ATOM MODULUS QUANTIZATION FOR MATCHING PURSUIT VIDEO CODING

Christophe De Vleeschouwer

Laboratoire de Télécommunications
Université catholique de Louvain
Louvain-la-Neuve, Belgium

Avideh Zakhor

Department of EECS
University of California
Berkeley CA, USA

ABSTRACT

We provide an analytical study of the selection and modulus quantization of matching pursuits (MP) coefficients. We demonstrate that an optimal rate-distortion trade-off is achieved by selecting the atoms up to a dead-zone threshold, and by defining the modulus quantizer in terms of that threshold. In doing so, we take into account quantization error re-injection resulting from inserting the modulus quantizer inside the MP atom computation loop. *In-loop* quantization affects the stepsize of the uniform quantizer, and results in a non-uniform optimal entropy constrained quantizer. Improvements larger than one dB are obtained for video coding.

1. INTRODUCTION

Matching pursuit (MP) is a greedy and iterative approximation algorithm that generates sparse representation of a signal with respect to an overcomplete set of basis functions. The MP expansion is defined in terms of index, sign, and modulus of a subset of basis functions. In this paper, we primarily deal with the modulus quantization of matching pursuits coefficients. Our study relies on two main assumptions. Firstly, linear reconstruction is assumed because complexity considerations dictate such reconstruction in most applications. Secondly, independent scalar modulus quantization is considered. This is applicable to coding scenarios where successively transmitted atom moduli are independent, i.e. where the atom coding order is dictated by entropy gain achieved on atom indices rather than by their modulus. This is the case for video coding applications, where a significant entropy gain is obtained by appropriate differential description of the atom positions [1]. This second assumption differentiates our work from [2], which is dedicated to systems coding the atoms in decreasing order of magnitude.

A previous work has already addressed the MP atom modulus quantization issue in the context of linear reconstruction and modulus-independent encoding order [3]. Our work not only provides an analytical derivation of the empirical results in [3], but also refines and completes them. Our study takes into account the quantization error re-injection resulting from insertion of the modulus quantizer inside the MP atom computation loop. Due to re-injection, the quantization error of an atom is likely to be corrected by subsequent iterations, resulting in improved coding efficiency. We show re-injection also impacts the quantizer design for both uniform quantization, and non-uniform optimal entropy constrained quantization.

The rest of the paper is organized as follows. Sections 2 and 3 express rate-distortion constraints, for the expansion process and for the quantization stage of the MP coder respectively. Section 2 shows that an optimal MP expansion must stop as soon as the moduli of the extracted atoms become smaller than a threshold. In Section 3, we assume an exponential distribution of the MP atom modulus, and review rate-constrained optimal quantization of such a random variable. We then consider the re-injection of the quantization error into the MP expansion loop, and show that re-injection reduces the impact of quantization error on the final reconstructed signal distortion. A formal analysis of this phenomenon reveals that the benefit of the re-injection is stronger for the atoms that are selected during the initial iterations of the MP expansion. As these atoms are also expected to have the largest moduli, we propose the use of a non-uniform quantizer, and present a method to design the R/D optimal non-uniform quantizer. Section 4 measures the impact of our work for video coding, and Section 5 includes conclusion.

2. EXPANSION IN A RATE/DISTORTION FRAMEWORK

MP is a greedy and iterative expansion process. In a coding context, a critical question is when to stop the process. This section shows that an expansion is optimal in the rate/distortion (R/D) sense if it captures all and only all atoms larger than a threshold. Let a dictionary $\mathcal{D} = \{\varphi_k\}_{k=1}^{S} \in R^M$ be a frame such that $\|\varphi_k\| = 1$ for all k. Given a source vector $x \in R^M$, MP is a greedy and iterative algorithm that approximates x by a linear combination of elements of $\mathcal{D}$, or atoms. At step i, the algorithm selects the dictionary function φ_{k_i} that maximizes $|< \varphi_{k_i}, R_i x >|$ and generates the residue for the next iteration, i.e. $R_{i+1} x = R_i x - < \varphi_{k_i}, R_i x > \varphi_{k_i}$. In the initial step, $R_0 x = x$. After n steps, the approximate signal is $x \approx \sum_{i=0}^{n-1}(s_i \cdot \alpha_i \cdot \varphi_{k_i})$, where $\alpha_i = | < \varphi_{k_i}, R_i x > |$ and $s_i = sign(< \varphi_{k_i}, R_i x >)$. To code the expansion, i.e. the index, sign, and modulus of the coefficients, it is desirable to achieve the highest quality at the lowest cost. A convenient way to formalize this problem is to use a Lagrange multiplier λ to define the relative importance of the distortion D and of the number of bits R. The optimal trade-off in terms of quality and coding cost is then the one that minimizes the Lagrangian cost function $\mathcal{L}(\lambda) = D + \lambda R$. For an MP expansion, define D_i to be the distortion or MP residue energy after i atoms have been selected and $\Delta D_i = D_i - D_{i-1}$ to be the decrease in residual energy due to the i^{th} atom. As atoms are roughly selected in decreasing order of magnitude, we have

