CONTEXT-BASED MULTIPLE DESCRIPTION WAVELET IMAGE CODING

DROR PORAT

CONTEXT-BASED MULTIPLE DESCRIPTION WAVELET IMAGE CODING

RESEARCH THESIS

SUBMITTED IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS
FOR THE DEGREE OF MASTER OF SCIENCE
IN ELECTRICAL ENGINEERING

DROR PORAT

SUBMITTED TO THE SENATE OF THE TECHNION — ISRAEL INSTITUTE OF TECHNOLOGY
TISHREI, 5770 HAIFA OCTOBER, 2009

THIS RESEARCH THESIS WAS DONE UNDER THE SUPERVISION OF PROF. DAVID MALAH IN THE DEPARTMENT OF ELECTRICAL ENGINEERING

ACKNOWLEDGMENT

I would like to express my deep and sincere gratitude to my supervisor, Prof. David Malah, for his excellent professional guidance, steady support and warm attitude. I would also like to thank Nimrod Peleg and the rest of the SIPL (Signal and Image Processing Lab) staff for their continuous technical assistance. I am heartily thankful to my parents, Shmuel and Zehava, for their endless support and encouragement throughout the way. Last but not least, I wish to thank my loved one Shira for her unconditional love and support.

THE GENEROUS FINANCIAL HELP OF THE TECHNION IS GRATEFULLY ACKNOWLEDGED

Contents

List of Figures

List of Tables

Abstract

Multiple description (MD) coding is a coding technique that produces several descriptions of a single source of information (e.g., an image), such that various reconstruction qualities are obtained from different subsets of the descriptions. The purpose of MD coding is to provide error resilience to information transmitted on lossy networks (i.e., networks that cannot avoid possible loss of packets). Since there is no hierarchy of descriptions in MD coding, representations of this type make all of the received descriptions useful (unlike, for example, layered coding, where a lost layer may also render other enhancement layers useless). Thus, MD coding is especially suitable for networks with no priority mechanisms for data delivery, such as the Internet.

Among previous works, MDs for image coding were generated via the utilization of a decomposition into polyphase-like components (a polyphase transform) and selective quantization, performed in the wavelet domain. In this research work, 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 a different polyphase component, to improve the coding efficiency of each redundant component (aimed to provide an acceptable reconstruction of a lost polyphase component in the case of a channel failure), 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 (performed in the general framework of Lagrangian optimization), and adaptive entropy coding.

Our experimental results clearly demonstrate the advantages of the proposed context-based MD image coder. Specifically, we also show that the proposed coder outperforms its predecessor—the original polyphase transform-based coder—across the entire redundancy range, and that the improvement in coding efficiency can indeed be attributed primarily to the effective utilization of contextual information.

List of Symbols

$A_{i,j}$	Activity of the wavelet coefficient $X_{i,j}$
B	Total number of classes from all subbands
$\mathcal{C}_{i,j}$	Context used for classification of the wavelet coefficient $X_{i,j}$
$C_0, \ldots, C_N$	Potential classes
D	Distortion
D_b	Distortion in encoding the class b
D'_b	The derivative of the distortion D_b
D_c	Central distortion
D_i	Distortion attained by Decoder i
D_s	Average side distortion
$D(R)$	Distortion rate function
$\hat{D}(R)$	Operational distortion rate function of a quantizer
$E\left[\cdot\right]$	Expectation
G	Classification gain
G_b	Synthesis gain associated with the subband to which class b belongs
G_θ	Givens rotation of angle θ
$H(\cdot)$	Entropy
H_Q	Output entropy of the quantizer Q

$I(\cdot\,;\,\cdot)$	Shannon mutual information
$J(\cdot)$	Lagrangian cost function
$\mathcal{L}(\cdot)$	Likelihood function
$\mathcal{L}^*(\cdot)$	Log-likelihood function
N	Number of pixels in the image
N_b	Number of coefficients in the class b
R	Rate
R_b	Rate for encoding the class b
R_i	Transmission rate over Channel i
R_T	Desired redundancy rate
$R(D)$	Rate distortion function
S	Number of subbands
$T_0, \ldots, T_{N-1}$	Classification thresholds
$X_{i,j}$	Wavelet coefficient at the array coordinates (i, j)
$\hat{X}$	Quantized value of X
$\{X_k\}_{k=1}^N$	Sequence of source symbols
$\{\hat{X}_k^{(i)}\}_{k=1}^N$	Reconstruction sequence produced by Decoder i
a_K	Approximation subband (at level K)
$\{b_j\}$	Bin boundaries (of the quantizer)
d_i^H	Horizontal detail subband (at level i)
d_i^V	Vertical detail subband (at level i)
d_i^D	Diagonal detail subband (at level i)
$d(\cdot, \cdot)$	Distortion measure
$\{e_1,\ e_2\}$	The standard basis of $\mathbb{R}^2$
l	Index assignment (in MDSQ)

$\{p_j\}$	Bin probablities (of the quantizer)
$\{q_j\}$	Reconstruction levels (of the quantizer)
α	Initial encoder (in MDSQ)
α_0	Encoder (in MDSQ)
β_0	Central decoder (in MDSQ)
$\beta_1,\ \beta_2$	Side decoders (in MDSQ)
Δ	Step size of the quantizer
δ	Reconstruction offset
η_b	The ratio N_b/N
λ	Laplacian parameter
$\hat{\lambda}$	Estimated Laplacian parameter
ξ	Lagrange multiplier
σ^2	Variance
$\langle\cdot,\ \cdot\rangle$	Inner product

List of Abbreviations

bpp	Bits per pixel
DPCM	Differential Pulse Code Modulation
DR	Distortion Rate
DWT	Discrete Wavelet Transform
ECSQ	Entropy-Constrained Scalar Quantization (or Quantizer)
EZW	Embedded Zerotree Wavelet (coder)
GGD	Generalized Gaussian Distribution
i.i.d.	Independent identically distributed
KLT	Karhunen-Loeve Transform
MD	Multiple Description
MDC	Multiple Description Coding
MDLVQ	Multiple Description Lattice Vector Quantization (or Quantizer)
MDSQ	Multiple Description Scalar Quantization (or Quantizer)
MDVQ	Multiple Description Vector Quantization (or Quantizer)
MLE	Maximum Likelihood Estimation (or Estimator)
MSE	Mean Squared Error
pdf	Probability density function
PSNR	Peak Signal-to-Noise Ratio

RD	Rate Distortion
SD	Single Description
SI	Subsample-Interpolation
SPIHT	Set Partitioning In Hierarchical Trees
SQ	Scalar Quantization (or Quantizer)
TCP	Transmission Control Protocol
UEP	Unequal Error Protection
UQ	Uniform Quantization (or Quantizer)
URURQ	Uniform Reconstruction with Unity Ratio Quantization (or Quantizer)
UTQ	Uniform Threshold Quantization (or Quantizer)
VQ	Vector Quantization (or Quantizer)

Chapter 1

Introduction

1.1 Multiple Description Coding

Multiple description (MD) coding is a coding technique that 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. 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. It is clear that in order to gain robustness to the possible loss of descriptions, some compression efficiency must be sacrificed (i.e., the representation is redundant). That is the reason why MD coding should be applied only if this disadvantage in compression efficiency is offset by the advantage of mitigating transport failures.

Many MD coding schemes focus on the case of two descriptions, mainly due to its relative tractability, and so does this work. The basic situation of MD coding, for the two-description case, is as follows [15]: Suppose that we wish to send a description of an information source over an unreliable communication network, namely, one with nonzero probability for the event of description loss. In order to deal with the possibility of a lost description, or due to constraints imposed by the network regarding the size of packets, we decide to send two descriptions instead and hope that at least one of them will get through. Obviously, each description should be individually good due to the possibility that only one description will get to the destination. However, in the case that both descriptions get through, we wish to maximize the combined descriptive information at the receiver.

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 [30], performed in the wavelet domain. Unlike many other MD coding techniques, such as MD scalar quantization [59] and MD correlating transforms [65], the technique of MD coding via polyphase transform and selective quantization explicitly separates description generation and redundancy addition, which significantly reduces the complexity of the system design and implementation. This technique also enables to easily generate descriptions of statistically equal rate and importance, a property which is well suited to communication systems with no priority mechanisms for data delivery (e.g., the Internet).

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 using selective quantization and added to this description. In case of a channel failure, this redundancy enables an acceptable reconstruction of the lost component.

1.2 Proposed Coding Scheme

In this research work, we develop and explore 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. 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 still produced 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). 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 both the UTQ and the URURQ.

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 intelligently and instantly adapt the arithmetic entropy encoder to the varying coefficients statistics.

Finally, we also suggest a way to determine, in practice, the optimal operating point for the proposed MD coder, based on the properties of the communication channel.

1.3 Thesis Outline

This thesis is organized as follows. Chapter 2 introduces the MD scenario, summarizes the main results in the theory of MD coding, and provides an overview of some existing MD coding techniques. Chapter 3 provides a detailed review of an efficient technique for MD image coding utilizing a polyphase transform and selective quantization, which forms the framework of our proposed context-based MD coding system. Chapter 4 provides the motivation for context-based MD coding in the wavelet domain, and also presents the outline of the proposed context-based MD wavelet image coder. A more detailed description of the main building blocks of the proposed coder is provided in the subsequent chapters.

In Chapter 5 we describe the context-based classification procedure, which enables the proposed coder to use a set of customized quantizers for improved coding efficiency. Chapter 6 provides a detailed description of the parametric model-based quantization and bit allocation schemes. In Chapter 7 we comprehensively address some implementation considerations, which are also generally related to efficient implementation of wavelet coders, including boundary extensions for the wavelet transform (especially referring to coding applications), and optimal bit allocation for the biorthogonal wavelet transform.

Chapter 8 provides exhaustive experimental results, which also include a comparison to other relevant MD image coding systems. It also provides a way to optimally determine, in practice, the operating point for the proposed coder, based on the properties of the communication channel.

Finally, Chapter 9 concludes this thesis and describes several possible directions for future research.

Chapter 2

Fundamentals of Multiple Description Coding

This chapter introduces the MD scenario, summarizes the main results in the theory of MD coding, and provides an overview of some existing MD coding techniques.

2.1 Introduction to Multiple Description Coding

As described in Section 1.1, the basic problem that multiple description coding deals with is how to represent a single source of information with several chunks of data ("descriptions") in such a way that the source can be approximated from any subset of the chunks [21].

