

Multiple Description Wavelet Based Image Coding

Sergio D. Servetto, *Member, IEEE*, Kannan Ramchandran, *Member, IEEE*, Vinay A. Vaishampayan, and Klara Nahrstedt, *Member, IEEE*

Abstract—We consider the problem of coding images for transmission over error-prone channels. The impairments we target are transient channel shutdowns, as would occur in a packet network when a packet is lost, or in a wireless system during a deep fade: when data is delivered it is assumed to be error-free, but some of the data may never reach the receiver. The proposed algorithms are based on a combination of multiple description scalar quantizers with techniques successfully applied to the construction of some of the most efficient subband coders. A given image is encoded into multiple *independent* packets of roughly equal length. When packets are lost, the quality of the approximation computed at the receiver depends only on the number of packets received, but does not depend on exactly which packets are actually received. When compared with previously reported results on the performance of robust image coders based on multiple descriptions, on standard test images, our coders attain similar PSNR values using typically about 50–60% of the bit rate required by these other state-of-the-art coders, while at the same time providing significantly more freedom in the mechanism for allocation of redundancy among descriptions.

Index Terms—Error resilience, image coding, joint source/channel coding, multiple description source coding.

I. INTRODUCTION

A. Need for Error Resilient Data Compression Algorithms

THE HIGH performance achieved by state-of-the-art image/video coding algorithms, combined with the sustained growth of the Internet and cellular networks over the last few years, has resulted in the emergence of a new family of communication services which involve the delivery of image/video data over error-prone channels. A most common form for these errors is that of transient channel shutdowns. For example, a typical network transmission might involve data

hopping over links of different capacities, and upon encounter of a low capacity link, packets must be dropped. Another common problem is network congestion: if upon arrival of a packet at a switching node the local buffer is full, that packet has to be dropped. In the case of wireless communication, during a deep fade, the probability of decoding error at the receiver becomes very high, and most of the received data is useless, resulting in discarding data frames received during the fade. In these (and other) transmission scenarios, it is reasonable to assume that there are periods of time during which the error correcting codes used at the lower layers of the system architecture ensure error-free delivery of the transmitted data, but that occasionally some data will be lost.

Besides the use of error correcting codes, many communication systems and storage devices provide *diversity* to combat possible channel impairments. For example, in the context of ATM networks, multipath connections are desirable because even if an intermediate switching node gets temporarily congested (thus resulting in a transient high rate of cell losses), it is less likely that such an event will occur on nodes along all paths simultaneously [25]. In the case of redundant arrays of inexpensive disks (RAID's), data is stored in multiple disks to provide resilience against individual disk failures [30]. In wireless systems, multipath and Doppler diversity are used to combat the fading effect [23]. In channels that undergo burst errors, interleaving is used to create an illusion of diversity, in which a number of lower capacity nonbursty channels are seen by the source.

Our main goal in this paper is to develop data compression algorithms capable of producing representations for images which are robust to the presence of errors of this nature. Not specifically for images, but in the context of coding an arbitrary information source, this problem has been thoroughly studied already in the Information Theory community, and is known as the problem of multiple descriptions (MD's). In this paper, we apply MD codes to the construction of our robust image coding algorithms. The main feature of our coders is that they do not produce a single bit stream, but instead multiple *equally important and independent* bit streams; this way, the quality of the reconstructed images depends only on the number of packets that arrive at the decoder, and not specifically on which of them did arrive. This idea is a generalization of the well-understood concept of successive refinements of an information source, where (say) an image can also be broken into (say) two descriptions p_1 , p_2 [9]. The difference is that in this case, description p_2 is a strict refinement of description p_1 , and it assumes knowledge of p_1 ; therefore, if p_1 is lost, p_2 cannot be used for decoding even if it arrives intact. With the framework based on multiple descriptions of this work, with *either* one of p_1 or p_2 we are able to

Manuscript received December 9, 1998; revised October 7, 1999. This work was supported in part by the Ray Ozzie Fellowship of the Department of Computer Science, University of Illinois at Urbana-Champaign, and by the National Science Foundation under Awards NSF MIP-97-03181 (CAREER) and CCR-96-23867 (CAREER), and Grant 9624396. Parts of this work were done while S. D. Servetto was a Summer Intern at AT&T Labs, and were presented at ICIP'98. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Yoshitaka Hashimoto.

S. D. Servetto is with the Laboratoire de Communications Audiovisuelles, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland (e-mail: servetto@lcvsun1.epfl.ch).

K. Ramchandran is with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA (e-mail: kannanr@eecs.berkeley.edu).

V. A. Vaishampayan is with the Information Sciences Research Center, AT&T Labs—Research, Florham Park, NJ 07932 USA (e-mail: vinay@research.att.com).

K. Nahrstedt is with the Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: klara@cs.uiuc.edu).

Publisher Item Identifier S 1057-7149(00)03863-X.

give a bound on the distortion in the reconstructed signal, while if both are received we can attain another bound that is lower.

B. Related Work

Multiple description codes have been used before in the context of robust image coding. The first MD image coder that we are aware of was proposed by Vaishampayan, consisting of an extension of a JPEG coder using MD scalar quantizers [33]. Wang *et al.* proposed another MD extension of a JPEG coder, using a class of pairwise correlating transforms to create the MD's [35]. Goyal *et al.* apply the more general transforms they proposed in [11], [13] to the development of another MD extension of JPEG [12]. Chung and Wang use lapped orthogonal transforms to construct yet another MD extension of JPEG [7].

Using more state-of-the-art image coding techniques, we proposed a MD wavelet coder, of which this paper is an extended version [27]. Along similar lines, Srinivasan and Chellappa propose a MD extension of a very high performance subband coder [15], using MD scalar quantizers [31]. Jiang and Ortega propose an MD extension to the SPIHT coder of Said and Pearlman [22], by separating Zerotrees into polyphase components [14]. Rogers and Cosman propose to rearrange bits at the output of one configuration of the SPIHT coder, in a way such that the loss of one packet results in an error that does not propagate beyond the image region contained in that packet [21]. Mohr *et al.* propose the use of error correcting codes of different strengths applied to different portions of a progressive bit stream such as that generated by the SPIHT coder [17].

The problem of breaking an image into pieces and then being able to reconstruct it from an arbitrary subset of these pieces is analogous to the problem of optical holography. For digital images, Bruckstein *et al.* first proposed two techniques for the computation of digital holograms: one of them is based on a pseudo-random pattern for sampling the image, in which arbitrary portions of the pattern contain roughly uniformly spaced image pixels, with a density proportional to the length of the pattern; the other is based on a frequency domain watermarking technique [5]. Ng and Kovčević propose an alternative solution to the same problem, based on the use of critically sampled filter banks [18]. Although our problem and the problem of digital holography are conceptually similar, none of the solutions proposed thus far for this problem deal with the issue of *efficiency* of the representation: they all work on raw image data only, and some do actually increase the number of bits in the representation.

C. Main Contributions and Paper Organization

The main contribution presented this work is the design and optimization of new and high-performance wavelet-based MD image coding algorithms with a number of remarkable features as follows.

- On standard test images, our coders typically attain equal PSNR values when compared to state-of-the-art MD coders, but using approx. only 50–60% of their required bit rate. Even when compared against state-of-the-art coders which are *not* robust to errors, our coders perform competitively too. And furthermore, these excellent

results are attained at very low computational complexity: our coders require only a single pass over the array of coefficients, in raster order.

- Our coders provide many more degrees of freedom in specifying the amount of redundancy to be allocated to given images, when compared to other state-of-the-art coders.