$$\Delta D_i < \Delta D_{i+1} \tag{1}$$

In the same way, define R_i the number of bits to encode i atoms and $\Delta R_i = R_i - R_{i-1}$ the incremental increase of the bit budget due to the i^{th} atom. As MP representations are sparse, the number

of bits devoted to an additional atom does not rapidly decrease with the iteration number. We have

$$\Delta R_i \approx \Delta R_{i+1} \tag{2}$$

For a Lagrangian multiplier λ, combining (1) and (2) gives

$$\Delta D_i + \lambda \Delta R_i < \Delta D_{i+1} + \lambda \Delta R_{i+1} \tag{3}$$

So, the incremental contribution $\Delta D + \lambda \Delta R$ of successive atoms to $\mathcal{L}(\lambda)$ can be assumed to be monotically increasing. This ensures the convexity of the (R, D) curve drawn along the expansion process. It also makes the selection of an additional atom worthwhile only until

$$\Delta D + \lambda \Delta R = 0, \tag{4}$$

defining a stopping criterion for the MP expansion in terms of incremental rate ΔR and decrease of MP residue energy ΔD due to the additional atom. Defining m to be the modulus of an additional atom and $\xi_{Q(.)}(m)$ to be its quantization error for a quantization method defined by $Q(.)$, the decrease ΔD of the MP residue energy is

$$\Delta D = -\left(m^2 - \xi^2_{Q(.)}(m) \right) \tag{5}$$

ΔR can be divided in two parts: the first one $R_{Q(.)}(m)$ corresponds to the quantized atom modulus; the second one ΔR_{index} corresponds to the rate for the sign and index of the additional atom, and is not directly affected by the quantization method. Based on Equations (4) and (5), the stopping criterion can be formulated in terms of the Lagrange multiplier λ, and of the modulus of the last selected atom Θ,

$$\lambda = \frac{-\Delta D}{\Delta R} = \frac{\Theta^2 - \xi^2_{Q(.)}(\Theta)}{\Delta R_{index} + R_{Q(.)}(\Theta)} \tag{6}$$

where Θ is the threshold modulus beyond which it is worthless to select an additional atom, and $\xi_{Q(.)}(\Theta)$ is the quantization error of modulus Θ. To achieve a rate constrained optimal representation, atoms have to be selected until no atom larger than the threshold Θ can be found on the residual signal. An important observation is that when the signal to expand is partitioned into smaller subspaces for complexity reasons, the stopping criteria has to be met in every subspace. Note also that, for a given λ, Θ depends on the quantizer $Q(.)$. A condition for overall optimality is to design a quantizer achieving R/D optimality for the same Lagrange multiplier λ. The quantizer design is investigated in the next section.

3. ATOM MODULUS QUANTIZER DESIGN

At constant number of atoms, a coarser modulus quantization increases the distortion D but decreases the bit budget R of the MP expansion. The goal of this section is to find the quantizer that minimizes the Lagrangian cost function $\mathcal{L}(\lambda) = D + \lambda R$ for a given multiplier λ. We consider independent scalar quantization of atoms, which is appropriate when atoms are coded in a random order of modulus. As the dead-zone subtracted atom modulus distribution closely fits an exponential model [3], we first review the entropy-constrained, or R/D optimal, quantization for exponentially distributed random variables. Then, we demonstrate that, due to the re-injection of the quantization error into the MP expansion loop, the quantization error of the atoms that are selected during the initial MP iterations are partly corrected by latter MP iterations. The benefit of the quantization error re-injection is taken into account to determine the uniform quantizer stepsize. The re-injection also motivates the design of a novel non-uniform quantization scheme.