Current systems that deliver content over packet networks typically generate the transmitted data with a progressive coder and use TCP (Transmission Control Protocol)—the standard protocol that controls retransmission of lost packets—in order to deliver the data over the network. The main drawback of such techniques is that the order in which packets are received is critical, which may cause intolerable delays when packets are lost. For example, suppose that an image is compressed using a progressive (layered) coder that generates L packets, numbered from 1 to L. The receiver reconstructs the image from the received packets and the quality of the reconstructed image is improved steadily as the number of consecutive packets received, starting from the first, increases. If the first two packets are received but the third packet, for instance, is not, the quality of the reconstructed image is proportionate to reception of only two packets, no matter how many of the other packets are received. Receiving the third packet may require retransmission of that packet, which introduces a delay. Such delays may be much longer than the inter-arrival times between received packets and therefore hurt the performance of the system. The above example shows that although progressive transmission works well when packets are received in order without loss, the performance of such a system may be degraded when packet losses occur.

In a constellation where losses are inevitable, it would be highly beneficial to create representations that make all of the received packets useful, not just those consecutive from the first. The purpose of MD coding is to create such representations. It is clear that in order to gain robustness to the possible loss of descriptions, some compression efficiency must be sacrificed. That is the reason why MD coding should be applied only if this disadvantage in compression efficiency is offset by the advantage of mitigating transport failures.

In the remainder of this chapter, unless specified otherwise, we focus on the two-description case for the sake of conciseness and clarity. Nevertheless, many of the ideas and results presented in the sequel can be extended to the case of more than two descriptions.

The basic situation of MD coding, for the two-description case, is as follows [15]: Suppose that we wish to send a description of a random process (information source) over an unreliable communication network—one with nonzero probability for the event of description loss. In order to deal with the possibility of a lost description, or due to constraints imposed by the network regarding the size of packets, we decide to send two descriptions instead and hope that at least one of them will get through. Obviously, each description should be individually good due to the possibility that only one description will get to the destination. However, in the case that both descriptions get through, we wish to maximize the combined descriptive information at the receiver.

The difficulty in generation of such descriptions is that good individual descriptions must be close to the process (by virtue of their goodness) and therefore must be highly dependent. Such a dependency implies that the second description will contribute only little extra information beyond one description alone. On the other hand, two descriptions that are independent from each other must be far apart and therefore cannot in general be individually good. Essentially, this demonstrates the fundamental tradeoff of MD coding: creating descriptions that are individually good and yet not too similar.

2.2 Information Theoretic Aspects

The following is known as the MD problem: if an information source is described with two separate descriptions, what are the concurrent limitations on the qualities of these descriptions taken separately and jointly?

The MD situation for the two-description case is shown schematically in Figure 2.1.

A sequence of source symbols $\{X_k\}_{k=1}^{N}$ is encoded and transmitted to three receivers

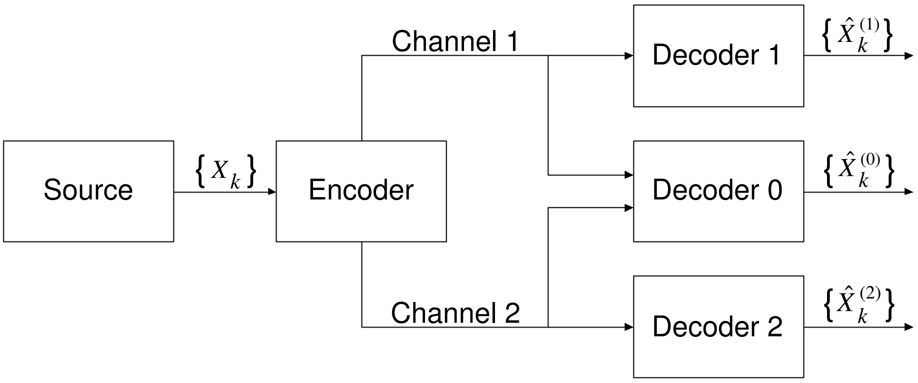

Figure 2.1: Scenario for MD coding with two channels and three receivers. In the general case there are M channels and $2^M - 1$ receivers.

over two noiseless (or error corrected) channels. One decoder, called the *central decoder*, receives information from both channels, while the other two decoders, called the *side decoders*, receive information only from their respective channels. We denote the transmission rate over Channel i by R_i, $i = 1, 2$, in bits per source sample. The reconstruction sequence produced by Decoder i is denoted by $\{\hat{X}_k^{(i)}\}_{k=1}^{N}$ with respective distortion

$$D_i = \frac{1}{N} \sum_{k=1}^{N} E\left[d\left(X_k, \hat{X}_k^{(i)} \right) \right], \; i = 0, 1, 2 \tag{2.1}$$

where $d(\cdot, \cdot)$ is a nonnegative, real-valued distortion measure[1].

It should be noted that the situation drawn in Figure 2.1 of three separate users or three classes of users (applicable to broadcasting over two channels, for example) can also serve as an abstraction to the case of a single user that can be in one of three

---

[1] Actually, a different distortion measure may be used for each of the reconstructed sequences. A single distortion measure is assumed here for simplicity.

states depending on which descriptions are received. We will usually adopt the latter interpretation.

Several interesting problems in rate distortion theory are introduced by the MD model. The central theoretical problem is the following [21]: given an information source and a distortion measure, determine the set of achievable values for the quintuple $(R_1, R_2, D_0, D_1, D_2)$. These theoretical bounds are of major importance and are discussed next.

Recall the definition of the rate distortion region in the one-description rate distortion problem: for a given source and distortion measure, the *rate distortion (RD) region* is the closure of the set of achievable rate distortion pairs (where a rate distortion pair (R, D) is called *achievable* if, for some positive integer n, there exists a source code with length n, rate R and distortion D). As an extension, the *MD rate distortion region*, or simply *MD region* is defined, for a given source and distortion measure, as the closure of the set of simultaneously achievable[2] rates and distortions in MD coding. Specifically, for the two-description case, the MD region is the closure of the set of achievable quintuples $(R_1, R_2, D_0, D_1, D_2)$.

Unfortunately, unlike RD regions in the one-description rate distortion problem [48], MD regions have not been completely characterized in terms of single-letter information theoretic quantities (entropy, mutual information, etc.). The following theorem of El Gamal and Cover [15] shows how achievable quintuples can be determined from joint distributions of source and reproduction random variables.

Theorem 2.1 [15] *Let $X_1, X_2, \ldots$ be a sequence of i.i.d. finite alphabet*[3] *random*

---

[2] The meaning of "achievable" in this context is the immediate extension from the one-description rate distortion problem.

[3] Although the proof of the theorem assumes a finite alphabet, it can be extended to the Gaussian source, as noted in [15].

variables drawn according to a probability mass function $p(x)$*. Let* $d(\cdot,\cdot)$ *be a bounded distortion measure. An achievable rate region for distortion* $\mathbf{D} = (D_1, D_2, D_0)$ *is given by the convex hull of all* (R_1, R_2) *such that*

$$\begin{aligned} R_1 &> I(X;\ \hat{X}^{(1)}) && (2.2)\\ R_2 &> I(X;\ \hat{X}^{(2)}) && (2.3)\\ R_1 + R_2 &> I(X;\ \hat{X}^{(0)}, \hat{X}^{(1)}, \hat{X}^{(2)}) + I(\hat{X}^{(1)};\ \hat{X}^{(2)}) && (2.4) \end{aligned}$$

for some probability mass function $p(x, \hat{x}^{(0)}, \hat{x}^{(1)}, \hat{x}^{(2)}) = p(x)p(\hat{x}^{(0)}, \hat{x}^{(1)}, \hat{x}^{(2)}|x)$ *such that*

$$\begin{aligned} D_1 &\geq E\left[d(X, \hat{X}^{(1)})\right] && (2.5)\\ D_2 &\geq E\left[d(X, \hat{X}^{(2)})\right] && (2.6)\\ D_0 &\geq E\left[d(X, \hat{X}^{(0)})\right] && (2.7) \end{aligned}$$

and where $I(\cdot\,;\ \cdot)$ *denotes Shannon mutual information [13].*

It should be noted that Theorem 2.1 does not specify *all* achievable quintuples, unfortunately. The points that can be obtained in this manner are therefore called *inner bounds* to the MD region. Other characterizations of achievable points in terms of information theoretic quantities for the two-description case [68] and for more than two descriptions [63] also generally do not give the entire MD region.

Points that are certainly not in the MD region are called *outer bounds.* Rate distortion functions provide the simplest outer bounds: In order to have distortion D_1, Decoder 1 must receive at least $R(D_1)$ bits per symbol, where $R(D)$ is the rate distortion function of the source (with the given distortion measure). By making similar arguments for the other two decoders we get the bounds

$$R_1 + R_2 \geq R(D_0) \tag{2.8}$$

$$R_i \geq R(D_i), \quad i = 1, 2. \tag{2.9}$$

The bounds (2.8) and (2.9) are usually loose due to the difficulty in making the individual and joint descriptions good.

Memoryless Gaussian sources and the squared error distortion measure[4] are of special importance (also) in MD coding. First, this is the only case for which the MD region is completely known. Ozarow showed in [41] that for this source and distortion measure, the MD region is exactly the largest set that can be obtained with the achievable region of El Gamal and Cover. Moreover, the MD region for any continuous-valued memoryless source with squared error distortion measure can be bounded using the MD region for Gaussian sources [67]. Useful intuition and insight into the limitations in MD coding can be gained from this special case and so we now turn to a more detailed analysis of it.

For a memoryless Gaussian source with variance σ^2 and squared error distortion measure, the MD region consists of $(R_1, R_2, D_0, D_1, D_2)$ that satisfy [21]

$$D_i \geq \sigma^2 2^{-2R_i}, \quad i = 1, 2 \tag{2.10}$$

$$D_0 \geq \sigma^2 2^{-2(R_1+R_2)} \cdot \gamma_D \tag{2.11}$$

where $\gamma_D = 1$ if $D_1 + D_2 > \sigma^2 + D_0$, and

$$\gamma_D = \frac{1}{1 - \left(\sqrt{(1-D_1)(1-D_2)} - \sqrt{D_1 D_2 - 2^{-2(R_1+R_2)}}\right)^2} \tag{2.12}$$

otherwise (obviously $\gamma_D \geq 1$ from (2.11) and the fact that $D(R) = \sigma^2 2^{-2R}$ is the distortion rate function for a Gaussian source with squared error distortion). The

---

[4] The squared error distortion measure is given by $d(x, \hat{x}) = (x - \hat{x})^2$.

relation (2.11) is pivotal as it indicates that there is a multiplicative factor γ_D by which the central distortion must exceed the distortion rate minimum. When at least one of the side distortions is sufficiently large, $\gamma_D = 1$ and the central reconstruction can be very good. Otherwise, there is a penalty in the central distortion, determined by the value of γ_D from (2.12).

