

Research Article

Multiple Description Coding with Redundant Expansions and Application to Image Communications

Ivana Radulovic and Pascal Frossard

LTS4, Swiss Federal Institute of Technology (EPFL), Signal Processing Institute, 1015 Lausanne, Switzerland

Received 15 August 2006; Revised 19 December 2006; Accepted 28 December 2006

Recommended by Béatrice Pesquet-Popescu

Multiple description coding offers an elegant and competitive solution for data transmission over lossy packet-based networks, with a graceful degradation in quality as losses increase. In the same time, coding techniques based on redundant transforms give a very promising alternative for the generation of multiple descriptions, mainly due to redundancy inherently given by a transform, which offers intrinsic resiliency in case of loss. In this paper, we show how partitioning of a generic redundant dictionary can be used to obtain an arbitrary number of multiple complementary, yet correlated, descriptions. The most significant terms in the signal representation are drawn from the partitions that better approximate the signal, and split to different descriptions, while the less important ones are alternatively distributed between the descriptions. As compared to state-of-the-art solutions, such a strategy allows for a better central distortion since atoms in different descriptions are not identical; in the same time, it does not penalize the side distortions significantly since atoms from the same partition are likely to be highly correlated. The proposed scheme is applied to the multiple description coding of digital images, and simulation results show increased performances compared to state-of-the-art schemes, both in terms of distortions and robustness to loss rate variations.

Copyright © 2007 I. Radulovic and P. Frossard. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

Efficient transmission of information over erasure channels has attracted a lot of efforts over the years, from different research communities. Such a problem becomes especially challenging when the coding block length is limited, or when the channel is not perfectly known, like in most typical image communication problems. It becomes therefore nontrivial to efficiently allocate the proper amount of channel redundancy, in order to ensure the robustness to channel erasures and, in the same time, to avoid wasting resources by overprotecting the information. When information losses are almost inevitable, and complexity or delay constraints limit the application of long channel codes, or retransmission, it becomes primordial to design coding schemes where all available bits can help the signal reconstruction.

An elegant solution to that kind of problems consists in describing the source information by several descriptions that can be used independently for the signal reconstruction. This is known as multiple description coding (MDC). The motivation behind multiple description coding is to encode the source information in such a way that high-quality recon-

struction is achieved if all the descriptions are available, and that the quality gracefully degrades in case of channel loss. As represented in Figure 1, the distortion depends on the number of descriptions available for the reconstruction, and typically decreases as the number of descriptions increases. Since multiple description coding offers several advantages, such as interesting graceful degradation in the presence of loss, and a certain robustness to uncertainty about channel characteristics, it has motivated the developments of numerous interesting coding algorithms. Some of these approaches completely rely on the redundancy present in the source, while others try to introduce a controlled amount of redundancy such that the distortion after reconstruction gracefully degrades in the presence of loss. The main challenge remains to limit the increase of rate compared to a single description case, and to trade off side and central distortions depending on the channel characteristics.

Redundant transforms certainly represent one of the most promising alternatives to generate descriptions with a controlled correlation, which nicely complement each other for efficient signal reconstruction. Recent advances in signal approximation have also demonstrated the benefits of

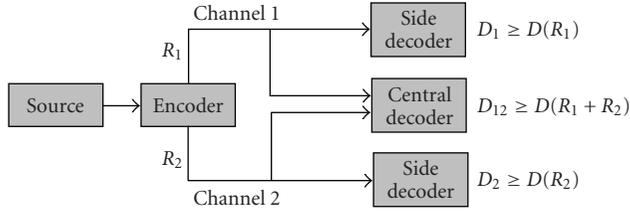


FIGURE 1: MDC with two descriptions, encoded with rates R_i . The distortion depends on the number of descriptions available at receivers.

flexible overcomplete expansion methods, particularly for multidimensional signals like natural images dominated by geometric features, where classical orthogonal transforms have shown their limitations. Transforms that build a sparse expansion of the signal, over a redundant dictionary of functions, are able to offer increased energy compaction, and design flexibility that often results in interesting adaptivity to signal classes. In addition, since the components of the signal are not orthogonal, they offer intrinsic resiliency to channel loss, which naturally makes redundant transforms interesting in multiple description coding schemes.

In this paper, we build on [1] and we present a method for the generation of an arbitrary number $N \geq 2$ of descriptions, by partitioning generic redundant dictionaries into coherent blocks of atoms. During encoding, atoms of the same dictionary partition are distributed in different descriptions. Since they are chosen from blocks of correlated atoms, such an encoding strategy does not bring an important penalty in the side distortion. In the same time, as they are still different, they all contribute to improvement of the reconstruction quality, and therefore decrease the central distortion as opposed to the addition of pure redundancy. The new encoding scheme is then applied to an image communication problem, where it is shown to outperform classical MDC schemes based on unequal error protection of signal components. The main contributions of this paper reside in the design of a flexible multiple description scheme, able to generate an arbitrary number of balanced descriptions, based on a generic dictionary. It additionally outperforms classical MDC schemes in terms of average distortion, and resilience to incorrect estimation of channel characteristics.

The paper is organized as follows. Section 2 presents an overview of the most popular multiple description coding strategies, with an emphasis on the redundant transforms and their potentials. In Section 3, we show how to partition redundant dictionaries, in order to generate multiple descriptions with a controlled correlation. Reconstruction of the signal with the available descriptions is discussed, and the influence of the distribution of atoms in the redundant dictionary is analyzed. Section 4 presents the application of the proposed multiple description coding scheme to a typical image communication scenario, while Section 5 finally provides simulation results that highlight the quality improvements compared to MDC schemes based on atom repetition, or unequal error protection. Finally, Section 6 concludes the paper.

2. RELATED WORK

This section presents a brief overview of multiple description coding techniques, with a particular emphasis on algorithms based on redundant transforms, and methods applied to multidimensional signals like images or video. The first and certainly simplest idea for the generation of multiple descriptions is based on information splitting [2], which basically distributes the source information between the different descriptions. This technique is quite effective if redundancy is present in the source signal, as it is typically the case in image and video signals. For example, wavelet coefficients of an image could be split into polyphase components [3]. Similarly, video information can be split into sequences of odd and even frames, which has been applied for the generation of two descriptions of video [4]. However, information splitting is generally limited to the generation of two descriptions due to drastic loss in coding efficiency when the number of descriptions increases.

Multiple descriptions can also be produced by extending quantization techniques with proper index assignment methods. These techniques lead to a refined quantization of the source samples, when the number of description increases. Multiple description coding based on both scalar and vector quantizations have been proposed [5, 6]. The multiple description scalar quantization (MDSQ) concept has also been successfully applied to the coding of images (see, e.g., [7] or [8]) or image sequences [9]. However, multiple description coding based on quantization techniques is mostly limited to two descriptions due to the rapid increase in complexity when the number of descriptions augments.

Transform coding has also been proposed to produce multiple descriptions [10], where it basically helps in reintroducing a controlled amount of redundancy to a source composed of samples with small correlation (as produced by typical orthogonal transforms). This redundancy becomes eventually beneficial to recover the information that has been lost due to channel erasures. The JPEG image coding standard can be modified to generate two descriptions by rotating the DCT coefficients [11, 12], and reintroducing a nonnegligible correlation between them. In practice however, the design of optimal correlating transforms is quite challenging. While solutions hold for a Gaussian source in the case of two descriptions, the generalization to a larger number of descriptions does not have yet any analytical solution.

Instead of implementing a transform that tries to provide uncorrelated coefficients, followed by a correlating transform to increase robustness to channel errors, redundant transforms can advantageously be used to provide a signal expansion with a controlled redundancy between components. Typical examples of redundant signal expansions are based on frames, or matching pursuit approximation. In [13], harmonic frames are used to generate multiple descriptions, and it was shown that this kind of expansion performs better than unequal error protection (UEP) schemes. Similar conclusions can be drawn from [14], where a frame expansion is applied to the wavelet coefficient zero trees to generate two or four descriptions. However, use of frames for the generation

of multiple descriptions is quite limited by the fact that not all subsets of received frame components enable a good signal reconstruction [13]. In [15], the authors compare four MDC schemes for video, based on redundant wavelet decompositions, and they give an insight in the tradeoff between the side and central distortions for the two schemes that perform best. Another related work is presented in [16], where a scheme for multiple description scalable video coding based on a motion-compensated redundant analysis was proposed.

