CONTEXT-BASED MULTIPLE DESCRIPTION WAVELET IMAGE CODING

Dror Porat and David Malah

Department of Electrical Engineering, Technion - Israel Institute of Technology
Haifa 32000, Israel
Phone: +(972)-4-8294745, fax: +(972)-4-8292793
Email: dporat@tx.technion.ac.il, malah@ee.technion.ac.il

ABSTRACT

Multiple description coding is a source coding technique that produces several descriptions of a single information source, such that various reconstruction qualities are obtained from different subsets of the descriptions. It thus can provide error resilience to information transmitted on lossy networks. Among previous works, MDs for image coding were generated via polyphase transform and selective quantization, performed in the wavelet domain. In this paper, we present an effective way to exploit the special statistical properties of the wavelet decomposition to provide improved coding efficiency, in the same general framework. We propose a novel coding scheme that efficiently utilizes contextual information, extracted from another polyphase component, to improve the coding efficiency of each redundant component. Our experimental results demonstrate that the proposed coder outperforms its predecessor across the entire redundancy range, and that the improvement in coding efficiency can indeed be attributed primarily to the effective utilization of contextual information.

1. INTRODUCTION

In *Multiple description (MD) coding*, one represents a single source of information (e.g., an image) with several chunks of data, called descriptions, in such a way that the source can be approximated from any (non-empty) subset of the descriptions [3]. The purpose of MD coding is to provide error resilience to information transmitted on lossy networks, such as the Internet, where inevitable loss of data may severely degrade the performance of conventional coding techniques. For example, in layered coding, where the information is represented hierarchically, a lost layer may also render other enhancement layers useless. MD coding, on the other hand, makes all of the received descriptions useful, and thus can better mitigate transport failures.

Among prior works, MDs for image coding were generated via the utilization of a decomposition into polyphase-like components (a polyphase transform) and selective quantization [4], performed in the wavelet domain. Unlike many other MD coding techniques, this technique explicitly separates description generation and redundancy addition, which significantly reduces the complexity of the system design and implementation. For description generation, in the two-description case discussed herein, the aforementioned technique employs a polyphase transform, and each of the two resulting polyphase components is coded independently at a source coding rate to constitute the primary part of information for its corresponding description. In order to explicitly add redundancy to each description, the other polyphase component is then coded at a (usually lower) redundancy coding rate and added to this description. In case of a channel failure, this redundancy enables an acceptable reconstruction of the lost component.

In this paper, we present an effective way to exploit the special statistical properties of the wavelet decomposition to provide improved coding efficiency, in the general framework of polyphase transform-based MD image coding. We propose a novel coding scheme that efficiently utilizes contextual information, extracted from the primary polyphase component of each description, to improve the coding efficiency of the corresponding redundant polyphase component, and thus enables the proposed MD coder to achieve improved overall performance. This is accomplished by means of various coding procedures, such as context-based classification of the wavelet coefficients, parametric model-based adaptive quantization, efficient optimal bit allocation, and adaptive entropy coding.

In order to efficiently utilize the statistical dependencies between neighboring wavelet coefficients, and avoid the need for an explicit characterization of these dependencies, we use an effective context-based classification procedure, inspired by that of [9]. To avoid the penalty of forward classification, the classification is based on contexts formed from quantized coefficients of the primary polyphase component of the description, which are also available at the decoder, and thus no transmission of side information is required. Nevertheless, a controlled amount of side information is added to the description and transmitted to the decoder, in order to improve the performance of the system. This side information includes the classification thresholds, allowing to select a class for a coefficient given its context, as well as the source statistics of each class, where each class is modeled using a parametric Laplacian distribution. The parametric modeling is also utilized by the bit allocation, quantization and entropy coding procedures that follow.

The context-based classification procedure enables the proposed coder to utilize a set of quantizers, each customized to an individual class. For this task, we examine two types of quantizers: the uniform threshold quantizer (UTQ) and the uniform reconstruction with unity ratio quantier (URURQ) [8]. Both of these quantizers well approximate the optimum entropy-constrained scalar quantizer (ECSQ) for the Laplacian distribution, assuming mean squared error (MSE) distortion, and are relatively simple to design and operate. In order to avoid the high complexity of entropy-constrained design algorithms for the quantizers, we propose an efficient design strategy, that is based on a pre-designed indexed array of MSE-optimized quantizers of different step sizes for the Laplacian distribution. To further reduce the complexity of the proposed design strategy, we also derive closed form expressions for the distortions attained by the quantizers.