In the special case that $R_1 = R_2$ and $D_1 = D_2$, called the *balanced* case in the terminology of MD coding, the following side distortion bound for a source with unit variance can be proved to hold [20] (in addition to the bound $D_1 \geq \sigma^2 2^{-2R_1}$):

$$D_1 \geq min \left\{ \frac{1}{2} \left[1 + D_0 - (1 - D_0) \sqrt{1 - 2^{-2(R_1+R_2)}/D_0} \right], 1 - \sqrt{1 - 2^{-2(R_1+R_2)}/D_0} \right\}. \tag{2.13}$$

Introducing the total rate $R_1 + R_2$ as the sum of a base rate $r = R(D_0)$ and a redundancy $\rho = R_1 + R_2 - R(D_0)$ gives the alternative formulation

$$D_1 \geq \begin{cases} \frac{1}{2} \left[1 + 2^{-2r} - (1 - 2^{-2r}) \sqrt{1 - 2^{-2\rho}} \right], & \text{for } \rho \leq r - 1 + \log_2(1 + 2^{-2r}) \\ 1 - \sqrt{1 - 2^{-2\rho}}, & \text{for } \rho > r - 1 + \log_2(1 + 2^{-2r}). \end{cases} \tag{2.14}$$

This bound is plotted in Figure 2.2 as a function of the redundancy ρ for several values of the base rate r (in bits/sample).

It is instructive to examine the behavior of the bound (2.14) in the low redundancy region. In this region, the partial derivative of the side distortion lower bound with respect to ρ is given by

$$\frac{\partial D_1}{\partial \rho} = -\frac{(1 - 2^{-2r}) \ln 2}{2} \cdot \frac{2^{-2\rho}}{\sqrt{1 - 2^{-2\rho}}} \tag{2.15}$$

and for any fixed $r > 0$ we get

$$\lim_{\rho \to 0^+} \frac{\partial D_1}{\partial \rho} = -\infty. \tag{2.16}$$

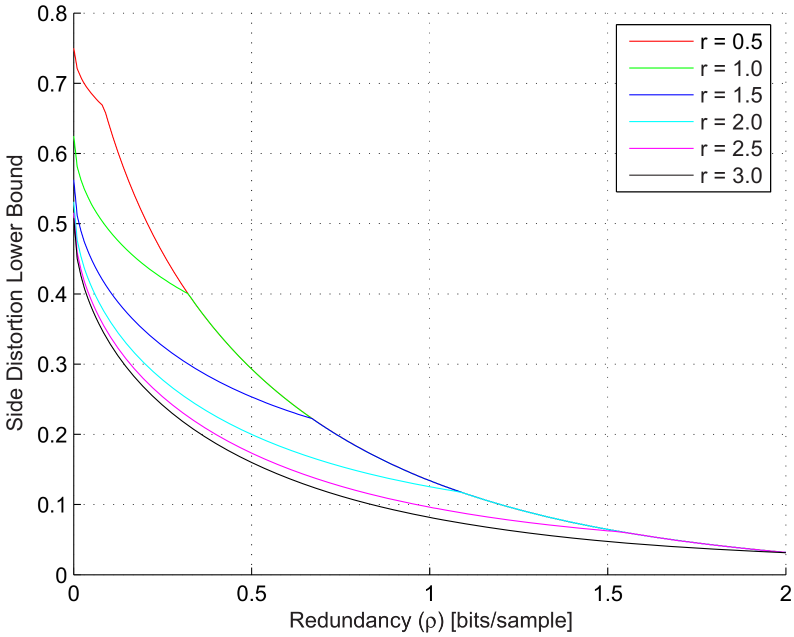

Figure 2.2: Side distortion lower bound vs. redundancy ρ for several values of base rate r.

For an interpretation of this result, consider an MD coding system that achieves the rate distortion bound for the central decoder. If a small additional increase in rate is allowed, what would be an intelligent way to spend it? A common performance measure in MD coding systems is some linear combination of central and side distortions (in which the weights usually correspond to probabilities of receiving certain combinations of descriptions). In such a case, the infinite slope at $\rho = 0^+$ implies that it will be much more efficient to dedicate the small additional rate to reducing the side distortion rather than to reducing the central distortion. Essentially, a nonzero redundancy should ideally be used by such a system.

The significant effect of a small amount of redundancy in MD coding systems is observed even when actual performance is considered, not merely bounds. For example, in the MD coding technique of correlating transforms (discussed in Section 2.3.3), an infinite side distortion slope is again observed at zero redundancy [23, 24].

Ozarow's result is often interpreted as an exponential lower bound on the product of central and side distortions. Assuming that $R_1 = R_2 \gg 1$ and $D_1 = D_2 \approx 2^{-2(1-\nu)R_1}$ with $0 < \nu \leq 1$, γ_D can be estimated as $(4D_1)^{-1}$. The bound in (2.11) then gets the form

$$D_0 \cdot D_1 \geq \frac{\sigma^2}{4} 2^{-4R_1} \tag{2.17}$$

which explains the reason for this interpretation of Ozarow's result. This optimal rate of exponential decay of the product of central and side distortions can be attained by MD quantization techniques [59, 60] (MD quantization is discussed in Section 2.3.2).

2.3 Generation of Multiple Descriptions

This section provides a concise description of some common MD coding techniques. Powerful MD codes can be built from combination of MD coding techniques, like MD quantizers and transforms [21], along with channel codes and basic components of modern compression systems, such as prediction, scalar quantization, decorrelating transforms and entropy coding [22].

One of the simplest ways to produce multiple descriptions is to partition the source data into several parts and then compress each part independently to produce the descriptions. Proper interpolation can be used to decode from any (non-empty) subset of the descriptions. For example, one can separate a speech signal into odd- and even-numbered samples to produce two descriptions, encode them independently and employ one-dimensional interpolation upon decoding if one of the descriptions is lost, as in Jayant's subsample-interpolation (SI) approach [28, 29]. This separation can also be adapted to multidimensional data, such as images [52]. Although this technique can be very powerful, it has a serious drawback, as it relies completely on the existence and amount of redundancy already present in the source (in the form of correlation or memory). Additionally, in order to be complementary to modern compression systems that initially reduce correlation by means of prediction or decorrelating transforms, an MD coding technique must work well on memoryless sources. We now turn to describe such techniques in detail.

2.3.1 Progressive Coding and Unequal Error Protection

Generation of two descriptions can be easily performed by duplication of a single description of the source. This description can be produced with the best (single-description) compression technique available. If only one description is received, the performance is the best possible in such a case. Nevertheless, receiving both descriptions provides no further improvement in performance.

A more flexible and controllable approach is to repeat only some fraction of the data. Ideally, the repeated data would be the most important, and thus this type of fractional repetition is naturally matched to progressive source coding. In order to produce two R-bit descriptions, first set $\zeta \in [0,1]$ and encode the source to a rate $(2-\zeta)R$ with a progressive source code. The first ζR bits are the most important and thus are repeated in both descriptions, while the remaining $(2-\zeta)R-\zeta R = 2(1-\zeta)R$ bits are split between the descriptions (in a meaningful way). Essentially, this strategy protects some of the bits with a rate-1/2 channel code[5] (repetition), while the other bits are left unprotected. It is therefore called an unequal error protection (UEP) strategy for MD coding. UEP can also be easily generalized to the case of more than two descriptions.

When designing an MD coding system based on UEP, one has to determine the amount of data to be coded at each channel code rate. Techniques for this assignment that are based on optimization of a scalar objective function are described in [38, 43].

In order to assess the performance of UEP as an MD coding technique, the MD rate distortion region (defined in Section 2.2) can be compared to the corresponding

---

[5] A block channel code that maps k input symbols to n output symbols, with $n > k$, is said to have rate k/n.

region attained with UEP. Both regions can be determined completely for the two-description case with memoryless Gaussian source and squared error distortion measure. Comparing the minimum attainable central distortions for fixed values of side distortions can be used to characterize the difference between these regions. The maximum gap in central distortion over all rates and side distortions is about 4.18 dB [20]. Although bounded, this gap is quite significant, and analogous comparisons show that it increases for more descriptions [63].

2.3.2 Multiple Description Scalar Quantization

We begin this section with a description of a natural way to use scalar quantizers to create multiple descriptions for communication over two channels—a purpose also known in this context as *channel splitting* due to its origin. We later introduce a more systematic and efficient approach to construct scalar quantizer pairs known as *multiple description scalar quantization (MDSQ).*

We start with an example demonstrating the difficulty in creating two useful descriptions of a source without added redundancy. Consider trying to communicate, using 4 bits, a single real number $x \in [-1, 1]$. One option, which is optimal if the source is uniformly distributed, is the four-bit uniform scalar quantizer shown in Figure 2.3(a) with black labels. There is no satisfactory way to split the four bits into two pairs for transmission over two channels since any reconstruction by the side decoder that does not receive the most significant bit (MSB) will, on average, be poor. One reasonable labeling under these circumstances is shown in Figure 2.3(a) with red and blue labels, but even this does not produce good estimates in a consistent manner.

A natural way to use scalar quantizers for channel splitting is to combine two

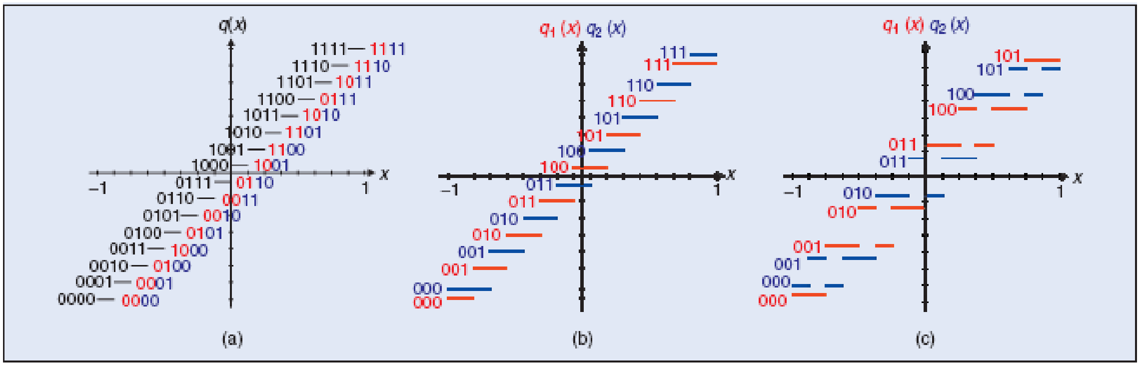

Figure 2.3: (a) A four-bit uniform scalar quantizer. (b) Two three-bit quantizers that complement each other so that both outputs together give about four-bit resolution. (c) More complicated quantizers together attain four-bit resolution while each having fewer output levels. (From [21])

uniform quantizers (one for each channel) with an offset between them, as shown in Figure 2.3(b) in red and blue. Due to this offset, combining the information from both channels gives one additional bit of resolution. To avoid clutter, we do not show the reconstructions computed when both channels are available; as an example, if $q_1(x) = 110$ and $q_2(x) = 101$, then x must be in the interval $[7/16, 9/16)$ and thus is reconstructed to $1/2$.