Signal processing techniques involving carefully designed transforms have been studied recently, as a means of constructing multiple descriptions of a source [11], [13], [19], [35]. While these techniques are known to have undesirable asymptotic properties in the high bit rate regime (e.g., the single-channel distortions do not go to zero in general as bit rate increases), they are generally believed to perform well at low rates. On the other hand, design techniques for MD scalar quantizers are based on making certain high rate approximations, thus raising the issue of whether these quantizers can be used in typical medium to low rate regimes. Further motivation for the use of MD scalar quantizers in the specific case of images is provided by the fact that, asymptotically at high rates, good transforms to use in a single description coder are also good transforms to use in a MD coder [3]. In this work we show the feasibility of MD scalar quantizers, for a “real” source (i.e., images), and in the medium to low bit rate regime, thus providing strong support for the usefulness of these quantizers even when the high rate assumptions made in their design are relaxed.

The rest of this paper is organized as follows. In Section II, we present the design and parameter optimization of a MD image coding algorithm. In Section III, we generalize the previous construction from the case of two descriptions to the case of an arbitrary number of descriptions. In Section IV we present thorough experimental results, and finally, in Section V, we present conclusions and discuss future research directions.

II. IMAGE CODING INTO MUTUALLY REFINABLE PACKETS

In this section, we consider the problem of encoding a still image into two mutually refinable packets.

A. Motivation and Plan Outline

1) *Review of Single Description Coding of Images:* There is a de-facto standard architecture for single description image coding algorithms, consisting of first applying a linear decorrelating transform to the input image, then performing scalar quantization of the transform coefficients, and finally performing entropy coding of the quantized bins. In this standard architecture, wavelets have been used as a linear decorrelating transform in order to attain some of the best known results in terms of coding efficiency.

Some of the most successful wavelet coders derive their high coding performance from their ability to identify sets of coefficients with different statistics within image subbands, and then coding each of these sets with respect to an appropriate statistical model [6], [15], [16], [22], [28], [26], [29], [36]. Since these sets typically are image dependent, this information is not known a priori, and therefore must be somehow conveyed to the decoder. This can be done either *explicitly* [15], [22], [28], [29],




Fig. 1. Architecture of the multiple description transform coder.

[36] or *implicitly* [16], [16], [26]. In the explicit case, “map” bits describing these sets are included in the bit stream; these are bits that do not convey information about the value of subband coefficients, but instead configure the decoder appropriately to decode such values. In the implicit case, the information regarding sets of coefficients is deduced only from data causally available at the decoder, so that no explicit map bits are required.

2) *Issues in the Generation of Two Descriptions:* To perform MD coding of images, we use the standard MD transform coding architecture proposed in [3], consisting of replacing the single description quantizer with a MD quantizer, and then performing entropy coding of the two output quantized streams independently. This architecture is illustrated in Fig. 1.

We decided to build our MD image coder based on this architecture mainly because of two reasons. One is that for coding Gaussian sources with memory using the MD transform coder, the optimal transform to use for decorrelation is the Karhunen–Loeve transform (KLT) [3]; in the case of single descriptions, this is often the justification used for the de-facto architecture, and for the use of decorrelating transforms that approximate the KLT. More practically however, the state-of-the-art on the design of single description coders based on that architecture is very mature, leading one to be optimistic about the performance attainable by extensions of these good coders to the MD case.

The issue of extending good single description image coders to support multiple descriptions using MD scalar quantizers [32] is not altogether straightforward: not any good classical coder will result in good performance under the MD constraints, just by inserting an index assignment and an extra entropy coder in the loop.

- In the context of MD coding, it must be possible to decode each description independently of whether other descriptions are available at the decoder or not. As a result, if the coding technique employed makes use of explicit map information, enough map bits must be spent within each description to ensure that each one of them can be decoded independently of the others. But map information is inherently different from basic data in that, in general, it does not admit approximate representations. While it makes perfect sense to talk about the accuracy to which a given wavelet coefficient is described, what is meant by “an approximate representation of map information” that could be used to build multiple descriptions is, at least, both unclear and dependent on specific coding frameworks. When this information is replicated in both descriptions and both arrive to the decoder, half of the map bits carry no information. This argument suggests that backward adaptive coders (such as [6], [16], [26]) are to be preferred over




Fig. 2. Basic architecture of a diversity system using a MD scalar quantizer.

coders that explicitly transmit map data (such as [15], [22], [28], [29], [36]).

- To further complicate things however, not any backward adaptive coder will perform well either. Some of these good coders derive part of their gains not only from their ability to adapt the entropy coder to local changes of statistics in image subbands, but also from adapting the quantizer applied to each coefficient. Now, in the context of MD coding, whereas descriptions are encoded independently of each other and must be decodable on their own, the requirement that when both arrive at the decoder they must be combined to produce a higher quality approximation imposes *synchronization* constraints among descriptions. Therefore, if adaptation of the MD quantizers is used, and a mismatch occurs in the choice of quantizers (due to the independent adaptation for each description), a decoding failure occurs when both descriptions arrive, since the information they carry is inconsistent; so, either no adaptation is used, or map data has to be included in each description to maintain synchronization, thus bringing us back to the problem pointed out above. This is not the case when entropy coders instead are adapted on a local basis, since in this case all that changes is the length of the codewords used, but not the encoded information itself.

B. Design of a Multiple Description Image Coder

MD systems dealing with an arbitrary number of descriptions involve a large number of parameters (rates for each description, and distortions for all subsets of descriptions), and therefore both optimization and performance comparisons become fairly involved. However, for two balanced descriptions, this is not the case: only one rate (the number of bits spent on each description) and two distortions (the distortion when either one of the channels fail or when both work) needs be considered. Next, we present the design of a high-performance MD coder that is amenable to performance analysis and optimization.

1) *Review of Multiple Description Scalar Quantization:* The first “practical”¹ results on multiple descriptions were presented by Vaishampayan, where a simple procedure is given to design MD scalar quantizers with some remarkable asymptotic properties [32], [34]. Fig. 2 illustrates the role of a MD quantizer in a two-channel diversity system.

A MD quantizer consists of two main components: a *scalar quantizer* (that maps continuous-valued random variables to points in a countable set), and an *index assignment* (that splits

¹Practical in the sense that previous work had focused on the characterization of the MD rate/distortion region under different conditions, but not on the effective construction of good MD codes.




Fig. 3. Two index assignments: (a) staggered quantization cells; (b) higher spread cells. The idea is that bins of the scalar quantizer are placed in a matrix, and then quantizer indices corresponding to row and column entries are sent over each channel. If both descriptions are available, the original quantization bin can be recovered; if not, the original quantization bin is known to be one of those in the received row/column.

the information about each sample into two complementary and possibly redundant descriptions of the same sample). An index assignment is an injection $I : \mathcal{N} \rightarrow \mathcal{N} \times \mathcal{N}$ (\mathcal{N} is the set of natural numbers). When the scalar quantizer maps the source to a finite number of points (say, n), the map I can be thought of as a matrix of size $n \times n$, in which only n locations are occupied. Clearly, there exist $n!C(n^2, n) = O(n^{3n})$ distinct such mappings. Two example index assignments are shown in Fig. 3.

The problem of designing good index assignments is thoroughly studied in [32], using Ozarow's characterization of the MD rate/distortion region for the Gaussian source as a guiding principle [10], [20]. Under the assumption of equal and high rates for both descriptions, and for a squared error distortion measure, a construction of a large class of index assignments is presented in [32], for which the exponential rate of decay of the MSE is exactly that predicted by Ozarow's result.²

2) Reconstruction Levels of the Single-Channel Quantizers: Image coders based on uniform quantizers typically perform inverse quantization by mapping bins to the midpoint of their cell. However, taking the same approach for inverting single-channel quantizers leads to remarkably poor performance if its cells are large (i.e., in the low excess rate regime). To overcome this problem it is necessary to reconstruct not to the midpoints of these cells, but to their centroids instead. However, in order to compute centroids, we need a model for the distribution of coefficients within each subband.

