

# A *Posteriori* Quantization of Progressive Matching Pursuit Streams

Pascal Frossard, *Member, IEEE*, Pierre Vandergheynst, *Member, IEEE*,  
Rosa Maria Figueras i Ventura, *Student Member, IEEE*, and Murat Kunt, *Fellow, IEEE*

**Abstract**—This paper proposes a rate-distortion optimal *a posteriori* quantization scheme for matching pursuit (MP) coefficients. The *a posteriori* quantization applies to an MP expansion that has been generated offline and cannot benefit of any feedback loop to the encoder in order to compensate for the quantization noise. The redundancy of the MP dictionary provides an indicator of the relative importance of coefficients and atom indices and, subsequently, on the quantization error. It is used to define a universal upper bound on the decay of the coefficients, sorted in decreasing order of magnitude. A new quantization scheme is then derived, where this bound is used as an Oracle for the design of an optimal *a posteriori* quantizer. The latter turns the exponentially distributed coefficient entropy-constrained quantization problem into a simple uniform quantization problem. Using simulations with random dictionaries, we show that the proposed exponentially upper bounded quantization (EUQ) clearly outperforms classical schemes. Stepping on the ideal Oracle-based approach, a suboptimal adaptive scheme is then designed that approximates the EUQ but still outperforms competing quantization methods in terms of rate-distortion characteristics. Finally, the proposed quantization method is studied in the context of image coding. It performs similarly to state-of-the-art coding methods (and even better at low rates) while interestingly providing a progressive stream that is very easy to transcode and adapt to changing rate constraints.

**Index Terms**—Compression, image coding, matching pursuit, progressive stream, quantization, redundancy.

## I. INTRODUCTION

COMPRESSION has now reached a turning point where new algorithms have to meet simultaneous new constraints as robustness, adaptivity, and sparsity of the representation. In this context, nonorthogonal transforms present several interesting properties that position them as an interesting alternative to orthogonal transforms like the discrete cosine transform (DCT) or wavelet-based schemes. Decomposing a signal over a redundant dictionary improves the compression efficiency, especially at low bit rate, where most of the signal energy is captured by only few elements. The main limitation of nonorthogonal transforms is, however, the encoding complexity, since the number of possible decompositions becomes infinite. Matching pursuit (MP) algorithms [1] provide an interesting way to iteratively decompose the signal

in its most important features with a limited complexity. They output a stream composed of both atoms or basis functions (more practically a reference index in a dictionary of functions) and their respective coefficients.

The aim of this paper is first to study the effects of quantization onto reconstruction of MP streams. Since the MP coefficients generally take on real values, quantization is necessary to reduce the bandwidth needed to transmit them. Quantization errors have been studied in [2]–[4] in the context of overcomplete frame expansions and consistent MP. This new study focuses on the transmission of multiresolution MP streams with a *posteriori* quantization. A *posteriori* quantization refers to a scheme that codes an MP expansion that has been generated offline. On the contrary, an *a priori* quantization scheme is included in the MP algorithm itself, which can thus compensate for quantization errors. The resulting stream is in this case targeted for a particular bit rate. In some practical cases, a genuine stream is however computed once and then quantized several times to satisfy possibly different rate constraints. This particular choice is driven by applied considerations: When using MP for coding images or videos, the decomposition itself is a bottleneck. The application simply cannot afford the price of running MP several times with different quantizer settings. Contrary to in-loop quantization schemes [5], [6], where the encoder uses the quantized coefficient to update the residual signal, the *a posteriori* quantized version of the coefficients indeed does not influence the MP expansion.

As usual, in overcomplete expansions, the set of functions that form the dictionary plays a crucial role in MP coding. A very redundant dictionary generally allows to capture the main features of the signal with only a few dictionary elements. However, the coding rate of the function parameters obviously increases with the size of the dictionary. Meanwhile, the relative importance of each component of the decomposition also directly depends on the dictionary. The chances to find a function that closely fits the input signal, and thus most of its energy, grows with the dictionary size. The redundancy of the dictionary therefore leads the energy decay rate of the residual signal, which has been proven to be upper bounded by an exponential curve [7], [8] in MP decompositions. The contribution of each MP coefficient therefore clearly depends on its position within the signal representation.

Based on the characterization of the energy decay curve, an exponentially bounded quantization (EUQ) scheme is proposed for the MP coefficients, which are sorted in decreasing order of magnitude. The EUQ scheme turns the exponentially distributed coefficient entropy-constrained quantization problem

Manuscript received October 17, 2002; revised May 5, 2003. The associate editor coordinating the review of this paper and approving it for publication was Dr. Vikram Krishnamurthy.

The authors are with the Signal Processing Institute, Swiss Federal Institute of Technology (EPFL), 1015 Lausanne, Switzerland.

Digital Object Identifier 10.1109/TSP.2003.821105

into a simple uniform quantization problem by dynamically reducing the coefficient quantization range along with increasing iteration numbers. This theoretically optimal scheme is shown to outperform previously proposed MP quantizers, especially at low bit rates. We refer to this scheme as *Oracle based* because it assumes that the upper bound is given for the dictionary in use. Unfortunately, the redundancy factor can be very difficult to compute in practical situations, i.e., for large dictionaries. We thus introduce a modified scheme that, although suboptimal, still achieves very good performances. This scheme adapts to the actual MP coefficient value to provide the quantization algorithm with heuristics parameters. It is shown to favorably compare with other quantization algorithms on random signals. Finally, in the practical case of image coding, it provides results similar to state-of-the-art coders [9], while additionally providing the advantage of generating a progressive stream, where coefficients are sent in decreasing order of magnitude. This worthy feature enables a very easy transcoding to adapt to changing rate constraints.

The paper is organized as follows: Section II first overviews the MP algorithm and its convergence properties. Section III studies the *a posteriori* quantization of MP coefficients and proposes a theoretically optimal exponentially upper bounded quantization (EUQ) algorithm that benefits from the properties of the encoding. Section IV builds on the optimal scheme to design an adaptive, but suboptimal, quantizer, which performs efficiently in practical algorithms. Section V provides comparisons between the adaptive quantization scheme and state-of-the-art coders, in the case of random signals, as well as natural images. Finally, concluding remarks are given in Section VI.

## II. MP EXPANSIONS

### A. MP Overview

In contrast to orthogonal transforms, overcomplete expansions of signals are not unique. The number of feasible decompositions is infinite, and finding the best solution under a given criteria is a NP-complete problem. In compression, one is interested in representing the signal to be coded with the smallest number of elements, that is, in finding the solution with most of the energy on only a few coefficients. MP is one of the suboptimal approaches that greedily approximates the solution to this NP-complete problem.