Indeed, the quantizers in Figure 2.3(b) produce good side reconstructions (i.e., the reconstructions from either channel alone are good), but this comes at the cost of a fairly high redundancy: the total rate over both channels is 6 bits/sample, while the reconstruction quality when both channels are received is comparable to that of the 4-bit quantizer in Figure 2.3(a). More generally, using two B-bit quantizers (offset from each other) to create the descriptions produces approximately $(B+1)$-bit resolution for the central decoder, while the total rate is $2B$ bits/sample. Due to the high redundancy, this solution is usually unsatisfactory (at least unless the side reconstructions are of extreme importance).

Reudink [44] was the first to propose efficient channel splitting techniques with a reasonable amount of redundancy that do not rely entirely on preexisting redundancy in the source sequence. The quantizers shown in Figure 2.3(c) in red and blue are based on one of Reudink's families of quantizer pairs. Although each of these quantizers looks somewhat strange due to the fact that its quantization cells are not connected, both quantizers complement each other to produce cells with small intersections. For example, knowing that $q_1(x) = 100$ only limits x to $[1/4, 3/8) \bigcup [1/2, 3/4)$; also knowing that $q_2(x) = 100$ then localizes x to the interval $[1/2, 5/8)$. Note that each quantizer in Figure 2.3(c) has only six outputs. Considering each of these quantizers as a $\log_2 6 \approx 2.6$-bit quantizer, the quantizer pair in Figure 2.3(c) attains 4-bit resolution with a total rate which is lower than that of the quantizers in Figure 2.3(b).

Reudink's work was archived as a technical report [44], but was not published. Many years later, Vaishampayan [59] independently developed a theory for designing MD scalar quantizers. Based on this theory, we now give a formal notation for MD scalar quantization and a systematic way to construct scalar quantizer pairs like those in Figure 2.3.

A fixed-rate MD scalar quantizer is comprised of an encoder α_0 and three decoders: a central decoder β_0 and two side decoders β_1 and β_2. For each scalar sample x of the source, the encoder α_0 produces a pair of quantization indices (i_1, i_2). The central decoder β_0 produces an estimate of x from (i_1, i_2) and the side decoders β_1 and β_2 produce estimates from i_1 and i_2, respectively. For example, in Figure 2.3(b) and (c) the i_1 and i_2 indices are shown in red and blue, respectively. The reconstructions by the side decoders β_1 and β_2 are indicated by vertical positions, while the reconstruction by the central decoder β_0 is implicit.

Vaishampayan [59] proposed a useful visualization technique for the encoding operation. First, the encoding operation α_0 is presented as $\alpha_0 = l \circ \alpha$—the composition of an index assignment l and an initial encoder α. The initial encoder α is an ordinary quantizer, i.e., it partitions the real line into cells that are intervals. The index assignment l, which must be invertible, produces a pair of indices (i_1, i_2) from the single index produced by the ordinary quantizer α. Vaishampayan's visualization technique is to write out l^{-1}, forming the index assignment matrix.

The index assignment matrix corresponding to the quantizer in Figure 2.3(b) is shown in Table 2.1, where the cells of the initial encoder α, taken in increasing values of x, are numbered from 0 to 14. The redundancy in the representation by this quantizer is indicated by the fact that the corresponding index assignment matrix has only 15 out of 64 cells occupied. The qualities of the side reconstructions are indicated by the relatively small range of values in any row or column of the matrix (a maximum difference of 1).

	000	001	010	011	100	101	110	111
000	0							
001	1	2						
010		3	4					
011			5	6				
100				7	8			
101					9	10		
110						11	12	
111							13	14

Table 2.1: Index assignment matrix for the quantizer in Figure 2.3(b). The indices in the left column and top row of the matrix correspond to the red and blue labels in Figure 2.3(b), respectively.

Lower redundancy can be achieved using an index assignment matrix with a higher fraction of occupied cells. Table 2.2 shows the index assignment matrix for the quantizer in Figure 2.3(c). A higher fraction of occupied cells (16 out of 36) in this example

	000	001	010	011	100	101
000	0	1				
001	2	3	5			
010		4	6	7		
011			8	9	11	
100				10	12	13
101					14	15

Table 2.2: Index assignment matrix for the quantizer in Figure 2.3(c). The indices in the left column and top row of the matrix correspond to the red and blue labels in Figure 2.3(c), respectively.

indicates lower redundancy, but the extended range of values in any row or column of the index assignment matrix (a maximum difference of 3) implies that the side distortions are higher.

If all the cells of the index assignment matrix are occupied, there is no redundancy. However, the side distortions in this case are necessarily high. For a four-bit quantizer α and zero redundancy, Table 2.3 shows the best possible index assignment, as also marked in Figure 2.3(a) with red and blue labels.

	00	01	10	11
00	0	1	5	6
01	2	4	7	12
10	3	8	11	13
11	9	10	14	15

Table 2.3: Index assignment matrix for the quantizer in Figure 2.3(a). The indices in the left column and top row of the matrix correspond to the red and blue labels in Figure 2.3(a), respectively.

In the design of an MD scalar quantizer, the optimization of the ordinary quantizer α and the decoders β_0, β_1 and β_2 is relatively easy. The optimization of the index assignment l is very difficult, however. Instead of addressing the exact optimal index assignment problem, Vaishampayan [59] gave several heuristic guidelines that are likely to give performance which is close to the best possible. Essentially, the cells of the index assignment matrix should be filled from the upper-left to the lower-right and from the main diagonal outward, as demonstrated in the examples above.

The theory of designing MD scalar quantizers is extended from fixed-rate quantization to entropy-constrained quantization in [61].

The design of an MD scalar quantizer offers convenient adaptation of the quantizer according to the relative importance of the central distortion and each side distortion. For the balanced case, where $R_1 = R_2$ and $D_1 = D_2$, and assuming high rates, the central and side distortions can be traded off while keeping their product constant [59, 61]. In addition, it can be shown that the exponential decay of this product as a function of the rate matches the optimal decay implied by [41], while the leading constant terms are consistent with what one would expect from high-resolution analysis of ordinary (single description) scalar quantizers [60].

The formalism of MD scalar quantization can be extended to MD vector quantization (MDVQ), but the actual quantizer design and encoding may become impractical if no additional constraints are placed. This is because the complexity of the initial quantization mapping α increases exponentially with the vector's dimension N and because the heuristic guidelines for the design of the index assignment l in MDSQ do not extend to MDVQ as there is no natural order on $\mathbb{R}^N$. These difficulties are avoided by the elegant technique of *MD lattice vector quantization (MDLVQ)* of Servetto et al. [47, 62], where the index assignment problem is simplified by lattice symmetries and the complexity is reduced due to the structure of the lattice [12]. Other related MD quantization techniques are described in [17, 18, 26].

Obviously, MD scalar quantization can also be applied to transform coefficients. For the case of high rates, it is shown in [3] that the transform optimization problem is essentially unchanged from conventional (single description) transform coding. The technique of MD correlating transforms described next also utilizes transforms and scalar quantizers, but, in contrast to MD scalar quantization, gets its MD character from the transform itself.

2.3.3 Multiple Description Correlating Transforms

This section provides a concise description of a useful technique for meeting MD coding objectives in the framework of standard transform-based coding through the design of correlating transforms [24]. The basic idea in standard transform coding is to eliminate statistical dependency and produce uncorrelated transform coefficients [22]. In the scenario of MD coding, however, statistical dependency between transform coefficients may prove useful as it can improve the estimation of transform coefficients that are in a lost description. The technique of MD coding using correlating transforms, proposed by Wang et al. [65], is based on explicit introduction of correlation between pairs of random variables through a linear transform (called a *correlating transform* in this context due to its purpose).

Relating to the two-description case, let X_1 and X_2 be statistically independent zero-mean Gaussian random variables with variances $\sigma_1^2 > \sigma_2^2$. In the case of conventional (single description) source coding, there would be no reason to use a linear transform prior to quantization. For a rate R (in bits per sample) and assuming high-rate entropy-coded uniform quantization, the MSE distortion per component would be given by [22]

$$D_0 = \frac{\pi e}{6}\sigma_1\sigma_2 2^{-2R}, \tag{2.18}$$

which is the best possible performance in this case.

In the two-description MD scenario, suppose that the quantized versions of X_1 and X_2 are sent on channels 1 and 2, respectively. Due to the statistical independence of X_1 and X_2, side decoder 1 cannot do better than estimating X_2 using its mean. Therefore

$$D_1 = \frac{\pi e}{12}\sigma_1\sigma_2 2^{-2R} + \frac{1}{2}\sigma_2^2 \tag{2.19}$$

and, similarly,

$$D_2 = \frac{\pi e}{12}\sigma_1\sigma_2 2^{-2R} + \frac{1}{2}\sigma_1^2. \tag{2.20}$$

Under the assumption that each channel is equally likely to fail, instead of considering D_1 and D_2 separately, it is reasonable to consider the average distortion when one description is lost

$$\overline{D}_1 = \frac{1}{2}(D_1 + D_2) = \frac{1}{4}(\sigma_1^2 + \sigma_2^2) + \frac{\pi e}{12}\sigma_1\sigma_2 2^{-2R}. \tag{2.21}$$

The average side distortion $\overline{D}_1$ could be reduced if side decoder i had some information about X_j, $j \neq i$. This can be achieved by sending correlated transform coefficients instead of the X_i's themselves, i.e., to send quantized versions of Y_1 and Y_2 given by

$$\begin{bmatrix} Y_1 \\ Y_2 \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}. \tag{2.22}$$

Specifically, let us consider the original transform proposed in [65], where

$$\begin{bmatrix} Y_1 \\ Y_2 \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}. \tag{2.23}$$

Since $[Y_1, Y_2]^T$ is obtained using an orthonormal transformation and MSE distortion is considered, the distortion in approximating the X_i's equals the distortion in approximating the Y_i's. Since the variance of both Y_1 and Y_2 is $(\sigma_1^2 + \sigma_2^2)/2$, the central distortion is given by

$$D_0 = \frac{\pi e}{6} \left(\frac{\sigma_1^2 + \sigma_2^2}{2} \right) 2^{-2R}, \tag{2.24}$$

which is worse than the performance attained without using the correlating transform by a constant factor of

$$\Gamma = \frac{(\sigma_1^2 + \sigma_2^2)/2}{\sigma_1 \sigma_2}. \tag{2.25}$$

Note that $\Gamma \geq 1$ as the ratio of an arithmetic and a geometric mean, and that $\sigma_1^2 > \sigma_2^2$ implies $\Gamma > 1$.