The statistical properties of subband data have been thoroughly studied before in the context of image coding, and different models have been proposed [15], [16], [24]. However, a feature common to all these models is that subbands can be assumed to be drawn from zero mean, unimodal, symmetric distributions. For this class of distributions we design a simple inverse quantization rule, whose operation is illustrated in Fig. 4.

Using this rule for inverse quantization instead of the midpoint rule resulted in an increase of 6–7 dB for the single-channel reconstructions in the low excess rate regime (i.e., when single-channel cells are composed of the union of a large number of two-channel cells), and a loss of 0.4–0.5 dB in the high excess rate regime.

²However, the sub-exponential terms in the MSE are *not* the same: the performance of the MD scalar quantizer is bounded away from the Gaussian MD rate/distortion bound, as one would intuitively expect.




Fig. 4. Inversion of single-channel quantizers. A quantization cell for the single-channel quantizers consists of a (possibly disjoint) union of intervals, representing cells from the scalar two-channel quantizer. To dequantize, we select the midpoint of the most likely interval (i.e., the interval closest to the origin).

3) Encoding: At a high level, the MD image coding algorithm works as follows. First, a given input image is decomposed into subbands, and then a uniform scalar quantizer (with a possibly wider zero bin) is applied to each of the subband coefficients, thus producing a quantized field. Two descriptions of this field are then created, by mapping each quantized coefficient to a pair of numbers, using the index assignment component of a MD quantizer (as in Fig. 3).

The choice of index assignment is the mechanism by which performance of the two-channel decoder is traded off for performance of the single-channel decoder, by means of controlling the amount of redundancy in a MD code. Now, it is well known that different subbands carry unequal weight in terms of overall signal energy content: whereas the loss of the pure lowpass projection is likely to render the entire reconstruction worthless, the loss of substantial portions of the high frequency subbands is much less significant. Therefore, it seems intuitively clear that there are gains to be had by allowing some degree of adaptation when choosing index assignments, instead of using a fixed one for all subbands. Our algorithm provides this by choosing one index assignment per subband, and explicitly encoding this choice as map bits into both descriptions: since the cost of the quantizer parameters are amortized over the entire subband, the penalty paid by doing so is more than compensated by the gains due to adaptation. An experiment to quantify these gains is presented in Section IV.

Next, subband descriptions are entropy coded, independently of each other. The coefficients in each description are dequantized using the single channel decoder, and a local variance estimate is formed based on causally available data (this is so that the estimate depends only on data available to the decoder). Based on comparing this estimate against a fixed threshold, each coefficient is classified into one of two possible classes; and each class is entropy coded separately, with respect to its own probability model. This classification step attempts to separate regions of locally large and locally small variances, a basic principle used in most state-of-the-art image coders. Table I presents a pseudocode description of our proposed algorithm.

4) Decoding: In order to decode, two steps have to be performed. First, individual descriptions have to be entropy decoded. If either one of the descriptions is lost, then the available description is dequantized using the single channel inverse quantizer (actually, this already happened in the entropy decoding loop), and the wavelet is inverted, thus yielding the single channel image estimate. If both descriptions arrive, then prior to inverse quantization the two descriptions have to be recombined

TABLE I
PSEUDOCODE DESCRIPTION FOR THE MD ENCODER

Input: image, scalar quantizer, array of index assignments (one per subband).
Output: two compressed bit streams.
Algorithm:
1. Take a wavelet transform of the input image, and apply the input quantizer to each coefficient.
2. Create two descriptions of each subband, by applying to each coefficient its subband-dependent index assignment.
3. For each description:
For each quantized coefficient \hat{c}_{ij} in each subband, start with two uniform density estimates $h^{(0)}$ and $h^{(1)}$, and then:
(a) Use the single-channel inverse quantizer to dequantize coefficients within a causal neighborhood.
(b) Compute a binary context word w_{ij} : 1 if a local variance estimate (computed based on the dequantized data) exceeds a fixed threshold; 0 otherwise.
(c) Use the density $h^{(w_{ij})}$ as the probability table in an arithmetic coder, to entropy code \hat{c}_{ij} .
(d) Update the corresponding density estimate $h^{(w_{ij})}$, based on the new observed value \hat{c}_{ij} .

by inverting the index assignment. Pseudocode for this is shown in Table II.

C. Optimization of System Parameters

Next, we present an algorithm to find optimal parameters of the proposed MD coding algorithm.

1) *Motivation:* In the single description case, available communications resources typically place constraints on a single variable affecting the quality of the reconstructed images: the number of bits available to encode them. In the MD case such a constraint is also meaningful: it is descriptions that actually get transmitted over communication channels, and therefore the number of bits used to encode each description cannot exceed the capacity of these channels. However, in the MD case, there is one extra independent variable that affects the quality of the reconstructed images: channel failures.

To illustrate this point, consider a setup in which at least one of the two channels is very likely to fail; in that case the reconstructed image quality is, most of the time, that achievable using only single-channel decoders. On the other hand, if channels fail rarely, then most of the time the reconstructed image quality is equal to that achievable using the (better) two-channel decoder. As a result, both the capacity of each channel *and* the frequency with which each of them fail affect image quality, and hence must be taken into account when choosing what parameters (i.e., what scalar quantizer and index assignments) to use for coding a particular image.

In our formulation, for a fixed number of bits for each description, we allow the user to specify a minimum quality for the images reconstructed using single-channel decoders. If the error rate of the channel is high, the user should request high quality single-channel reconstructions, at the expense of

TABLE II
PSEUDOCODE DESCRIPTION OF TWO-CHANNEL DECODER

Input: Two compressed bit streams.
Output: A reconstructed image.
Algorithm:
1. Entropy decode each description, to obtain two quantized fields $\hat{c}^{(1)}$ and $\hat{c}^{(2)}$ of subband coefficients.
2. Use the pair $\hat{c}_{ij}^{(1)}$ and $\hat{c}_{ij}^{(2)}$ to recover the central bin \hat{c}_{ij} , by inverting the index assignment corresponding to that subband.
3. Dequantize the field of \hat{c}_{ij} 's, using the uniform scalar quantizer; then apply the inverse wavelet transform, to produce the output reconstructed image.

the quality achievable by the two-channel reconstructions. We formalize these concepts next.

2) *Formulation:* We formulate now a discrete optimization problem, yielding optimal performance in the presence of application-specified constraints. Consider the following definitions.

- Let δ denote an arbitrary scalar quantizer, and let I denote an arbitrary index assignment.
- Let f denote an image, and $(f_m)_{m:1 \dots M}$ its decomposition into M wavelet subbands.
- Let $\mathcal{D}_0(f, \delta, I_S)$ denote the mean squared error (MSE) in the central reconstruction of f using the given central quantizer δ and the set of index assignments

$I_{\mathbf{s}}, \mathbf{s} = (s_1 \cdots s_M)$; i.e., one index assignment per image subband ($s_i \geq 2$ is the number of diagonals in the index assignment applied to the i th subband). Let $\mathcal{D}_1(f, \delta, I_{\mathbf{s}}), \mathcal{D}_2(f, \delta, I_{\mathbf{s}})$ denote the single-channel MSE's.

- Let $\mathcal{R}_1(f, \delta, I_{\mathbf{s}}), \mathcal{R}_2(f, \delta, I_{\mathbf{s}})$ denote the number of bits required to encode each description of f using the given central quantizer and index assignments.

Then, our goal is to find a pair $(\delta, I_{\mathbf{s}})$ to solve:

$$\min_{\delta, I_{\mathbf{s}}} \mathcal{D}_0(f, \delta, I_{\mathbf{s}}) \quad (1)$$

subject to

$$\mathcal{R}_1(f, \delta, I_{\mathbf{s}}) = \mathcal{R}_2(f, \delta, I_{\mathbf{s}}) \leq R_{\text{budget}} \quad (2)$$

$$\mathcal{D}_1(f, \delta, I_{\mathbf{s}}) = \mathcal{D}_2(f, \delta, I_{\mathbf{s}}) \leq D_{\text{budget}} \quad (3)$$

where the user-specified parameters are R_{budget} (the available bit rate to encode each description), and D_{budget} (the maximum distortion acceptable for single-channel reconstructions).