MP is an adaptive algorithm that iteratively decomposes any function  $f$  in the Hilbert space  $\mathcal{H}$  in a possibly redundant dictionary of functions called *atoms* [1]. Let  $\mathcal{D} = \{g_\gamma\}_{\gamma \in \Gamma}$  be such a dictionary with  $\|g_\gamma\| = 1$ , and let  $\Gamma$  represent the set of possible indices. The function  $f$  is first decomposed as follows:

$$f = \langle g_{\gamma_0} | f \rangle g_{\gamma_0} + \mathcal{R}f \quad (1)$$

where  $\langle g_{\gamma_0} | f \rangle g_{\gamma_0}$  represents the projection of  $f$  onto  $g_{\gamma_0}$ , and  $\mathcal{R}f$  is a residual component. Since all elements in  $\mathcal{D}$  have by definition a unit norm, it is easy to see from (1) that  $g_{\gamma_0}$  is orthogonal to  $\mathcal{R}f$ , and this leads to

$$\|f\|^2 = |\langle g_{\gamma_0} | f \rangle|^2 + \|\mathcal{R}f\|^2. \quad (2)$$

To minimize  $\|\mathcal{R}f\|$ , one must choose  $g_{\gamma_0}$  such that the projection coefficient  $|\langle g_{\gamma_0} | f \rangle|$  is maximum. The pursuit is carried out by applying iteratively the same strategy to the residual component. After  $N$  iterations, one has the following decomposition for  $f$ :

$$f = \sum_{n=0}^{N-1} \langle g_{\gamma_n} | \mathcal{R}^n f \rangle g_{\gamma_n} + \mathcal{R}^N f \quad (3)$$

where  $\mathcal{R}^N$  is the residual of the  $N^{\text{th}}$  step with  $\mathcal{R}^0 f = f$ . Similarly, the energy  $\|f\|^2$  is decomposed into

$$\|f\|^2 = \sum_{n=0}^{N-1} |\langle g_{\gamma_n} | \mathcal{R}^n f \rangle|^2 + \|\mathcal{R}^N f\|^2. \quad (4)$$

When the dictionary is complete, the original function  $f$  can be exactly decomposed as

$$f = \sum_{n=0}^{+\infty} \langle g_{\gamma_n} | \mathcal{R}^n f \rangle g_{\gamma_n}. \quad (5)$$

Designing complete dictionaries is a critical task in the MP-based algorithm [10]. Although MP places very few restrictions on the dictionary, the latter is strongly related to convergence speed and, thus, to coding efficiency. In this paper, convergence speed has to be understood as the ability for MP to capture most of the input signal energy in just a few iterations. Any collection of arbitrarily sized and shaped functions can generally be used as dictionary, as long as completeness is respected, but the convergence of the MP expansion depends on the signal characteristics and the size and structure of the dictionary.

### B. Convergence of MP

The convergence speed of MP corresponds to its ability to extract the maximum signal energy in a few iterations. In other words, it corresponds to the decay rate of the residue. Since the convergence speed directly depends on the dictionary, the decay rate of the residual energy can be bounded once the dictionary is known, even without *a priori* information about the input signal.

The approximation error decay rate in MP has been shown to be bounded by an exponential [7], [8]. In other words, the decay of the residue norm is faster than an exponential decay curve, whose rate depends on the dictionary only. From [7], there exists a decay parameter  $\lambda > 0$  such that, for all  $N \geq 0$

$$\|\mathcal{R}^N f\| \leq 2^{-\lambda N} \|f\| \quad (6)$$

or, equivalently

$$\|\mathcal{R}^{m+1} f\| \leq 2^{-\lambda} \|\mathcal{R}^m f\|, \forall m. \quad (7)$$

The decay rate<sup>1</sup> can be written as [7]

$$2^{-\lambda} = (1 - \alpha^2 \beta^2)^{1/2} \quad (8)$$

<sup>1</sup>The decay parameter  $\lambda$  decreases when the size of the signal space increases. However, at the limit of infinite dimensional spaces, the convergence is no longer exponential [11].

where  $\beta$  is the redundancy factor, and  $\alpha \in (0, 1]$  represents an optimality factor. This latter factor depends on the algorithm that, at each iteration, searches for the best atom in the dictionary. The optimality factor  $\alpha$  is set to one when MP browses the complete dictionary at each iteration. The parameter  $\beta$  depends on the dictionary construction. It represents the ability of the dictionary to capture features of any input function  $f$ . The upper bound on the coefficient norm is reached in the worst case, where the input function is the farthest from any dictionary vector. Hence,  $\beta$  is defined by

$$\sup_{\gamma_m} |\langle f, g_{\gamma_m} \rangle| \geq \beta \|f\|. \quad (9)$$

It can be interpreted as the cosine of the maximum angle between any direction in  $\mathcal{H}$  and the closest element of the dictionary, which we simply assume to be a  $L$ -dimensional subspace of  $\mathcal{H}$  in the finite-dimensional case [1]. This parameter characterizes the redundancy of the dictionary and tends to one when the size of the complete dictionary increases. A geometrical formulation of  $\beta$  is provided in [12].

Finally, the parameter  $\beta$  characterizes the redundancy of the dictionary for uniformly distributed atoms, independently from the function  $f$  to approximate. At the same time,  $\beta$  characterizes the importance of the atom itself in carrying information. Indeed, the importance of the atom index increases with the redundancy of the dictionary (i.e., the  $\beta$  value), whereas the relative importance of the coefficient decreases. In contrast, most of the information is carried by the coefficients in orthogonal transforms.

### III. QUANTIZATION OF MP COEFFICIENTS

#### A. A *Posteriori* Coefficients Quantization

Recall that the interest of the *a posteriori* quantization scheme lies in the context of asymmetric applications. Due to the encoding complexity, the MP stream is computed only once and then differently quantized to match various receiver requirements. An optimal coefficient quantization cannot be included in the MP decomposition, in contrast to [2]–[6] and [13], since this would imply dedicated encoding for each end-user. While *a posteriori* quantization provides flexibility regarding the coding rate, the associated distortion is, however, not compensated for by further MP iterations, contrary to in-loop quantization. The quantization error propagation can nevertheless be kept small for sufficiently fine quantization. Fig. 1 shows that distortion resulting from *a posteriori* quantization is slightly higher than that of the in-loop quantization, depending on the quantization step size, however. The bitstream size is directly proportional to the number of MP iterations since all iterations have the same size, for a uniform quantization scheme and random signal and dictionary. With in-loop quantization, MP adapts to the quantized coefficient, and the residue corresponds to the stream used for reconstruction. It therefore generally performs better than the *a posteriori* quantization, even if the MP decay can be slightly slower. The difference between both quantization schemes is, however, very small for large coefficients (where the relative quantization error is kept small). For very small coefficients




Fig. 1. Comparison of *a priori* and *a posteriori* uniform coefficient quantization of the MP expansion of ten-sample random real signals over a dictionary of 128 random vectors. The distortion have been averaged over 10 independent encodings.

(i.e., high stream size), the quantization step size is too large to capture the coefficient value. The *a posteriori* quantization error can thus even increase with the number of iterations (i.e., very small coefficients), as in the case of a mid-step uniform quantizer used in Fig. 1. However, an in-loop quantization scheme stays less flexible than the *a posteriori* quantization since the stream is generally targeted for a given rate.

The aim of quantization is to offer the best possible reconstruction quality for a given bit budget. We propose a rate-distortion optimal solution taking benefit from the exponential decay of the residual energy. This interesting property directly drives the quantization of the coefficients in two ways. Intuitively, the quantization error on an MP element depends first on the iteration number. The highest iteration number elements indeed bring less energy than the first elements. Their quantization can thus be relatively coarser than for the high energy coefficients. Second, the number of MP elements can also be adapted to the available bandwidth by dropping the lowest energy elements.