In [17, 18], the authors propose to generate two descriptions of video sequences with a matching pursuit algorithm. In their implementation, the elements of a redundant dictionary (the so-called atoms) that best approximate a signal are repeated in both descriptions, while the remaining atoms are alternatively split between the descriptions. The redundancy between the descriptions is controlled with a number of shared atoms. The same principle, combined with multiple description scalar quantization, can also be found in [19, 20], where the authors used the orthogonalized version of matching pursuit. However, the problem with these solutions is that they do not exploit the redundancy inherently offered by the transform, but they rather introduce channel redundancy by repeating the most important information. If no loss occurs, such a repetition results in an obvious waste of resources. This is exactly what we propose to avoid in the multiple description coding algorithm that relies on partitioning of the redundant dictionary into coherent blocks of atoms. In this way, descriptions can be made similar to being robust to channel erasures, yet different enough to improve the signal reconstruction when the channel is good.

3. MULTIPLE DESCRIPTION CODING WITH REDUNDANT DICTIONARIES

3.1. Motivations

While most modern image compression algorithms, such as the JPEG standard family, have been designed following the classical coding paradigm based on orthogonal transform and scalar quantization, new representation methods have recently been proposed in order to improve the shortcomings inherent to classical algorithms. Even if important improvements have been offered by different types of wavelet transforms, optimality of the approximation is only reached for specific cases. In particular, it has been shown that wavelet transforms are suboptimal for the approximation of multidimensional signals like natural images, which are dominated by edges and geometric features. Adaptive and nonlinear approximations over redundant dictionaries of functions have emerged as an interesting alternative for image coding, and have been proven to be highly effective, especially at low bit rate [21].

In addition to increased design flexibility, and improved energy compaction properties, redundant dictionaries also offer some intrinsic resiliency to loss of information, due to channel erasures, for example. Since the components of the signal expansion are not orthogonal, efficient reconstruction strategies can be derived in order to estimate lost elements, and to improve the quality of the signal reconstruction. That

quality can yet be dramatically improved by a careful signal encoding strategy, where information has to be arranged in such a way that the simultaneous loss of important correlated components becomes unlikely. This naturally leads to the concept of multiple description coding that exactly pursues this objective. Instead of introducing redundancy in the signal expansion to fight against channel loss, one can exploit the redundancy of the dictionary and partition it, such that multiple complementary yet correlated descriptions can be built by proper distribution of the signal components.

The inherent redundancy present in the transform step and the good approximation properties offered by overcomplete expansions obviously motivate the use of redundant dictionaries in the design of joint source and channel coding strategies. Multiple description image coding stands as a typical application where the benefits of properly designed redundant dictionary are particularly advantageous. While previous works mostly use complex frame construction, or unequal protection based on forward error correction mechanisms [13, 14], we propose in this paper to build multiple descriptions with a dictionary partitioning algorithm, and a greedy signal approximation based on a modified matching pursuit algorithm.

3.2. Definitions

Before going more deeper into the construction of descriptions, we now fix the notations and definitions that are used in the rest of this paper. We consider a scenario with N descriptions that are denoted by D_i , with $1 \leq i \leq N$. Each description contains M signal components, and descriptions are balanced in terms of size and importance. The distortion induced by the signal reconstruction with only one description is called the *side* distortion, while the distortion after reconstruction from several descriptions is called *partial* distortion. Finally, if all the descriptions are used for the signal reconstruction, the distortion is called the *central* distortion. In the case where all descriptions have approximately equal size, and all the side distortions are similar, we say that the descriptions are *balanced*.

We now briefly recall a few definitions that allow to characterize redundant dictionaries. First, we consider a set of signals s that lay in a real d -dimensional vector space \mathcal{R}^d endowed with a real-valued inner product. We further assume that any of these signals is to be represented with a finite collection of unitary norm elementary signals called the *atoms*. Denote by $\mathcal{D} = \{a_i\}_{i=1}^{|\mathcal{D}|}$ such a collection of $|\mathcal{D}|$ atoms that we call a *dictionary*. Redundant dictionaries are such that the number of atoms in the dictionary is usually much bigger than the dimensionality of the signal, that is, $|\mathcal{D}| \gg d$. There is no particular constraint regarding the dictionary, except that it should span the entire signal space.

Several metrics have been proposed to characterize the redundant dictionary \mathcal{D} . For example, the *structural redundancy* β reports the distribution of atoms in the dictionary, and is written as

$$\beta = \inf_{a, \|a\|=1} \sup_{p \in \mathcal{D}} |\langle a, a_p \rangle|. \quad (1)$$

Basically, it measures the cosine of maximum possible angles between any direction of the signal s , and its closest direction among the atoms in \mathcal{D} . The structural redundancy β obviously depends on the dictionary construction and controls the approximation rate for overcomplete signal expansions over the dictionary \mathcal{D} .

Another metric, which is often simpler to compute, reflects the worst-case correlation between any two atoms in the dictionary. It is defined as the *coherence* of the dictionary, and is written as

$$\mu = \max_{\{a_p, a_q\} \in \mathcal{D}} |\langle a_p, a_q \rangle|. \quad (2)$$

Obviously, orthogonal basis has a coherence $\mu = 0$, while highly redundant dictionaries have coherence close to 1. Since the coherence only reflects an extreme property of the dictionary, the *cumulative coherence* $\mu_1(m)$ has been proposed to measure the maximum total correlation between a fixed atom with m distinct atoms. It is written as

$$\mu_1(m) = \max_{|\Lambda| \leq m} \max_{a_p \in \Lambda} \sum_{\lambda \in \Lambda} |\langle a_p, a_\lambda \rangle|, \quad (3)$$

where $\Lambda \subset \mathcal{D}$. In general, the cumulative coherence gives more information about the dictionary, but it is more difficult to compute. In the worst case, we can bound it as $\mu_1(m) \leq m\mu$.

Finally, it is often useful to partition redundant dictionaries into groups of atoms, for tree-based search algorithms [22], for example, or for controlling the construction of multiple descriptions, as detailed later. In this case, the dictionary \mathcal{D} is partitioned into blocks or subdictionaries $\{\mathcal{D}_i\}$ such that $\bigcup_i \mathcal{D}_i = \mathcal{D}$ and $\mathcal{D}_i \cap \mathcal{D}_j = \emptyset$ for $i \neq j$. It then becomes interesting to characterize the distance between these subdictionaries. The block coherence μ_B is therefore defined by

$$\mu_B = \max_{i \neq j} \max_{a_p \in \mathcal{D}_i, a_q \in \mathcal{D}_j} |\langle a_p, a_q \rangle|. \quad (4)$$

A special class of redundant dictionaries represents the dictionaries that can be partitioned into independent groups of correlated atoms, which are called *block-coherent* dictionaries.

3.3. MDC with partitioned dictionaries

Multiple description coding is an efficient strategy to fight against channel erasures, and redundant dictionaries of functions certainly offer interesting properties for the construction of correlated descriptions. Descriptions, which typically represent sets of signal components, should be built in such a way that they are complementary in providing a good signal approximation, and yet correlated to provide robustness to channel erasures. We propose to achieve this construction by partitioning the dictionary into blocks of similar atoms. Each atom of a block is then put in a different description, which ensures that descriptions are correlated. In the same time, since atoms in a block are different, they all contribute in improving the approximation of the signal.

In more details, recall that our objective is to generate an arbitrary number N of descriptions of the signal s , which are balanced in size and distortion. Each description contains a subset of atoms drawn from the dictionary \mathcal{D} , along with their respective coefficients that represent the contribution of the atom in the signal approximation. We first partition the dictionary into clusters of N similar atoms. Each of these clusters is represented by a particular function that we call a *molecule*. A molecule is representative of the characteristics of the atoms within a cluster, and can be computed, for example, as a weighted sum of the N atoms of the cluster.

Then, instead of searching for the atoms that best approximate the signal s , the signal expansion is performed at the level of molecules. When the best representative molecules are identified, the atoms that compose the corresponding cluster in the dictionary are distributed between the different descriptions. This strategy first does not penalize considerably the side distortion, resulting from the reconstruction of the signal with one description only, since the atoms in dictionary clusters are likely to be very correlated. Second, proper reconstruction strategies are able to exploit the information brought by the different atoms of a cluster, in order to increase the quality of the signal approximation. Finally, it is interesting to note that a search performed on the molecules typically decreases the computational complexity of the signal expansion (e.g., a typical speedup factor of $\log_2 N$ can be achieved with respect to a full search on the dictionary).