3) *Approximation of the Optimal Solution:* Motivated by asymptotic properties of MD quantizers at high rates, we develop an efficient algorithm to compute an *approximate* solution of the problem defined by (1)–(3). We do not attempt to find the optimal solution because, although we do not have a proof yet, we conjecture that our problem is NP-Hard. If this were true it would imply that, with high likelihood, the most efficient algorithm to solve our problem would not be substantially better than just trying out repeatedly each one of the $O(n^{3n})$ possible index assignments, a clearly intractable task.

Define the *spread* of an index assignment to be the number of diagonals occupied by the central bins in the matrix representation; for example, the spread of the assignments in Fig. 3(a) and (b) is 2 and 3, respectively. To derive an efficient algorithm, we restrict the class of index assignments we consider to be of the form shown in Fig. 3, i.e., to the set of bandwidth-constrained matrices. We are motivated to consider this particular subset because, at high rates, the performance of a MD scalar quantizer is entirely determined by the spread [32].

Another reason why we need to resort to an approximation is because, at practical middle/low bit rates, exact equality for $\mathcal{R}_1, \mathcal{R}_2$, and for $\mathcal{D}_1, \mathcal{D}_2$, cannot be attained in general. Using the index assignments designed in [32], we have found the difference in the length of the descriptions to be less than 0.1% (e.g., two descriptions of length 16 380 and 16 341 bytes), and the difference in PSNR of each description to be less than 0.02 dB. Since we regard these approximations to be good enough, we do not further pursue this issue. However, we feel it is important to emphasize that exact equality, as prescribed by the constraints in the optimization problem, cannot be achieved in general.

4) *Computation of Optimal Index Assignments:* Consider first the following simpler problem: for a given quantizer step size δ , find a vector of index assignments $I_{\mathbf{s}}^*(\delta)$, such that

$$I_{\mathbf{s}}^*(\delta) = \arg \min_{I_{\mathbf{s}}} \mathcal{D}_1(f, \delta, I_{\mathbf{s}})$$




Fig. 5. Trellis for the dynamic programming solution: in the horizontal axis, the different subbands are shown; in the vertical axis, the possible index assignments for each subband are shown. The cost $C(k, i, j)$ of an arrow is the single-channel distortion that results from using an index assignment with i diagonals on subband k . DP is used to find a path of minimum cost in this trellis, corresponding to an optimal choice of index assignments.

(this problem will be a key element of the solution to the more general problem, presented below).

Now, when the choice of each index assignment I_{s_i} in the optimal vector $I_{\mathbf{s}}$ can be made independently of all other choices, dynamic programming (DP) can find the optimal $I_{\mathbf{s}}$ efficiently [8]; this assumption clearly holds for our coder, since subbands are encoded independently of each other. We solve this problem efficiently by searching for a minimum cost path satisfying the given constraints, in the trellis shown in Fig. 5.

Further reductions in complexity can be obtained by using Lagrangian or hybrid DP-Lagrangian methods, but are not considered here [37].

5) *Computation of the Approximate Solution:* Using the algorithm presented above, here we present an algorithm to find parameters solving (1)–(3), when restricted to the class of bandwidth-constrained matrices.

The basic idea is to take advantage of the monotonicity of both \mathcal{R} and \mathcal{D} as a function of δ . We search for the smallest quantizer δ such that the rate constraint R_{budget} is met (with equality), because this guarantees that \mathcal{D}_0 will be minimized, since \mathcal{D}_0 is independent of $I_{\mathbf{s}}$. Then, for that particular value of δ we verify the constraint on \mathcal{D}_1 : if it is satisfied, we are done. But if it is not, then the constraints are inconsistent: increasing δ can only result in larger values of \mathcal{D}_1 , and δ cannot be decreased without violating the constraint on \mathcal{R} . At each step, for each candidate δ considered, the best vector of index assignments $I_{\mathbf{s}}^*(\delta)$ is found using the algorithm presented above. A pseudocode description of the proposed algorithm is presented in Table III.

III. IMAGE CODING INTO ANY NUMBER OF MUTUALLY REFINABLE PACKETS

In this section, we consider the problem of compressing an image into not just two, but *any* number of mutually refinable packets. Our goal is to come up with an image representation based on multiple packets, in a way such that approximations to the original image can be obtained from *arbitrary* subsets of packets, and such that the quality of the approximation depends only on the size of the subset but *not* on which specific packets are selected. This problem is a generalization of the classical problem of multiple descriptions, to deal with more than two descriptions.

TABLE III
PSEUDOCODE DESCRIPTION OF THE OPTIMIZATION PROCEDURE

```

Input: image,  $R_{\text{budget}}$ ,  $D_{\text{budget}}$ , an interval  $[\delta_0, \delta_1]$  known to contain  $\delta^*$ .
Output: optimal parameters  $\delta^*$  and  $I_s^*$ , if they exist.
Algorithm:

1.  $\delta_L := \delta_0$ ;  $\delta_H := \delta_1$ ;

2. for  $n \geq 2$ :
    (a)  $\delta_n := \frac{1}{2}(\delta_L + \delta_H)$ .
    (b) Compute  $I_s^*(\delta_n) = \arg \min_{I_s} \mathcal{D}_1(f, \delta_n, I_s)$ .
    (c) if  $\frac{1}{2}(\mathcal{R}_1(f, \delta_n, I_s^*(\delta_n)) + \mathcal{R}_2(f, \delta_n, I_s^*(\delta_n))) < R_{\text{budget}}$ , then  $\delta_L := \delta_n$ ;
        else, if  $\frac{1}{2}(\mathcal{R}_1(f, \delta_n, I_s^*(\delta_n)) + \mathcal{R}_2(f, \delta_n, I_s^*(\delta_n))) > R_{\text{budget}}$ , then  $\delta_H := \delta_n$ ;
        else exit loop;

3. if  $\frac{1}{2}(\mathcal{D}_1(f, \delta_n, I_s^*(\delta_n)) + \mathcal{D}_2(f, \delta_n, I_s^*(\delta_n))) \leq D_{\text{budget}}$ , then return  $\delta^* := \delta_n$  and  $I_s^* := I_s^*(\delta_n)$ ;
    else the constraints cannot be satisfied.

```

A. Motivation and Plan Outline

1) *Robust Transmission Over IP Networks:* We are motivated to consider this more general version of the problem of multiple descriptions essentially because of constraints imposed by an application that is of great interest to us: robust transmission of compressed data over IP networks.

- Packets transmitted over IP networks are limited in size. For example, over an ethernet, the size of the largest packet that can be transmitted³ is 1500 bytes, and suppose we encode an image into two descriptions of length 3000 bytes each. Although applications are allowed to transmit packets longer than the MTU limit, the transport and/or IP protocols will fragment a large packet into smaller ones, transmit all of these small packets separately, and reassemble them into a large packet at the receiving end. However, should a single small packet get lost, even if all other small packets arrive these are discarded. To prevent this “domino” packet drop effect, one would rather create four descriptions of size 1500 bytes each, rather than two descriptions of size 3000 bytes each.
- IP networks do not support the notion of priority classes. All packets of the same length are *a priori* equally likely to be dropped. Hence, a coder capable of producing multiple packets of roughly equal length and equal importance is better suited for this application than classical coders based on multiresolution principles.

2) *Extensions to Multiple Packets:* There are significant differences between the standard problem of multiple descriptions (two packets), and the more general case we are interested in here (many packets). There are a number of possible approaches to designing these more general coders, and in this work we focus on approaches based on MD quantization ideas.