Let  $c_{\gamma_m}$  represent the scalar product  $\langle g_{\gamma_m} | \mathcal{R}^n f \rangle$ . From (6), its norm is upper bounded by an exponential function, which can be written as

$$|c_{\gamma_m}| \leq (1 - \alpha^2 \beta^2)^{n/2} \|f\|. \quad (10)$$

The upper bound depends on both the energy of the input function and the construction of the dictionary. Since the coefficient obviously cannot bring more energy than the residual function, the norm of the coefficient is strongly related to the residual energy decay curve. Fig. 2 represents the energy of the MP coefficients for the expansion of random signals over a 128-random vector dictionary ( $\alpha = 1$ ). It clearly illustrates that the coefficients norm can be upper bounded by an exponentially decaying curve. This qualitative result depends neither on the signal nor on the dictionary. However, the decay rate directly depends on the redundancy factor  $\beta$ , as presented in Section II.

#### B. Redundancy-Driven Uniform Coefficients Quantization

The exponential upper bound on the coefficients is now used to design an efficient quantization scheme. In the following




Fig. 2. Coefficient norm versus the iteration number for the MP expansion of three random signals of length 10 over a 128-random vector dictionary ( $\alpha = 1$ ).

discussion, we assume that an Oracle has computed the structural redundancy factor of the MP dictionary, and we use it to derive an efficient quantizer. Later on, we will see how one can drop this sometimes illusory requirement through an adaptive algorithm. Our starting point is the observation we already made: There is clearly no need to quantize all coefficients on the same range since their values decrease exponentially. In other words, the quantization applied to the first coefficients is certainly not efficient on the last ones, as their range is (exponentially) smaller. Bits can thus be saved by simply limiting the quantization region between 0 and the exponential decay curve given by the parameters (i.e.,  $\beta$  and  $\|f\|$ ), after possible coefficient reordering. The quantized coefficients are then sent to the decoder together with an additional bit of sign.

The following paragraph proposes to compute the optimal number of coefficients, as well as the optimal number of bits per coefficient in EUQ. The distortion at decoding can be bounded by the sum of the quantization error and the approximation error due to the limit on the number of MP iterations. Let  $\xi_j = c_{\gamma_j} - \tilde{c}_{\gamma_n}$  denote the error on the coefficient. The total distortion  $D$  can thus be written as

$$\begin{aligned} D &\leq \sum_{j=0}^{N-1} |\xi_j|^2 + \|\mathcal{R}^N f\|^2 \\ &\leq \sum_{j=0}^{N-1} |\xi_j|^2 + (1 - \alpha^2 \beta^2)^N \|f\|^2 \end{aligned} \quad (11)$$

where the energy of the residue at iteration  $N$  is bounded thanks to (6).

Assume now that the distribution of the coefficients norm is uniform between 0 and the exponential upper bound given by (10). This is clearly an oversimplified hypothesis, but, as can be seen on Fig. 3, the exact distribution depends on the iteration number, and its properties dramatically change for different coefficients. The first coefficients are quite narrowly distributed close to the upper bound, whereas the others tend to be increasingly more compactly distributed near small values. This behavior is mainly due to the very conservative upper bound defined in (10), which assumes no *a priori* knowledge on the input signal. The simplified uniform distribution model therefore re-




Fig. 3. Distribution of coefficients  $n = \{0, 4, 9, 19\}$  relatively to the exponential upper bound for the MP expansion of 1000 ten-sample random signals over a random dictionary of 128 atoms ( $\alpha = 1$ ).

flects the lack of a real structure of coefficients distribution in this particular context. The uniform quantization is, moreover, justified in a generic algorithm, where the only available parameters are the signal energy and the dictionary redundancy, without considering neither the particular relation between the signal and the dictionary or dependency among successive coefficients. For example, one could assume a Gaussian coefficient distribution around an average value decaying faster than the exponential upper bound. However, the evolution of this average value clearly depend on the efficiency of the dictionary to decompose a particular signal  $f$ , and such an analysis is left for a future study. Finally, it can be observed that successive coefficient norms are likely to be close, and the distance to the upper bound increases with the coefficients order since the upper bound is clearly conservative. The adaptive algorithm proposed in the next section will take benefit from these observations.

Under the previous assumption, the coefficient  $c_j$  is uniformly quantized within the exponentially decaying quantization range  $I_j = \nu^j \|f\|$ , where  $\nu = (1 - \alpha^2 \beta^2)^{1/2}$ . Let  $n_j$  be the number of quantization steps within  $[0, I_j]$  for the quantization of the  $j^{\text{th}}$  coefficient. Similarly, let  $a_j$  represent the number of bits needed to code the atom index. The atom indices can be entropy coded to improve the compression and  $a_j$  represents in this case the average length of atom codewords without loss of generality. The coding rate for  $N$  MP atoms is therefore given by

$$R = \sum_{j=0}^{N-1} \log_2(n_j) + \sum_{j=0}^{N-1} a_j. \quad (12)$$

In addition, in the case of fine uniform quantization, the distortion can be written as

$$D_Q = \sum_{j=0}^{N-1} |\xi_j|^2 = \sum_{j=0}^{N-1} \frac{\nu^{2j} \|f\|^2}{12 n_j^2}. \quad (13)$$

The optimal quantization problem, which minimizes the distortion for a given bit budget  $R_{\text{budget}}$ , can now be formulated as follows.

*Problem 1:* Given an MP expansion represented by a set of coefficients  $c_j$  and their respective atoms  $g_{\gamma_j}$  coded on code-words of length  $a_j$ . Find  $N$ , which is the number of atoms, and  $\{n_j\}$ , with  $0 \leq j \leq N-1$ , which is the number of bits needed to code their respective coefficients, such that the distortion  $D$  is minimized, while the coding rate is smaller than the bit budget, i.e.,  $R \leq R_{\text{budget}}$ .

The Lagrangian multiplier method [14], [15] is well suited for this kind of constrained optimization problem. It defines a cost function  $\mathcal{L}(\lambda)$  as the sum of the objective distortion function and the constraint on the rate, weighted by the Lagrangian multiplier  $\lambda$ . This formulation provides a solution to the hard constrained problem of finding the optimal set of  $n_j$  and  $N$  by converting it to a set of unconstrained problems driven by  $\lambda$ . In our case, the cost function can be written as

$$\begin{aligned} \mathcal{L}(\lambda) &= D + \lambda R \\ &= \sum_{j=0}^{N-1} \frac{\nu^{2j} \|f\|^2}{12 n_j^2} + \nu^{2N} \|f\|^2 \\ &\quad + \lambda \left( \sum_{j=0}^{N-1} \log_2(n_j) + \sum_{j=0}^{N-1} a_j \right). \end{aligned} \quad (14)$$

The optimal quantization is obtained by differentiating  $\mathcal{L}(\lambda)$  with respect to both  $n_j$  and  $N$ . First, solving

$$\frac{\partial \mathcal{L}(\lambda)}{\partial n_j} = 0, \forall j \quad (15)$$

for  $n_j$  positive and finite yields

$$n_j = \sqrt{\frac{\|f\|^2 \nu^{2j} \log 2}{6 \lambda}}. \quad (16)$$

The solution of (16) is a minimum of the Lagrangian since the second derivative is positive at this point, regardless of the value of  $j$ . Hence, the optimal quantization imposes an exponential law on the number of quantization levels

$$\frac{n_{j+1}}{n_j} = \frac{\nu^{j+1}}{\nu^j} = (1 - \alpha^2 \beta^2)^{1/2}. \quad (17)$$

Interestingly, this previous relation leads to an equivalent participation of each iteration to the total distortion. Indeed, the average distortion per coefficient is equal to

$$\frac{\lambda}{2 \log(2)}, \forall n_j > 1 \quad (18)$$

independently of the iteration. Notice, however, that the Lagrangian formulation provides only an approximation to the optimal quantization. Indeed, in a practical case,  $n_j$  can only take integer values, which are often limited to integer powers of two.