For bit allocation between the various classes in the different subbands, we develop an optimal and efficient model-based bit allocation scheme, in the general framework of Lagrangian optimization, which also takes into account the non-energy preserving nature of the biorthogonal wavelet transform. Our bit allocation scheme, which is based on variance scaling and on the aforementioned pre-designed indexed array of MSE-optimized quantizers, enables the proposed coder to avoid complex on-line bit allocation procedures, as well as to effectively and instantly adapt the arithmetic entropy encoder to the varying coefficients statistics.

Our experimental results show that the proposed coder outperforms its predecessor [4] across the entire redundancy range, and that the improvement in coding efficiency can indeed be attributed primarily to the effective utilization of contextual information.

The paper is organized as follows. In Section 2 we provide a detailed description of the main building blocks of the proposed coder. Section 3 provides some experimental results. Finally, concluding remarks follow in Section 4.

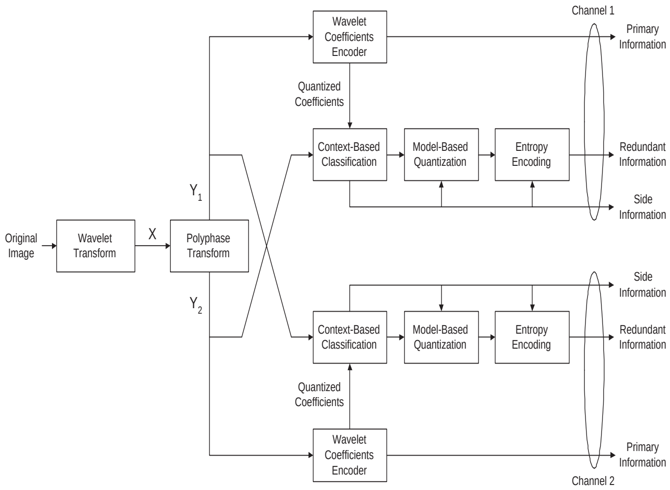

Figure 1: Simplified block diagram of the proposed MD encoder.

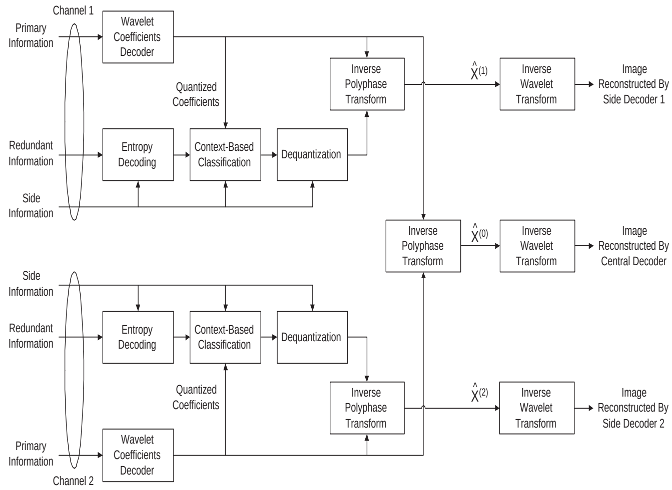

Figure 2: Simplified block diagram of the proposed MD decoder.

2. CONTEXT-BASED MULTIPLE DESCRIPTION WAVELET IMAGE CODER

A simplified block diagram of the proposed MD encoder is shown in Figure 1. As shown, the input image is first wavelet transformed to produce the dyadic wavelet decomposition denoted by X. The resulting decorrelated transform coefficients are decomposed into the two polyphase components Y_1 and Y_2 via a polyphase transform. Each of the two polyphase components Y_1 and Y_2 is then encoded independently by a wavelet coefficients encoder to form the primary part of information for its corresponding description. Depending on the exact type of polyphase transform, this can generally be a standard wavelet encoder (e.g., SPIHT [7] encoder), operating directly on the wavelet coefficients. In order to enable an acceptable reconstruction of the lost component in the case of a channel failure, each description also carries information about the other component. This redundant information, useful only in the case of a lost description, is produced by efficient algorithms for classification, bit allocation, quantization, and adaptive entropy encoding, that use contextual information extracted from quantized coefficients of the complementary polyphase component. These quantized coefficients are obtained directly from the wavelet coefficients encoder that produces the primary part of information for the description.