The case of two descriptions based on MD scalar quantizers is essentially reduced to two instances of the classical quantization problem. This is so because a MD quantizer is essentially a pair of quantizers such that, individually, these are good quantizers; but then they have the added constraint that when

considering the quantizer whose set of Voronoi cells is obtained as the pairwise intersection of the cells of the two quantizers, this new quantizer is also good. A natural mechanism to extend this construction to generate $N \geq 2$ mutually refinable packets is to design N good quantizers, in a way such that the quantizer that results from intersecting the Voronoi cells of any k out of the N quantizers is also good. This is the approach taken by Berger-Wolf and Reingold [4], using design criteria based on the same high-rate quantization principles applied in [32]. In that approach, each packet carries information about every single subband coefficient: N descriptions are created by quantizing each coefficient with the N quantizers, and putting the quantizer outputs into each description.

MD scalar quantizers have been found to perform extremely well in the low bit rate regime, even when the high-rate assumptions made in their design are clearly violated [27]. However, no such corroboration has been found yet when the number of descriptions is allowed to increase.

3) *Design Principles:* If the condition that each packet carries information about every single coefficient is relaxed, it becomes possible to explore interesting tradeoffs in the generation of more than two descriptions. To understand the nature of these tradeoffs, consider Fig. 6.

From the point of view of pure source coding, it is desirable to collect all coefficients within a Zerotree structure and encode them in a single packet: high compression efficiency will be attained by doing so. However, if the packet carrying that particular Zerotree is lost, the damage to the signal will be severe: observe in Fig. 6 how, since all the memory in the process is lost with the erased packet, no form of signal processing can be applied to the reconstructed image to recover the missing eye. Therefore, robustness requirements dictate that spatially clustered coefficients be spread among multiple packets and that they be coded independently of each other. Hence, we see that memory in the source can be used either to improve compression efficiency, or to combat channel impairments:⁴ this is the source of tradeoffs to explore when designing robust coders.

³This is known as maximum transmission unit (MTU).

⁴This is a manifestation of the tension between the goals of source and channel coding.




Fig. 6. To illustrate problems arising when attempting to spread subband coefficients across multiple packets; left: a typical Zero-tree, right: the reconstruction that results from completely losing a Zero-tree.

Based on these observations, these are the design principles that we use in the construction of our robust coder:

a) Decompose into spatio-temporal subbands using a classical transform: The idea in this case is to exploit the standard energy-compaction gains attained by classical transforms, by coding each subband separately. However, within subbands, an iid model is used so that the damage caused by a lost packet can be diffused over the entire image (rather than concentrating it in a localized spatial section).

b) Implement a subband-dependent unequal error protection scheme: On average, significantly more damage will be caused to a signal by the complete erasure of a low-frequency subband coefficient than by the loss of a high-frequency coefficient. Therefore, it seems natural to design a coder in which the number of lost packets it takes to erase all the information related to a coefficient be dependent on the importance of the subband to which that coefficient belongs.

B. Design of a Multiple Description Image Coder

Next, we present the design of a high-performance algorithm to compress images into any number of mutually refinable packets.

1) Encoding: Consider a family $\{\mathcal{C}_n\}_{n=1}^E$ of good erasure-resilient codes, indexed by a parameter which indicates the maximum number of erasures that the code can tolerate and still recover the message. The simplest example of such codes is the family of $(n, 1, n)$ repetition codes, with trivial encoding and decoding algorithms, but with poor asymptotic performance. With slightly more complexity, maximum-distance separable (MDS) codes such as Reed–Solomon codes, or “almost” MDS codes⁵ such as those of Alon and Luby provide more examples [2]. Using these codes, we define a mapping $\rho : \{1 \dots M\} \rightarrow$

$\{1 \dots E\}$ (M is the number of subbands in the wavelet expansion of the image, E is the number of codes in the family). The interpretation of this function ρ is that it specifies the number of lost packets required to completely erase all the information on the description of a subband coefficient: this is how the UEP property above discussed is implemented.

Finally, once appropriate codes are applied to each description of each subband, and given a target number N of independent packets desired, the components of each codeword are entropy coded and placed into each of the packets. The operation of the proposed encoding algorithm is illustrated in Fig. 7.

2) Decoding: The decoding steps are straightforward. First, each of the received packets are entropy decoded. Then, for each subband, if the number of packets received is enough to decode the coefficients of the description of that subband, the appropriate erasure-resilient code is applied, and the description is recovered; else the description is lost. Finally, an approximation to the original image is obtained by MD dequantizing each subband, and inverting the transform.

3) Remarks: Except for minor modifications, the technique we propose to allocate redundancy and create N descriptions of a given image is essentially the same as the priority encoding transmission scheme [1]. However, there is one factor which, without further research, prevents us from using the more general erasure-resilient codes computed by PET systems: it is the fact that in our construction, we need to deal with the—somewhat bizarre—concept of having to entropy code channel codewords.

In our experiments, we used a very simple form of MDS codes to add redundancy: the family of $(n, 1, n)$ repetition codes. For these codes, compression of the resulting channel codewords is not difficult, because whatever statistical model was available to encode source outcomes can also be used to encode the channel codewords (these are nothing but multiple identical copies of the original source outcomes). A very

⁵Also known as $(1 + \epsilon)$ -MDS codes.

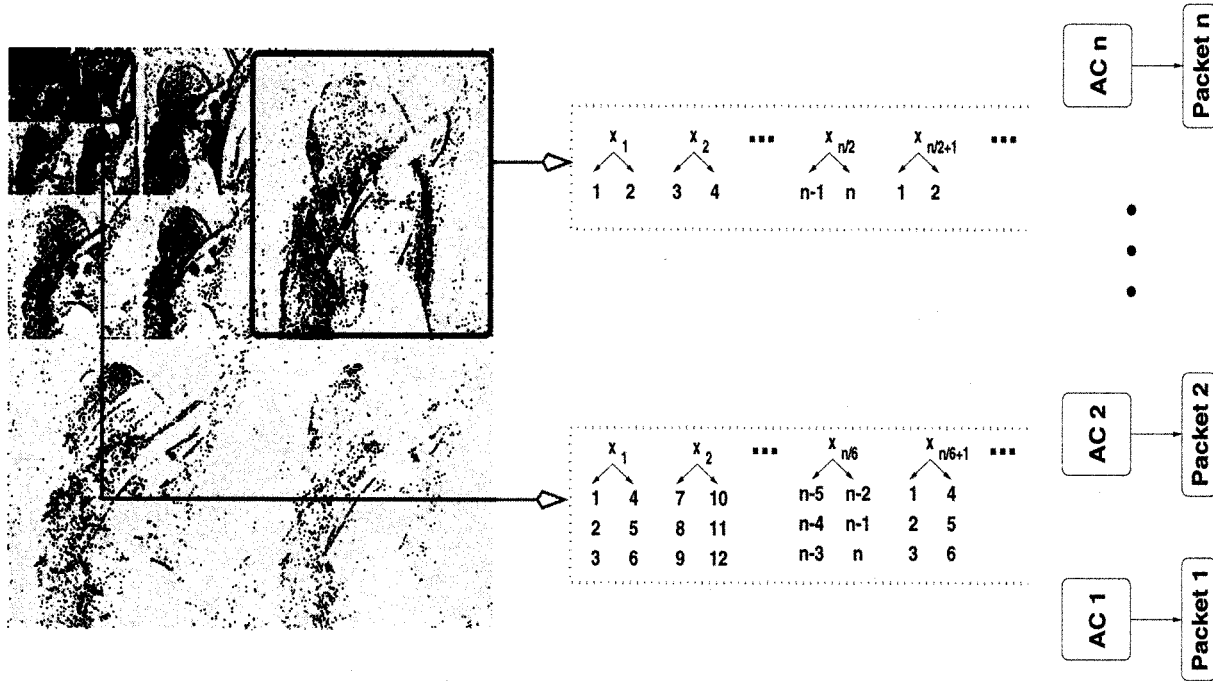


Fig. 7. Robust compression of subband data. In a first step, a MD quantizer is applied to each subband coefficient, creating two descriptions of the original value. Then, depending on the subband to which a coefficient belongs, an appropriate error correcting code is applied: in this example, no code is applied on the highlighted high frequency subband, and a simple (3, 1, 3) code is applied on each description of the highlighted low frequency subband. Then, the components of each codeword are distributed among multiple arithmetic coders, one per packet to be transmitted.

interesting problem opened up by our work is the study of dependencies between the statistical properties of source outcomes and those of the codewords obtained using more general MDS or $(1 + \epsilon)$ -MDS codes.