More formally, suppose that a set of M molecules $\{m_j\}$ are selected as the best representative features of the signal s . The multiple description coding scheme allocates the child a_i^j of molecule m_j to the description i , where $i = 1, 2, \dots, N$. The atoms that compose the description i can subsequently be represented by a generating matrix Φ_i , with $\Phi_i = \{a_i^j\}$ and $j = 1, 2, \dots, M$. In addition to atoms, the descriptions also carry coefficients that reflect the relative contribution of each atom in the signal reconstruction. Coefficients are simply given by the projection of the signal s onto the generating matrix Φ_i as

$$\Phi_i s^T = C_i, \quad (5)$$

where $C_i = \langle s, a_i \rangle$ gives the contribution of each atom in Φ_i . C_i 's are continuous-valued vectors, which obviously need to be quantized before coding and transmission. We assume in this paper that they are uniformly quantized into \tilde{C}_i , with the same scalar quantizer and the same quantization step size Δ for all the coefficients. Even if that quantization strategy may not be optimal, it consists on a very common model used for the quantization of coefficients obtained by frame expansions (e.g., [13, 23]), and we use it also in this paper, for the sake of simplicity. We additionally assume that all the coefficients are quantized to the next lower quantization level, and that Δ is small enough. The quantization noise then becomes independent of the signal, and we can write

$$\tilde{C}_i = C_i + \eta, \quad (6)$$

where η denotes the quantization noise. The quantized coefficients \tilde{C}_i 's together with indices of atoms from Φ_i finally form the description i .

3.4. Signal reconstruction

The signal is eventually reconstructed with the descriptions that are available at the decoder, after possible erasures on a lossy channel. The redundant signal expansion proposed in the previous section obviously does not conserve the energy of the signal, which cannot be reconstructed by a simple linear combination of vectors \widetilde{C}_i 's and the atoms from the generating matrices Φ_i , obtained from the available descriptions. We therefore need to design a decoding process that removes the redundancy that has been introduced in the encoding stage, and we distinguish between two cases, based on the number of available descriptions.

If only one description i is available, the signal is simply reconstructed by determining the best approximation r_i of the signal s in a least-mean-square sense. It is given by

$$r_i = (\Phi_i^\dagger \cdot \widetilde{C}_i)^T = (\Phi_i^\dagger \cdot (C_i + \eta))^T, \quad (7)$$

where T and \dagger , respectively, denote the transpose and pseudoinverse operations. Such a reconstruction induces an MSE distortion D_i that can be expressed as

$$D_i = \frac{\|s - r_i\|^2}{S} = \frac{\|s - (\Phi_i^\dagger \cdot (C_i + \eta))^T\|^2}{S}. \quad (8)$$

Here S denotes the signal size. The distortion is composed of the distortion D_i^a due to the approximation of s over Φ_i , and the distortion due to quantization D_i^q . Recall that these two terms can be separated due to the high-rate approximation assumption that leads to the independency of the signal and the quantization noise. The source distortion can be further expressed as

$$\begin{aligned} D_i^a &= \frac{\|s\|^2 - \text{tr}(ss^T \Phi_i^T (\Phi_i \Phi_i^T)^{-1} \Phi_i)}{S} \\ &= \frac{\|s\|^2 - \text{tr}(C_i^T (\Phi_i \Phi_i^T)^{-1} C_i)}{S}, \end{aligned} \quad (9)$$

where S corresponds to the signal size and $\text{tr}(\cdot)$ denotes the matrix trace. In order to bound the distortion D_i^a , we consider the worst-case scenario where the correlation between any atoms in Φ_i is equal to μ_B , which is the maximal possible correlation between any two partitions in the dictionary \mathcal{D} . In this case, $(\Phi_i \Phi_i^T)^{-1}$ is a matrix having elements $(1 + \mu_B(M - 2))/(1 - \mu_B)(1 + \mu_B(M - 1))$ on the main diagonal, and $-\mu_B/(1 - \mu_B)(1 + \mu_B(M - 1))$ elsewhere. Therefore, we have

$$\begin{aligned} D_i^a &\leq \frac{\|s\|^2}{S} - \frac{\sum_i C_i^2}{S(1 - \mu_B)} + \frac{\mu_B (\sum_i C_i)^2}{S(1 - \mu_B^2(M - 1) + \mu_B(M - 2))} \\ &\leq \frac{\|s\|^2}{S} - \frac{\sum_i C_i^2}{S(1 + \mu_B(M - 1))}. \end{aligned} \quad (10)$$

Similarly, the quantization distortion can be written as

$$D_i^q = \frac{\Delta^2}{3S} \text{tr}(\Phi_i \Phi_i^T)^{-1}. \quad (11)$$

An upper bound on the quantization distortion can be derived by assuming the worst-case scenario, where the correlation between any pair of atoms is given by μ_B :

$$\begin{aligned} D_i^q &\leq \frac{M\Delta^2}{3S} \frac{1 + \mu_B(M - 2)}{(1 + \mu_B(M - 1))(1 - \mu_B)} \\ &\leq \frac{M\Delta^2}{3S} \frac{1}{1 - \mu_B}. \end{aligned} \quad (12)$$

We can note that the application of scalar quantization on correlated components induces a distortion that is inversely proportional to $1 - \mu_B$. Note that the quantization error could be reduced by orthogonalization of Φ_i at encoder, or by using vector quantization, for example. The design of an optimal quantization strategy for redundant signal expansions is however beyond the scope of the present paper.

Finally, if $k \geq 2$ descriptions are available for the signal reconstruction, we can proceed in a similar way. Denote by K the set of received k descriptions, and by r_K and D_K the corresponding reconstruction and distortion, respectively. The best signal approximation in a least-mean-squares sense r_K is obtained by grouping the generating matrices and coefficient vectors of the available descriptions i , with $1 \leq i \leq k$. Denote by Φ_K the set of k received matrices and by \widetilde{C}_K the corresponding set of received coefficients. The reconstruction can therefore be expressed as

$$r_K = (\Phi_K^\dagger \cdot \widetilde{C}_K)^T. \quad (13)$$

Since the matrix Φ_K has dimensions $kM \times M$, computing its pseudoinverse is quite involving. However, the computational complexity can be drastically reduced using the fact that $\Phi_K^\dagger = \Phi_K^T (\Phi_K \Phi_K^T)^{-1}$. Namely, instead of computing a pseudoinverse of Φ_K , we simply compute the inverse of $\Phi_K \Phi_K^T$ that is a symmetric $M \times M$ matrix.

The MSE distortion after signal reconstruction D_K again contains two components, the distortion due to the signal approximation D_K^a , and the distortion due to quantization D_K^q . The distortion due to the signal approximation can be written as

$$D_K^a = \frac{\|s\|^2 - \text{tr}(ss^T \Phi_K^T (\Phi_K \Phi_K^T)^{-1} \Phi_K)}{S}. \quad (14)$$

Similarly to the single-description case, it can be bounded as

$$D_K^a \leq \frac{\|s\|^2}{S} - \frac{\sum_{i=1}^{kM} C_i^2}{S(1 + \mu(kM - 1))}, \quad (15)$$

where we consider the worst-case scenario with any two atoms having a correlation μ that is the maximal correlation between any pair of atoms in the dictionary \mathcal{D} . The quantization distortion is given by

$$D_K^q = \frac{\Delta^2}{3S} \text{tr}(\Phi_K \Phi_K^T)^{-1}. \quad (16)$$

Under similar assumptions, it can be bounded by

$$D_K^q \leq \frac{M\Delta^2}{3S} \frac{1 + \mu(M - 2)}{(1 + \mu(M - 1))(1 - \mu)} \leq \frac{kM\Delta^2}{3S} \frac{1}{1 - \mu}. \quad (17)$$

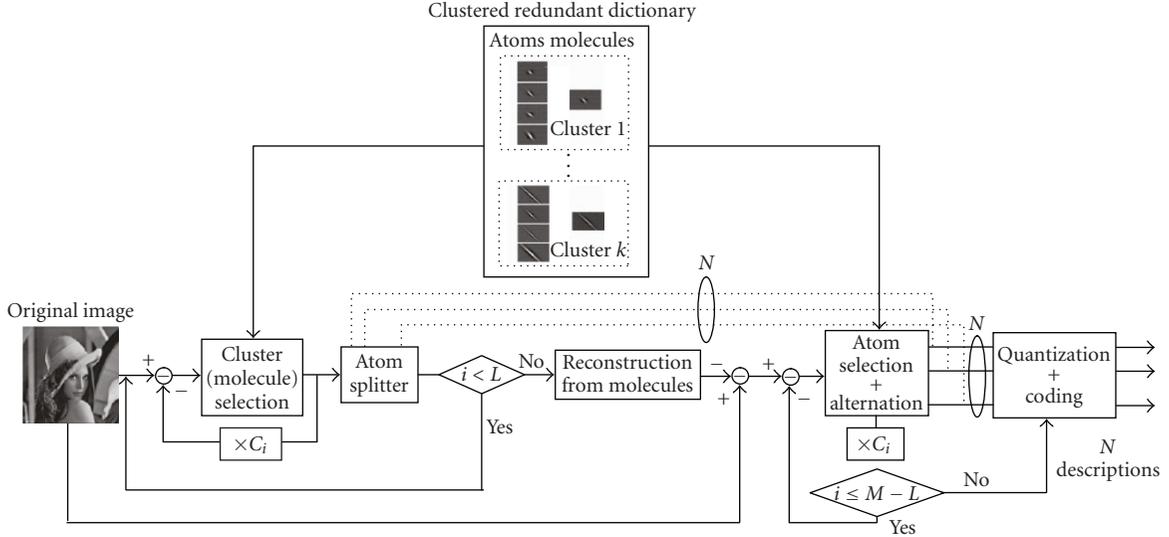


FIGURE 2: Block diagram of the multiple description image coding algorithm.

