Multi-resolution Source Coding Using Entropy Constrained Dithered Scalar Quantization *

Qian Zhao Hanying Feng Michelle Effros

Abstract

In this paper, we build multi-resolution source codes using entropy constrained dithered scalar quantizers. We demonstrate that for n-dimensional random vectors, dithering followed by uniform scalar quantization and then by entropy coding achieves performance close to the n-dimensional optimum for a multi-resolution source code. Based on this result, we propose a practical code design algorithm and compare its performance with that of the Set Partitioning in Hierarchical Trees (SPIHT) algorithm on natural images.

I Introduction

The apparent divergence in the lossy coding literature between theoretical arguments and practical algorithms may perhaps be attributed to the difficulty of the problem. Entropy constrained dithered quantization (ECDQ) [1, 2], which uses simple methods to achieve performance that is provably close to the theoretical optimum, represents a rare tool in the struggle to close that divide. In this paper, we build a multi-resolution ECDQ and demonstrate its potential by deriving theoretical performance bounds and giving practical implementation results. We begin with some definitions.

Rate-distortion function: The (per-symbol) rate-distortion function for an n-dimensional random source X^n is

$$R_n(D) = \inf_{\{f_{U^n|X^n}(u^n|x^n):\ (1/n)Ed(X^n,U^n)\leq D\}} \frac{1}{n} I(X^n;U^n),$$

where $f_{U^n|X^n}(u^n|x^n)$ is the conditional probability density function (pdf) of the representation u^n given the source sample x^n, $I(X^n;U^n)$ is the mutual information between X^n and U^n, and $(1/n)Ed(X^n,U^n)$ is the per-symbol expected distortion between X^n and U^n. We use $d(x,\hat{x}) = (x-\hat{x})^2$ throughout.

Entropy constrained dithered quantization: Let $Q : \mathcal{R} \rightarrow \mathcal{R}$ be a uniform scalar quantizer with step size Δ (giving $Q(x) = i\Delta$ for all $x \in (i\Delta - \Delta/2, i\Delta + \Delta/2]$) and Z be a "dither" random variable uniformly distributed on $(-\Delta/2, \Delta/2]$. For the

*Q. Zhao (qian.zhao@oracle.com) is with Oracle Corp., Redwood City, CA 94065. H. Feng (fhy@stanford.edu) is with the Department of Electrical Engineering at Stanford University, Stanford, CA 94305. M. Effros (effros@z.caltech.edu) is with the Department of Electrical Engineering (MC 136-93) at the California Institute of Technology, Pasadena, CA 91125. This material is based upon work partially supported by NSF Grant No. CCR-0220039 and the Caltech's Lee Center for Advanced Networking.

Proceedings of the Data Compression Conference (DCC'04)
1068-0314/04 $ 20.00 © 2004 IEEE

theory, we assume that the encoder and decoder both know Z. For each sample X_i of source X, the encoder describes $Q(X_i + Z)$ using a universal conditional entropy code conditioned on Z. The decoder decodes the description of $Q(X_i + Z)$ and builds reconstruction $\hat{X}_i = Q(X_i + Z) - Z$. (We here describe scalar ECDQ; vector ECDQ replaces scalar dither and quantization with vector equivalents. We treat only the scalar case.) From [3], we have:

1. Given $\hat{X}_i = Q(X_i + Z) - Z$, $\hat{X}_i - X_i$ is independent of X_i and is distributed as $-Z$, giving expected distortion $D = E(\hat{X}_i - X_i)^2 = EZ^2 = \Delta^2/12$.

2. Using techniques like block coding, the rate of an ECDQ can be made arbitrarily close to the corresponding conditional entropy. Hence we measure the rate of the ECDQ as $H(Q(X_i+Z)|Z) = I(X_i; X_i - Z) = I(X; X - Z)$ when $X_1, X_2, \ldots$ come from a stationary source.

Multi-resolution source code: An M-resolution source code has M encoder maps $f_m : \mathcal{R}^n \to \{1, \ldots, L_m\}$ and M decoder maps $g_m : \{1, \ldots, L_1\} \times \cdots \times \{1, \ldots, L_m\} \to \mathcal{R}^n$. We use R_m and D_m to denote the distortion and (total) rate at resolution m. Thus $D_m = (1/n)Ed(X^n, g_m(f_1(X^n), \ldots, f_m(X^n)))$; rate R_m depends on the techniques (e.g., fixed-rate or entropy coding) used to describe the encoder output.

We here generalize ECDQ to multi-resolution source coding. We define the *rate redundancy* at resolution m of an MR-ECDQ as the difference between the rate R_m achieved by the MR-ECDQ in resolution-m and the rate-distortion function $R_n(D_m)$ at the same distortion D_m. (Notice that we cannot generally achieve $R_m - R_n(D_m) = 0$ for all m simultaneously [4].) In Section II, we show that for any source and any distortions $D_1 \geq \cdots \geq D_M$, MR-ECDQ achieves $R_m - R_n(D_m) < 1.755$ for all m. We also give tighter bounds for a restricted class of distortions and a code that allows time sharing. In Section III, we give a practical algorithm for implementing MR-ECDQ and compare its empirical performance with that of SPIHT.