IV. EXPERIMENTAL RESULTS

In this section we report coding results. In our experiments, we use the 10–18 Daubechies wavelet, and target coding rates in the range 0.25–1 bpp. Sample image reconstructions, matlab files used to generate the plots presented below, and C source code, can be found at <http://lcavwww.epfl.ch/~servetto/>.

A. MD Image Coding Performance

In a first experiment, we show performance results for our optimized MD coder. For comparison, we also present results for the case when the same index assignment is fixed for all subbands, thus eliminating the need to search for optimal combinations. The corresponding plots are shown in Fig. 8 and sample image reconstructions are shown in Fig. 9.

Observe how in Fig. 8 the convex hull corresponding to subband adaptation lies strictly below that of the fixed image-wide choices. In the low excess-rate regime (high single-channel distortion and low two-channel distortion, the bottom part of this plot), an increase in PSNR of +3 dB in the performance of the single-channel decoder occurs when index assignments are freely chosen for each subband. In the high excess-rate regime, the gap is negligible.

It is also interesting to observe how, in this example, in order to improve the quality of the two-channel decoder by 0.76 dB (from 38.69 dB up to 39.45 dB) it is necessary to tolerate a




Fig. 8. Subband adaptation of the index assignments. The dots marked $d = 2, 3, 5, 7$ correspond to the single-channel/two-channel MSE tradeoff achieved by fixed, image-wide MD quantizers; all other dots correspond to the performance achieved by different combinations of index assignments, one per subband. In all cases, the central quantizer is adjusted so that the total rate would be 0.5 bpp/description (total 1 bpp).

single-channel decoded image that is 7.08 dB worse (from 35.53 dB down to 28.45 dB). We see therefore that, for practical applications, there is little gain to be had by using index assignments other than the staggered case ($d = 2$). This is because in order to be able to gain less than one dB in PSNR for the two-channel decoder, the quality of the images obtained by the single-channel decoder has to be reduced by 7–8 dB. Hence, the optimization process is very useful to understand the limitations of our proposed coding framework, as a benchmarking tool, and



Fig. 9. Sample image reconstructions, for Lena, at 0.5 bpp/channel. (a) Two-channel decoder, high redundancy (PSNR: 38.69 dB), (b) single-channel decoder, high redundancy (PSNR: 35.53 dB), (c) two-channel decoder, low redundancy (PSNR: 39.45 dB), and (d) single-channel decoder, low redundancy (PSNR: 28.45 dB).

is certainly a very interesting intellectual exercise. But in practice, the extra complexity involved in choosing optimal coding parameters is just not worth the marginal improvements in perceptual image quality, especially in the high excess-rate regime.

B. Comparison Against other Multiple Description Image Coders

1) *Two Descriptions:* As a reference, we compare the performance attained by our MD coder (first in its two-packet configuration), against MD image coders of Wang, Orchard and Reibman [35], and of Srinivasan and Chellappa [31]. The resulting PSNR's are shown in Fig. 10.

We would like to point out that the comparison with the JPEG-type coder of [35] is not entirely fair, in that it is not clear what part of the gains we present come from using a superior single description wavelet-based coder (wavelets versus JPEG) and what from a different technique for constructing MD's.⁶ However, we strongly suspect that the gains we obtain cannot be attributed *only* to the use of a better transform and better post-transform processing/modeling of transform data. Performance comparisons for the two-channel decoder are certainly not useful, since in this case all we are comparing is

⁶However, we include the comparison anyway because this is one of the few existing good MD image coders we know of.




Fig. 10. PSNR values achieved by our coder and by the two reference coders (left: Lena, compared to [35]; right: Barbara, compared to [31]). Similar PSNR values are obtained by our coder using about 50%–60% of the bit rate required by these other coders.

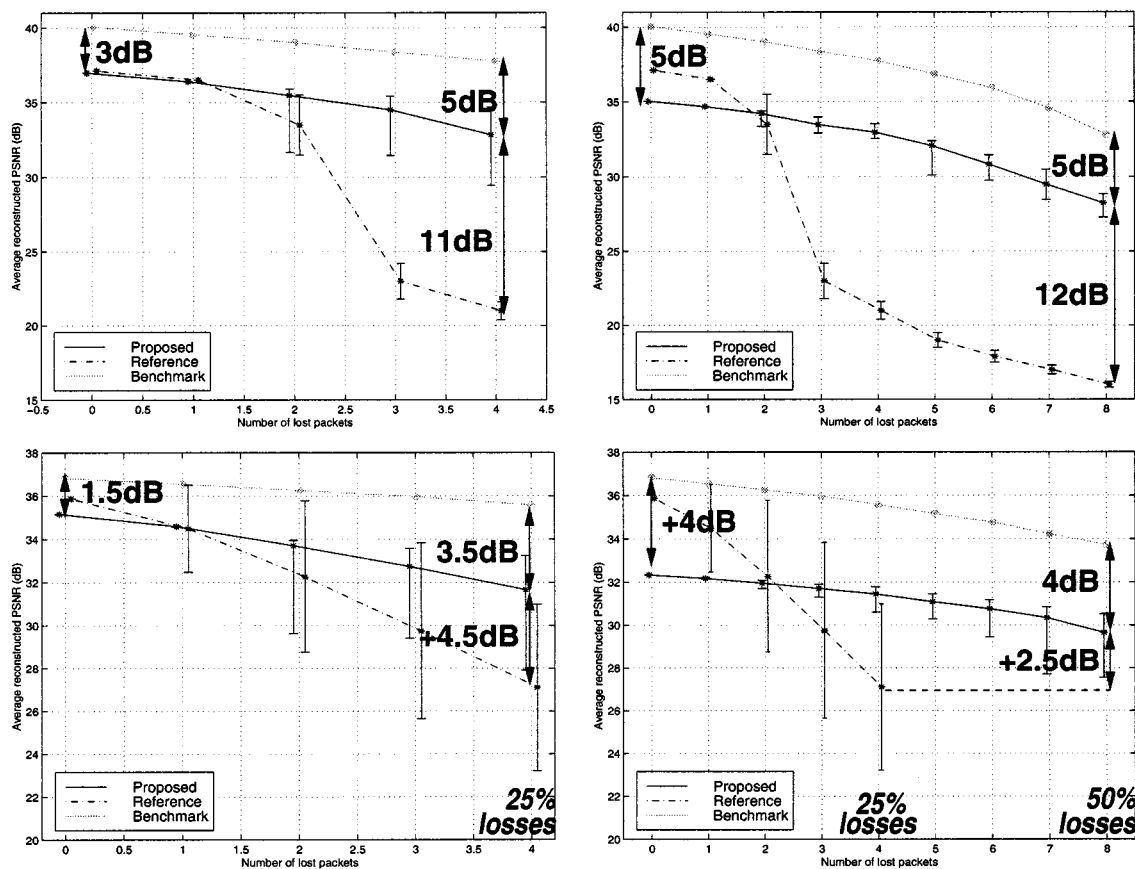


Fig. 11. PSNR values achieved by our coder and by the two reference coders. Top: Lena coded at 1 bpp into 10 packets (0.1 bpp/packet), compared to [12]; bottom: Lena coded 0.5 bpp into 16 packets (1/32 bpp/packet), compared to [14]. In both comparisons, left plots show results for low-redundancy allocations in our coder, right plots for high-redundancy allocations. The system denoted “Benchmark” refers to an ideal, genie-aided encoder which knows prior to transmission which packets will be lost, and hence uses a single-description coder and reduces its transmission rate accordingly.

a standard JPEG coder against a good wavelet coder, and there the differences are known to be, on average, 2–2.5 dB in favor of wavelet coders, which our experiments verify. However, for the single-channel decoders, which is when the technique used to generate MD’s comes into play, the gains we observe are in excess of 4 dB; no wavelet coder is known to have achieved such dramatic gains over a JPEG coder, and hence we do not

think that these gains can be explained only in terms of a DCT-versus-Wavelet argument.