For each side decoder, the total distortion is approximately equal to the quantization error plus the distortion in estimating the missing coefficient from the received one. Using the fact that Y_1 and Y_2 are jointly Gaussian, it can be shown that [24]

$$\overline{D}_1 \approx \frac{\sigma_1^2 \sigma_2^2}{\sigma_1^2 + \sigma_2^2} + \frac{\pi e}{12} \left(\frac{\sigma_1^2 + \sigma_2^2}{2} \right) 2^{-2R}. \tag{2.26}$$

By comparing (2.21) and (2.26), one can see that the use of the transform given by (2.23) reduced the constant term in $\overline{D}_1$ by the factor Γ^2, while the exponential term was increased by the factor Γ. Along with (2.18) and (2.24), this shows that the technique of correlating transforms enables the central and side distortions to be traded off. Intermediate tradeoffs can be obtained by using other orthogonal transforms. As also discussed later in this section, additional operating points, including more extreme tradeoffs, can be attained by using nonorthogonal transforms [39]. Note that the factor Γ approaches unity as $\sigma_1/\sigma_2 \to 1$; thus, the pairwise correlation has no effect when the joint density of X_1 and X_2 has spherical symmetry (this fact holds equally well for both orthogonal and nonorthogonal transforms).

We now give a useful geometric interpretation to the use of correlating transforms for MD coding. As before, let $X = [X_1,\ X_2]^T$ where X_1 and X_2 are statistically independent zero-mean Gaussian random variables with variances $\sigma_1^2 > \sigma_2^2$. Denoting the standard basis of $\mathbb{R}^2$ by $\{e_1,\ e_2\}$, we note that the level curves of the joint probability density function (pdf) of X are ellipses with principal axis aligned with e_1 and secondary axis aligned with e_2. Without using a correlating transform, X is essentially represented by $(\langle X,\ e_1\rangle,\ \langle X,\ e_2\rangle)$, where $\langle\cdot,\ \cdot\rangle$ denotes inner product. This is shown in Figure 2.4(a).

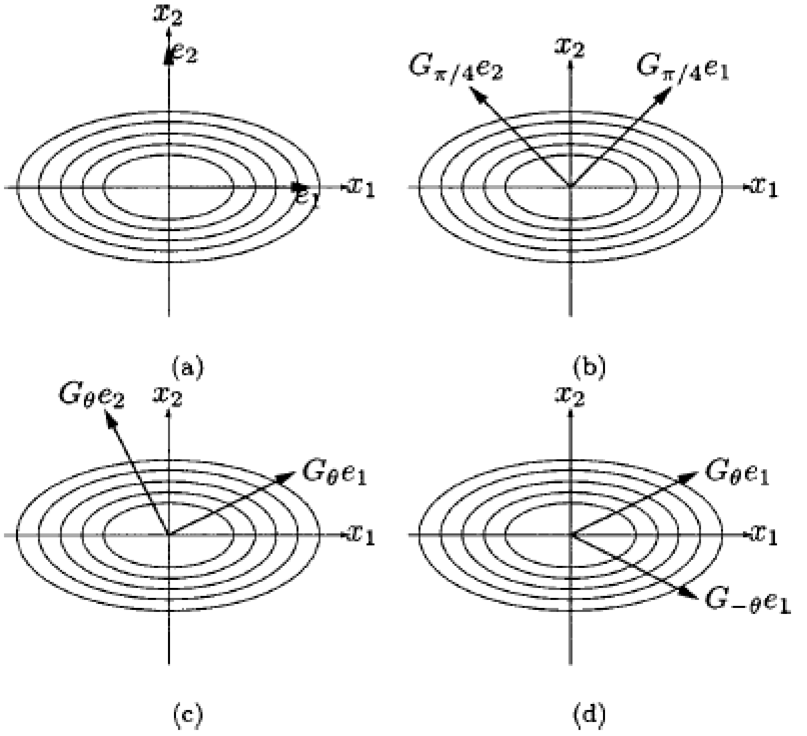

Figure 2.4: Basis configurations for correlating transforms. (a) The standard basis. (b) Basis for the original correlating transform of [65]. (c) Generalization to arbitrary orthogonal bases. (d) Bases that are symmetric with respect to the principal axis of the source density. (From [24])

Assume that uniform scalar quantized versions of $\langle X,\, e_1\rangle$ and $\langle X,\, e_2\rangle$ are used as descriptions. For a given total rate, it was demonstrated earlier in this section that the central distortion D_0 and the average side distortion $\overline{D}_1 = (D_1 + D_2)\,/\,2$ can be traded off by using the correlating transform of (2.23), which corresponds to using the representation

$$\left(\langle X,\, \frac{1}{\sqrt{2}}[1,1]^T\rangle,\, \langle X,\, \frac{1}{\sqrt{2}}[-1,1]^T\rangle\right). \tag{2.27}$$

This is only a single operating point, however, and it may be desirable to trade off D_0 and $\overline{D}_1$ in a continuous manner.

The representation in (2.27) is immediately recognized as

$$\big(\langle X,\, G_{\pi/4}e_1\rangle,\, \langle X,\, G_{\pi/4}e_2\rangle\big) \tag{2.28}$$

where G_θ is a Givens rotation of angle θ (see Figure 2.4(b)). Thus, a natural extension is to consider all representations of the form (see Figure 2.4(c))

$$\left(\langle X,\, G_\theta e_1\rangle,\, \langle X,\, G_\theta e_2\rangle\right), \text{ for } 0 \le \theta \le \pi/4 \tag{2.29}$$

which offers a continuous tradeoff between D_0 and $\overline{D}_1$. This representation has an undesirable asymmetry, however, as it produces unequal side distortions (except for $\theta = \pi/4$). A geometric reasoning for this asymmetry, based on Figure 2.4(c), is as follows. D_1 is the variation of X which is *not* captured by $\langle X,\, G_\theta e_1\rangle$, or the variation perpendicular to $G_\theta e_1$ (neglecting the quantization error as D_1 and D_2 are equally affected by quantization). Similarly, D_2 is the variation perpendicular to $G_\theta e_2$. Due to the fact that $G_\theta e_1$ and $G_\theta e_2$ are not symmetrically situated with respect to the pdf of X (except for $\theta = \pi/4$), the side distortions D_1 and D_2 are unequal.

The equal weighting of D_1 and D_2 in $\overline{D}_1$ suggests equal importance of both descriptions, and so equal side distortions may be desirable. If so, based on the geometric observation above, it is reasonable to represent X by (see Figure 2.4(d))

$$(\langle X,\, G_\theta e_1\rangle,\, \langle X,\, G_{-\theta} e_1\rangle)\,,\ \text{for } 0 < \theta < \pi/2. \tag{2.30}$$

Additionally, note that in order to capture most of the principal component of the source, the basis should be skewed toward e_1 (see Figure 2.4(d)). Although equal side distortions are attained using this representation (i.e., $D_1 = D_2$), its use introduces a new problem.

The representation of X using (2.30) is a nonorthogonal basis expansion (for $\theta \neq \pi/4$). The uniform scalar quantization of such a representation produces non-square partition cells. These non-square partition cells are undesirable as they have normalized second moments that are higher than those of square cells, and thus may lead to higher distortions [22]. The insight attributed to Vaishampayan is that a correlating transform can be applied *after* quantization has been performed in an orthogonal basis representation, which ensures that the partition cells are square regardless of the transform. The advantage of this approach over the original method of pairwise correlating transforms was demonstrated in [39]. Although the transform in this case maps from a discrete set to a discrete set, it can be designed to approximate a linear transform [23, 24].

From the geometric view of Figure 2.4, one can easily understand the significance of the ratio σ_1/σ_2 in MD coding using correlating transforms. When $\sigma_1/\sigma_2 \rightarrow 1$, the level curves become circles and the variation perpendicular to either basis vector is then invariant to the basis; the side distortions are thus unaffected by the correlating transform and this holds equally well for both orthogonal and nonorthogonal bases.

The MD coding technique of correlating transforms can also be generalized to the case of more than two descriptions. Such an extension is described and analyzed in [24].

Chapter 3

Multiple Description Image Coding via Polyphase Transform

This chapter reviews an efficient technique for MD image coding utilizing a polyphase transform and selective quantization [30]. This method forms the framework of our proposed context-based MD coding system, introduced in Chapter 4.

3.1 System Outline

Unlike many other MD coding techniques, such as MD scalar quantization and MD correlating transforms (both are described in Chapter 2), the technique of MD coding via polyphase transform and selective quantization explicitly separates *description generation* and *redundancy addition.* This separation can significantly reduce the complexity of the system design and implementation, especially in the case of more than two descriptions. The polyphase transform-based technique described in this chapter also enables to easily generate descriptions of statistically equal rate and importance, a property which is well suited to communication systems with no priority mechanisms for data delivery, such as the Internet.

For description generation, this approach employs a polyphase transform (i.e., decomposition to polyphase-like components; specific options are described below), and each of the 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, other polyphase components are then coded at a (usually lower) redundancy coding rate using selective quantization and added to this description. In case of channel failures, this redundancy enables an acceptable reconstruction of the lost components. This is shown in Figure 3.1 for the two-description case.

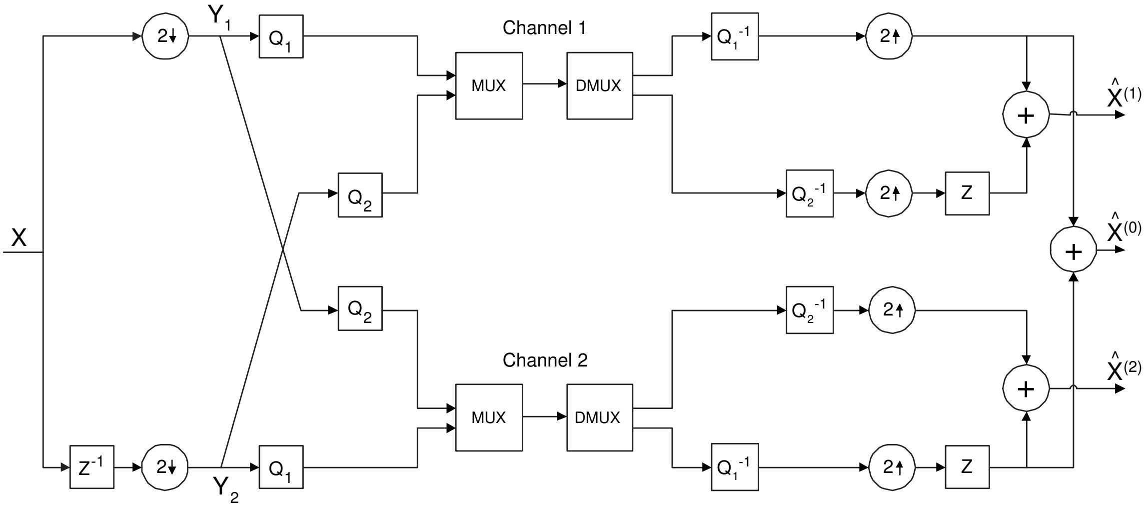

Figure 3.1: MD coding system based on polyphase transform and selective quantization.