In order to utilize the statistical dependencies between neighboring wavelet coefficients, for improving the coding efficiency of the redundant information, the context used for the classification of a given coefficient is formed from neighboring quantized coefficients, obtained from the complementary polyphase component. To avoid the need for an explicit characterization of the statistical dependencies between neighboring coefficients, we use an adaptive classification procedure that also produces a controlled amount of side information, which is transmitted to the decoder. This side information includes the classification thresholds, allowing to select a class for a coefficient given its context, as well as the source statistics of each class, where each class is modeled using a parametric distribution. The parametric modeling is also utilized by the bit allocation, quantization and entropy coding procedures that follow.

A simplified block diagram of the corresponding decoder is shown in Figure 2. If both descriptions arrive, the primary part of information from each description is decoded using a wavelet coefficients decoder (corresponding to the wavelet coefficients encoder used for encoding the primary part of information). The central decoder then composes the resulting polyphase components to produce the reconstructed wavelet decomposition $\hat{X}^{(0)}$. Finally, an inverse wavelet transform is applied to the recovered wavelet coefficients to produce the reconstructed image.

If only description i is received ($i \in \{1,2\}$), the primary part of information from the received description is decoded, as before, to recover the quantized polyphase component Y_i. The complementary polyphase component is then recovered from the redundant part of information. Side decoder i composes the recovered polyphase components to produce the reconstructed wavelet decomposition $\hat{X}^{(i)}$. An inverse wavelet transform follows to produce the reconstructed image.

The following sections provide a more detailed description of the main building blocks of the proposed coder, which include the classification, quantization and bit allocation schemes.

2.1 Context-Based Classification

The proposed MD coder utilizes classification for efficient coding of the redundant information, which enables an acceptable reconstruction of the lost polyphase component in the case of a channel failure. The context-based classification procedure, inspired by that of [9], enables the use of a set of quantizers, each customized to an individual distribution component. Compared to the use of a single average quantizer fitted to the overall input statistics, this offers a potential increase in coding efficiency [9].

For description $i \in \{1,2\}$, define the polyphase component Y_i to be the *primary polyphase component*, and define the complementary component to be the *redundant polyphase component*. The *context* of a given wavelet coefficient in the redundant polyphase component, is then defined as a set of quantized coefficients from the primary polyphase component, which is used to characterize the given coefficient. Note that since the coefficients that form the context are already quantized, the decoder can use the exact same context used by the encoder. As detailed in [1], significant statistical dependencies exist between neighboring wavelet coefficients, and these can be utilized for efficient context-based classification. For a given coefficient in the redundant polyphase component, we therefore form a context from neighboring quantized coefficients that belong to the primary polyphase component. Assuming a plain polyphase transform, which, for each column in each subband, simply groups all the odd-numbered coefficients into one polyphase component and all the even-numbered coefficients into the other component, Figure 3 shows the context of a given wavelet coefficient $X_{i,j}$. For ease of illustration, the primary and redundant polyphase components are interleaved in the figure. As shown, the classification of the wavelet coefficient $X_{i,j}$ is based on the following context:

$$\mathscr{C}_{i,j} = \left\{\hat{X}_{i-1,j-1}, \hat{X}_{i-1,j}, \hat{X}_{i-1,j+1}, \hat{X}_{i+1,j-1}, \hat{X}_{i+1,j}, \hat{X}_{i+1,j+1}\right\}$$