2) *Many Descriptions*: As another reference, we compare the performance attained by our MD coder (now in its many-packets configuration), against MD image coders of Goyal *et al.* [12] and of Jiang and Ortega [14]. The resulting PSNR’s are shown in Fig. 11.

In this case we would like to point out that, using a simple form of the PET system (based on trivial repetition codes), our coder is able to attain a much wider range of tradeoffs between redundancy allocations and robustness. For certain choices of redundancy parameters (MD quantizer and channel codes to combat erasures), we are able to attain very high performance when only a small number of packets is lost, at the expense of poor performance and rapid decay when the number of lost packets increases. Alternatively, we can sacrifice some compression performance when the number of lost packets is small, and in return get a much slower decay as well as greatly reduced uncertainty in performance when the number of packets increases. To the best of our knowledge, ours is the first system to provide such functionality.

C. Remarks on Computational Complexity

Excluding the DP optimization step which, as argued above, is not critical from a practical point of view, the encoder/decoder pair is of a remarkable low complexity for the performance it delivers. To give a practical sense of this, running on an Ultra Sparc 1 workstation with a 140 Mhz CPU, it takes 6.8 s to produce a 512×512 image having two descriptions each encoded at 0.5 bpp, and 4.8 s to produce the same image based on a single description. These times are all-inclusive: network I/O, and computation of the direct and inverse wavelet transforms are included. And our implementation is *far* from optimized for speed.

Furthermore, after the wavelet calculation, all processing is done in a single pass, and scanning each subband in raster order: the flow of control does not depend on the particular data set being coded, it is a simple raster-oriented algorithm, particularly well suited for implementation on a DSP chip. And all steps except the final entropy coding admit a parallel implementation: this is of particular interest mainly for the wavelet transform, which is by far the most cpu-intensive part of the encoder/decoder, and one we have little control over in designing our algorithms.

V. CONCLUSIONS

In this work, we presented the design, implementation and optimization of error resilient data compression algorithms. Unlike most of the recent developments in the field (based on DCT's and subspace methods), our algorithms are based on the use of wavelets, MD scalar quantizers, and erasure-resilient codes.

Our coding results are a significant improvement over the state-of-the-art in the field: on average, we can match the image quality achieved by most other coders using about 50%–60% of their bit rate. And furthermore, we attain this using a remarkably simple coding structure, both conceptually as well as computationally. However, we believe there is still room for improvement: when removing the MD constraints, as a single description image coder, our coder falls approx. 0.8–0.9 dB below the performance of the best published coder we are aware of [15]. As argued, the MD constraints impose restrictions on the set of tools available to build a good coder. Yet it remains to be seen whether “smarter” coders within our framework, or even coders

based on entirely different frameworks, are able to close this gap.

Further work is required to complete the proof of the conjectured intractability result mentioned in Section II-C. If our conjecture proves true, an approximation algorithm for generating index assignments based on a source description, capable of generalizing asymptotically optimal MD quantizers [32], [34] (and hopefully improving upon their performance at low rates) would be extremely useful.

Perhaps the single most important problem yet to be resolved in the context we set up in this work is that of redundancy allocation (in the $N > 2$ case). Recall that we motivated the problem of more than two descriptions based on constraints imposed by the application of transmission over IP networks (Section III). In this scenario, the choice of priority function ρ specifying which codes are to be applied on which subbands is a major factor determining the overall performance of such systems; and that choice should depend on channel conditions (high losses \rightarrow more redundancy, etc). We are currently studying this problem.

ACKNOWLEDGMENT

The authors would like to thank T. Berger-Wolf, for many useful discussions on the structure and properties of index assignments, W. Jiang for providing them a copy of his paper [14] prior to its publication, and the anonymous reviewers, for multiple suggestions which resulted in a significantly improved manuscript.

REFERENCES


- [1] A. Albanese, J. Blömer, J. Edmonds, M. Luby, and M. Sudan, “Priority encoding transmission,” *IEEE Trans. Inform. Theory*, vol. 42, pp. 1737–1744, Nov. 1996.
- [2] N. Alon and M. Luby, “A linear time erasure-resilient code with nearly optimal recovery,” *IEEE Trans. Inform. Theory*, vol. 42, pp. 1732–1736, Nov. 1996.
- [3] 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.
- [4] T. Y. Berger-Wolf and E. M. Reingold, “Optimal index assignment for multichannel communications,” in *Proc. SIAM Symp. Discrete Algorithms*, Baltimore, MD, 1999, (long version available from <http://emr.cs.uiuc.edu/~reingold/algorithms.shtml>).
- [5] A. M. Bruckstein, K. J. Holt, and A. N. Netravali, “Holographic representation of images,” *IEEE Trans. Image Processing*, vol. 7, pp. 1583–1597, Nov. 1998.
- [6] C. Chrysafis and A. Ortega, “Efficient context-based entropy coding for lossy wavelet image compression,” in *Proc. Data Compression Conf.*, Snowbird, UT, 1997.
- [7] D. Chung and Y. Wang, “Multiple description image coding based on lapped orthogonal transforms,” in *Proc. IEEE Conf. Image Processing*, Chicago, IL, 1998.
- [8] T. Cormen, C. Leiserson, and R. Rivest, *Introduction to Algorithms*. Cambridge, MA: MIT Press, 1990.
- [9] W. Equiritz and T. Cover, “Successive refinement of information,” *IEEE Trans. Inform. Theory*, vol. 37, pp. 269–275, Mar. 1991.
- [10] A. A. El Gamal and T. Cover, “Achievable rates for multiple descriptions,” *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 851–857, Nov. 1982.
- [11] V. K. Goyal and J. Kovačević, “Optimal multiple description transform coding of Gaussian vectors,” in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, 1998.
- [12] V. K. Goyal, J. Kovačević, R. Arian, and M. Vetterli, “Multiple description transform coding of images,” in *Proc. IEEE Int. Conf. Image Proc.*, Chicago, IL, 1998.
- [13] V. K. Goyal, J. Kovačević, and M. Vetterli, “Multiple description transform coding: Robustness to erasures using tight frame expansions,” in *Proc. IEEE Int. Symp. Inform. Theory*, Cambridge, MA, 1998.