We can note that the distortion at reconstruction is clearly linked to the properties of the dictionary, as expected. In particular, partial and central distortions are influenced by the coherence within the dictionary, while the side distortion depends on the block coherence. The design of an optimal dictionary has therefore to trade off correlation within dictionary partition, and correlation between dictionary partitions. The compromise between side and central distortions is typical in multiple description coding, and the best working point depends on the quality of the communication channel. In the next section, we present an application of the above scheme to a typical image communication scenario.

4. MULTIPLE DESCRIPTION IMAGE CODING

4.1. Overview

This section proposes the application of multiple description coding with redundant dictionaries, to a typical image communication problem. The overall description of the algorithm is given in Figure 2. The redundant dictionary is partitioned into blocks of similar atoms, and each partition is represented by the molecules. The image is first decomposed into a series of L molecules, which are iteratively selected with a modified matching pursuit algorithm. The children atoms are distributed into the different descriptions. Each description is later refined by the addition of $M-L$ atoms. The residual signal, after subtraction of the approximation obtained with the molecules, is decomposed with a typical matching pursuit algorithm. The selected atoms are distributed in a round-robin fashion, to the different descriptions. Finally, coefficients are computed by projection of the signal on the set of atoms that compose each description. Eventually, they are uniformly quantized and entropy-coded along with the atom indexes, to form the final descriptions. The next subsections describe in more details the key parts of the multiple description image coding algorithm.

4.2. MDC with modified matching pursuit

Even if redundant dictionaries present interesting advantages for the approximation of multidimensional signals like images, searching for the sparsest (shortest) signal representation in a redundant dictionary of functions is in general an NP-hard problem [24]. Fortunately, it is usually sufficient to find a nearly optimal solution that would reduce the search complexity in a great manner, and very simple algorithms like matching pursuit [25] have been shown to provide very good approximation performance.

Matching pursuit is a simple greedy algorithm that iteratively decomposes any function s in the Hilbert space \mathcal{H} with atoms from a redundant dictionary. Let all the atoms, denoted by a_i , have a unit norm $\|a_i\|_2 = 1$ and let $\mathcal{D} = \{a_i\}$, $i = 1, 2, \dots, |\mathcal{D}|$. By setting $R_0 = s$, the signal is first decomposed as

$$R_0 = \langle a_0, R_0 \rangle a_0 + R_1, \quad (18)$$

where a_0 is chosen so as to maximize the correlation with R_0 :

$$a_0 = \arg \max_{\mathcal{D}} |\langle a_i, R_0 \rangle|, \quad (19)$$

and R_1 is the residual signal after the first iteration. The algorithm proceeds iteratively, by applying the same procedure to the residual signal. It can be shown that the energy of the residual after M iterations satisfies

$$\|R_M\|^2 = \|s\|^2 - \sum_{i=0}^{M-1} |\langle R_i, a_i \rangle|^2. \quad (20)$$

The approximation performance of matching pursuit is tightly linked to the structure of a dictionary, and it has been demonstrated that the norm of the residual after M iterations can be bounded by [26]

$$\|R_M\|^2 \leq (1 - \alpha^2 \beta^2) \|R_{M-1}\|^2 \leq (1 - \alpha^2 \beta^2)^M \|s\|^2, \quad (21)$$

where β is the structural redundancy defined in (1) and $\alpha \in (0, 1]$ is an optimality factor. This factor depends on the algorithm that searches for the best atom in the dictionary, at each iteration (e.g., $\alpha = 1$ for a full-search strategy). Matching pursuit represents a simple, flexible, yet efficient algorithm for signal expansion over redundant dictionaries. We therefore choose to use a modified matching pursuit algorithm to decompose the image in a series of molecules.

We propose to generate N descriptions by distributing similar, but not identical, atoms in different descriptions. As explained in the previous section, this can be achieved by computing the representation of the signal on the level of molecules, instead of the atoms themselves. The L molecules m_j , $j = 1, \dots, L$, that best approximate the signal s are selected by running matching pursuit on the set of molecules, which yields

$$s = \sum_{j=0}^{L-1} \langle R_j, m_j \rangle m_j + R_L. \quad (22)$$

The multiple descriptions are then built by distributing each atom from the blocks corresponding to these molecules, into different descriptions. Formally, if a molecule m_j is chosen in the j th stage of MP, we attribute its child a_i^j to description i , with $i = 1, 2, \dots, N$.

Redundant expansions offer the possibility of capturing most of the signal energy in a few atoms. That property is typically observed also for matching pursuit expansions, where the first selected atoms are the most important ones for the signal approximation (see (21)). In the same time, atoms that are selected after in later iterations only bring a small contribution to the signal reconstruction. We therefore propose to adopt a two-stage algorithm, where the first iterations are run on molecules, which capture most of the image energy. It offers us the possibility to put similar, and high energy atoms in the different descriptions. However, it may be wasteful to code with redundancy the molecules that only bring a small contribution. Therefore, the second stage of the encoding runs a classical matching pursuit algorithm on the atoms themselves, and distribute them in the different descriptions without any added redundancy. The most efficient joint source and channel coding schemes proceed by unequal error protection, and we basically pursue the same idea here.

After the L most significant molecules have been identified, a residual signal is built by subtracting the reconstructed signal with all the selected molecules, from the original image. A matching pursuit expansion of the residual signal is then performed on the level of atoms. The atoms are simply distributed alternatively between descriptions, to eventually generate descriptions with a total of M atoms. Upon completing both stages, the M atoms in description i are gathered in a generating matrix $\Phi_i = \{a_i^j\}$, with $j = 1, 2, \dots, M$, where the first L rows of Φ_i are children of the L selected molecules, and the remaining $M - L$ rows correspond to atoms that are alternatively distributed between descriptions. To generate description i , the signal is finally projected onto Φ_i , $C_i = \Phi_i s^T$. C_i s are uniformly quantized into \tilde{C}_i . Together with indices of atoms in Φ_i , \tilde{C}_i are attributed to description i .

Note finally that the choice of the number of molecules L depends on the transmission channel properties, and directly trades off the side and central distortions. We will see below how one can choose optimal L based on losses in the network.

4.3. Dictionary

A great amount of research has focused on the construction of “good” dictionaries. Some examples include spikes and sinusoids [27], wavelet packets [28], frames [29], or Gabor atoms [25], for example. We propose to use here an overcomplete dictionary composed of edge-like functions, as proposed in [21]. The structured dictionary is built on two mother functions. First, an isotropic Gaussian 2D function is responsible for efficient representation of the low-frequency characteristics of an image:

$$g_1(x, y) = \frac{1}{\sqrt{\pi}} e^{-(x^2+y^2)}. \quad (23)$$

The second mother function is an anisotropic function that consists of Gaussian along one direction and a second derivative of a Gaussian along another direction:

$$g_2(x, y) = \frac{2}{\sqrt{3}\pi} (4x^2 - 2) e^{-(x^2+y^2)}. \quad (24)$$

Such a shape is chosen in order to capture the contours that represent most of the content of natural images. Geometric transforms (translation, rotation, and scaling) are then applied to the mother functions to build a structured redundant dictionary. We allow the translation parameters to be any integers smaller than the image size. The scaling is isotropic and varies from 1/32 to 1/4 of the image size on a logarithmic scale with a resolution of one third of octave. As for the second function, we use the same translation parameters and the scaling parameters are uniformly distributed on a logarithmic scale from one to 1/8 of the image size, with a resolution of one third of octave. We also allow the rotation parameter to vary in increments of $\pi/18$.