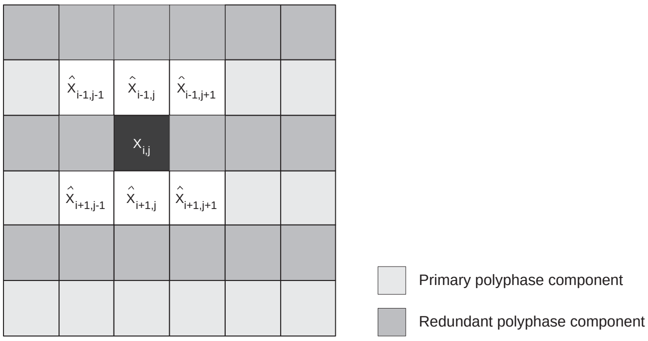

Figure 3: Context of a given wavelet coefficient $X_{i,j}$. The coefficient belongs to the redundant polyphase component, and is shown in black. The quantized coefficients that form the context belong to the primary polyphase component, and are shown in white.

where $\hat{X}$ denotes the value of X after quantization by the wavelet coefficients encoder (see Figure 1), and the coordinate system refers to the whole subband (i.e, prior to the polyphase transform).

We now describe the classification rule, which assigns one of a finite number of classes to a coefficient $X_{i,j}$, given its context $\mathscr{C}_{i,j}$. Due to the difficulty of characterizing the full multidimensional relationship between the magnitudes of neighboring wavelet coefficients, the classification rule is based on a weighted average of the magnitudes of coefficients in $\mathscr{C}_{i,j}$. Define the *activity* $A_{i,j}$ of the coefficient $X_{i,j}$, predicted from its context $\mathscr{C}_{i,j}$, as

$$\begin{aligned} A_{i,j} &= a_1|\hat{X}_{i-1,j-1}| + a_2|\hat{X}_{i-1,j}| + a_3|\hat{X}_{i-1,j+1}| + \\ &\quad + a_4|\hat{X}_{i+1,j-1}| + a_5|\hat{X}_{i+1,j}| + a_6|\hat{X}_{i+1,j+1}| \end{aligned}$$

where the fixed relative weights $\{a_k\}$ satisfy $\sum_k a_k = 1$. In the proposed coder we set the weights $\{a_k\}$ to be inversely proportional to the Euclidean distances between the positions of the corresponding coefficients in $\mathscr{C}_{i,j}$ and the position of the coefficient $X_{i,j}$.

In order to classify a coefficient $X_{i,j}$, its context $\mathscr{C}_{i,j}$ is first used to predict its activity $A_{i,j}$. The (nonnegative) activity $A_{i,j}$ is then compared with a set of pre-computed adaptive classification thresholds to determine the class assigned to $X_{i,j}$. Specifically, assuming $N+1$ potential classes $C_0, C_1, \ldots, C_N$, and given N monotonically increasing classification thresholds

$$T_0 = 0 < T_1 < T_2 < \cdots < T_{N-1} < \infty ,$$

the coefficient $X_{i,j}$ is assigned to the class C_k if $A_{i,j} \in (T_{k-1}, T_k]$, where $T_{-1} = -\infty$ and $T_N = \infty$. This is shown in Figure 4.

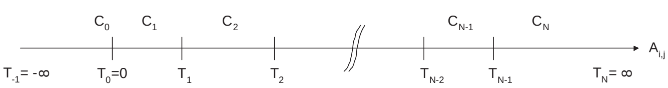

Figure 4: Classification of the coefficient $X_{i,j}$ based on its activity $A_{i,j}$.

Our class-adaptive quantization scheme is based on a parametric distribution model. Specifically, we assume that the coefficients in each class of each subband are drawn from a (zero-mean) Laplacian distribution, i.e., according to the probability density function (pdf) $f_\lambda(x) = \lambda/2 \cdot e^{-\lambda|x|}$. Since we employ a predictive (DPCM-like) quantization scheme to encode the coefficients in the approximation subband, we assume a Laplacian distribution model (for the prediction errors) in the approximation subband as well.