- [14] W. Jiang and A. Ortega, "Multiple description coding via polyphase transform and selective quantization," in *Proc. SPIE Visual Commun. Image Processing Conf.*, San Jose, CA, 1999.
- [15] R. Joshi, H. Jafarkani, J. Kasner, T. Fischer, N. Farvardin, M. Marcellin, and R. Bamberger, "Comparison of different methods of classification in Subband coding of images," *IEEE Trans. Image Processing*, vol. 6, pp. 1473–1486, Nov. 1997.
- [16] S. LoPresto, K. Ramchandran, and M. Orchard, "Image coding based on mixture modeling of wavelet coefficients and a fast estimation-quantization framework," in *Proc. Data Compression Conf.*, Snowbird, UT, 1997.
- [17] A. E. Mohr, E. A. Riskin, and R. E. Ladner, "Graceful degradation over packet erasure channels through forward error correction," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, 1999.
- [18] F. Ng and J. Kovačević, "Nonredundant image representations," in *Proc. IEEE Int. Conf. Image Processing*, CA, 1997.
- [19] M. T. Orchard, Y. Wang, V. A. Vaishampayan, and A. Reibman, "Redundancy rate-distortion analysis of multiple description coding using pairwise correlating transforms," in *Proc. IEEE Int. Conf. Image Processing*, Santa Barbara, CA, 1997.
- [20] L. H. Ozarow, "On a source coding problem with two channels and three receivers," *Bell Syst. Tech. J.*, vol. 59, no. 10, pp. 1909–1921, 1980.
- [21] J. Rogers and P. Cosman, "Robust wavelet zerotree image compression with fixed-length packetization," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, 1998.
- [22] A. Said and W. Pearlman, "A new, fast, and efficient image codec based on set partitioning in hierarchical trees," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 6, pp. 243–250, Mar. 1996.
- [23] A. Sayeed and B. Aazhang, "Joint multipath-doppler diversity in mobile wireless communications," *IEEE Trans. Commun.*, vol. 47, pp. 123–132, Jan. 1999.
- [24] S. D. Servetto, "Compression and reliable transmission of digital image and video signals," Ph.D. dissertation, Univ. Illinois, Urbana-Champaign, 1999.
- [25] S. D. Servetto, K. Ramchandran, K. Nahrstedt, and A. Ortega, "Optimal segmentation of a VBR source for its parallel transmission over multiple ATM connections," in *Proc. IEEE Int. Conf. Image Processing*, Santa Barbara, CA, 1997.
- [26] S. D. Servetto, K. Ramchandran, and M. T. Orchard, "Image coding based on a morphological representation of wavelet data," *IEEE Trans. Image Processing*, vol. 8, pp. 1161–1174, Sept. 1999.
- [27] S. D. Servetto, K. Ramchandran, V. A. Vaishampayan, and K. Nahrstedt, "Multiple-description wavelet based image coding," in *Proc. IEEE Int. Conf. Image Processing*, Chicago, IL, 1998.
- [28] S. D. Servetto, J. M. Rosenblatt, and K. Ramchandran, "A binary Markov model for the quantized wavelet coefficients of images and its rate/distortion optimization," in *Proc. IEEE Int. Conf. Image Processing*, Santa Barbara, CA, 1997.
- [29] J. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. Signal Processing*, vol. 41, pp. 3445–3462, Dec. 1993.
- [30] A. Silberschatz and P. Gavrilin, *Operating Systems Concepts*, 4th ed. Reading, MA: Addison-Wesley, 1994.
- [31] M. Srinivasan and R. Chellappa, "Multiple description subband coding," in *Proc. IEEE Int. Conf. Image Processing*, Chicago, IL, 1998.
- [32] V. A. Vaishampayan, "Design of multiple description scalar quantizers," *IEEE Trans. Inform. Theory*, vol. 39, pp. 821–834, May 1993.
- [33] —, "Application of multiple description codes to image and video transmission over lossy networks," in *Proc. 7th Int. Workshop Packet Video*, Brisbane, Australia, 1996.
- [34] 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.
- [35] Y. Wang, M. T. Orchard, and A. Reibman, "Multiple description image coding for noisy channels by pairing transform coefficients," in *Proc. 1st Workshop IEEE Signal Processing Soc. Multimedia Signal Processing*, Princeton, NJ, 1997.
- [36] Z. Xiong, K. Ramchandran, and M. Orchard, "Space quantization for wavelet image coding," *IEEE Trans. Image Processing*, vol. 6, pp. 677–693, May 1997.
- [37] Y. Yoo, A. Ortega, and K. Ramchandran, "A novel hybrid technique for discrete rate-distortion optimization with applications to fast codebook search for SVQ," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, Atlanta, GA, 1996.

Sergio D. Servetto (S'95–M'99) was born in Argentina in 1968. He received the Lic. Inf. degree from the Universidad Nacional de La Plata (UNLP), Argentina, in 1992, and the M.Sc. degree in electrical engineering and Ph.D. degree in computer science from the University of Illinois, Urbana-Champaign (UIUC), in 1996 and 1999, respectively.

From 1991 to 1994, he was a Permanent Member of Technical Staff of the Computer Research and Advanced Applications Group (CRAAG), IBM, Argentina, where he was involved in the development of optical character recognition systems designed to operate in noisy industrial environments. From 1994 to 1999, he was a Graduate Research Assistant with the Image Formation and Processing Group, Beckman Institute, UIUC, and at the Multimedia Operating Systems and Networking Group, Department of Computer Science, UIUC. He held part-time teaching appointments at UNLP during 1992–1993 and UIUC during 1996–1997. He was with Bell Labs, Murray Hill, NJ, and AT&T Labs, Florham Park, NJ, during the Summers of 1997 and 1998, respectively. Since the Fall of 1999, he has been the First Assistant to Prof. Martin Vetterli, Laboratoire de Communications Audiovisuelles (LCAV), Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland. His research interests include coding and transmission over networks of image and video signals, as well as fundamental problems in algorithms, signal processing, communications, control, quantization and coding theory that arise in that context.

Dr. Servetto was the recipient of the 1998 Ray Ozzie Fellowship for outstanding graduate students in computer science, and of the 1999 David J. Kuck Outstanding Thesis Award for the best doctoral dissertation of the year, both from the Department of Computer Science, UIUC.




Kannan Ramchandran (S'93–M'93) received the B.S. degree from the City College of New York, and the M.S. and Ph.D. degrees from Columbia University, New York, all in electrical engineering, in 1982, 1984, and 1993, respectively.

From 1984 to 1990, he was a Member of Technical Staff at AT&T Bell Labs, working in telecommunications R&D. From 1990 to 1993, he was a Graduate Research Assistant with the Center for Telecommunications Research at Columbia University. From 1993 to 1999, he was on the faculty of the Electrical and

Computer Engineering Department, University of Illinois, Urbana-Champaign, and was affiliated with the Beckman Institute and the Coordinated Science Laboratory. Since the Fall of 1999, he has been an Associate Professor in the Electrical Engineering and Computer Science Department, University of California, Berkeley. His research interests include image and video compression and communications, multirate signal processing and wavelets, fast algorithms for signal and image processing, and unified algorithms for signal processing, communications, and networking.

Dr. Ramchandran was the recipient of the 1993 Elaihu I. Jury Award at Columbia University for the best doctoral dissertation in the area of systems, signal processing, or communications. He received an NSF Research Initiation Award in 1994, an Army Research Office Young Investigator Award in 1996, an NSF CAREER Award in 1997, and an Office of Naval Research Young Investigator Award in 1997. He was the co-recipient of the 1996 Senior Best Paper Award from the IEEE Signal Processing Society. He is a Member of the IEEE IMDSP Technical Committee, and serves as an Associate Editor for the IEEE TRANSACTIONS ON IMAGE PROCESSING.



Vinay A. Vaishampayan is with AT&T Shannon Laboratory, Florham Park, NJ. His research interests are in signal processing, communications, and information theory.



Klara Nahrstedt (S'93–M'95) received the B.A. degree in mathematics from Humboldt University, Berlin, Germany, in 1984, and the M.Sc. degree in numerical analysis from the same university in 1985. She received the Ph.D. degree from the Department of Computer and Information Science, University of Pennsylvania, Philadelphia, in 1995.

She is an Assistant Professor with the Computer Science Department, University of Illinois, Urbana-Champaign. Her research interests are directed toward multimedia services, protocols, quality of service provision and resource management in real-time multimedia distributed systems. She is coauthor of the widely used book *Multimedia: Computing, Communications and Applications* (Englewood Cliffs, NJ: Prentice-Hall).

Dr. Nahrstedt is the recipient of the Early NSF Career Grant and Xerox Award for research achievements. She is a member of the ACM and IEEE Computer Society.