-length codes," *IEEE Trans. Inform. Theory*, vol. IT-8, pp. 292–304, July 1962.
- [29] T. J. Ferguson and J. H. Rabinowitz, "Self-synchronizing Huffman codes," *IEEE Trans. Inform. Theory*, vol. IT-30, pp. 687–693, July 1984.
- [30] N. Farvardin, "A study of vector quantization for noisy channels," *IEEE Trans. Inform. Theory*, vol. 36, pp. 799–809, July 1990.
- [31] G. Buch, F. Burkert, J. Hagenauer, and B. Kukla, "To compress or not to compress?," in *Proc. IEEE GLOBECOM*, London, U.K., Nov. 1996, pp. 196–203.
- [32] B. Hochwald and K. Zeger, "Tradeoff between source and channel coding," *IEEE Trans. Inform. Theory*, vol. 43, pp. 1412–1424, Sept. 1997.
- [33] V. K. Goyal, *Single and Multiple Description Transform Coding with Bases and Frames*. Philadelphia, PA: Soc. Industr. Appl. Math., 2001.
- [34] R. Zamir, "Gaussian codes and Shannon bounds for multiple descriptions," *IEEE Trans. Inform. Theory*, vol. 45, pp. 2629–2635, Nov. 1999. See [59] for generalizations to locally quadratic distortion measures.
- [35] C. E. Shannon, "Coding theorems for a discrete source with a fidelity criterion," *IRE Int. Conv. Rec.*, pt. 4, vol. 7, pp. 142–163, 1959.
- [36] T. Berger, *Rate Distortion Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1971.
- [37] V. A. Vaishampayan, "Design of multiple description scalar quantizers," *IEEE Trans. Inform. Theory*, vol. 39, pp. 821–834, May 1993.
- [38] D. O. Reudink, "The channel splitting problem with interpolative coders," Bell Labs., Tech. Rep. TM80-134-1, Oct. 1980.
- [39] V. A. Vaishampayan and J. Domaszewicz, "Design of entropy-constrained multiple-description scalar quantizers," *IEEE Trans. Inform. Theory*, vol. 40, pp. 245–250, Jan. 1994.
- [40] V. A. Vaishampayan and J.-C. Batllo, "Asymptotic analysis of multiple description quantizers," *IEEE Trans. Inform. Theory*, vol. 44, pp. 278–284, Jan. 1998.
- [41] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Trans. Inform. Theory*, vol. 44, pp. 2325–2383, Oct. 1998.
- [42] S.-M. Yang and V. A. Vaishampayan, "Low-delay communication for Rayleigh fading channels: An application of the multiple description quantizer," *IEEE Trans. Commun.*, vol. 43, no. 11, pp. 2771–2783, Nov. 1995.
- [43] J.-C. Batllo and V. A. Vaishampayan, "Asymptotic performance of multiple description transform codes," *IEEE Trans. Inform. Theory*, vol. 43, pp. 703–707, Mar. 1997.
- [44] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices and Groups*. New York: Springer-Verlag, 1988, vol. 290, Grundlehren der mathematischen Wissenschaften.
- [45] V. K. Goyal, J. Zhuang, and M. Vetterli, "Transform coding with backward adaptive updates," *IEEE Trans. Inform. Theory*, vol. 46, pp. 1623–1633, July 2000.
- [46] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [47] V. K. Goyal and M. Vetterli, "Manipulating rates, complexity, and error-resilience with discrete transforms," in *Conf. Rec. 32nd Asilomar Conf. Signals, Systems, and Computers*, vol. 1, Pacific Grove, CA, Nov. 1998, pp. 457–461.
- [48] V. K. Goyal, "Transform coding with integer-to-integer transforms," *IEEE Trans. Inform. Theory*, vol. 46, pp. 465–473, Mar. 2000.
- [49] —, "Beyond traditional transform coding," Ph.D. dissertation, Univ. Calif., Berkeley, 1998. Published as Univ. California, Berkeley, Electron. Res. Lab. Memo. UCB/ERL M99/2, Jan. 1999.
- [50] D. Sinha, J. D. Johnston, S. Dorward, and S. Quackenbush, "The perceptual audio coder (PAC)," in *The Digital Signal Processing Handbook*. Piscataway, NJ: IEEE Press, 1998, ch. 42, pp. 42.1–42.18.
- [51] R. Areal, J. Kovačević, and V. K. Goyal, "Multiple description perceptual audio coding with correlating transforms," *IEEE Trans. Speech Audio Processing*, vol. 8, pp. 140–145, Mar. 2000.
- [52] M. Effros and A. Goldsmith, "Capacity definitions and coding strategies for general channels with receiver side information," in *Proc. IEEE Int. Symp. Information Theory*, Cambridge, MA, Aug. 1998, p. 39.
- [53] J. Hong, "Discrete Fourier, Hartley, and Cosine transforms in signal processing," Ph.D. dissertation, Columbia Univ., New York, 1993.
- [54] A. Zandi, J. D. Allen, E. L. Schwartz, and M. Boliek, "CREW: Compression with reversible embedded wavelets," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 1995, pp. 212–221.
- [55] R. Calderbank, I. Daubechies, W. Sweldens, and B.-L. Yeo, "Wavelet transforms that map integers to integers," *Appl. Comput. Harm. Anal.*, vol. 5, no. 3, pp. 332–369, July 1998.
- [56] G. Strang, *Linear Algebra and Its Applications*, 3rd ed. San Diego, CA: Harcourt Brace Jovanovich, 1988.
- [57] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 1985. Reprinted with corrections 1987.
- [58] W. H. R. Equitz and T. M. Cover, "Addendum to 'Successive refinement of information'," *IEEE Trans. Inform. Theory*, vol. 39, pp. 1465–1466, July 1993.
- [59] T. Linder, R. Zamir, and K. Zeger, "The multiple description rate region for high resolution source coding," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar.–Apr. 1998, pp. 149–158.
- [60] N. J. A. Sloane and A. D. Wyner, Eds., *Claude Elwood Shannon: Collected Works*. Piscataway, NJ: IEEE Press, 1993.