The coding rate can also be limited by transmitting only part of the MP expansion. On the one hand, even if small coefficients can be efficiently approximated by the exponential upper bound, the quantization scheme may decide not to code them because the cost of the atom index is too expensive. On the other hand, it can be seen from (14) that atoms may be transmitted alone, even if no coefficient is coded ( $n_j \leq 1$ ). Depending on the redundancy of the dictionary, atom indices often

carry more information than coefficients. With efficient entropy coding, the atom indices may, moreover, become very cheap to code. One can, therefore, imagine a scheme where the coefficients are simply interpolated from the exponential decay curve, especially for high order iterations (i.e., small coefficients). The optimal number of iterations  $N$  is thus given by minimizing the Lagrangian cost function of (14), where  $n_j$  has been replaced by its optimal value from (16):

$$\mathcal{L}(\lambda) = \sum_{j=0}^{N-1} \mathcal{L}_j(\lambda) + \nu^{2N} \|f\|^2 \quad (19)$$

with

$$\mathcal{L}_j(\lambda) = \begin{cases} \frac{\nu^{2j} \|f\|^2}{12 n_j^2} + \lambda (\log_2 n_j + a_j), & \text{if } n_j \geq 1 \\ \frac{\nu^{2j} \|f\|^2}{12} + \lambda a_j, & \text{otherwise.} \end{cases} \quad (20)$$

Because indices can be transmitted without coefficients, the Lagrangian is defined as a piecewise function. The breakpoint occurs at  $j = N_B = \lfloor -\log(n_0)/\log(\nu) \rfloor$  (i.e.,  $n_j = 1$ ), where  $n_0 = \sqrt{\|f\|^2 \log(2)/6 \lambda}$  from (16). Notice that  $N_B$  is positive only if  $\lambda \leq \|f\|^2 \log(2)/6$ . Otherwise, the weight on the rate in the Lagrangian cost function becomes much more important than the distortion, and the best scheme would be not to transmit any coefficient. Without loss of generality, let us now assume that  $a_j = A$  is constant, where  $A$  corresponds to the average size of the atom indices after possible entropy coding. The optimal number of iterations  $N$  is given by the following theorem.

*Theorem 1:* The Lagrangian cost function of (19) has a unique minimum in terms of the number of MP iterations, and this optimal solution is given by

$$N = \begin{cases} N_1, & \text{if } A > -\frac{23 \log(\nu)+1}{2 \log(2)} \\ N_2, & \text{otherwise} \end{cases} \quad (21)$$

where

$$\begin{aligned} N_1 &= -\frac{1 + 2A \log(2) - \log(\nu) + \log\left(\frac{\|f\|^2 \log(2)}{6 \lambda}\right)}{2 \log(\nu)} \\ &\quad - \frac{W_{-1}\left(3 \cdot 2^{3-2A} \nu \log(\nu) e^{-1}\right)}{2 \log(\nu)} \end{aligned} \quad (22)$$

$$N_2 = \frac{\log\left(-\frac{\lambda A}{2\|f\|^2 \log(\nu) \left(1 + \frac{1}{12 \nu^2 - 12}\right)}\right)}{2 \log(\nu)} \quad (23)$$

are the optimal values of  $N$  on both sides of the piecewise function (i.e., for, respectively,  $N \leq N_B$  and  $N > N_B$ ). In (22),  $W_{-1}(x)$  represents the second branch of the Lambert  $W$  function [16].

The choice between  $N_1$  and  $N_2$  depends only on the relation between  $A$  and  $\nu$ . Indeed, the atom index size and the decay rate of the coefficient norm drive the decision of coding an atom without coefficient. Finally, the solution of (21) is a minimum of the Lagrangian cost function since the second derivative is positive at  $N$ . A detailed proof is given in the Appendix for completeness.




Fig. 4. Average R-D curve for the MP expansion of ten ten-sample random real signals over a 50-vector random dictionary in the case of both exponential and uniform quantization.

Now that the unconstrained problem of (14) has been solved optimally for an arbitrary  $\lambda$ , the next step is to find the optimal  $\lambda_{\text{opt}}$  that guarantees a bit budget  $R_{\text{budget}}$ . The solutions of the optimization problem form a convex hull of the achievable rate-distortion curve, and the bit budget constraint imposes a  $\lambda_{\text{opt}}$  that represents the slope of the rate-distortion function at  $R = R_{\text{budget}}$ . Several bisection methods have been proposed to solve this second problem [17]–[19]. Since, in our case, the bit rate is expressed as a function of the Lagrangian multiplier  $\lambda$ , the constraint on the bit budget directly imposes an approximated value for  $\lambda_{\text{opt}}$  from

$$R_{\text{budget}} = \begin{cases} \sum_{j=0}^{N_1-1} \log_2(n_j) + N_1 A, & \text{if } A > -\frac{23}{2} \frac{\log(\nu)+1}{\log(2)} \\ \sum_{j=0}^{N_2-1} \log_2(n_j) + N_2 A, & \text{otherwise} \end{cases} \quad (24)$$

where  $N_1$ ,  $N_2$ , and  $n_j$  are functions of  $\lambda$  from the above equations. Recall that the approximation is due to the constraint of integer number of bits for each coefficient. The values of  $\lambda_{\text{opt}}$  can finally easily be computed through numerical methods.

Fig. 4 shows the rate-distortion characteristic of the MP expansion of a ten-sample random signal over a 50-vector dictionary. The redundancy-driven quantization scheme is compared with different lower step uniform quantization schemes, which simply quantize the coefficients over a range defined by the energy of the signal. The proposed EUQ clearly outperforms the uniform quantization since it adapts to the rate and to the range of the coefficients to provide the best approximation for the available rate. Uniform quantization schemes provide good results for low rate since the first coefficients are finely quantized. For high rates, however, the coefficients become too small compared with the quantization steps, and they are set to zero. The error may even increase in the case of mid-step quantization since the quantized coefficient becomes larger than the true value computed by the MP (see Fig. 1).

Finally, it can be noted that the exponentially upper bounded uniform quantization is equivalent to the division of the coeffi-

cients by an exponential quantization table factor, multiplied by a quantizer scale factor, which is given by the bit budget. Practically speaking, the only parameters to pass to the decoder are  $\nu$ ,  $\|f\|$ , and  $n_0$  or, equivalently, the quantizer scale factor. The optimal number of iterations thus only depends on the bit budget and the design of the input dictionary. The redundancy-driven quantization becomes particularly interesting for highly redundant dictionaries, where the coefficient value decreases very rapidly. In this case, the distribution of bits among coefficients is particularly efficient compared to uniform quantization.