In order to adaptively determine the classification thresholds $T_1, T_2, \ldots, T_{N-1}$, with the goal of maximizing the coding gain due to classification, we utilize the classification thresholds design procedure proposed in [9] (where additional details can be found). The outputs of this design procedure are the classification thresholds $T_1, \ldots, T_{N-1}$ and the maximum likelihood Laplacian parameter estimates $\hat{\lambda}_0, \hat{\lambda}_1, \ldots, \hat{\lambda}_N$ for each of resulting classes. The design procedure is performed separately for each of the descriptions, and is repeated independently for each of the subbands. For each description, the corresponding classification thresholds and Laplacian parameter estimates are transmitted to the decoder as side information. Denoting the number of subbands by S, this involves transmission of $S \cdot [(N-1) + (N+1)] = 2NS$ real numbers (after proper quantization), per description. The parametric source statistics of the various classes (i.e., the Laplacian parameter estimates), enable us to further utilize efficient bit allocation and quantization schemes, as described next.

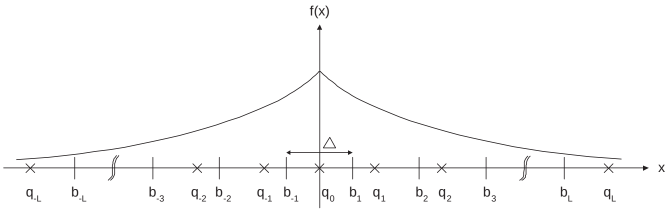

Figure 5: Uniform threshold quantizer (UTQ) with step size Δ and an odd number of levels $N = 2L+1$. The bin boundaries are denoted by $\{b_j\}$ and the reconstruction levels by $\{q_j\}$. The Laplacian pdf is also shown for illustration.

2.2 Parametric Model-Based Quantization

The context-based classification enables the proposed coder to utilize a set of quantizers, each customized to an individual class. As demonstrated in [2], the optimum ECSQ for the Laplacian distribution, assuming MSE distortion, can be well approximated by a uniform threshold quantizer (UTQ). The UTQ is a scalar quantizer with a fixed step size Δ, and is used effectively by various image subband quantization schemes. Since the reconstruction levels affect only the quantizer distortion, with no influence on the rate, they can be optimized for distortion alone. In the case of MSE distortion, this leads to the well-known centroid condition (i.e., conditional expectation) for the position of the reconstruction levels. We focus on UTQs with odd number of levels ("mid-tread" UTQs) to be able to obtain low bit rates (i.e., low redundancy), if desired. This is due to the fact that if the number of levels is even and the pdf is symmetric (e.g., Laplacian), no symmetric quantizer can yield an output entropy lower than 1 bit/sample [2]. Figure 5 shows a UTQ with an odd number of levels.

A general entropy-constrained design algorithm of UTQ is given in [2]. For the Laplacian (and exponential) distribution, a non-iterative entropy-constrained design algorithm of UTQ is provided by Sullivan [8]. Due to the relatively high complexity of the aforementioned design algorithms, we employ a more efficient design strategy: As detailed in the following section, the proposed coder utilizes an effective optimal bit allocation scheme, which is based on a pre-designed array of optimized UTQs of different step sizes (where the optimization is for minimum MSE distortion). This enables us to greatly simplify the design of each UTQ, as each UTQ is optimized for a given step size, with no constraint on its output entropy. We also store the pre-computed bin probabilities of each UTQ, which enable us to instantly adapt the entropy encoder to the varying coefficients statistics.

For a UTQ with a given step size Δ, the reconstruction levels $\{q_j\}$ and the bin probabilities $\{p_j\}$ can be easily computed by integration. To further reduce the complexity of our proposed design strategy, we derive a closed form expression for the MSE distortion D attained by the UTQ. For conciseness, we only provide here the general guidelines and the obtained expression: First, we prove that the reconstruction offsets $q_j - b_j$, for $j = 1, \ldots, L-1$, do not depend on j. Namely, we show that

$$q_j = b_j + \delta, \;\; j = 1, \ldots, L-1 ,$$

where

$$\delta = \frac{1}{\lambda} - \frac{\Delta}{e^{\lambda\Delta} - 1}$$

and λ is the parameter of the Laplacian distribution. Next, using integration by parts and appropriate substitutions of variables, we obtain the following closed form expression for the average distortion of the UTQ:

$$\begin{aligned} D \;\; = \;\; & \frac{2}{\lambda^2} - e^{-\lambda\frac{\Delta}{2}} \left(\left(\frac{\Delta}{2}\right)^2 + \frac{\Delta}{\lambda} + \frac{2}{\lambda^2} \right) + \\ & + \left(\sum_{j=1}^{L-1} e^{-\lambda q_j} \right) \cdot \left[e^{\lambda\delta} \left(\delta^2 - \frac{2\delta}{\lambda} + \frac{2}{\lambda^2} \right) + \right. \\ & \left. - e^{\lambda(\delta-\Delta)} \left((\delta-\Delta)^2 - \frac{2(\delta-\Delta)}{\lambda} + \frac{2}{\lambda^2} \right) \right] + \\ & + \frac{1}{\lambda^2} e^{1-\lambda q_L} . \end{aligned}$$

The rate R is estimated by the entropy of the quantizer output, given by

$$H_Q = - \sum_{j=-L}^{j=L} p_j \log_2 p_j .$$

As an alternative to the UTQ, we also examined the utilization of the uniform reconstruction with unity ratio quantizer (URURQ) [8], defined by its step size Δ—the distance between adjacent *reconstruction* levels. Additional details, including the derivation of a closed form expression for the distortion attained by the URURQ, are given in [6].

For brevity, in the sequel we generally refer to the quantizers used by the proposed coder (to quantize the coefficients of the redundant polyphase component of each description) only as UTQs. Nevertheless, the same guidelines hold for URURQs as well.

2.3 Optimal Model-Based Bit Allocation

The context-based classification enables the proposed coder to utilize a set of quantizers, each customized to an individual class, for coding the redundant polyphase component of each description. Define the *base rate* of a given description as the average bit rate (in bits/pixel) at which the primary polyphase component of the description is encoded. Similarly, define the *redundancy rate* of the description as the average bit rate (in bits/pixel) at which the redundant polyphase component of the description is encoded. Given the desired redundancy rate for a description, the encoder needs to determine the rate at which each of the customized quantizers operates. To this end, the proposed coder utilizes an optimal model-based bit allocation scheme. The bit allocation, which is based on the Laplacian model, is performed in the general framework of Lagrangian optimization [5]. It is carried out for each description separately, given the desired redundancy rate for the description, and performed over all classes of redundant coefficients from all subbands simultaneously.

Given the desired redundancy rate for the description, we wish to find the optimal rates $\{R_b\}_{b=1}^{B}$ (assuming a total of B classes) for each of the different classes (from all subbands), such that the resulting MSE distortion in the image domain is minimized, subject to the redundancy rate constraint. As before, we assume that the coefficients in each class are drawn from a Laplacian distribution (the Laplacian parameter estimate $\hat{\lambda}$ for each class is obtained from the classification thresholds design procedure). We use Lagrangian optimization to solve this constrained optimization problem, and seek to minimize the Lagrangian cost function $J(\xi) = D + \xi R$ (for a fixed Lagrange multiplier ξ), where D is the MSE distortion in the image domain resulting from quantization of coefficients in the redundant polyphase component, and R is the redundancy rate of the description. Presenting the distortion D as the weighted sum of individual distortions (due to quantization of coefficients) in the different classes and the rate R as the weighted sum of individual rates, and differentiating with respect to R_b, yields the following rate allocation equations

$$D'_b(R_b) = -\frac{\xi}{G_b}, \; b = 1, \ldots, B , \tag{1}$$

where $D'_b(R_b)$ is the derivative of the transform domain distortion D_b of the class b, evaluated at R_b bits/coefficient, and where G_b is the synthesis (inverse transform) energy gain factor associated with the subband to which the class b belongs[1].

We utilize variance scaling to efficiently solve the rate allocation equations (1), and thus avoid computing the operational distortion rate function of the optimal UTQ of each class b (for the class-specific Laplacian parameter estimate $\hat{\lambda}_b$, or variance estimate $\hat{\sigma}_b^2 = 2/\hat{\lambda}_b^2$). Hence, we only compute off-line (and also index by the slope) the operational distortion rate function of the UTQ for the unit-variance Laplacian distribution, and infer the required operating point of the UTQ of each class b from that normalized operational distortion rate function and the normalized slope $-\xi/(G_b\hat{\sigma}_b^2)$. That is, the actual quantizer for the class b is simply the scaled version, by $\hat{\sigma}_b$, of the quantizer corresponding to the normalized slope on the normalized operational distortion rate function (which is practically stored as an indexed array of MSE-optimized UTQs). The quantizer's bin probabilities are also used to instantly adapt the arithmetic entropy encoder that follows the quantization stage to the statistics of the class b, as captured by the model-based classification algorithm.