II Code Design and Theoretical Performance Analysis

We build an MR-ECDQ using a single dither random variable Z_1 and a family of nested uniform scalar quantizers. Let Z_1 be uniformly distributed on $(-\Delta_1/2, \Delta_1/2]$. For each $m \in \{1, \ldots, M\}$, let K_m be a positive integer, $\Delta_m = \Delta_{m-1}/K_m$, and Q_m be the uniform scalar quantizer with basic cell $(-\Delta_m/2, \Delta_m/2]$; the result is a collection of quantizers with nested encoding regions. In resolution m, the MR-ECDQ uses a universal conditional entropy code to describe $Q_m(X_i + Z_1)$ given $Q_1(X_i + Z_1), \ldots, Q_{m-1}(X_i + Z_1)$ and Z_1. To achieve rates close to the entropy, we again allow multiple symbols to be entropy coded together (e.g., with an arithmetic code). In resolution m, the decoder builds reconstruction $\hat{X}_i = Q_m(X_i + Z_1) - Z_1$. We begin by bounding the rate redundancy for the resulting nested MR-ECDQ.

Theorem 1 *Let $(D_1, \ldots, D_M)$ and $(R_1, \ldots, R_M)$ be the distortions and total rates at resolutions $1, \ldots, M$ of a nested MR-ECDQ, where $D_m = D_{m-1}/K_m^2$ for some integer K_m. Then $R_m - R_n(D_m) < 0.755$ for all $m \in \{1, \ldots, M\}$ and any source.*

Proceedings of the Data Compression Conference (DCC'04)
1068-0314/04 $ 20.00 © 2004 IEEE

Proof. Assume we use block entropy coding with block length n. Only nested distortion values $D_m = D_{m-1}/K_m^2$ can be obtained via the nested MR-ECDQ,

$$D_m = \frac{1}{12}\Delta_m^2 = \frac{1}{12}\frac{\Delta_{m-1}^2}{K_m^2} = \frac{1}{K_m^2}D_{m-1} = \frac{1}{\prod_{j=2}^m K_j^2}D_1 = \frac{1}{12\prod_{j=2}^i K_j^2}\Delta_1^2.$$

Define $Z_m = Z_1 - Q_m(Z_1)$ $(m = 2,\ldots,M)$; then Z_m is uniformly distributed on $(-\Delta_m/2, \Delta_m/2]$. Let $Z_m^n = (Z_m,\ldots,Z_m)$, and use $Q_m(X^n + Z_1^n) = (Q_m(X_1 + Z_1),\ldots,Q_m(X_n + Z_1))$ to denote a vector of n quantized values; then by induction,

$$\begin{aligned}
R_1 &= (1/n)H(Q_1(X^n + Z_1^n)|Z_1^n) = (1/n)I(X^n; X^n - Z_1^n) && (1)\\
R_m &= R_{m-1} + (1/n)H(Q_m(X^n + Z_1^n)|Z_1^n, Q_1(X^n + Z_1^n),\ldots,Q_{m-1}(X^n + Z_1^n))\\
&= (1/n)H(Q_m(X^n + Z_1^n), Q_{m-1}(X^n + Z_1^n),\ldots,Q_1(X^n + Z_1^n)|Z_1^n)\\
&= (1/n)H(Q_m(X^n + Z_1^n)|Z_1^n) && (2)\\
&= (1/n)I(X^n; X^n - Z_m^n) && (3)
\end{aligned}$$

where (1) and (3) come from Property 2 in Section I, and (2) holds since $Q_m(X^n+Z_1^n)$ uniquely determines $Q_j(X^n + Z_1^n)$ for any j such that $1 \le j < m$. Let U_m^n be the n-dimensional vector achieving $R_n(D_m)$ for source X^n. Then by [2, Theorem 2], we can bound the rate redundancy of the nested MR-ECDQ at resolution m as

$$R_m - R_n(D_m) = (1/n)I(X^n; X^n - Z^n) - (1/n)I(X^n; U^n) < 0.755. \quad \square$$

The following results treat the case where $D_1,\ldots,D_M$ do not satisfy the above structure. Lemma 1 [5] is useful in obtaining these results.