The dictionary is finally partitioned into blocks of similar atoms, represented by molecules. In general, such partitions can be obtained by either a *top-down* or a *bottom-up* clustering approach. The former method tries to segment the initial dictionary into a number of subdictionaries, each of them consisting of atoms that satisfy some similarity constraints. Alternatively, the bottom-up approach groups the atoms as long as similarity constraints are satisfied. Since the bottom-up approach becomes rapidly complex when each cluster has to contain a fixed number N of atoms, we propose to use a top-down approach in this paper.

The top-down approach recursively segments our dictionary, to eventually generate a tree structure whose leaves are the atoms from \mathcal{D} . We use a top-down tree based pursuit algorithm [30], which implements a clustering strategy based on segmentation, where a fixed number N of similar atoms are grouped together. The trees were constructed using the *k-means* algorithm. Each of the nonleaf nodes in the tree is associated with the list of the atoms it represents. A

molecule can be computed as a simple weighted sum of the atoms it spans, taking into account the distance from the corresponding atoms. Different metrics can be used for the distance measure; one of the most popular ones is $d(a_i, a_j) = 1 - |\langle a_i, a_j \rangle|^2$. If the atoms are strongly correlated, their distance is close to 0, while in the case of orthogonal atoms this distance is 1.

4.4. Distortion model

We have previously derived the upper bounds on both reconstruction and quantization errors based on some dictionary properties as well as number of descriptions and number of atoms per description. However, since these bounds are computed in the worst-case scenario in terms of atom correlation, they are generally too loose in practical applications like image coding.

In order to define tighter bounds for the encoding scheme proposed above, we bound its behavior by the performance of a classical matching pursuit algorithm. Indeed, the signal reconstruction (13) leads to the best approximation in a least-squares sense, which is not necessarily the case in classical reconstructions with simple linear combinations of atoms selected by matching pursuit. Therefore, we can always bound the distortion due to our least-mean-squares approximation, by the matching pursuit distortion given in (21).

Finally, we can model the distortion due to signal approximation as the sum of two terms, corresponding to the two coding steps of the proposed scheme. The first one refers the distortion due to the approximation with L molecules, while the second one describes the distortion due to the refinement stage of $M - L$ atoms. We can approximate it in the following manner:

$$D_K^q = c + a \cdot b^L + c_K + a_K \cdot b_K^{k(M-L)}. \quad (25)$$

The shape of D_K^q fits the shape given by (21), up to an additive constant. The distortion decay is captured by terms a, b, c, a_K, b_K, c_K , that are chosen to fit best the real distortion values. Similarly, the quantization distortion is modelled as

$$D_K^q = k \cdot d_K \cdot \Delta^2. \quad (26)$$

This model keeps the shape of derived upper bounds in (12) and (17), up to multiplicative constants that are again chosen to fit the real quantization distortion values.

This distortion model can now be used to find the optimal number of molecules L and the optimal number of descriptions for a given communication channel, such that the average distortion is minimized. The average distortion D_{av} is given as

$$D_{av} = \sum_{|K|=0}^N \binom{N}{|K|} p^{N-|K|} (1-p)^{|K|} D_K, \quad (27)$$

where p is the channel loss probability and $D_\emptyset = \|s\|^2/S$. Figure 3 finally illustrates the model accuracy. It shows the minimal achievable average distortion for three descriptions for loss probabilities of $p \in [10^{-4}, 0.05]$. We can see that the model provides a very good approximation of the actual distortion values.

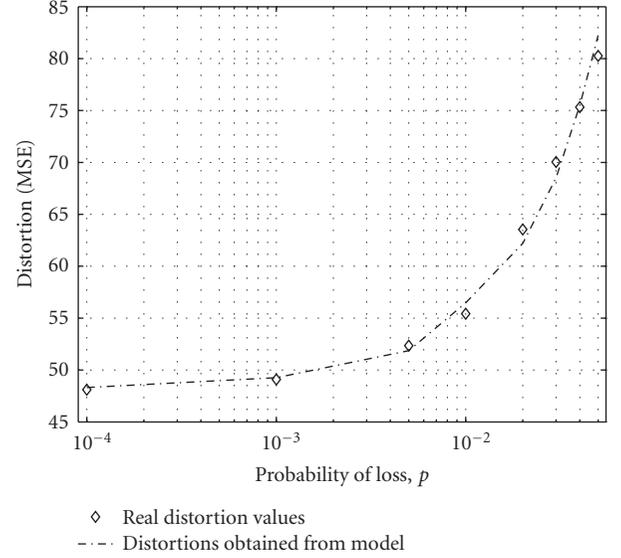


FIGURE 3: Minimal achievable average distortions for the case of three descriptions: real values versus model.

5. SIMULATION RESULTS

5.1. Settings

This section analyzes the performance of the proposed coding scheme, in typical image communication scenario. We assume that each description corresponds to one packet, and therefore is either received error-free or completely lost. We show the results for Lena and Peppers images, both of size 128×128 , obtained by averaging over 1000 simulations of random packet losses. The distortion of the reconstructed signal is the mean square error (MSE). Finally, we do not implement any concealment or postfiltering strategy at the decoder.

We first show the behavior of the proposed scheme as a function of number of descriptions and network losses. We then analyze in more details the performance of our scheme in the case where the number of descriptions is limited to 2 and, respectively, 3 descriptions. We compare these performances to two MDC schemes that implement simple atom repetition [17], and unequal error protection (UEP) [31]. These two schemes are illustrated in Figure 4. The atom sharing scheme repeats a certain number of the most important atoms a_i in all the descriptions, while the remaining atoms are alternatively split between descriptions. On the other side, FEC scheme applies a systematic code, column-wise across the N -packet block. Here, atoms are protected according to their importance.

Finally, we analyze the performance of our scheme compared to an MDC scheme based on unequal error protection, when the number of descriptions can be optimized with respect to the transmission channel characteristics. Overall, the results demonstrate that the proposed scheme is competitive with state-of-the-art MDC schemes that are able to generate any number of descriptions. Moreover, the proposed

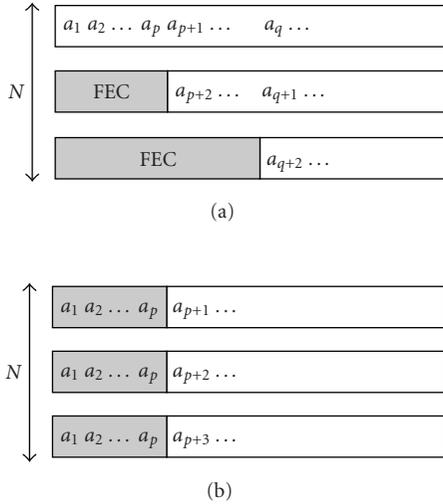


FIGURE 4: (a) FEC scheme and (b) atom sharing scheme.

scheme is less sensitive to bad estimation of the loss probability, which clearly penalizes optimized unequal error protection schemes.

5.2. Optimal number of descriptions

In the first experiment, we observe the behavior of the proposed MDC scheme, when the overall bit rate is fixed and the number of descriptions varies. We fix the total number of atoms to 600 and vary the number of descriptions between 2 and 4, as well as the number of atoms per description. We use 11 bits to code the atom indexes, and all the coefficients are quantized uniformly with the step size 1, which results in the total rate of 1.35 kB. We choose the optimal number of molecules L in each of the cases, in such a way that the average distortion is minimized. The minimal achievable average distortions are computed as a function of packet loss probability p , where $p \in [10^{-4}, 0.05]$. The results are illustrated in Figure 5.

When the losses are very low (i.e., $p < 10^{-3}$), a small number of descriptions are generally the best choice, as they allow for efficient redundancy and good approximation performance since the number of closely related atoms is small. As the losses increase, the optimal number of descriptions also augments, as expected. However, the significant difference in performance can only be observed when the loss rate exceeds 1%. At a loss rate of 5%, four descriptions improve the performance of 1.7 dB, respectively, 0.2 dB, with respect to the cases with 2 and 3 descriptions only. Note that similar observations have already been reported in other MDC schemes (e.g., [32, 33]). It confirms that the case of two descriptions, which is the most frequently studied, is not necessarily optimal, and that the ability to generate more descriptions is certainly beneficial at high loss rates. Finally, we can conjecture that in realistic cases, building more than four descriptions only brings negligible improvements, and this is the limit we will use in our simulations.