3.1. Entropy-constrained scalar quantization

It has been shown that the dead-zone subtracted MP atom modulus distribution closely matches an exponential distribution [3]. Formally, if the MP expansion selects atoms up to a threshold Θ, the Θ-substracted modulus random variable M_Θ has an exponentially decreasing probability distribution function $f_{M_\Theta}(.)$. As an exponentially distributed random variable is memoryless, it has a uniform optimal quantizer [4], i.e. all its steps have the same size Ω. For simplicity and also because a more precise derivation has not improved the results [5], we approximate the R/D optimal Ω with a high-resolution analysis, i.e. by considering quantization with small bin width Ω. This high-resolution assumption is made all along the paper. In this case, the quantizer becomes a mid-quantizer, and the squared-error distortion and entropy are approximately given by [6]

$$D_{M_\Theta}(\Omega) \approx \Omega^2/12 \tag{7}$$

$$H_{M_\Theta}(\Omega) \approx -\int f_{M_\Theta}(x) \log_2 f_{M_\Theta}(x)dx - \log_2 \Omega \tag{8}$$

Using a Lagrange multiplier λ to define the incremental relative importance of distortion and entropy, the optimal entropy-constrained scalar quantizer (ECSQ) is the one which minimizes the objective function $J_{M_\Theta} = D_{M_\Theta} + \lambda H_{M_\Theta}$ [4], i.e. that solves

$$\frac{\partial D_{M_\Theta}(\Omega)}{\partial \Omega} + \lambda \frac{\partial H_{M_\Theta}(\Omega)}{\partial \Omega} = 0 \tag{9}$$

Inserting (7) and (8) into (9) defines the R/D optimal stepsize Ω to quantize an exponentially random variable as a function of λ, i.e.

$$\lambda = -\frac{\partial D_{M_\Theta}}{\partial H_{M_\Theta}} = \frac{\ln(2).\Omega^2}{6} \tag{10}$$

3.2. In-loop quantization of MP atom modulus

To improve coding efficiency, *in-loop* quantization re-injects the quantization error in the MP expansion process so that it can be corrected by subsequent iterations. Here, we analyze the evolution of the energy of the MP residue along the quantized expansion process in order to understand the impact of the re-injection on the distortion of the quantized moduli. Using the notations introduced in Section 2, x is a given source vector in R^M and $R_i x$ is the expansion residue after i MP iterations. At the initial step, $R_0 x = x$ is the signal to expand. At step i, the algorithm selects the dictionary function φ_{k_i} that best matches $R_i x$ and generates the residue for the next iteration. If $\hat{\alpha}_i$ is defined to be the quantized value of $\alpha_i = <R_i x, \varphi_{k_i}>$, then the residue based on the quantized modulus is given by $R_{i+1} x = R_i x - \hat{\alpha}_i \varphi_{k_i}$. Defining σ_i^2 as the quantization distortion on α_i, and observing that $R_i x - \alpha_i \varphi_{k_i}$ is orthogonal to φ_{k_i}, with $||.||$ referring to the quadratic norm in R^M, we have

$$||R_{i+1}x||^2 = ||R_i x||^2 - \alpha_i^2 + \sigma_i^2 \tag{11}$$

Without loss of generality, we can assume that at step i a fraction γ_i of the residual signal energy is captured by the MP expansion, i.e. $\alpha_i^2 = \gamma_i ||R_i x||^2$. Then, we have

$$
\begin{aligned}
||R_1 x||^2 &= ||R_0 x||^2 - \alpha_0^2 + \sigma_0^2 \ with \ \alpha_0^2 = \gamma_0 ||R_0 x||^2 \\
&= (1 - \gamma_0)||R_0 x||^2 + \sigma_0^2 \\
||R_2 x||^2 &= ||R_1 x||^2 - \alpha_1^2 + \sigma_1^2 \ with \ \alpha_1^2 = \gamma_1 ||R_1 x||^2 \\
&= (1 - \gamma_0)(1 - \gamma_1)||R_0 x||^2 + (1 - \gamma_1)\sigma_0^2 + \sigma_1^2 \\
||R_{n+1} x||^2 &= \Pi_{i=0}^n (1 - \gamma_i)||R_0 x||^2 + \sum_{i=0}^n \left(\Pi_{j=i+1}^n (1 - \gamma_j) \right) \sigma_i^2
\end{aligned} \tag{12}
$$

In the above equation, the factor $\left(\Pi_{j=i+1}^n (1 - \gamma_j) \right)$ multiplying σ_i^2 is a direct consequence of re-injecting the quantization error into the MP expansion loop. This factor is smaller than one, showing that re-injection reduces the impact of quantization error on the final distortion $||R_{n+1} x||^2$.