Finally, we note that determining the value of the Lagrange multiplier ξ such that the total rate constraint is satisfied can be performed using the bisection method. The resulting value of ξ, which specifies the bit allocation, is also transmitted to the decoder (as part of the side information for the description).

3. EXPERIMENTAL RESULTS

The proposed MD coder aims to improve the performance of the original polyphase transform-based coding scheme of [4], by utilizing the special statistical properties of the wavelet decomposition to improve the coding efficiency of the redundant polyphase component of each description. In other words, for a given quality of central reconstruction, the proposed coder aims to provide improved side reconstructions. As an example, Figure 6 shows the performance of the proposed coder, compared to that of the original polyphase transform-based coder, utilizing either the plain polyphase transform or the vector-form polyphase transform [4], for the image Lena (of size 512×512 pixels) and a total rate of 1 bpp. As shown, the proposed coder indeed attains improved side reconstructions (i.e, higher side PSNR), for a given quality of central reconstruction, across the entire redundancy range. We also note that this it true whether UTQs or URURQs are utilized by the proposed coder, since both types of quantizers have proved to yield comparable performances in our simulations.

The proposed coder exploits contextual information to improve the coding efficiency of the redundant polyphase component of each description. To this end, the proposed coder utilizes various coding procedures, such as context-based classification, parametric model-based adaptive quantization, efficient optimal bit allocation and adaptive entropy coding. It is interesting to examine whether

[1]This energy gain factor represents the amount of squared error in the synthesized image introduced by a unit error in a transform coefficient of the subband. For an orthonormal transform $G_b = 1, \forall b$.

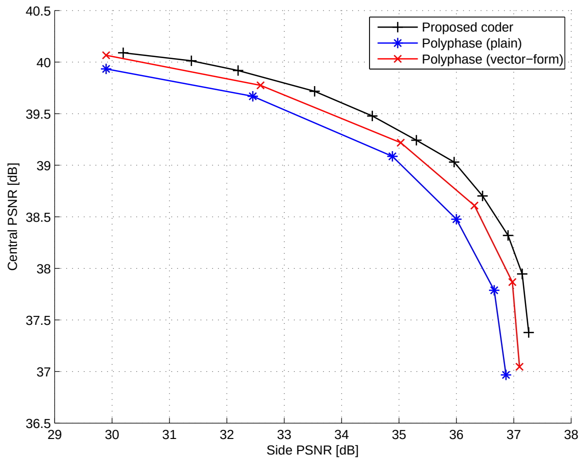

Figure 6: Performance of the proposed MD coder (in black), compared to that of the original polyphase transform-based MD coder [4], utilizing either the plain polyphase transform or the vector-form polyphase transform (for the image Lena and a total rate of 1 bpp).

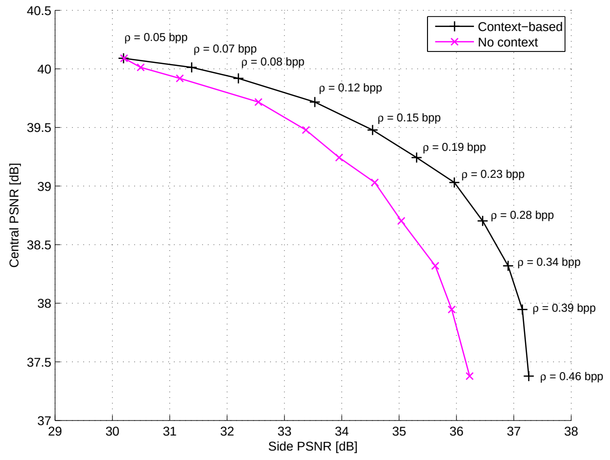

Figure 7: Performance of the "no-context coder", compared to that of the proposed context-based coder in its standard configuration (for the image Lena and a total rate of 1 bpp). Also shown is the redundancy rate ρ corresponding to various points on the performance curves.

the achieved improvement in coding efficiency can indeed be attributed primarily to the effective utilization of contextual information. We refer to the improvement in coding efficiency due to the utilization of contextual information as the *context gain*.