Lemma 1 *[5, Lemma 1] Let $R_n(D)$ be the rate-distortion function, then for any $0 < D_1 \le D_2$, $R_n(D_1) - R_n(D_2) \le (1/2)\log(D_2/D_1)$.*

Theorem 2 *Let $(D_1,\ldots,D_M)$ and $(R_1,\ldots,R_M)$ be the target expected distortions and total rates at resolution $1,\ldots,M$ of an MR-ECDQ. Then $R_m - R_n(D_m) < 1.755$ for all $m \in \{1,\ldots,M\}$ and any source.*

Proof. Given target distortions $D_1 \ge D_2 \ge \ldots \ge D_M$ at resolutions $1,\ldots,M$, we design a nested MR-ECDQ code with distortion $D'_m \le D_m$ at each resolution m. Set $k(1) = 1$ and for each $m > 1$ choose integer $k(m)$ to satisfy $D_1/4^{k(m)} \le D_m < D_1/4^{k(m)-1}$. Let $D'_m = D_1/4^{k(m)}$. From Theorem 1, we can achieve $R_m - R_n(D'_m) < 0.755$ for each $m \in \{1,\ldots,M\}$. Thus by Lemma 1

$$\begin{aligned}
R_m - R_n(D_m) &= R_m - R_n(D'_m) + R_n(D'_m) - R_n(D_m)\\
&< 0.755 + \frac{1}{2}\log\frac{D_m}{D'_m}\\
&< 0.755 + \frac{1}{2}\log 4,
\end{aligned}$$

since $D'_m = D_1/4^{k(m)} \le D_m < D_1/4^{k(m)-1} = 4D'_m$. $\square$

We can tighten the above bound by allowing the time-sharing between the resolutions of a nested MR-ECDQ. Lemma 2 is useful for that bound.

Lemma 2 *Suppose* $0 < D_2/4 \le D_1 < D_2$ *and* $D = \alpha D_1 + (1-\alpha)D_2$*, where* $0 \le \alpha \le 1$*, then* $\alpha R_n(D_1) + (1-\alpha)R_n(D_2) - R_n(D) < C_1 = 0.3414$.

Proof:

$$\begin{aligned}
&\alpha R_n(D_1) + (1-\alpha)R_n(D_2) - R_n(D) \\
&= \alpha(R_n(D_1) - R_n(D)) + (1-\alpha)(R_n(D_2) - R_n(D)) \\
&\le \alpha(R_n(D_1) - R_n(D)) && (4)\\
&\le \frac{\alpha}{2}\log\frac{D}{D_1} && (5)\\
&= \frac{\alpha}{2}\log\left(\alpha + (1-\alpha)\frac{D_2}{D_1}\right) \le \frac{\alpha}{2}\log(4-3\alpha) && (6)\\
&< 0.3414, && (7)
\end{aligned}$$

where (4) follows since $D_2 \ge D$ and the rate-distortion function is a non-increasing function; (5) follows from Lemma 1; (6) follows from the definition of D since $D_2/4 \le D_1$; and (7) comes from the maximization of $\alpha\log(4-3\alpha)$ over $\alpha \in [0,1]$. □

Theorem 3 *Let* $(D_1, \ldots, D_M)$ *and* $(R_1, \ldots, R_M)$ *be the target expected distortions and total rates at resolution* $1, \ldots, M$ *of an MR-ECDQ that allows time-sharing. Then* $R_m - R_n(D_m) < 1.097$ *for all* $m \in \{1, \ldots, M\}$ *and any source.*

Proof. For $1 \le m \le M$, choose integer $k(m)$ and constant $\alpha_m \in (0,1]$ such that

$$\frac{D_1}{4^{k(m)}} < D_m \le \frac{D_1}{4^{k(m)-1}}, \quad \text{and} \quad \alpha_m\frac{D_1}{4^{k(m)-1}} + (1-\alpha_m)\frac{D_1}{4^{k(m)}} = D_m.$$

Since $D_1 \ge \ldots \ge D_M$, $1 = k(1) \le \ldots \le k(M)$. We use time sharing between the resolutions of an $(k(M)+1)$-resolution nested MR-ECDQ. At resolution m, the nested MR-ECDQ achieves distortion $D'_m = D_1/(4^{m-1})$ and rate R'_m satisfying $R'_m - R_n(D'_m) < 0.755$ by Theorem 1. For notational simplicity, we use Q_m to denote the output indices of this MR-ECDQ's scalar quantizer at resolution m.

First suppose that $M = 2$. We use Q_1 in the first resolution description. In resolution 2, we time-share between $Q_{k(2)}$ and $Q_{k(2)+1}$ with proportions α_2 and $1-\alpha_2$, respectively. The resulting rate and distortion at resolution 1 are $(R_1, D_1) = (R'_1, D'_1)$; hence, $R_1 - R_n(D_1) = R'_1 - R_n(D'_1) < 0.755$. The rate and distortion at resolution 2 are $(R_2, D_2) = (\alpha_2 R'_{k(2)} + (1-\alpha_2)R'_{k(2)+1}, \alpha_2 D'_{k(2)} + (1-\alpha_2)D'_{k(2)+1})$; hence

$$\begin{aligned}
R_2 - R_n(D_2) &= \alpha_2 R'_{k(2)} + (1-\alpha_2)R'_{k(2)+1} - R_n(D_2)\\
&= \alpha_2(R'_{k(2)} - R_n(D'_{k(2)})) + (1-\alpha_2)(R_{k(2)+1} - R_n(D'_{k(2)+1}))\\
&\quad + \alpha_2 R_n(D'_{k(2)}) + (1-\alpha_2)R_n(D'_{k(2)+1}) - R_n(D_2)\\
&< 0.755 + \alpha_2 R_n(D'_{k(2)}) + (1-\alpha_2)R_n(D'_{k(2)+1}) - R_n(D_2) < 1.097,
\end{aligned}$$

where the last inequality follows from Lemma 2.