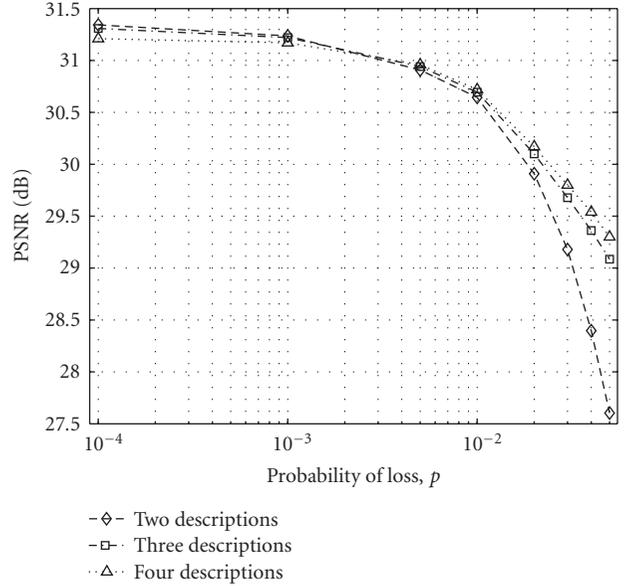


FIGURE 5: Comparison of minimal achievable distortions for two, three, and four descriptions, when the total rate is fixed (Lena image).

5.3. Two descriptions

We now compare the performance of our scheme for $N = 2$ descriptions with other MDC strategies (when $N = 2$, the UEP scheme is equivalent to the atom sharing scheme). We first observe the evolution of the minimal achievable average distortion with respect to the packet loss probability p . Similar to the previous experiments, we build descriptions with $M = 300$ atoms, of 18 bits each (i.e., the total bit rate is again around 1.35 kB). The number of shared atoms in the atom sharing scheme, and the number of molecules L in the proposed scheme are optimized. The results are shown in Figure 6. We can see that our scheme provides improvement of up to 0.6 dB compared to the atom sharing (and UEP) scheme. This is due to the fact that our scheme takes advantage from all the received atoms, while the existing schemes cannot use the redundant atoms, which are a waste of resources when no loss occurs.

Next, we compare both schemes optimized for a given loss ratio p , but when the actual channel characteristics are somewhat different (as it may happen in practical scenarios when channel status changes). Figure 7 shows the performance of both schemes optimized for $p = 10^{-3}$, while the actual loss probability covers the range $[10^{-4}, 0.1]$. We can see that our scheme always gives better results and the improvement is up to 1.4 dB. While the atom sharing scheme seems to work well in the very narrow range around the loss, it is optimized for our scheme tends to be more robust in much wider range of losses, and thus more resilient to bad estimation of the channel characteristics.

We finally observe the images reconstructed with different numbers of descriptions. Both encoding schemes have been optimized for $p = 10^{-3}$, and a total rate of 1.35 kB. The

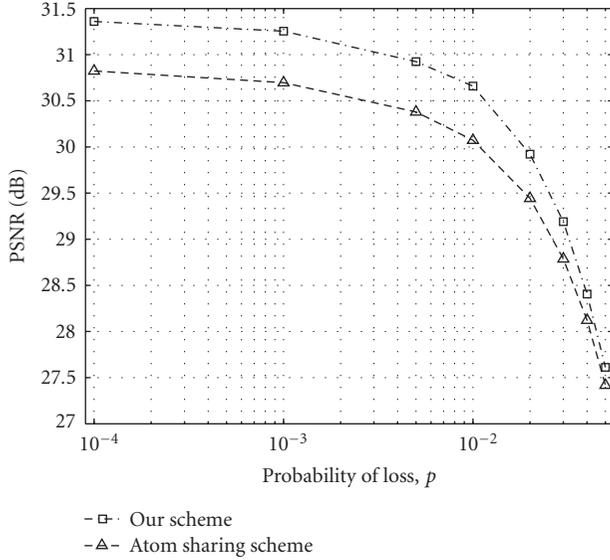


FIGURE 6: PSNR versus loss probability for the proposed scheme, and the atom sharing scheme, optimized for two description and a total rate of 1.35 kB (Lena image).

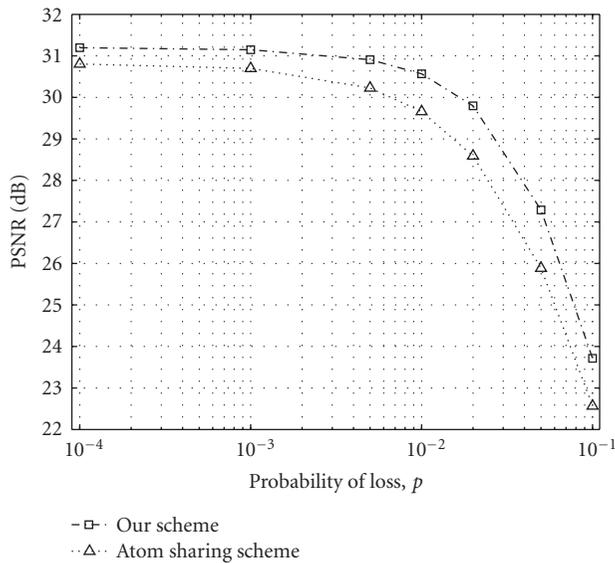


FIGURE 7: PSNR versus actual loss probability, for the proposed scheme, and the atom sharing scheme, optimized for two descriptions and a total rate of 1.35 kB, and a loss probability of 10^{-3} (Lena image).

images are given in Figure 8, for our scheme, and the atom sharing scheme. We can observe that the side reconstruction is better for the proposed MDC scheme (i.e., 3.5 dB improvement), while the central reconstruction gives an improvement of 0.4 dB. The difference in side distortion is mostly due to the fact that the number of repeated atoms is very small in the atoms sharing scheme optimized for low loss probability ($p = 10^{-3}$). Better central distortion is expected, since the important atoms are not repeated in our scheme, and cor-



FIGURE 8: Reconstructed Lena images, as a function of a number of received descriptions, from 1 description on the left column, to 2 descriptions on the right column. (Top row: our scheme, Bottom row: atom sharing scheme.)

related, yet different atoms bring more information for the reconstruction.

5.4. Three descriptions

We now consider the case of $N = 3$ descriptions, and propose a similar analysis as above. The minimal average distortion as a function of p for the proposed scheme, an MDC scheme based on atom sharing, and an unequal error protection scheme is given in Figures 9 and 10 for the Lena and Peppers images, respectively. We see that our scheme outperforms the existing schemes in a wide range of losses, especially at low packet loss ratios, where the advantage in the central distortion becomes predominant (i.e., improvement of about 0.6 dB in the case of Lena). As the losses exceed 2%, the FEC scheme tends to slightly outperform our scheme, and at $p = 5\%$ the improvement reaches almost 0.3 to 0.5 dB. This can be explained by the fact that the FEC scheme protects different atoms according to their importance, and therefore is more flexible to protect the strongest atoms, which is beneficial at high loss rate. It is also interesting to notice that the FEC and atom sharing scheme perform similarly at low losses, while there is an increasing gain in favor of FEC scheme as the loss ratio increased, since redundancy is allocated more efficiently with an unequal error protection strategy.

Figures 11 and 12 show the behavior of the three schemes, when the actual loss probability is different from the expected one. The schemes have all been optimized for a loss

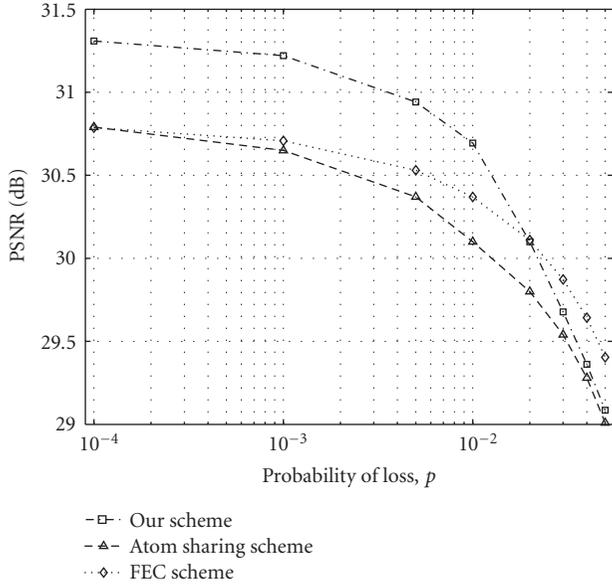


FIGURE 9: PSNR versus loss probability, for the proposed scheme, the UEP FEC scheme, and the atom sharing scheme, optimized for three descriptions and a total rate of 1.35 kB (Lena image).