The context gain can be measured experimentally by comparing the performance of the proposed MD coder, in its standard configuration (e.g., with four classes per subband), to that of an almost identical coder, in which the classification procedure assigns all the coefficients in each subband to a single class. Such a modified coder, which we refer to as the "no-context coder", thus uses a single average quantizer for each subband (which is fitted to the overall statistics in the subband), and ignores any contextual information. Figure 7 shows the performance of the "no-context coder", compared to that of the proposed coder in its standard configuration (for the image Lena and a total rate of 1 bpp). For a given quality of central reconstruction, the context gain is the gap between the side PSNR attained by the proposed coder to that attained by the "no-context coder". As shown, the context gain increases significantly with the redundancy (up to more than 1 dB improvement in side PSNR). This proves that the improvement in coding efficiency offered by the proposed MD coder can indeed be attributed primarily to the effective utilization of contextual information and the special statistical properties of the wavelet decomposition.

4. CONCLUSION

Among prior works, MDs for image coding were generated via the utilization of a polyphase transform and selective quantization, performed in the wavelet domain. In this paper we presented an effective way to exploit the special statistical properties of the wavelet decomposition to provide improved coding efficiency, in the same general framework. We have proposed a novel coding scheme that efficiently utilizes contextual information, extracted from the primary polyphase component of each description, to improve the coding efficiency of the corresponding redundant polyphase component, and thus enables the proposed MD coder to achieve improved overall performance. This is accomplished by means of various coding procedures, such as context-based classification of the wavelet coefficients, parametric model-based adaptive quantization (using either UTQs or URURQs), efficient optimal bit allocation, and adaptive entropy coding. Our experimental results show that the proposed coder outperforms its predecessor across the entire redundancy range, and that the improvement in coding efficiency can indeed be attributed primarily to the effective utilization of contextual information and the special statistical properties of the wavelet decomposition.

REFERENCES

[1] R. W. Buccigrossi and E. P. Simoncelli, "Image compression via joint statistical characterization in the wavelet domain," *IEEE Trans. Image Process.*, vol. 8, no. 12, pp. 1688–1701, Dec 1999.

[2] N. Farvardin and J. Modestino, "Optimum quantizer performance for a class of non-Gaussian memoryless sources," *IEEE Trans. Inf. Theory*, vol. 30, no. 3, pp. 485–497, May 1984.

[3] V. K. Goyal, "Multiple description coding: compression meets the network," *IEEE Signal Process. Mag.*, vol. 18, no. 5, pp. 74–93, Sep 2001.

[4] W. Jiang and A. Ortega, "Multiple description coding via polyphase transform and selective quantization," in *Proc. Visual Communic. and Image Proc.*, San Jose, CA, Jan 1999, pp. 998–1008.

[5] A. Ortega and K. Ramchandran, "Rate-distortion methods for image and video compression," *IEEE Signal Process. Mag.*, vol. 15, no. 6, pp. 23–50, Nov 1998.

[6] D. Porat, "Context-based multiple description wavelet image coding," Master's thesis, Technion—IIT, Oct 2009, http://sipl.technion.ac.il/siglib/FP/Dror-Porat_Thesis.pdf.

[7] A. Said and W. A. Pearlman, "A new, fast, and efficient image codec based on set partitioning in hierarchical trees," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 6, no. 3, pp. 243–250, Jun 1996.

[8] G. J. Sullivan, "Efficient scalar quantization of exponential and Laplacian random variables," *IEEE Trans. Inf. Theory*, vol. 42, no. 5, pp. 1365–1374, Sep 1996.

[9] Y. Yoo, A. Ortega, and B. Yu, "Image subband coding using context-based classification and adaptive quantization," *IEEE Trans. Image Process.*, vol. 8, no. 12, pp. 1702–1715, Dec 1999.