We use an inductive argument to go from $M = 2$ to general multi-resolution codes. Given a code that meets the desired distortion constraints on resolutions $1, \ldots, M-1$, we want to modify that code to meet the desired constraints on resolution M as well. The $(M-1)$-resolution code time-shares between $Q_{k(M-1)}$ and $Q_{k(M-1)+1}$ in resolution $M-1$. If $k(M) > k(M-1)$, then time-sharing between $Q_{k(M)}$ and $Q_{k(M+1)}$ in resolution M gives $R_M < R_n(D_M) + 1.097$. If $k(M) = k(M-1)$, then increasing the fraction of the data on which the code uses $Q_{k(M+1)}$ to give a total fraction $(1-\alpha_M)$ achieves the same effect. □

III Code Implementation and Experimental Results

Implementing the MR-ECDQs described in the previous section presents several difficulties. First, by requiring that Z_1 be available at both the encoder and the decoder, the MR-ECDQ algorithm (like the ECDQ algorithm before it) requires shared randomness not generally available at the encoder and decoder of a data compression system. Second, universal conditional entropy codes are not currently available, and even good low-complexity non-universal conditional entropy codes without the universality constraint are difficult to achieve. We next address each of these issues in turn and propose a practical MR-ECDQ implementation.

Code Implementation

In theory, dither random variable Z_1 is a real random variable uniformly distributed on $(-\Delta_1/2, \Delta_1/2]$ and known to both the encoder and decoder. To approximate this random variable in practice, we can generate a pseudo-random value at the encoder and describe that value explicitly to the decoder. Only one pseudo-random number is required since we use the same value of Z_1 for all data samples $X_1, X_2, \ldots$. The cost of an explicit description of that value is small when amortized, for example, over the size of a typical image. (Note that we need perform no explicit quantization of Z_1. Only the implicit quantization inherent in representing a floating point value digitally is required.) While this approach is simple, we here use another method.

Theorems 1–3 bound the expected rate-distortion performance over the distributions of source X and dither Z. If we encapsulate the performance of an MR-ECDQ by the Lagrangian performance measure $\sum_{m=1}^{M}[\alpha_m D_m + \beta_m R_m]$ with constraint $\sum_{m=1}^{M}\beta_m = 1$, then the above theorems apply directly to show that if $J(\alpha^M, \beta^M)$ is the optimal Lagrangian performance theoretically achievable for source X^n and Lagrangian constants (α^M, β^M), then there exists an MR-ECDQ with $\sum_{m=1}^{M}[\alpha_m D_m + \beta_m R_m] \leq c\sum_{m=1}^{M}\beta_m = c$, where $c \in \{.755, 1.097, 1.755\}$ (depending on the derivation assumptions). Given that the expected value of the Lagrangian performance (with respect to the distribution on Z_1) can be made to lie within c of its optimal performance, there must exist a single value of Z that achieves Lagrangian performance at least as good. We therefore fix the value of Z_1 for each image. Through experiments, we develop a rule of thumb for choosing Z_1, as discussed shortly.

Given this strategy, the implementation proceeds as follows. Let $[X_{\min}, X_{\max}]$ be the source alphabet. Set dither value $Z \equiv z_o \in [-X_{\max}, X_{\min}]$. For any length-$L$ data sequence $\{X_\ell\}_{\ell=1}^{L}$, we add z_o from each sample (see Figure 1(a)) and then $X_\ell + z_o$ using uniform scalar quantizer $Q_1(.)$ with basic cell size $\Delta_1 = X_{\max} - X_{\min}$.

The index of the quantized value is $I_1(X) = \lfloor X/\Delta_1 \rfloor$. We use universal entropy code (e.g., Lempel-Ziv-Welch code [6, 7] or Burrows-Wheeler Transform (BWT) coding [8]) on $I_1(X_1+z_o), I_1(X_2+z_o), \ldots, I_1(X_L+z_o)$. The decoder decodes $I_1(X_\ell+z_o)$ and reconstructs $Q_1(X_\ell+z_o)$ as the mid-point of the quantization cell, i.e., $Q_1(X_\ell+z_o) = I_1(X_\ell+z_o)\Delta_1 + \Delta_1/2$. The reconstruction for X_ℓ at the 1st resolution is $\hat{X}_{1,i} = Q_1(X_\ell+z_o) - z_o$, $i = 1, \ldots, L$.