Referring to Figure 3.1, the input X is first decomposed into two sub-sources Y_1 and Y_2 via a polyphase transform. Each of the two polyphase components Y_1 and Y_2

is then quantized independently by Q_1 to form the primary part of information for its corresponding description. 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—a coarsely[1] quantized version produced by the quantizer Q_2. The outputs of the appropriate quantizers are combined together by a multiplexer to form the descriptions, and each description is transmitted on its corresponding channel. At the receiver, if both descriptions arrive, the finely quantized versions of the two polyphase components (produced by Q_1 during encoding) are then used for reconstruction to produce $\hat{X}^{(0)}$. If one of the descriptions is lost, one finely quantized and one coarsely quantized polyphase component (the primary and redundant part of the received description, respectively) are used for reconstruction (the reconstruction from description i only is denoted by $\hat{X}^{(i)}$, $i = 1, 2$).

Note that in the encoded bit stream the redundant information (coarsely quantized polyphase components) is explicitly separated from the primary information (finely quantized polyphase components), which simplifies the design and implementation of the encoder and decoders in the system. Compared to the technique of MD correlating transforms, no use of an additional correlating transform is needed here, while careful index assignments required by MD scalar quantizers are also avoided by this approach. Despite the similarity to Jayant's subsample-interpolation (SI) approach [28, 29] (described in Section 2.3), the polyphase transform-based technique does not rely on the existence and amount of redundancy already present in the source and thus can also be applied to memoryless sources.

---

[1] The quantizers Q_1 and Q_2 are referred to as "fine" and "coarse", respectively, in order to demonstrate the concept of selective quantization. The actual relation between the resolutions of the different quantizers is generally determined according to the probability of description loss (i.e., channel failure).

The polyphase transform-based technique can also be adapted to the case of correlated input data, as shown in Figure 3.2. The correlated input data (e.g., an image) is first transformed by a decorrelating transform (e.g., KLT [22]). In the *description generation* stage, a polyphase transform is applied to the transform coefficients and each resulting polyphase component forms the primary part of a single description. The *redundancy addition* stage identifies for each description which of the other descriptions it will protect and introduces redundancy among descriptions accordingly. Each description is then quantized and entropy-coded independently, where the primary information is coded at a source coding rate and the redundant information is coded at a (usually lower) redundancy coding rate. The coded descriptions are independently packetized and transmitted over the network (e.g., the Internet).

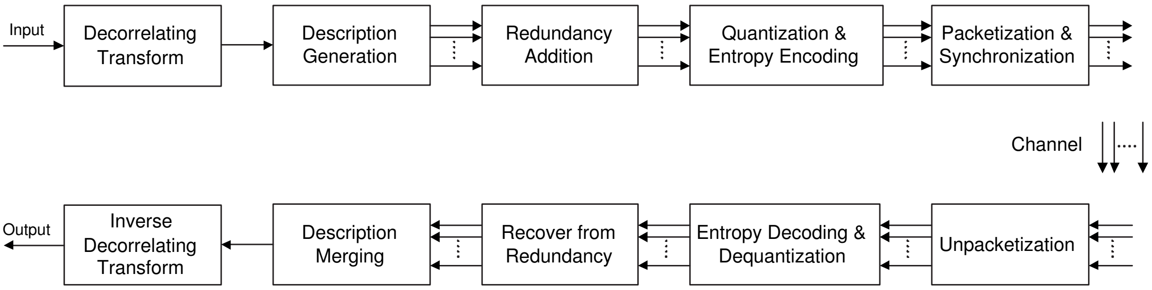

Figure 3.2: MD coding system based on polyphase transform and selective quantization for correlated input.

For decoding, the decoder first identifies which descriptions are available and which are lost. It then decodes the available descriptions and uses the redundancy information to reconstruct the polyphase components whose corresponding descriptions are missing. In the description merging stage the reconstructed polyphase components are recomposed together and finally the inverse of the decorrelating transform is applied to obtain the reconstructed output.

3.2 Experimental Results

In this section we review some of the experimental results for the polyphase transform-based MD image coder, as reported in [30]. It is aimed to demonstrate the performance of the polyphase transform-based coder and compare it with that of an alternative MD wavelet image coder based on multiple description scalar quantization (MDSQ).

For the two-description case, in the polyphase transform-based MD image coder of [30], the input image is first wavelet transformed and a polyphase transform is then applied to the wavelet coefficients to decompose them into two polyphase components. Two different types of polyphase transform are examined. One is the plain polyphase transform, which, for each row in each subband, simply groups all the even-numbered coefficients into one polyphase component and all the odd-numbered coefficients into the other component. The second, which can be viewed as a generalized polyphase transform in vector form, groups wavelet coefficients in different subbands corresponding to the same spatial location into a block structure similar to the zerotree structure [49]. One polyphase component then consists of all the even-numbered blocks and the other consists of all the odd-numbered blocks.

Finely quantized and coarsely quantized versions of the polyphase components then constitute the two descriptions, as explained in Section 3.1. The actual quantization and entropy coding of the different polyphase components is performed using the SPIHT wavelet image coder [45].

For the grayscale image Lena of size 512 × 512 pixels, Figure 3.3 shows[2] the PSNR (Peak Signal-to-Noise Ratio) values corresponding to the central and (average) side

---

[2] Figure 3.3 is based on Figure 5(a) from [30], after a swap between the curves referring to the plain and vector-form polyphase transforms, as implied by the accompanying text in [30].

distortions obtained for different amounts of redundancy and a total rate of 1 bpp (bits per pixel). Comparing the results for the two types of polyphase transforms, the vector-form polyphase transform yields slightly better results than the plain polyphase transform. This can be explained by the fact that the SPIHT coder is suited to the zerotree structure used by the vector-form polyphase transform and thus more efficient coding is obtained for this type of transform. A comparison with the MDSQ-based

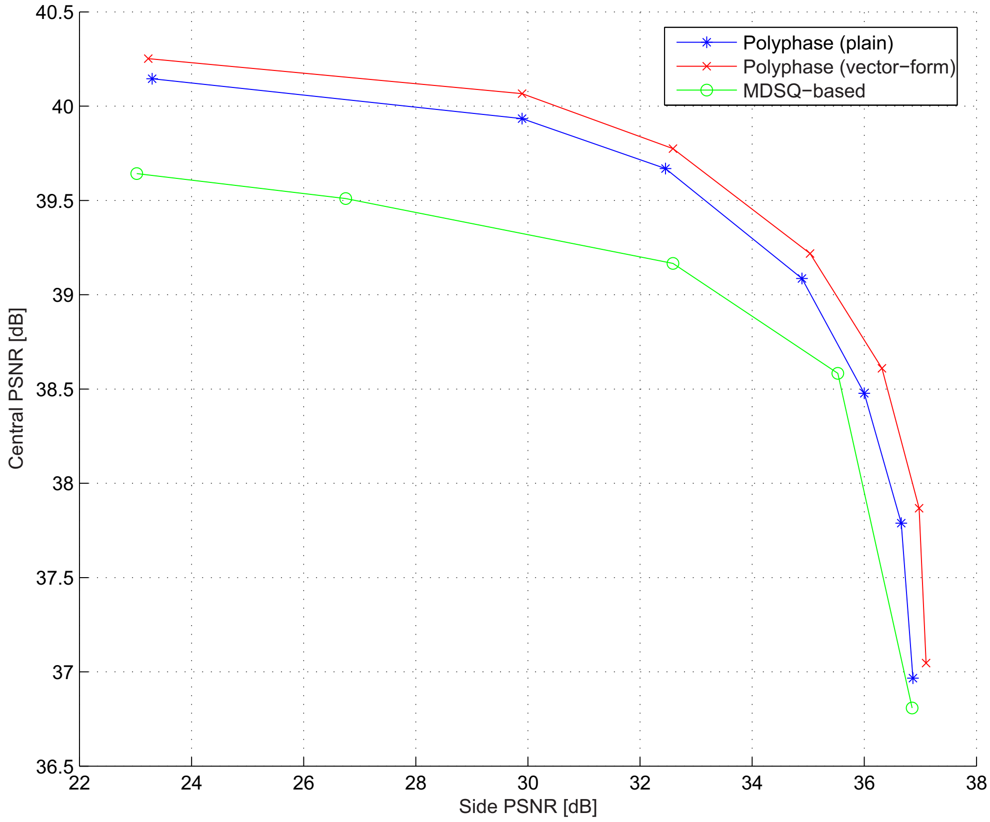

Figure 3.3: Central PSNR vs. side PSNR for image Lena; total rate of 1 bpp. The polyphase transform-based MD wavelet coder [30] (with two options for the polyphase transform) is compared with the MDSQ-based coder [46].

MD wavelet image coder of Servetto et al. [46] (results are obtained from [30]) is also shown in Figure 3.3 and demonstrates the effectiveness of the polyphase transform-based MD coding technique.

Chapter 4

Context-Based Multiple Description Wavelet Image Coding: Motivation and Proposed System Outline

This chapter provides the motivation for context-based MD coding in the wavelet domain, which enables the proposed coder to utilize some useful properties of the wavelet decomposition for meeting MD coding objectives, and also presents the outline of the proposed context-based MD wavelet image coder. The subsequent chapters will provide a more detailed description of the main building blocks of the proposed coder. These include the context-based classification procedure, which enables us to use a set of customized quantizers for improved coding efficiency, as well as the efficient parametric model-based quantization and bit allocation schemes.

4.1 Motivation

One of the first applications of wavelet theory was image coding. Early wavelet coders used coding techniques similar to those already used by subband coders, with the only possible difference being the choice of filters (wavelet filters were typically designed to satisfy certain smoothness constraints [14], while subband filters were designed to approximately satisfy the criteria of nonoverlapping frequency responses). In contrast, modern wavelet coders usually utilize the multiresolution nature of the wavelet decomposition in order to achieve improved coding efficiency.

Our proposed MD coder also operates in the wavelet domain and exploits the special properties of the wavelet decomposition for meeting MD coding objectives. This is accomplished by means of various coding procedures adapted for this task, such as context-based classification of subband coefficients, parametric model-based quantization and efficient bit allocation. This section provides the motivation for context-based MD coding in the wavelet domain by presenting a brief review of the wavelet transform, followed by a description of some useful statistical properties of wavelet coefficients. Later, we will describe in detail the various procedures that enable the proposed MD coder to utilize the useful properties of the wavelet decomposition for improved coding efficiency.