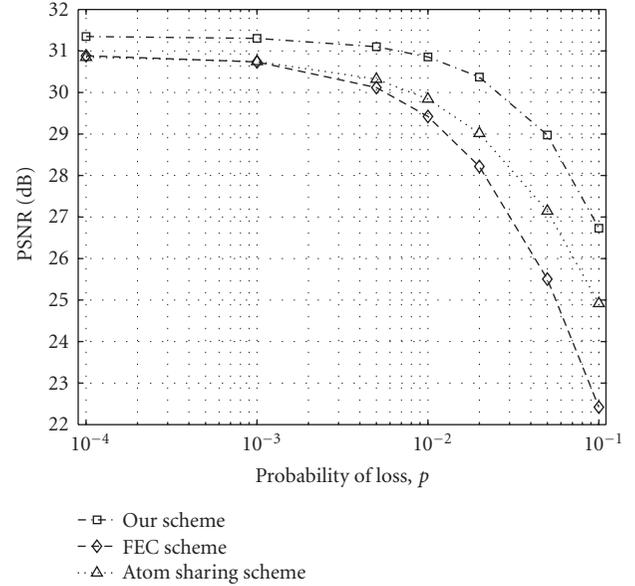


FIGURE 11: PSNR versus actual loss probability, for the proposed scheme, the UEP FEC scheme, and the atom sharing scheme, optimized for three descriptions and a total rate of 1.35 kB, and a loss probability of 10^{-3} (Lena image).

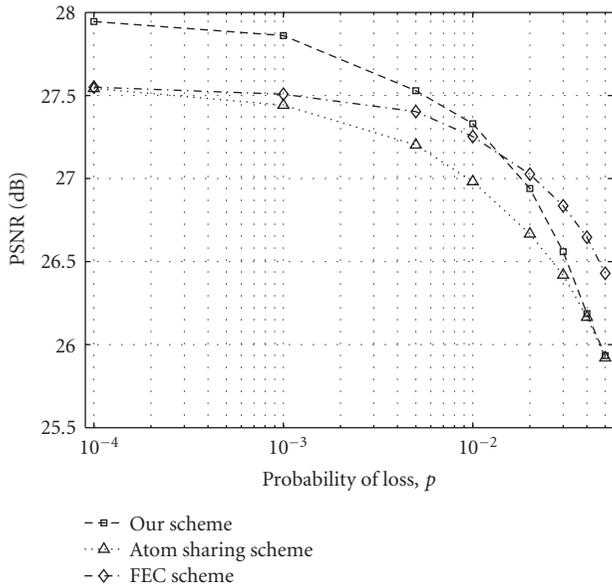


FIGURE 10: PSNR versus loss probability, for the proposed scheme, the UEP FEC scheme, and the atom sharing scheme, optimized for three descriptions and a total rate of 1.35 kB (Peppers image).

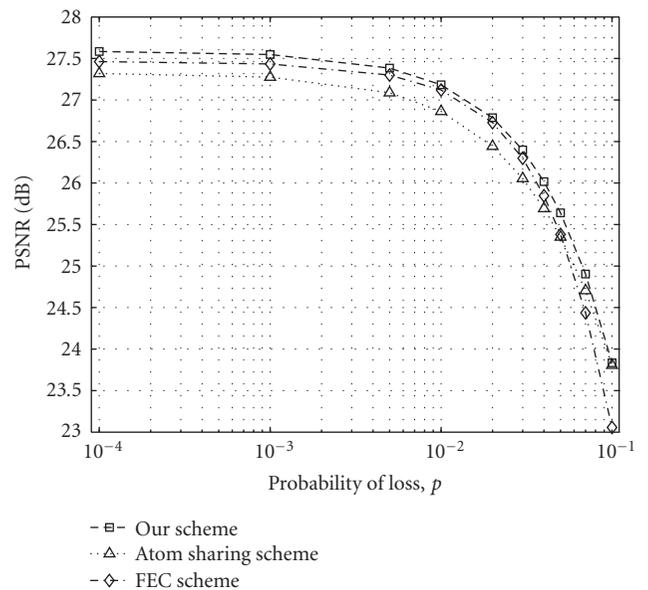


FIGURE 12: PSNR versus actual loss probability, for the proposed scheme, the UEP FEC scheme, and the atom sharing scheme, optimized for three descriptions and a total rate of 1.35 kB, and a loss probability of $5 \cdot 10^{-3}$ (Peppers image).

probability of $p = 10^{-3}$, respectively, $p = 5 \cdot 10^{-3}$, and we compute the average distortion when the actual loss probability varies. It can be seen again that the FEC scheme works well in a very narrow range of losses. Namely, when the loss probability increases, the FEC scheme becomes very vulnerable, giving the sharpest decrease in quality out of all compared schemes. It is also interesting to note that even if the

atom sharing scheme performs worse than the FEC scheme in average, it tends to be more robust to changing channel characteristics. Our scheme is the most resilient to such changes, and this is mostly visible at high loss rates (i.e., up to 1.8 dB improvement, resp., 4.3 dB compared to the atom sharing scheme, and the FEC scheme, for the Lena image).



FIGURE 13: Reconstructed Lena images, as a function of a number of received descriptions, from 1 description on the left, to 3 descriptions on the right column.

Finally, we represent the decoded images, reconstructed with different numbers of descriptions, for the three schemes that have been optimized for loss probabilities $p = 10^{-3}$ and $p = 5 \cdot 10^{-3}$, respectively. Figures 13 and 14 illustrate this comparison for two images, respectively, Lena and Peppers. We observe that since the UEP-based FEC is clearly optimized for low loss probability, the side and even partial distortions are generally quite unimportant. On the other side, the proposed scheme and the atom sharing scheme are more conservative in allocation of redundancy, and therefore more resilient to changes in the actual loss probability. Finally, we

can observe that our scheme, as expected, always performs best when all descriptions are available, since it does not send pure redundancy for important components, but rather correlated information that still improve the central distortion.

5.5. Improved FEC reconstruction

We have considered so far comparisons with state-of-the-art schemes that use ordinary reconstruction strategy based on a simple linear combination of the atoms available at decoder. The reconstruction can however be improved in the



FIGURE 14: Reconstructed Peppers images, as a function of a number of received descriptions, from 1 description on the left, to 3 descriptions on the right column.

case of MDC based on UEP protection, by using a similar projection method as in the MDC scheme proposed in this paper. This projection method will then optimize, in a least-mean-square sense, the approximation that can be constructed from the available atoms. For the sake of completeness, we provide here a comparison between the proposed MDC scheme and an FEC scheme whose reconstruction is improved by the projection method. We keep the same simulations settings as before, with a total rate equivalent to 600 atoms, and we vary the number of descriptions and number of FEC packets in order to reach the optimal working point for different channel loss probabilities. Results are depicted in Figure 15 for the Lena image. While the reconstruction is

slightly improved in the FEC scheme, the performance does not change significantly. We can see that our scheme still provides better results in the range of losses $[10^{-4}, 10^{-2}]$, mainly due to an improved central distortion. As the losses increase, the FEC scheme will tend to perform better, since it provides a high protection to the most important components, and tends to provide a slightly better side distortion, when optimized for high loss rates.

6. CONCLUSIONS

This paper has presented a multiple description coding scheme, which exploits the redundancy present in redundant

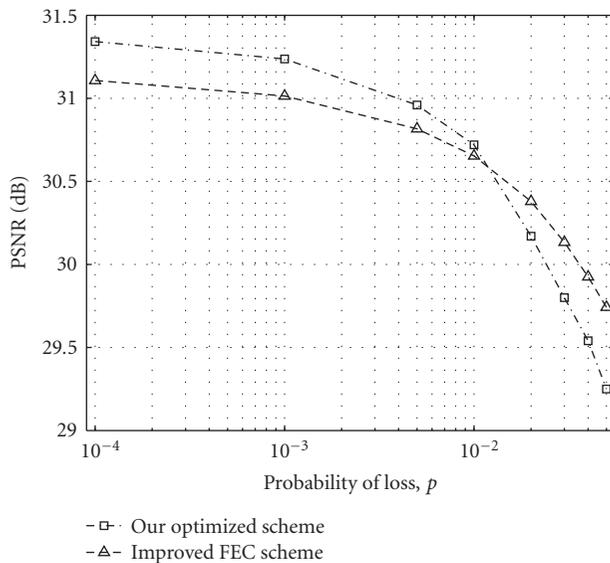


FIGURE 15: PSNR quality versus loss probability, for the proposed scheme and an FEC scheme that uses the projection method (Lena image).

dictionaries. Instead of repeating signal components, or adding pure redundancy to the signal decomposition, redundant transforms with partitioned dictionaries allow to control the correlation between the description, and put different, yet correlated atoms, in different descriptions. This allows for improving the central distortion, as descriptions nicely complement each other, without important penalty on the side distortion. Besides its flexibility, the proposed scheme presents the advantage by allowing for the generation of an arbitrary number N of balanced descriptions, while most of the schemes are generally limited to 2 descriptions.