At resolution $m > 1$, we first re-order the symbols in the sequence, so that symbols with identical index values $I_1, I_2, \ldots, I_{m-1}$ are placed at adjacent positions (see Figure 1). The index value $I_m(X_\ell+z_o)$ conditioned on $Q_{m-1}(X_\ell+z_o)$ is

$$I_m(X_\ell+z_o) = \begin{cases} 0 & \text{if } X_\ell+z_o < Q_{m-1}(X_\ell+z_o), \\ 1 & \text{if } X_\ell+z_o \geq Q_{m-1}(X_\ell+z_o). \end{cases}$$

The binary sequence $I_m(X_{\ell(m,1)}+z_o)I_m(X_{\ell(m,2)}+z_o)\ldots I_m(X_{\ell(m,L)}+z_o)$ is then coded with a universal entropy code, where $X_{\ell(m,1)}, \ldots, X_{\ell(m,L)}$is the data sequence after re-ordering. The decoder already knows I_k and Q_k for all X and all $k < m$; hence the decoder can recover $I_m(X+z_o)$ and re-order the data sequence appropriately. The reproduction value is the mid-point of the new quantization cell: $Q_m(X_\ell+z_o) = Q_{m-1}(X_\ell+z_o) + I_m(X_\ell+z_o)\Delta_m - \Delta_m/2$. The reconstruction for X_ℓ at the m-th resolution is $\hat{X}_{\ell,m} = Q_m(X_\ell+z_o) - z_o$, $l = 1, \ldots, L$.

The above procedure is equivalent to quantizing $X_\ell + z_o$ using uniform scalar quantizer $Q_m(.)$ with basic cell size $\Delta_m = \Delta_{m-1}/2$, then coding $Q_m(X_\ell+z_o)$ using a universal conditional entropy code conditioned on $Q_1(X_\ell+z_o), \ldots, Q_{m-1}(X_\ell+z_o)$. By rearranging the positions of pixels according to $I_1, \ldots, I_{m-1}$ (or equivalently $Q_1(X_\ell+z_o), \ldots, Q_{m-1}(X_\ell+z_o)$), a simple universal entropy code can capture the statistics of $Q_m(X_\ell+z_o)$ given $Q_1(X_\ell+z_o), \ldots, Q_{m-1}(X_\ell+z_o)$ and achieve performance close to the optimum.

We experiment with a few variations to further improve the performance of our algorithm.

1. We can scan the image in zigzag (rather than standard raster) scan order (see Figure 2).

2. Set $Q_m(\cdot)$ as the centroid (rather than the mid-point) of the quantization cell and use non-uniform scalar quantization for the m-th ($m \geq 1$) resolution.

3. Experimentally we find $z_o = (X_{\max} - X_{\min})/2 - \sum_{m=1}^{L} X_m/L$ to be a good value for the dither.

Experimental Results

We implement the algorithm of previous section and test it on a variety of images. We use both the Lempel-Ziv-Welch (LZW) code and BZIP (an algorithm based on the BWT) as lossless codes. We compare the rate-distortion performance of the variations on the above algorithm with each other and with the performance of the SPIHT algorithm [9]. The images we test include Barbara, Crowd, Airport, Lena, and a set of medical brain-scan images.