4.1.1 Wavelet Background

A typical one-dimensional (1-D) discrete wavelet transform (DWT) is shown in Figure 4.1. The one-dimensional input signal $x(n)$ is passed through a lowpass filter h and a highpass filter g, followed by downsampling by a factor of two, to constitute

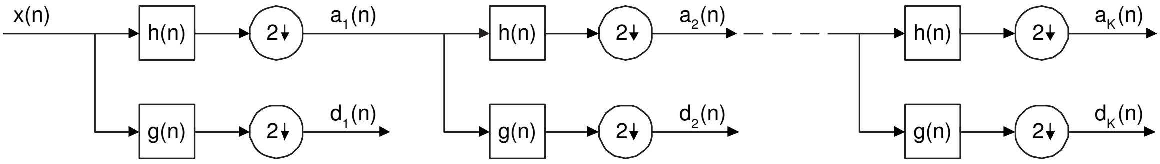

Figure 4.1: K-level one-dimensional discrete wavelet transform.

one level of the transform [57]. The outputs $a_1(n)$ and $d_1(n)$ are called the *approximation* and *detail* subbands of the signal, respectively. Multiple levels, or *scales*, of the wavelet transform are formed by repeating the filtering and decimation process, usually on the lowpass branch outputs only. To produce a K-level wavelet decomposition, as in Figure 4.1, the process is carried out for a total number of K times. The resulting approximation coefficients $a_K(n)$ and detail coefficients $d_i(n)$, $i \in \{1, \ldots, K\}$ are called wavelet coefficients.

The extension of a one-dimensional wavelet transform to a two-dimensional (2-D) wavelet transform for the processing of images can be performed using separable wavelet filters [27, 37]. Namely, a two-dimensional transform can be computed by applying a one-dimensional transform to all the rows of the input image, followed by a one-dimensional transform applied to all the resulting columns. This is illustrated in Figure 4.2, where $x(m, n)$ denotes the input image (m and n are the row and column indices, respectively) and h and g are the lowpass and highpass filters, respectively. Note that downsampling is performed in two stages, with one of them preceding the second filtering operation, to reduce the overall number of computations. The outputs $a_1(m, n)$, $d_1^H(m, n)$, $d_1^V(m, n)$ and $d_1^D(m, n)$ are called the *approximation*, *horizontal detail*, *vertical detail* and *diagonal detail* subbands of the image, respectively.

Multiple levels (scales) of the two-dimensional wavelet decomposition are formed

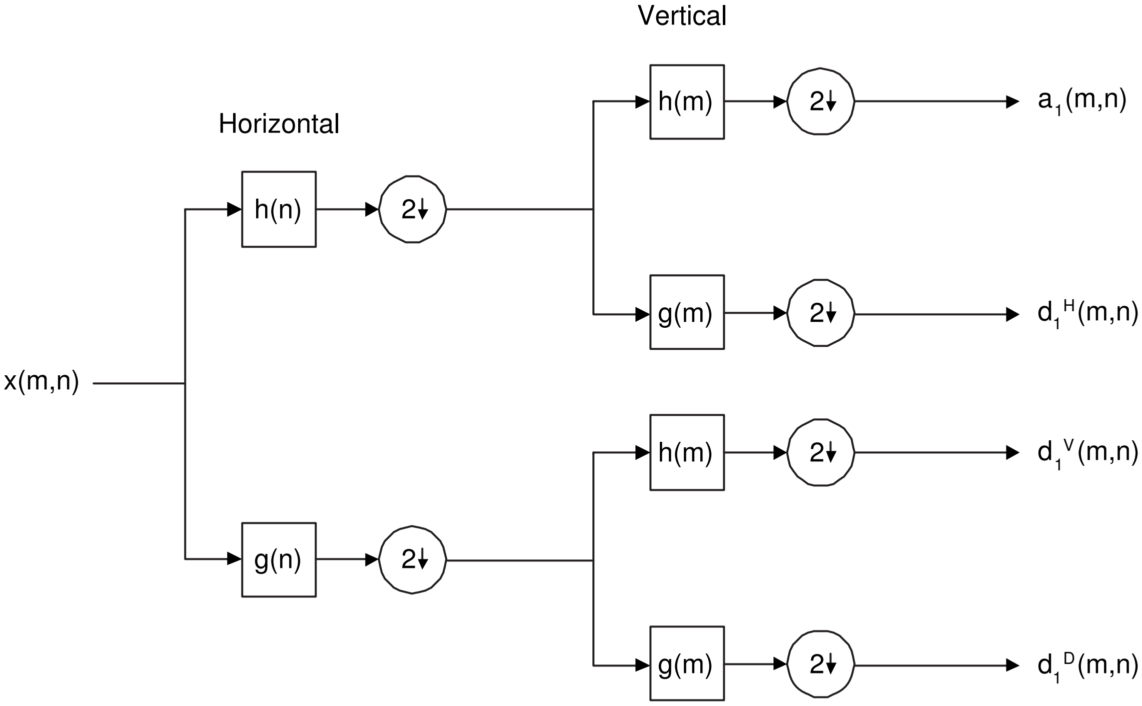

Figure 4.2: Two-dimensional discrete wavelet transform. Separable filters are first applied in the horizontal dimension and then in the vertical dimension to produce a one-level two-dimensional transform.

by repeating the process, typically on the approximation subband only. As an example, consider the original image of size 512 × 512 pixels shown in Figure 4.3. A three-level two-dimensional wavelet transform of the original image is shown in Figure 4.4. Note that the detail subbands (all subbands except the approximation subband a_3 in the upper left corner of the decomposition) have been enhanced to make their underlying structure more visible.

The octave-band frequency decomposition, obtained by recursively splitting the lowpass branch outputs, provides good spatial selectivity at higher frequencies, and good frequency selectivity at lower frequencies. This tradeoff is very useful for the analysis of natural images, as these often exhibit short, high-frequency events (e.g., edges), in addition to long, low-frequency events (e.g., typical scene's background).

4.1.2 Statistical Characterization of Wavelet Coefficients

For the purpose of image coding, one of the most attractive properties of the wavelet transform is its tendency to compact the energy of the input image into a relatively small number of wavelet coefficients [57]. This energy compaction property can also be shown to be equivalent to reducing the correlation amongst wavelet coefficients [58]. For natural images, as an example, much of the energy in the wavelet transform is concentrated in the approximation subband. Moreover, the energy in detail subbands is concentrated in a relatively small number of coefficients.

The first order statistics (distribution marginals) of detail[1] wavelet coefficients

---

[1] Approximation coefficients are usually modeled using a uniform distribution [9].

Figure 4.3: Original image used for demonstrating the two-dimensional wavelet transform.

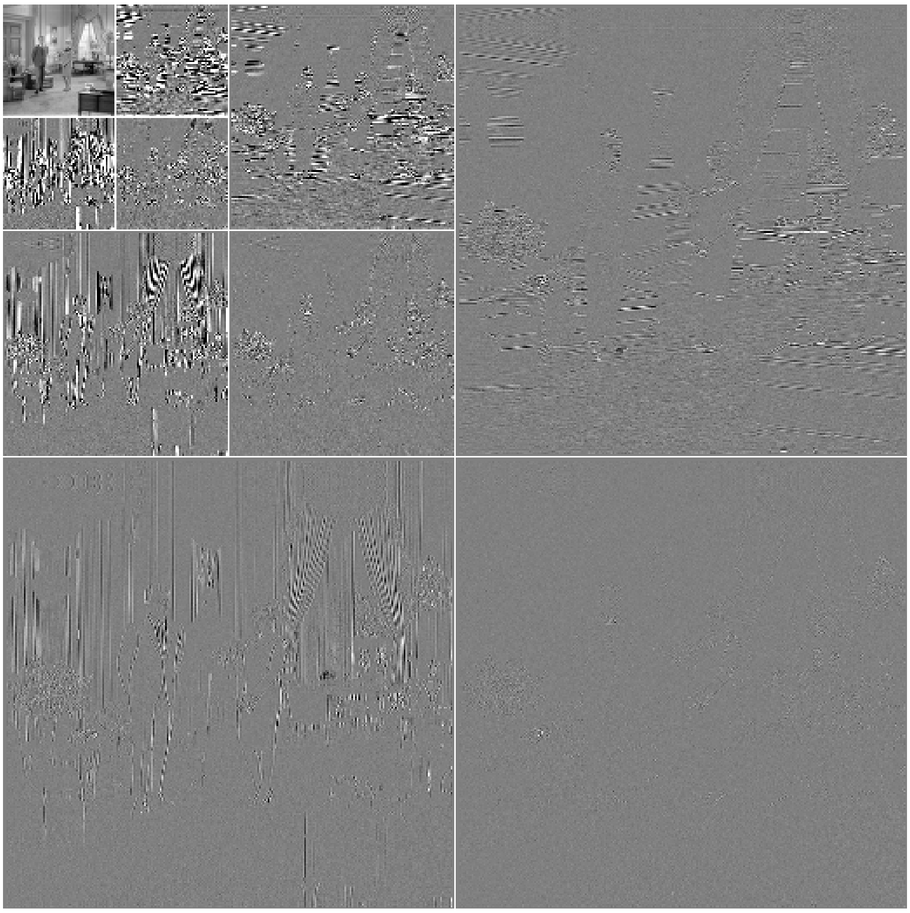

Figure 4.4: Example of a three-level two-dimensional wavelet transform of the original image shown in Figure 4.3 (the detail subbands have been enhanced to make their underlying structure more visible).

have been previously modeled [31, 33, 36, 37] using a two-parameter *generalized Gaussian* (also called *generalized Laplacian*) probability density function (pdf) of the form

$$f_{s,r}(x) = \frac{1}{N(s,r)}\, e^{-|x/s|^r} \tag{4.1}$$

where $N(s,r) = 2s\Gamma(1/r)/r$ and $\Gamma(a) = \int_0^\infty t^{a-1}e^{-t}\,dt$ is the Gamma function. The shape parameter r determines the exponential rate of decay and the distribution variance σ^2 is given by

$$\sigma^2 = \frac{s^2\Gamma\left(\frac{3}{r}\right)}{\Gamma\left(\frac{1}{r}\right)}. \tag{4.2}$$

The generalized Gaussian distribution (GGD) is a family of zero-mean symmetric distributions that include the Laplacian distribution (for $r = 1$) and the Gaussian distribution (for $r = 2$). In addition, the density function of (4.1) converges pointwise to a uniform density on $(-s, s)$ as $r \to \infty$.

For various types of images, excellent fits of the generalized Gaussian probability density function of (4.1) have been demonstrated in [9]. For the images in the rich sample set used in [9], the shape parameter r in (4.1) corresponding to the best fit is typically in the range $[0.5, 1]$, which implies that the first order statistics of detail wavelet coefficients can be reasonably modeled using a Laplacian distribution.

An example of the generalized Gaussian probability density function of (4.1) with $r = 1$ (i.e., Laplacian density) and $s = 1.5$ is shown in Figure 4.5. As shown, the density function is symmetric, peaked at zero and has relatively long tails. The long tails indicate the existence of a few subband coefficients with large magnitudes that concentrate a large portion of the energy in the subband. The high peak around zero indicates that most coefficients in the subband have magnitudes which are close to zero and thus have low energy.