3.3. Uniform quantization

For uniform quantization of exponentially distributed modulus, the quantization distortion is the same for all coefficients [4]. So, in (12), the distortion σ_i^2 is independant of i and, following the notations of Section 3.1, can be denoted by $D_{M_\Theta}(\Omega)$. When a total of N atoms are selected, the distortion D_{UQ} resulting from the *in-loop* uniform quantization of these atoms is the second term of the sum in (12) and can be written as

$$D_{UQ}(\Omega) = \sum_{i=0}^{N-1} \left(\Pi_{j=i+1}^{N-1}(1 - \gamma_j) \right) D_{M_\Theta}(\Omega) \qquad (13)$$

The entropy H_{UQ} for N atoms is N times the Θ-subtracted modulus entropy, i.e. $H_{M_\Theta}(\Omega)$ in (8). For a given N and λ, the R/D optimal quntizer stepsize Ω solves

$$\lambda = -\frac{\partial D_{UQ}}{\partial H_{UQ}} = -\frac{\sum_{i=0}^{N-1} \left(\Pi_{j=i+1}^{N-1}(1-\gamma_j) \right)}{N} \cdot \frac{\partial D_{M_\Theta}}{\partial H_{M_\Theta}} \qquad (14)$$

Combining (14) with (10) and introducing $\beta(N)$ to denote the first factor in (14), we get

$$\lambda = \beta(N) \cdot \frac{\ln(2) \cdot \Omega^2}{6} \qquad (15)$$

Equation (15) specifies the relationship between λ and Ω assuming $\beta(N)$ is known. The multiplication by a factor $\beta(N)$ that is smaller than one is due to re-injection. (15) shows that Ω increases as $\beta(N)$ decreases, indicating that taking re-injection into account results in larger quantization step than otherwise. To estimate $\beta(N)$, we note that by definition, all γ_i are much smaller than one, and we approximate them by their mean value, denoted by $\overline{\gamma}$. $\beta(N)$ then becomes

$$\beta(N) \approx \frac{\sum_{i=0}^{N-1}(1 - \sum_{j=i+1}^{N-1}\gamma_i)}{N}$$
$$\approx (1 - N\overline{\gamma}/2) \approx \sqrt{(1 - N\overline{\gamma})} \qquad (16)$$

In (12), by neglecting the quantization distortion in comparison with the energy of the residual signal, we also have

$$||R_N x||^2 \approx \Pi_{i=0}^{N-1}(1 - \gamma_i)||R_0 x||^2$$
$$||R_N x||^2 \approx (1 - N\overline{\gamma})||R_0 x||^2 \qquad (17)$$

From (16) and (17), $\beta(N)$ is estimated in terms of the ratio between the initial and final MP residual energy, i.e. $\beta(N) \approx \sqrt{||R_N x||^2 / ||R_0 x||^2}$. Since an optimal expansion selects the atoms up to a threshold Θ, the final residual energy $||R_N x||^2$ is also the energy E_Θ after all atoms larger than Θ have been selected. In that sense, $\beta(N)$ is a function of Θ and we have

$$\beta(N) \approx \sqrt{\frac{E_\Theta}{||R_0 x||^2}} \triangleq \beta'(\Theta) \qquad (18)$$

In practice, the final residual energy is only known once the expansion has been performed, which means that $\beta'(\Theta)$ can only be estimated a posteriori. However, as *in-loop* quantization performs the quantization along the expansion process, the stepsize Ω has to be known ahead of the expansion. Thus, we face a chicken-and-egg problem. Experiments have shown that for an arbitrary signal, $\beta'(\Theta)$ can be estimated from an initial expansion that is computed without quantization. For video coding applications, $\beta'(\Theta)$ can be estimated from the expansion obtained on the previous frame [5].