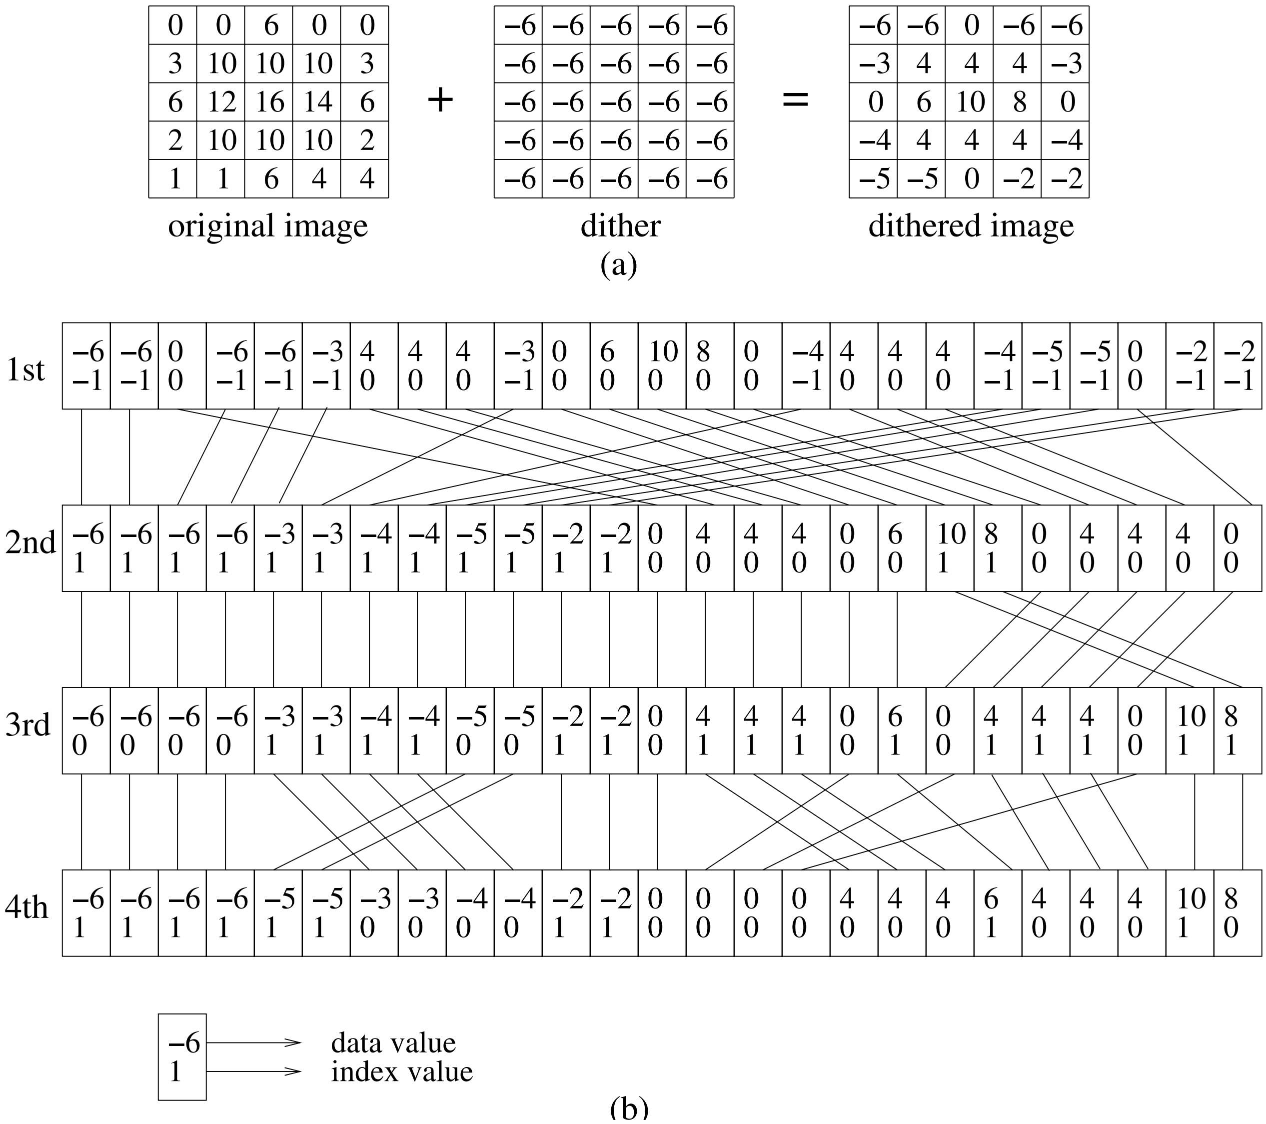

Figure 1: MR-ECDQ implementation: (a) original image and dithered image; (b) rearranging pixels positions in multi-resolution coding.

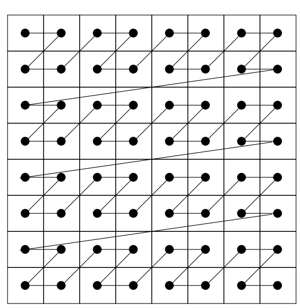

Figure 2: The zigzag scanning pattern.

Proceedings of the Data Compression Conference (DCC'04)
1068-0314/04 $ 20.00 © 2004 IEEE

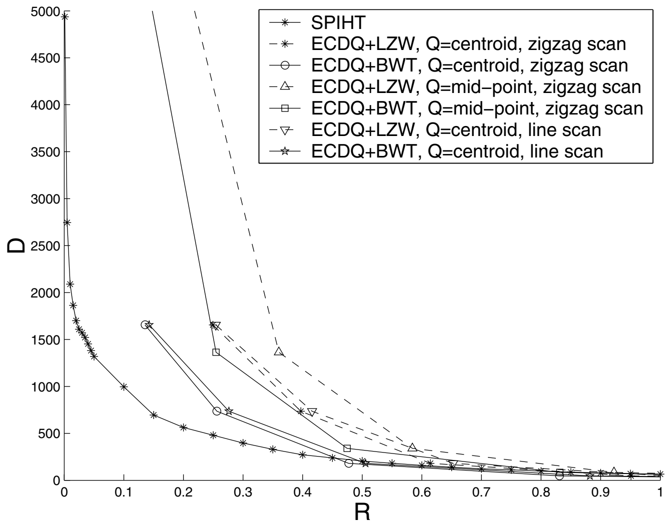

Figure 3: Variations on the MR-ECDQ algorithm.