The scheme proposed above heavily relies on the knowledge of the structural redundancy factor  $\beta$  or, equivalently,  $\nu$ . We call this scheme Oracle-based because, in practical situations, it can be too demanding to compute this parameter. Moreover, its influence is mixed with the suboptimality factor  $\alpha$  brought in by MP implementation choices. This additional factor is, moreover, very difficult to control. In the next section, we develop an adaptive algorithm that uses the previous theory to dynamically estimate the quantization parameters, therefore relaxing assumptions on the *a priori* knowledge of the redundancy factor.

#### IV. PRACTICAL SCHEME: ADAPTIVE QUANTIZATION

The previous optimal quantization scheme has several limitations, particularly in the practical case of large dictionaries. Moreover, the accuracy of the exponential upper bound is highly dependent on the MP algorithm implementation. To overcome these limitations, we now propose a suboptimal but very efficient algorithm based on the development of the previous section. The key idea lies in a dynamic computation of the redundancy factor  $\beta$  (i.e., the parameter  $\nu$ ) from the quantized data. Since this information is also available at the decoder, it will be able to perform the inverse quantization without any additional side information. As in the previous section, we only focus on the *magnitude* of coefficients, reporting their sign on an separate bit.

The MP coefficients are first reordered and sorted in the decreasing order of magnitude. This operation may be necessary since the MP algorithm does not guarantee a strict decay of the coefficient energy, but rather an upper bound. The adaptive quantization scheme then performs as follows. Let  $Q[c_k]$ ,  $k = 1, \dots, j-1$  denote the quantized counterparts of the  $j-1$  first coefficients. Due to the rapid decay of the magnitude, coefficient  $c_j$  is very likely to be smaller than  $Q[c_{j-1}]$ . It can thus be quantized in the range  $[0, Q[c_{j-1}]]$ . The number of quantization levels at step  $j$  is theoretically driven by the redundancy factor, as given by (17). The adaptive quantization uses an estimate of the redundancy factor to compute the number of quantization levels as

$$n_j = \tilde{\nu}_{j-1} n_{j-1}. \quad (25)$$

The estimate of the redundancy factor  $\tilde{\nu}$  is recursively updated, replacing the quantization range  $I_j$  of the optimal quantization scheme (see Section III-B) with the quantized coefficient  $Q[c_j]$  as

$$\tilde{\nu}_j = \left( \frac{Q[c_j]}{\|f\|} \right)^{1/j}. \quad (26)$$

The adaptive quantization is described by Algorithm 1. The quantization is completely determined by the choice of  $n_0$ , the number of bits for the first coefficient, and a positive value of  $N$ . The algorithm iteratively quantizes each of the coefficients and updates, at each step, the estimated value of the redundancy factor. It finally stops after a predefined number of atoms  $N$  has been reached or when it determines that the number of bits to code a coefficient becomes null (i.e.,  $n_j = 0$ ).

**Algorithm 1** Adaptive Quantization  
**Require:**  $I_0 = \|f\|$ ,  $n_0 = \sqrt{\|f\|^2 \log 2/6\lambda}$ ,  $\tilde{\nu}_0 = 1$   
**while**  $j < \max$  iterations ( $N$ ) **do**  
 $n_j = \tilde{\nu}_{j-1} n_{j-1}$   
 $Q[c_j] \leftarrow$  Quantize  $|c_j|$  on  $n_j$  levels with range  $I_{j-1}$   
 $I_j \leftarrow Q[c_j]$   
 $\tilde{\nu}_j \leftarrow (Q[c_j]/\|f\|)^{1/j}$  {update  $\tilde{\nu}$ }  
 $j \leftarrow j + 1$   
**end while**

When the bit stream has to conform to a given bit budget, the parameter  $n_0$  is computed as follows. First,  $\nu$  is estimated with (26) by training the dictionary on a large set of signals (e.g., images), encoded with the adaptive quantization algorithm. It can be shown empirically that the estimation quite rapidly tends to the asymptotic value of the redundancy factor. The estimation of  $\nu$  is then used to compute  $\lambda$  as a function of the given bit budget with (24), where  $N$  is given by (21) (for highly redundant dictionaries, we generally have  $N = N_1$ ). The value of  $\lambda$  determines the number of bits of the first coefficient  $n_0$ , and the number of atoms to be coded,  $N$ , using appropriately (22) or (23). The adaptive quantization algorithm is finally completely determined and generally yields bit rates very close to the bit budget. For high bit rates, it sometimes underestimates the quantizer efficiency, due to the fact that the quantization range estimated from the quantized coefficient values is more accurate than the one obtained from the theoretical exponential upper bound. In this case, the actual bit rate becomes smaller than the bit budget, which can anyway be reached by coding additional coefficients.

Notice that a loose estimation of  $\nu$  or  $\lambda$  will not impair the efficiency of the quantization algorithm but rather displace the resulting encoding on the working rate-distortion curve. Finally, several coefficients could be used for the computation of  $\tilde{\nu}$  in Algorithm 1. Such a modification will improve the accuracy of the estimation of quantization parameters and avoid potential oscillatory effects.

V. EXPERIMENTAL RESULTS

A. Random Signals

In this section, we now compare the adaptive entropy-constrained quantization scheme with an exponential quantization scheme [20] used in MP coding [9]. The exponential quantization is clearly expected to provide better results than uniform quantization [21] due to the particular distribution of the MP




Fig. 5. Rate-distortion curve for adaptive and exponential quantization of the MP expansion of ten-sample random real signals over a 128-atom dictionary.

coefficients. Both the exponential and the adaptive entropy-constrained quantization schemes, which are comparable in terms of complexity, are used for *a posteriori* MP coefficient quantization. In the particular case of coding of random signal coding, the distortion is reported to the coding rate that is estimated as the sum of the quantized coefficients entropy and the index average size. Note that the exponential quantization is similar to the one proposed in [9], where the dead zone is adapted to the statistics of the coefficients.

Fig. 5 shows the evolution of the MSE distortion versus the coding rate for both the adaptive entropy-constrained and the exponential quantization. The curve has been averaged on 100 MP expansions of a random  $\mathbb{R}^{10}$  source. The adaptive scheme clearly provides better results since it optimally distributes bits among MP coefficients. Moreover, it allows us to reach much lower coding rates since the number of coefficients is adaptively chosen according to the bit budget.

B. Image Coding

This section finally proposes the analysis of the behavior of the adaptive quantization scheme in the practical case of MP image coding. To this aim, an encoder has been built that uses either the adaptive EUQ scheme proposed in the previous section or the quantization scheme proposed in [6]. Note that in both cases, the quantization is performed *a posteriori* on a stream generated by an offline MP expansion.

The encoder used in these experiments is composed as follows: A low-frequency part that is 1/16th the size of the image is first coded through a combined DPCM and entropy coding scheme. This low-frequency part, whose coding rate is equivalent to 0.03 b/pixel in the experiments below, is then subtracted from the image. The residual high-frequency image is coded through MP, using a dictionary built on anisotropic refinement and rotation of two-dimensional functions [22]. Unlike most of the proposed coders, the proposed MP scheme does not search atoms on a block-by-block basis but rather in the whole image in order to improve the localization of the signal energy.