To sum up, for a given Lagrange multiplier λ, (15) specifies the stepsize to be used for quantizing atoms. In Section 2, for the same λ, (6) indicates when to stop the iterative MP expansion. Here, we combine these results to derive an explicit relation between the expansion stopping threshold Θ, and the quantizer stepsize Ω. We

begin by simplifying (6). As a direct consequence of the quantizer design, the quantization error on the last atom, i.e. $\xi_{Q(.)}(\Theta)$, is equal to $\Omega/2$. Experimentaly we have also observed that the cost in bits of the last selected atom, i.e. $R_{last} = \Delta R_{index} + R_{Q(.)}(\Theta)$, does not significantly depend on the stopping threshold Θ. Under these simplifications, combining (6) and (15) results in

$$\frac{\Theta^2 - \Omega^2/4}{R_{last}} = \beta'(\Theta) \cdot \ln(2) \cdot \Omega^2/6 \qquad (19)$$

For a given Θ, (19) gives the Ω that results in a R/D optimal expansion. For video coding, $R_{last} \in [17 \dots 19]$, and hence

$$\frac{\Omega}{\Theta} \approx 0.66 \sqrt{\beta'(\Theta)^{-1}} \qquad (20)$$

For $\beta'(\Theta) = 1$, this result is very close to the 0.6 empirical ratio found in [3]. The analytical derivation helps us understand the assumptions under which this result is valid and refines it when the impact of re-injection becomes significant, i.e. $\beta'(\Theta) >> 1$.

3.4. Non-uniform quantization

In the previous paragraph we have shown that re-injection affects the optimal uniform quantizer stepsize. Now we consider the design of the entropy constrained quantizer in presence of re-injection but in absence of uniformity constraint. Non-uniformity is justified by the observation that atoms are roughly selected in decreasing order of magnitude. As a consequence, the initial atoms, which have more chance to be corrected by subsequent MP iterations, are also the largest ones. This suggests increasing the quantization stepsize with the atom magnitude. Designing the R/D optimal non-uniform quantizer consists of fixing its bin boundaries, or equivalently the sequence of stepsizes $\{\Omega_i\}_{i>0}$, so that all bins have the same incremental benefit in distortion for a given incremental cost in rate, the ratio between them being defined by the Lagrangian multiplier λ. We propose a recursive approach: at each step, given the lower boundary of a quantizer bin, the upper boundary is determined. From Section 2, the lower boundary of the first bin is the stopping threshold Θ. Let us now consider the selection of the upper boundary of the i^{th} quantizer bin, assuming its lower boundary is known. The problem is first solved in an abstract way in order to formulate the conditions for optimality in terms that are manageable by a practical algorithm. Let n_1 denote the hypothetical iteration index for which the atom modulus α_{n_1} equals the lower boundary of the i^{th} quantizer bin. For any hypothetical iteration index $n_2 < n_1$, define α_{n_2} to be the selected atom modulus. As atoms are selected in decreasing order of magnitude, $\alpha_{n_2} > \alpha_{n_1}$ and the set of atoms selected between the n_2^{th} and n_1^{th} iterations of the MP expansion belong to $[\alpha_{n_1}, \alpha_{n_2}[$. Without loss of generality our problem is to find α_{n_2}, or equivalently n_2, so that the best Lagrangian R/D trade-off is achieved for the quantization of the $(n_1 - n_2)$ atoms belonging to $[\alpha_{n_1}, \alpha_{n_2}[$. Letting N be the hypothetical total number of atoms selected by the MP expansion up to the Θ threshold, the distortion $D_i(N, n_1, n_2)$ due to the quantization of the $(n_1 - n_2)$ atoms belonging to the i^{th} bin $[\alpha_{n_1}, \alpha_{n_2}[$ with a stepsize $\Omega_i = \alpha_{n_2} - \alpha_{n_1}$ is an immediate consequence of (12) and can be written

$$D_i(N, n_1, n_2) = \sum_{k=n_2+1}^{n_1} \left(\Pi_{j=k+1}^{N-1}(1 - \gamma_j) \right) D_{M_\Theta}(\Omega_i) \qquad (21)$$