First, we test the influence of the quantization rule, entropy code, and image scan order on the performance. Figure 3 shows the average results on a brain-scan image averaged over a range of Z values. The results are similar for other images that we tried. Here centroids outperform mid-points (counting the overhead needed to describe those centroids to the decoder), BWT outperforms LZW, and zigzag outperforms line scan.

In Figure 4, we plot the rate-distortion curves for different dither values for a brain image using BWT in MR-ECDQ. The dither value z_o has much greater influence on low rate performance than on high rate performance. For $R < 0.2$, $z_o = 96$ achieves the best performance over all tested dither values and its rate performance differs from that of SPIHT by no more than 0.08 bits/symbol at the same distortion. For $R \geq 0.5$, most z_o values outperform SPIHT, in some cases yielding rate differences larger than 0.5 for the same distortion.

Figure 5 compares the performance of MR-ECDQ in the image domain to MR-ECDQ applied on the image's wavelet coefficients. The latter achieves performance very close to that of SPIHT at low rates (the rate difference is at most 0.02). At higher rates, application of MR-ECDQ in the spatial domain outperforms both SPIHT and application of MR-ECDQ to the wavelet coefficients. Figure 6 gives similar results on another test image. Our conclusions from these figures are: at very low rates, the performance of MR-ECDQ on wavelet coefficients can be very close to or even the same as that of SPIHT; at higher rates MR-ECDQ on the original image consistently outperforms SPIHT on the images tested.

IV Summary

We design multi-resolution source codes using entropy constrained dithered scalar quantization and demonstrate constant rate redundancy bound at all resolutions, for all distortions, and for all sources. We further provide a practical implementation of

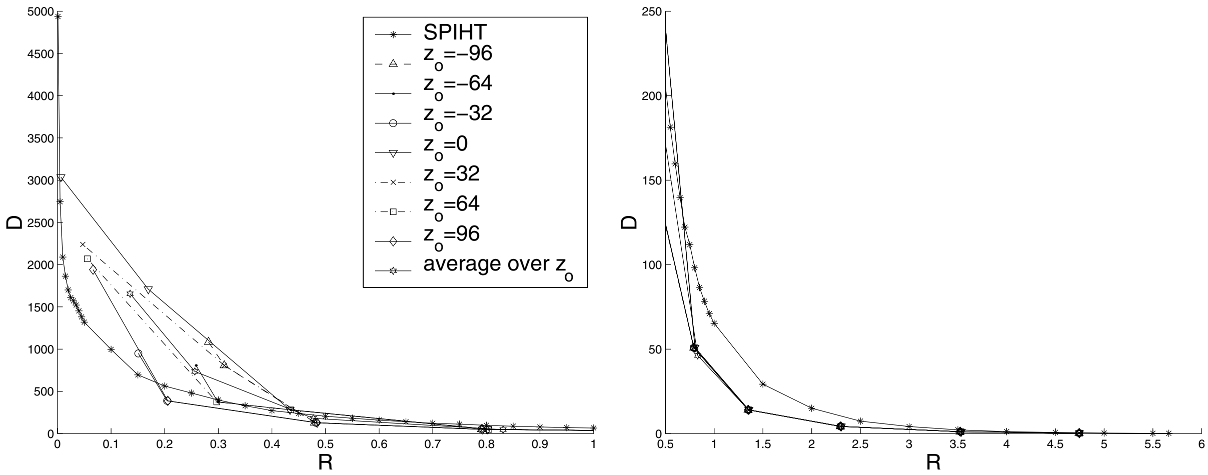

(a) At low resolution. (b) At high resolution.

Figure 4: MR-ECDQ+BWT's performance on a brain image at different z_o values.