The main difference between image coding and the MP expansion of random signal presented above resides in the coding




Fig. 6. PSNR-rate characteristic of the MP coding of *Lena* with the adaptive exponentially upper bounded and the exponential quantization.

of the atom positions. In order to limit the size of the dictionary, the atom positions have to be coded separately from the atom indices. Since the adaptive EUQ scheme heavily relies on the coefficient ordering along their decreasing magnitude, the coding of the atom position is, in this case, quite expensive. There are indeed no obvious statistical properties in the distribution of the atom positions. A small coding gain is, however, obtained by coding the position coordinates  $x$  and  $y$  with an arithmetic coder.

On the other hand, the uniform threshold quantizer proposed in [6] is totally independent of the atom order. The atoms are ordered along their position in the image.<sup>2</sup> The atom positions are then differentially coded on an image row through an arithmetic coder. A *end-of-line* symbol is then sent, followed by the number of empty lines to the next atom. Such a differential coding is very efficient when the number of atoms becomes large.

Fig. 6 represents the evolution of the PSNR as a function of the coding rate for the two MP encoders described here and above. Recall that the only differences between them are the coefficient quantization and the position coding. It can be seen that both schemes perform very similarly. The coder based on the adaptive exponentially upper bounded quantizer performs slightly better at low rates, where the penalty of raw position coding is not too important compared with differential coding. At high rates, the proposed method provides a slightly lower PSNR quality than the coder based on a uniform threshold quantizer. In this case, the benefit of the adaptive quantization is compensated by an expensive position coding. The visual quality is also very similar for both encoders, as shown in Fig. 7.

While the coding performance is comparable in both coders, the adaptive exponentially upper bounded quantizer offers a very important advantage in terms of scalability. The MP stream is intrinsically progressive and can be simply truncated to meet lower rate constraints by discarding the least important coefficients at the end of the stream. The quantization can also be easily adapted by changing the quantizer parameters. On the contrary, the differential coding of the position in the second encoder imposes decoding/recoding operations in the case of bit rate adaptation. In addition, the design of the corresponding

<sup>2</sup>The coding of the positions does not take into account block limits, whereas [6] codes the atom positions per block



Fig. 7. MP coding of a  $128 \times 128$  pixel *Lena* image, at 0.19 and 0.37 b/pixel with the adaptive exponentially upper bounded (AEUQ) and the uniform threshold (UTQ) quantizers. (a) AEUQ 0.19 b/pixel, PSNR = 21.85 dB. (b) UTQ 0.19 b/pixel, PSNR = 21.62 dB. (c) AEUQ 0.37 b/pixel, PSNR = 23.94 dB. (d) UTQ 0.37 b/pixel, PSNR = 23.89 dB

quantizer (e.g., the dead-zone, see [6]) has to be modified according to different bit budget constraints.

## VI. CONCLUSIONS

A new quantization for MP expansion has been proposed in this paper. It describes an *a posteriori* quantization of the MP stream to meet the constraints of asymmetric applications targeting heterogeneous decoders. The encoder takes advantage of the exponential decay of the MP coefficients, which has been shown to be directly driven by the dictionary structure. Depending on the parameters of this exponential upper bound, an optimal redundancy-driven coefficient quantization scheme has been proposed. This scheme clearly outperforms uniform quantization schemes since it adapts to the coefficient decay to provide the best possible approximation with the lowest coding rate. Following what has been taught by the Oracle-based approach, we derived an approximate entropy-constrained scheme that outperforms previously proposed MP quantizers. The performance of the adaptive scheme has finally been shown to favorably compare to state-of-the-art quantization in low bit-rate MP image coding, while additionally providing an advantageous progressive bitstream.

## APPENDIX OPTIMAL NUMBER OF ITERATIONS

### *Proof of Theorem 1*

The Lagrangian cost function of (19) has two pieces connected at a breakpoint  $N_B$  defined by the first noncoded co-



efficient. The first derivative of the right-hand side part of the Lagrangian (i.e., for  $N > N_B$ ) can be written as

$$\frac{\partial \mathcal{L}(\lambda)}{\partial N} = \lambda A + 2 \|f\|^2 \left(1 + \frac{1}{12 \nu^2 - 12}\right) \nu^{2N} \log \nu. \quad (27)$$

This expression has a zero given by

$$N_2 = \frac{\log \left( -\frac{\lambda A}{2 \|f\|^2 \log(\nu) \left(1 + \frac{1}{12 \nu^2 - 12}\right)} \right)}{2 \log(\nu)}. \quad (28)$$

However, (28) is by definition valid only for  $N > N_B$ . The condition for  $N_2$  to be larger than  $N_B$  is given by

$$\frac{\log \left( -\frac{\lambda A}{2 \|f\|^2 \log(\nu) \left(1 + \frac{1}{12 \nu^2 - 12}\right)} \right)}{2 \log(\nu)} > \frac{-\log \left( \sqrt{\frac{\|f\|^2 \log(2)}{6 \lambda}} \right)}{\log(\nu)} \quad (29)$$

or, for  $\nu \in ]0, 1[$  and  $A$  positive

$$A < -\frac{12 \log(\nu) + \frac{\log(\nu)}{\nu^2 - 1}}{\log(2)}. \quad (30)$$

The second derivative of the right-hand side part of the Lagrangian cost function is given by

$$\frac{\partial^2 \mathcal{L}(\lambda)}{\partial N^2} = 4 \|f\|^2 \left(1 + \frac{1}{12 \nu^2 - 12}\right) \nu^{2N} \log^2 \nu \quad (31)$$

and is positive for  $\nu \in ]0, 1[$ . Hence,  $N_2$  is a minimum of the Lagrangian cost function, provided that the constraint of (30) is satisfied. The constraint indicates whether it is worth transmitting the atom indices without coefficients and depends only on the size  $A$  of the atom codewords and the coefficient energy decay rate  $\nu$ . The first derivative of the left-hand side part of the Lagrangian cost function (i.e., for  $N < N_B$ ) is given by

$$\begin{aligned} \frac{\partial \mathcal{L}(\lambda)}{\partial N} &= \frac{\lambda \log(\nu)}{\log(2)} N + 2 \nu^{2N} \log(\nu) \|f\|^2 \\ &+ \frac{\lambda}{\log(2)} \left( \frac{1}{2} + \log \sqrt{\frac{\|f\|^2 \log(2)}{6 \lambda}} - \frac{\log(\nu)}{2} + A \log(2) \right). \end{aligned} \quad (32)$$

For the sake of clarity, let us rewrite (32) as

$$\frac{\partial \mathcal{L}(\lambda)}{\partial N} = a + b x + d e^x \quad (33)$$

with

- 
- $a = \frac{\lambda}{2 \log(2)} + \lambda A - \frac{\lambda \log(\nu)}{2 \log(2)} + \frac{\lambda}{2} \log \left( \frac{\|f\|^2 \log(2)}{6 \lambda} \right);$
- $b = \lambda/2 \log(2);$
- $d = 2 \|f\|^2 \log(\nu);$
- $x = 2 N \log(\nu).$