Similar to (14) and (15), given λ, the optimal i^{th} bin quantizer stepsize Ω_i solves

$$\lambda = -\frac{\sum_{k=n_2+1}^{n_1} \left(\Pi_{j=k+1}^{N-1}(1 - \gamma_j) \right)}{n_1 - n_2} \cdot \frac{\partial D_{M_\Theta}}{\partial H_{M_\Theta}}$$
$$= \beta_i(N, n_1, n_2) \cdot \frac{\ln(2) \cdot \Omega_i^2}{6} \qquad (22)$$

In (22), Ω_i is the width of the i^{th} bin. Given the lower bound of the bin, it defines the upper bound as a function of λ. For small γ_i, $\beta_i(N, n_1, n_2)$ can be approximated as

$$\beta_i(N, n_1, n_2) \approx (1 - \sum_{j=n_1+1}^{N-1} \gamma_j) \cdot \frac{\sum_{i=n_2+1}^{n_1}(1 - \sum_{j=i+1}^{n_1}\gamma_j)}{n_1 - n_2} \qquad (23)$$

Developments similar to the ones in (16) and (17) give

$$(1 - \sum_{j=n_1+1}^{N-1} \gamma_j) \quad \approx \quad \frac{||R_N x||^2}{||R_{n_1} x||^2} \qquad (24)$$

$$\frac{\sum_{k=n_2+1}^{n_1}(1 - \sum_{j=k+1}^{n_1}\gamma_j)}{n_1 - n_2} \quad \approx \quad \sqrt{\frac{||R_{n_1} x||^2}{||R_{n_2} x||^2}} \qquad (25)$$

and

$$\beta_i(N, n_1, n_2) \approx \frac{||R_N x||^2}{||R_{n_1} x||^2} \sqrt{\frac{||R_{n_1} x||^2}{||R_{n_2} x||^2}} = \frac{||R_N x||^2}{\sqrt{||R_{n_1} x||^2 \cdot ||R_{n_2} x||^2}} \qquad (26)$$

This factor depends on the residual signal energies after respectively n_1, n_2 and N MP iterations, with $N > n_1 > n_2$. To design the non-uniform quantizer, we use an alternative interpretation of these energy values. We note that each of them is the energy of the residual signal after all atoms larger than a particular threshold have been selected. $||R_{n_1} x||^2$ and $||R_{n_2} x||^2$ are measured once the atoms larger than the lower and higher boundary of the i^{th} bin have been selected respectively. $||R_N x||^2$ measures the energy of the final residue, i.e. once all atoms larger than the stopping threshold Θ have been selected. Formally, let θ_i be the lower boundary of the i^{th} bin, i.e.

$$\theta_i = \Theta + \sum_{l=1}^{i-1} \Omega_l^{j+1} \qquad (27)$$

Define E_θ to be the energy of the residue after all atoms larger than θ have been selected. E_Θ denotes the final residual energy. E_{θ_i} denotes the energy after all atoms larger than the lower boundary of the i^{th} bin have been selected. So, (26) can be written as

$$\beta_i(N, n_1, n_2) \approx \frac{E_\Theta}{\sqrt{E_{\theta_i} \cdot E_{\theta_{i+1}}}} \triangleq \beta_i'(\Theta, \theta_i, \theta_{i+1}) \triangleq \beta_i'(\Theta) \qquad (28)$$

For notation convenience, $\beta_i'(\Theta, \theta_i, \theta_{i+1})$ is simply refered to as $\beta_i'(\Theta)$ in the following. Similar to arguments used in deriving (20), (6) can be combined with (22) to arrive at

$$\frac{\Omega_i}{\Theta} \approx 0.66 \cdot \sqrt{\beta_i'(\Theta)^{-1}} \qquad (29)$$