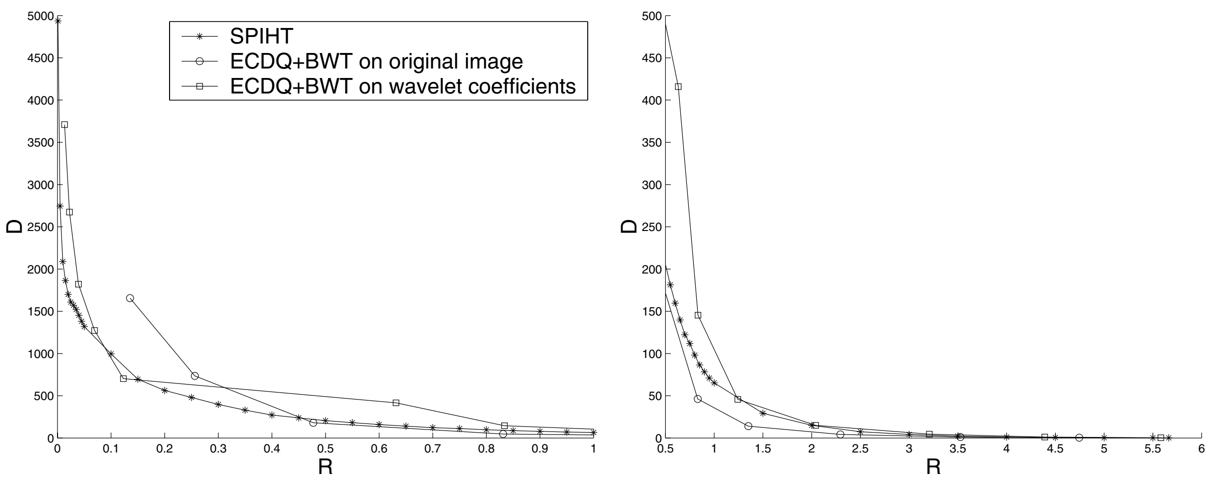

(a) At low resolution. (b) At high resolution.

Figure 5: Comparison of the average performance of different MR-ECDQ algorithms with SPIHT on a brain-scan image.

Proceedings of the Data Compression Conference (DCC'04)
1068-0314/04 $ 20.00 © 2004 IEEE

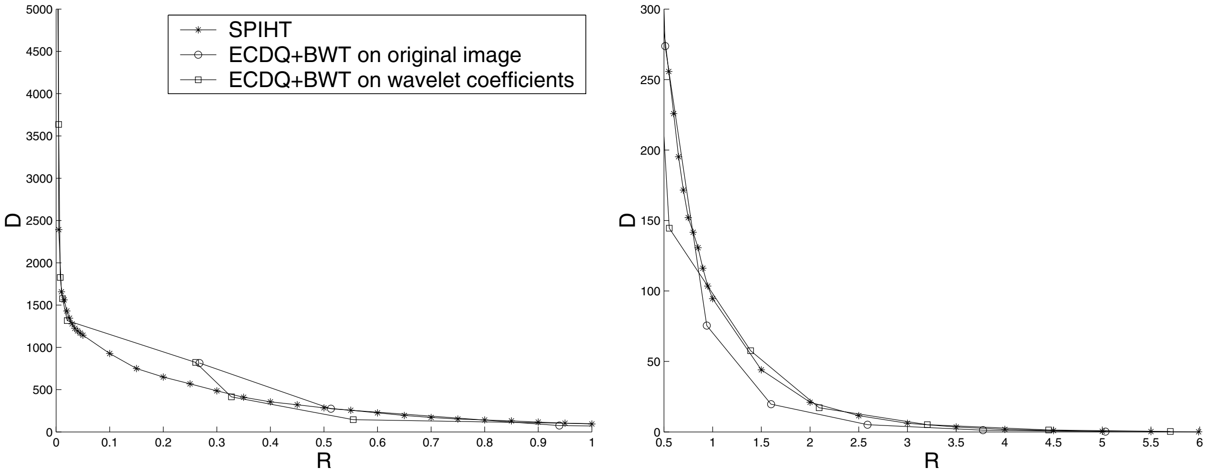

(a) At low resolution. (b) At high resolution.

Figure 6: Comparison of the performance of MR-ECDQ with SPIHT on image Lena.

MR-ECDQ and show its competitive rate-distortion performance on images.

References

[1] J. Ziv. On universal quantization. *IEEE Transactions on Information Theory*, IT-31(3):344–347, May 1985.

[2] R. Zamir and M. Feder. On universal quantization by randomized uniform/lattice quantizers. *IEEE Transactions on Information Theory*, 38(2):428–436, March 1992.

[3] Y. Frank-Dayan and R. Zamir. Dithered lattice-based quantizers for multiple descriptions. *IEEE Transactions on Information Theory*, 48(1):192–204, January 2002.

[4] W. H. R. Equitz and T. M. Cover. Successive refinement of information. *IEEE Transactions on Information Theory*, 37(2):269–275, March 1991.

[5] H. Feng and M. Effros. Improved bounds for the rate loss of multi-resolution source codes. *IEEE Transactions on Information Theory*, 49(4):809–821, April 2003.

[6] J. Ziv and A. Lempel. A universal algorithm for sequential data compression. *IEEE Transactions on Information Theory*, 23(3):337–343, May 1977.

[7] T. A. Welch. A technique for high-performance data compression. IEEE Computer, 17(6):8–19, June 1984.

[8] M. Burrows and D. J. Wheeler. A block-sorting lossless data compression algorithm. Technical Report SRC 124, Digital Systems Research Center, Palo Alto, CA, May 1994.

[9] A. Said and W. A. Pearlman. A new, fast, and efficient image codec based on set partitioning in hierarchical trees. *IEEE Transactions on Circuits and Systems for Video Technology*, 6(3):243–250, June 1996.

Proceedings of the Data Compression Conference (DCC'04)
1068-0314/04 $ 20.00 © 2004 IEEE