The application of the new multiple description coding scheme to a typical image communication scenario demonstrates that it outperforms other MDC schemes based on atom repetition, or unequal FEC protection, especially for low-loss probability. In addition, the proposed scheme presents an increased resilience to wrong estimation of the communication channel characteristics, while unequal error protection schemes are very sensitive to differences between expected and actual loss probabilities. While the proposed method can interestingly be implemented for any generic dictionary, the definition of optimally distributed dictionaries, for typical MDC scenarios, is still under investigation.

ACKNOWLEDGMENT

This work has been supported by the Swiss National Science Foundation Grant PP-002-68737.

REFERENCES

- [1] I. Radulovic and P. Frossard, "Multiple description image coding with block-coherent redundant dictionaries," in *Proceedings of Picture Coding Symposium*, Beijing, China, April 2006.
- [2] N. S. Jayant, "Subsampling of a DPCM speech channel to provide two 'self-contained' half-rate channels," *The Bell System Technical Journal*, vol. 60, no. 4, pp. 501–509, 1981.
- [3] W. Jiang and A. Ortega, "Multiple description coding via polyphase transform and selective quantization," in *Visual Communications and Image Processing*, vol. 3653 of *Proceedings of SPIE*, pp. 998–1008, San Jose, Calif, USA, January 1999.
- [4] J. G. Apostolopoulos, "Reliable video communication over lossy packet networks using multiple state encoding and path diversity," in *Visual Communications and Image Processing*, vol. 4310 of *Proceedings of SPIE*, pp. 392–409, San Jose, Calif, USA, January 2001.
- [5] V. A. Vaishampayan, "Design of multiple description scalar quantizers," *IEEE Transactions on Information Theory*, vol. 39, no. 3, pp. 821–834, 1993.
- [6] V. A. Vaishampayan, N. J. A. Sloane, and S. D. Servetto, "Multiple-description vector quantization with lattice codebooks: design and analysis," *IEEE Transactions on Information Theory*, vol. 47, no. 5, pp. 1718–1734, 2001.
- [7] S. D. Servetto, K. Ramchandran, V. A. Vaishampayan, and K. Nahrstedt, "Multiple description wavelet based image coding," *IEEE Transactions on Image Processing*, vol. 9, no. 5, pp. 813–826, 2000.
- [8] M. Srinivasan and R. Chellappa, "Multiple description subband coding," in *Proceedings of IEEE International Conference on Image Processing (ICIP '98)*, vol. 1, pp. 684–688, Chicago, Ill, USA, October 1998.
- [9] A. Jagmohan, A. Sehgal, and N. Ahuja, "Two-channel predictive multiple description coding," in *Proceedings of IEEE International Conference on Image Processing (ICIP '05)*, vol. 2, pp. 670–673, Genova, Italy, September 2005.
- [10] V. K. Goyal and J. Kovačević, "Generalized multiple description coding with correlating transforms," *IEEE Transactions on Information Theory*, vol. 47, no. 6, pp. 2199–2224, 2001.
- [11] Y. Wang, M. T. Orchard, and A. R. Reibman, "Multiple description image coding for noisy channels by pairing transform coefficients," in *Proceedings of 1st IEEE Workshop on Multimedia Signal Processing*, pp. 419–424, Princeton, NJ, USA, June 1997.
- [12] V. K. Goyal, J. Kovačević, R. Arean, and M. Vetterli, "Multiple description transform coding of images," in *Proceedings of IEEE International Conference on Image Processing (ICIP '98)*, vol. 1, pp. 674–678, Chicago, Ill, USA, October 1998.
- [13] V. K. Goyal, J. Kovačević, and J. A. Kelner, "Quantized frame expansions with erasures," *Applied and Computational Harmonic Analysis*, vol. 10, no. 3, pp. 203–233, 2001.
- [14] S. S. Channappayya, J. Lee, R. W. Heath Jr., and A. C. Bovik, "Frame based multiple description image coding in the wavelet domain," in *Proceedings of IEEE International Conference on Image Processing (ICIP '05)*, vol. 3, pp. 920–923, Genova, Italy, September 2005.
- [15] T. Petrișor, C. Tillier, B. Pesquet-Popescu, and J.-C. Pesquet, "Comparison of redundant wavelet schemes for multiple description coding of video sequences," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP '05)*, vol. 5, pp. 913–916, Philadelphia, Pa, USA, March 2005.
- [16] T. Petrișor, C. Tillier, B. Pesquet-Popescu, and J.-C. Pesquet, "Redundant multiresolution analysis for multiple description video coding," in *Proceedings of 6th IEEE Workshop on Multimedia Signal Processing (MMSP '04)*, pp. 95–98, Siena, Italy, September–October 2004.
- [17] X. Tang and A. Zakhor, "Matching pursuits multiple description coding for wireless video," *IEEE Transactions on Circuits*

- and Systems for Video Technology*, vol. 12, no. 6, pp. 566–575, 2002.
- [18] T. Nguyen and A. Zakhor, “Matching pursuits based multiple description video coding for lossy environments,” in *Proceedings of IEEE International Conference on Image Processing (ICIP '03)*, vol. 1, pp. 57–60, Barcelona, Spain, September 2003.
- [19] H.-T. Chan, C.-M. Fu, and C.-L. Huang, “A new error resilient video coding using matching pursuit and multiple description coding,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 15, no. 8, pp. 1047–1052, 2005.
- [20] G. Karabulut and A. Yongacoglu, “Multiple description coding using orthogonal matching pursuit,” in *Proceedings of the 3rd Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net '04)*, pp. 529–534, Bodrum, Turkey, June 2004.
- [21] R. M. Figueras i Ventura, P. Vandergheynst, and P. Frossard, “Low-rate and flexible image coding with redundant representations,” *IEEE Transactions on Image Processing*, vol. 15, no. 3, pp. 726–739, 2006.
- [22] P. Jost, P. Vandergheynst, and P. Frossard, “Tree-based pursuit: algorithm and properties,” *IEEE Transactions on Signal Processing*, vol. 54, no. 12, pp. 4685–4697, 2006.
- [23] G. Rath and C. Guillemot, “Frame-theoretic analysis of DFT codes with erasures,” *IEEE Transactions on Signal Processing*, vol. 52, no. 2, pp. 447–460, 2004.
- [24] G. Davis, S. Mallat, and M. Avellaneda, “Adaptive greedy approximations,” *Constructive Approximation*, vol. 13, no. 1, pp. 57–98, 1997.
- [25] S. G. Mallat and Z. Zhang, “Matching pursuits with time-frequency dictionaries,” *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [26] S. G. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, San Diego, Calif, USA, 2nd edition, 1999.
- [27] S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic decomposition by basis pursuit,” *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 33–61, 1999.
- [28] R. R. Coifman and M. V. Wickerhauser, “Entropy-based algorithms for best basis selection,” *IEEE Transactions on Information Theory*, vol. 38, no. 2, part II, pp. 713–718, 1992.
- [29] I. Daubechies, *Ten Lectures on Wavelets*, SIAM, Philadelphia, Pa, USA, 1992.
- [30] P. Jost, P. Vandergheynst, and P. Frossard, “Tree-based pursuit,” Tech. Rep. TR-ITS-2004.13, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, July 2004.
- [31] P. Frossard and P. Vandergheynst, “Unequal error protection of atomic image streams,” Tech. Rep. TR-ITS-2005.007, Signal Processing Institute, Lausanne, Switzerland, January 2005.
- [32] I. Radulovic and P. Frossard, “Fast index assignment for balanced N-description scalar quantization,” in *Proceedings of Data Compression Conference (DCC '05)*, p. 474, Snowbird, Utah, USA, March 2005.
- [33] J. Østergaard, J. Jensen, and R. Heusdens, “n-Channel symmetric multiple-description lattice vector quantization,” in *Proceedings of Data Compression Conference (DCC '05)*, pp. 378–387, Snowbird, Utah, USA, March 2005.