We note that $\beta_i'(\Theta)$ decreases as i increases, which indicates that the stepsize Ω_i increases with the quantizer bin index i. This is to be expected as quantization error of large moduli atoms in early MP iterations is corrected by subsequent iterations. The energy values needed to estimate $\{\beta_i'(\Theta)\}_{i>0}$ in (28) are derived either from a non-quantized expansion of the current signal or from the previous frame of the video sequence, thus avoiding excessive computation. Specifically, the R/D optimal expansion is computed as follows. Let $\{\beta_i'^j(\Theta)\}_{i>0}$ be the β' sequence derived for the j^{th} frame of the video sequence. The process is initialized with $\beta_i'^0(\Theta) = 1$ for all i. At frame $j + 1$, for a given Θ, $\{\Omega_i^{j+1}\}_{i>0}$ is computed from (29), using $\{\beta_i'^j(\Theta)\}_{i>0}$, and subsequently the signal expansion is performed with *in-loop* quantization. In order to estimate the sequence $\{\beta_i'^{j+1}(\Theta)\}_{i>0}$ for the next frame, the final residual energy and the residual energy after all atoms larger than the boundaries of each quantization bin have been selected are computed. This computation is performed by accumulating the atom contribution to the residue energy decrease on a quantizer bin basis [5]. For the k^{th} atom, this contribution is $\alpha_k^2 - (\alpha_k - \hat{\alpha}_k)^2$, where α_k and $\hat{\alpha}_k$ are the non-quantized and quantized moduli.

4. APPLICATION TO VIDEO CODING

In Fig.1, three uniform and one non-uniform quantizers are compared. All uniform quantizer stepsizes have been selected according to equation (20), i.e. $\Omega/\Theta = 0.66 \sqrt{\beta'(\Theta)^{-1}}$. For both UQ, IS and UQ, the quantization error re-injection is neglected, i.e. $\beta' = 1$. They are different in that UQ, IS stops the search as soon as an atom smaller than Θ has been encountered, while UQ continues the search until no atoms larger than Θ can be found in any subspace. UQ, IS is the approach used in [3] while UQ conforms to the stopping criteria expressed in Section 2. The third uniform quantizer ULQ considers the quantization error re-injection and sets β' according to (18). The non-uniform quantizer, $NULQ$, is designed based on (29) and (28). As expected, $NULQ$ outperforms all uniform quantizers. However, at low bitrates the improvement is negligible. This is because most of the bits are devoted to motion vector coding, and few atoms are encoded. On the contrary, at high bitrates, the non-uniform quantizer results in 0.5 to 2 dBs improvement over [3].
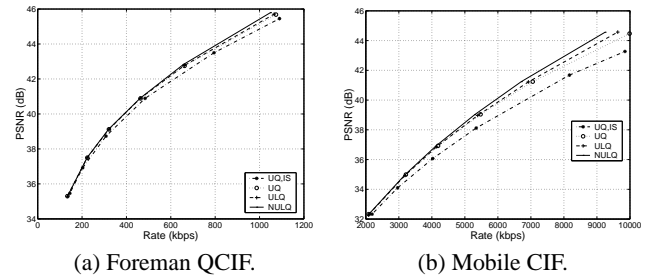

(a) Foreman QCIF.   (b) Mobile CIF.

Fig. 1. R/D curves for video sequences encoded at 30 fps.

5. CONCLUSIONS

We have presented a R/D analysis of the MP expansion and quantization for systems that transmit atoms in random order of magnitude. Our study shows that, at optimality, atoms are selected up to a quality-dependent threshold. It also validates the characterization of the quantizer proposed in [3], in terms of the ratio between its stepsize and the expansion stopping threshold. Moreover, we have shown that due to the re-injection of the quantization error during the MP expansion, the R/D optimal quantizer is non-uniform.

6. REFERENCES

[1] O. Al-Shaykh, E. Miloslavsky, T. Nomura, R. Neff, and A. Zakhor, "Video compression using matching pursuits," *IEEE Transactions on CSVT*, vol. 9, no. 1, pp. 113–143, February 1999.

[2] P. Frossard and P. Vandergheynst, "A posteriori quantized matching pursuit," in *Data Compression Conference*, Snowbird, UT, March 01.

[3] R. Neff and A. Zakhor, "Modulus quantization for MP video coding," *IEEE Trans. on CSVT*, vol. 10, no. 6, pp. 895–912, September 2000.

[4] G.J. Sullivan, "Efficient scalar quantization of exponential and laplacian random variables," *IEEE Transactions on Information Theory*, vol. 42, no. 5, pp. 1365–1374, September 1996.

[5] C. De Vleeschouwer and A. Zakhor, "In-loop atom modulus quantization for MP and its application to video coding," *IEEE Transactions on Information Theory*, To be submitted.

[6] R.M. Gray and D.L. Neuhoff, "Quantization," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 1–63, October 1998.