The zeroes of (33) are given by

$$x = -\frac{a + b W \left( \frac{d}{b} e^{-a/b} \right)}{b} \quad (34)$$

where  $W(z)$  is the Lambert W function [16]. It is a well-defined function of the complex variable  $z$  and has two real branches

for  $z \in [-e^{-1}, 0]$ ,  $W_0(z)$ , and  $W_{-1}(z)$ . When  $A > 1$ , which simply means that an atom index is specified by more than one bit, the argument of  $W$  in (34) satisfies

$$-e^{-1} \leq (2^{3-2A} 3 \nu \log(\nu) e^{-1}) < 0 \quad (35)$$

for  $\nu \in ]0, 1[$ . This ensures the existence of two solutions, and hence, the Lagrangian cost function presents two extrema, which are further studied below. Let  $x_0$  and  $x_{-1}$  denote the solutions of (34) for, respectively, the first and second branches of the Lambert W function. Respectively, this yields  $N_{10} = x_0/2 \log(\nu)$  and  $N_{11} = x_{-1}/2 \log(\nu)$ . Since the definition of the Lambert W function implies that  $-1 \leq W_0(z) \leq 0$  and  $W_{-1}(z) < -1$ , the solution given by the first branch of the Lambert function is, moreover, strictly larger than the solution given by the second branch, i.e.,  $N_{10} > N_{11}$ .

The second derivative of the left-hand side part of the Lagrangian cost function is given by

$$\frac{\partial^2 L}{\partial N^2} = 2 b \log(\nu) + 2 d \log(\nu) e^{2 \log(\nu) N} \quad (36)$$

where  $b \geq 0$  and  $d < 0$  for  $\nu \in ]0, 1[$ , and  $\lambda \geq 0$ . The second derivative is therefore positive when

$$2 d \log(\nu) e^{2 \log(\nu) N} > -2 b \log(\nu) \quad (37)$$

or equivalently

$$e^{2 \log(\nu) N} > -\frac{b}{d}. \quad (38)$$

From (34), we easily check that

$$e^{2 \log(\nu) N_1} = e^{-a/b} e^{-W(d/b e^{-a/b})}. \quad (39)$$

Since, by definition,  $e^{-W(z)} = W(z)/z$  for  $z \neq 0$ , we have

$$e^{2 \log(\nu) N_1} = \frac{b}{d} W \left( \frac{d}{b} e^{-a/b} \right). \quad (40)$$

From (38)

$$e^{2 \log(\nu) N_{10}} = \frac{b}{d} W_0 \left( \frac{d}{b} e^{-a/b} \right) < -\frac{b}{d} \quad (41)$$

and

$$e^{2 \log(\nu) N_{11}} = \frac{b}{d} W_{-1} \left( \frac{d}{b} e^{-a/b} \right) > -\frac{b}{d} \quad (42)$$

since, by definition,  $-1 \leq W_0(z) \leq 0$  and  $W_{-1}(z) < -1$ . The second derivative is therefore positive for  $N = N_{11}$  and negative for  $N = N_{10}$ , and the Lagrangian cost function presents a minimum at  $N = N_{11}$  and a maximum at  $N = N_{10}$ . Since  $N_{11} < N_{10}$ , the Lagrangian cost function is first a decreasing function of  $N$ , for  $N$  small. Moreover, the maximum of the left-hand side part of the function (i.e.,  $N = N_{10}$ ) is always after the breakpoint  $N_B$  since

$$-W_0 \left( \frac{d}{b} e^{-a/b} \right) < 1 + 2 A \log(2) - \log(\nu) \quad (43)$$

for  $\nu \in ]0, 1[$  and  $A$  positive. It can thus be omitted in further analysis. It remains to be shown that the optimal solution for the left part of the Lagrangian function is positive and falls before

$N_B$ . Positivity at  $N_{11}$  is guaranteed if and only if the following condition is respected:

$$\lambda < \frac{4 \|f\|^2 \log(2) \log(\nu)}{W_{-1}(3 \cdot 2^{3-2A} \nu \log(\nu) e^{-1})}. \quad (44)$$

This means that in the worst case for the Lambert function, we have a limit on the rate in the Lagrangian cost function

$$\lambda < 4 \|f\|^2 \log(2) \log\left(\frac{1}{\nu}\right). \quad (45)$$

In addition,  $N_{11}$  is smaller than  $N_B$  if

$$-W_{-1}(3 \cdot 2^{3-2A} \nu \log(\nu) e^{-1}) > 1 - \log(\nu) + 2 \log(2) A. \quad (46)$$

Since  $W(z) = a$  is equivalent to  $a e^a = z$  by definition, the previous condition is equivalent to

$$A > -\frac{23 \log(\nu) + 1}{2 \log(2)}. \quad (47)$$

Conditions of (30) and (47) are mutually exclusive. If one constraint is respected, the other cannot be satisfied, and this ensures the existence of one and only one minimum to the Lagrangian cost function. This can be shown by proving that the product of both conditions is always positive

$$\left(-\frac{23 \log(\nu) + 1}{2 \log(2)} - A\right) \left(-\frac{12 \log(\nu) + \frac{\log(\nu)}{\nu^2 - 1}}{\log(2)} - A\right) \geq 0. \quad (48)$$

Indeed, if

$$A + \frac{12 \log(\nu) + \frac{\log(\nu)}{\nu^2 - 1}}{\log(2)} \leq 0 \quad (49)$$

then

$$A + \frac{23 \log(\nu)}{2 \log(2)} \leq -\frac{\log(\nu)}{2 \log(2)} - \frac{2 \log(\nu)}{2 \log(2) (\nu^2 - 1)}. \quad (50)$$

Furthermore, if (49) is respected, then

$$A + \frac{23 \log(\nu) + 1}{2 \log(2)} \leq 0. \quad (51)$$

From (50), the last relation can be rewritten as

$$\frac{2 \log(\nu)}{\nu^2 - 1} \leq 1 - \log(\nu). \quad (52)$$

Setting  $\nu = e^{-t/2}$ , for  $t \in ]0, \infty[$  is equivalent to proving that

$$t \geq (e^{-t} - 1) \left(1 + \frac{t}{2}\right) \quad (53)$$

where both left and right terms are zero for  $t = 0$ . The left term increases with  $t$  with a slope 1. The right term also increases with  $t$  since its first derivative is always positive, but the increasing rate is always smaller or equal to 1. This proves

the inequality of (53). It can be shown in the same way that if one of the factor of (48) is negative, the other one is negative as well. Altogether, we thus have a unique solution for the optimal number of iterations.

## REFERENCES

- [1] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Processing*, vol. 41, pp. 3397–3415, Dec. 1993.
- [2] V. K. Goyal, M. Vetterli, and N. T. Thao, "Quantized overcomplete expansions in  $R^N$ : analysis, synthesis and algorithms," *IEEE Trans. Inform. Theory*, vol. 44, pp. 16–31, Jan. 1998.
- [3] Z. Cvetkovic, "Source coding with quantized redundant expansions: Accuracy and reconstruction," in *Proc. IEEE Data Compression Conference*, 1999, 279, pp. 344–358.
- [4] Z. Cvetkovic and M. Vetterli, "Overcomplete expansions and robustness," in *Proc. IEEE-SP Int. Symp. Time-Frequency Time-Scale Anal.*, 1996, 297, pp. 325–328.
- [5] V. K. Goyal and M. Vetterli, "Dependent coding in quantized matching pursuit," *Proc. SPIE Visual Commun. Image Process.*, vol. 3024, pp. 2–12, 1997.
- [6] R. Neff and A. Zakhor, "Adaptive modulus quantizer design for matching pursuit video coding," in *Proc. IEEE Int. Conf. Image Processing*, vol. 2, 1999, 354, pp. 81–85.
- [7] S. Mallat, *A Wavelet Tour of Signal Processing*, 2nd ed. New York: Academic, 1999.
- [8] G. Davis, S. Mallat, and M. Avellaneda, "Adaptive greedy approximations," *J. Constructive Approx.*, vol. 13, pp. 57–98, 1997.
- [9] R. Neff and A. Zakhor, "Modulus quantization for matching-pursuit video coding," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 10, pp. 895–912, Sept. 2000.
- [10] P. Frossard, "Robust and multiresolution video delivery: From H.26x to matching pursuit based technologies," Ph.D. dissertation, Swiss Fed. Inst. Technol., Lausanne, Switzerland, 2000.
- [11] L. K. Jones, "On a conjecture of Huber concerning the convergence of projection pursuit regression," *Ann. Statist.*, vol. 15, no. 2, pp. 880–882, 1987.
- [12] P. Frossard and P. Vandergheynst, "Redundancy in nonorthogonal transforms," in *Proc. IEEE Int. Symp. Inform. Theory*, Washington, DC, June 2001, p. 196.
- [13] V. K. Goyal, M. Vetterli, and N. T. Thao, "Quantization of overcomplete expansions," in *Proc. IEEE Data Compression Conf.*, 1995, 298, pp. 13–22.
- [14] H. Everett, "Generalized Lagrange multiplier method for solving problems of optimum allocation of resources," *Oper. Res.*, vol. 11, pp. 399–417, 1963.
- [15] A. Ortega and K. Ramchandran, "Rate-distortion methods for image and video compression," *IEEE Signal Processing Mag.*, vol. 15, pp. 23–50, Nov. 1998.
- [16] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth, "On the Lambert W function," *Adv. Comput. Math.*, vol. 5, pp. 329–359, 1996.
- [17] G. M. Schuster and A. K. Katsaggelos, "An optimal quadtree-based motion estimation and motion-compensated interpolation scheme for video compression," *IEEE Trans. Image Processing*, vol. 7, pp. 1505–1523, Nov. 1998.
- [18] Y. Shoham and A. Gersho, "Efficient bit allocation for an arbitrary set of quantizers," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 36, pp. 1445–1453, Sept. 1988.
- [19] K. Ramchandran and M. Vetterli, "Best wavelet packet bases in a rate-distortion sense," *IEEE Trans. Image Processing*, vol. 2, pp. 160–175, Apr. 1993.
- [20] G. J. Sullivan, "Efficient scalar quantization of exponential and Laplacian random variables," *IEEE Trans. Inform. Theory*, vol. 42, pp. 1365–1374, Sept. 1996.
- [21] P. Frossard and P. Vandergheynst, "A posteriori quantized matching pursuit," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 27–29, 2001.
- [22] P. Vandergheynst and P. Frossard, "Efficient image representation by anisotropic refinement in matching pursuit," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Processing*, vol. 3, Salt Lake City, UT, May 2001, pp. 1757–1760.

**Pascal Frossard** (S'96–M'01) received the M.S. and Ph.D. degrees, both in electrical engineering, from the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland, in 1997 and 2000, respectively. In 1998, he graduated from the Graduate School of the Communications System Division of the EPFL.

From 1998 to 2000, he worked with the Signal Processing Laboratory, EPFL, as a research and teaching assistant under a grant from Hewlett-Packard. Between 2001 and 2003, he was a member of the research staff at the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he worked on media compression and streaming technologies. Since April 2003, he has been an assistant professor at EPFL, supported by the Swiss National Science Foundation. He works on nonlinear signal transforms applied to visual information and adaptive media streaming algorithms. In 2002, he was co-editor of a Special Issue on new visual coding methods for *Signal Processing*.

Dr. Frossard is co-guest editor of a special issue of the IEEE TRANSACTIONS ON MULTIMEDIA on media streaming. In 2002, he was the general chairman of the IEEE International Conference and Exposition on Multimedia (ICME 2002).

**Pierre Vanderghenst** (M'01) received the M.S. degree in physics and the Ph.D. degree in mathematical physics from the Université Catholique de Louvain, Louvain, Belgium, in 1995 and 1998, respectively.

From 1998 to 2001, he was a Postdoctoral Researcher with the Signal Processing Laboratory, Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland. He is now an Assistant Professor of visual information representation theory at EPFL, where his research focuses on computer vision, image and video analysis, and mathematical techniques for applications in visual information representation.

Dr. Vanderghenst is co-Editor-in-Chief of *Signal Processing*.

**Rosa Maria Figueras i Ventura** (S'00) was born in Catalonia, Spain, in 1977. She received the M.Sc. degree in telecommunication engineering from the Escola Tècnica Superior d'Enginyers de Telecomunicació de Barcelona (ETSETB), Technical University of Catalonia (UPC), Barcelona, Spain, in 2000. Since 2001, she has been pursuing the Ph.D. degree with the Signal Processing Institute, EPFL, Switzerland, on the domain of image processing and representation.

In 2000, she was an exchange student with the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland.

**Murat Kunt** (S'70–M'74–SM'80–F'86) was born in Ankara, Turkey, on January 16, 1945. He received the M.S. degree in physics and the Ph.D. degree in electrical engineering, both from the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland, in 1969 and 1974, respectively.

From 1974 to 1976, he was a visiting scientist at the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, where he developed compression techniques for X-ray images and electronic image files. In 1976, he returned to EPFL, where, presently, he is Professor of electrical engineering and Director of the Signal Processing Institute, which is one of the largest at EPFL. He conducts teaching and research in digital signal and image processing with applications to modeling, coding, pattern recognition, scene analysis, industrial developments, and biomedical engineering. His laboratory participates in a large number of European projects under various programs, such as ESPRIT, EUREKA, RACE, HCM, COMMETT, and COST. He is the author or the co-author of more than 200 research papers and 15 books and holds seven patents. He consults for governmental offices including the French General Assembly.

Dr. Kunt is the Editor-in-Chief of the *Signal Processing Journal* and a founding member of the European Association for Signal Processing (EURASIP). He serves as a chairman and/or a member of the scientific committees of several international conferences and on the editorial boards of the PROCEEDINGS OF THE IEEE, *Pattern Recognition Letters*, and *Traitement du Signal*. He was the co-chairman of the first European Signal Processing Conference, Lausanne, 1980, and the General Chairman of the International Image Processing Conference (ICIP'96), Lausanne, 1996. He was the President of the Swiss Association for Pattern Recognition from its inception until 1997. He received the gold medal from EURASIP for meritorious service, the IEEE ASSP Technical Achievement Award, the IEEE Third Millennium Medal, an honorary doctorate degree from the Catholic University of Louvain, Louvain, Belgium, the Technical Achievement Award from EURASIP, and the Imaging Scientist of the Year Award from the IS&T and SPIE in 1983, 1997, 2000, 2001, 2002, and 2003, respectively.