Although the coefficients produced by the wavelet transform are approximately

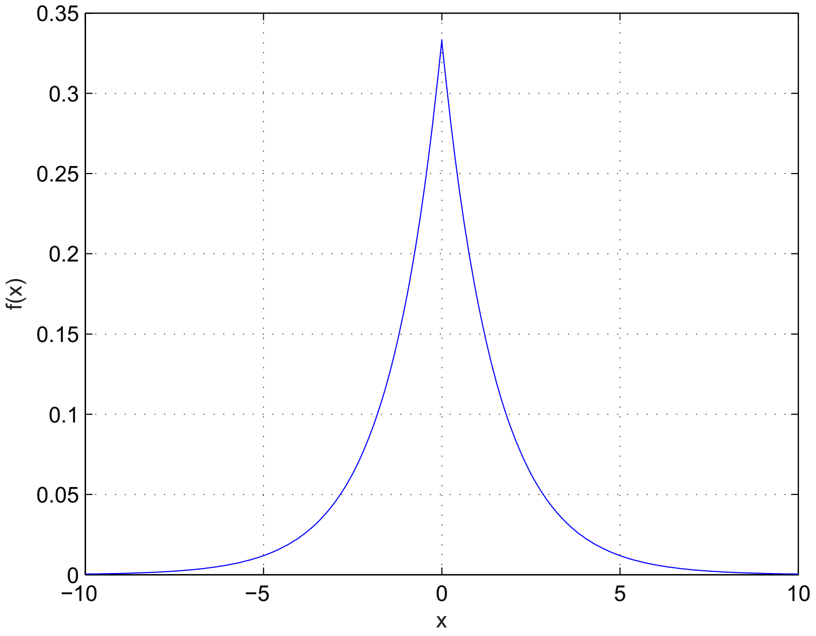

Figure 4.5: The generalized Gaussian probability density function of (4.1) with $r = 1$ and $s = 1.5$.

decorrelated [49], visual inspection of the wavelet decomposition (see Figure 4.4, for example) reveals that wavelet coefficients are *not* statistically independent. In fact, large magnitude coefficients are usually gathered at neighboring spatial locations, as well as at the same relative spatial locations of subbands at adjacent scales and orientations [9]. Experiments show that the correlation coefficient between the square of a wavelet coefficient and the square of the coefficient corresponding to the same spatial location at the next coarser scale (also called *parent*; see Figure 4.6 for illustration) tends to be between 0.2 and 0.6 with a concentration around 0.35 [49].

Numerous image compression schemes implicitly utilize the spatial and scale-to-scale dependencies present in the wavelet decomposition. The embedded zerotree wavelet (EZW) coder of Shapiro [49], whose introduction in 1993 marked the beginning of the era of modern lossy wavelet coding, exploits the likelihood of a coefficient to have small magnitude if the coefficients corresponding to the same spatial location at coarser scales have small magnitudes. In this technique, entire trees of zeros are encoded with a single symbol, thus capturing a portion of the conditional distribution of a coefficient given the coefficients corresponding to the same spatial location at coarser scales. Said and Pearlman [45] generalized Shapiro's zerotrees and introduced set partitioning techniques to effectively encode these generalized trees of zeros. The resulting coding scheme, known as SPIHT (Set Partitioning In Hierarchical Trees), outperforms EZW at even lower levels of complexity. LoPresto et al. [36] use a generalized Gaussian model for the wavelet coefficients and estimate the model parameters from local neighborhoods.

In order to explicitly examine and utilize the statistical relationship between the magnitudes of wavelet coefficients, and due to the difficulty of characterizing the full

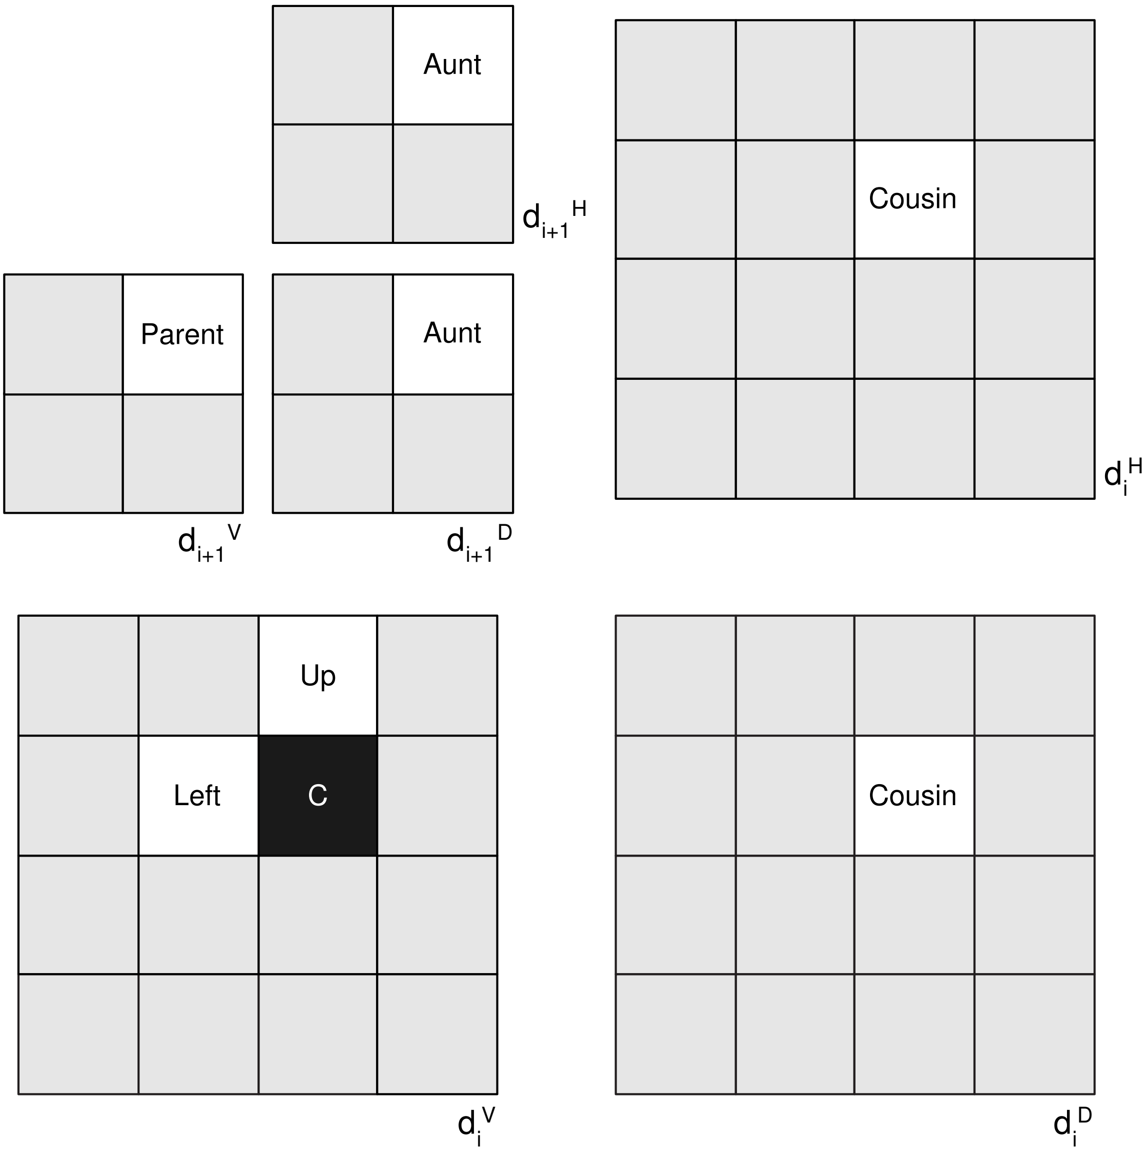

Figure 4.6: Potential conditioning neighbors (also at adjacent scales and orientations) for a given wavelet coefficient C.

multidimensional density, Buccigrossi and Simoncelli [9] proposed to use a linear predictor for the magnitude of a coefficient, based on the magnitudes of its neighbors (also at adjacent scales and orientations), as shown in Figure 4.6. In order to determine which coefficients to include in the conditioning set, rather than exhaustively explore all possible subsets of candidate neighbors, a greedy algorithm was applied. Specifically, the conditioning set was constructed incrementally: at each step, the remaining neighbor whose inclusion maximized the mutual information between the coefficent and its predictor was joined to the conditioning set. Using this analysis based on three sample images, and imposing causality (assuming a standard scanline ordering of the coefficients), it was found that the local neighbors within the subband (labeled "Left" and "Up" in Figure 4.6) contributed the most to the mutual information. The parent and the cousins (in that order) followed with smaller contributions, while joining any of the remaining potential neighbors to the conditioning set offered only a marginal increase in the resulting mutual information.

As described above, the utilization of the statistical dependencies present in the wavelet decomposition enables modern wavelet coders to achieve improved coding efficiency. The proposed context-based MD wavelet image coder, introduced in the following section, is designed to utilize the fundamental properties of the wavelet decomposition in the framework of multiple description coding. In the following sections, we will show how these useful properties of the wavelet decomposition can be exploited for meeting MD coding objectives.

4.2 Proposed System Outline

Chapter 3 described a useful technique for MD coding via the utilization of a polyphase transform and selective quantization [30]. Section 4.1 introduced the wavelet decomposition and its fundamental statistical properties. Our proposed context-based MD wavelet image coding system is aimed to provide improved coding efficiency by means of effective utilization of the statistical properties of the wavelet decomposition, in the general framework of polyphase transform-based MD coding. We focus on the two-description case, mainly due to its tractability, as is common in the MD coding literature.

For image coding (i.e., correlated input data), the polyphase transform-based technique of [30] first employs a decorrelating transform, such as wavelet transform. Referring to Figure 3.1 in Section 3.1 (which does not show the decorrelating transform itself), the resulting decorrelated transform coefficients, denoted by X, are decomposed into two sub-sources Y_1 and Y_2 via a polyphase transform. Each of the two polyphase components Y_1 and Y_2 is then quantized independently by Q_1 to form the primary part of information for its corresponding description. 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—a coarsely quantized version produced by the quantizer Q_2. At the receiver, if both descriptions arrive, the finely quantized versions of the two polyphase components (produced by Q_1 during encoding) are used for reconstruction. If one of the descriptions is lost, one finely quantized and one coarsely quantized polyphase component (the primary and redundant part of the received description, respectively) are used for reconstruction.

As also pointed out in [30], the structural similarities existing between the polyphase components Y_1 and Y_2 can potentially improve the quantization efficiency, if utilized effectively. In the case that only description i is received ($i \in \{1, 2\}$), side decoder i uses the finely quantized version of the polyphase component Y_i (produced by Q_1) and the coarsely quantized version of the other polyphase component (produced by Q_2) for reconstruction. Since this coarsely quantized component needs to be decoded only in that case and the finely quantized component is available to both the encoder and the relevant side decoder, the finely quantized component can also serve to provide contextual information that will improve the quantization efficiency of Q_2. This is shown in Figure 4.7.

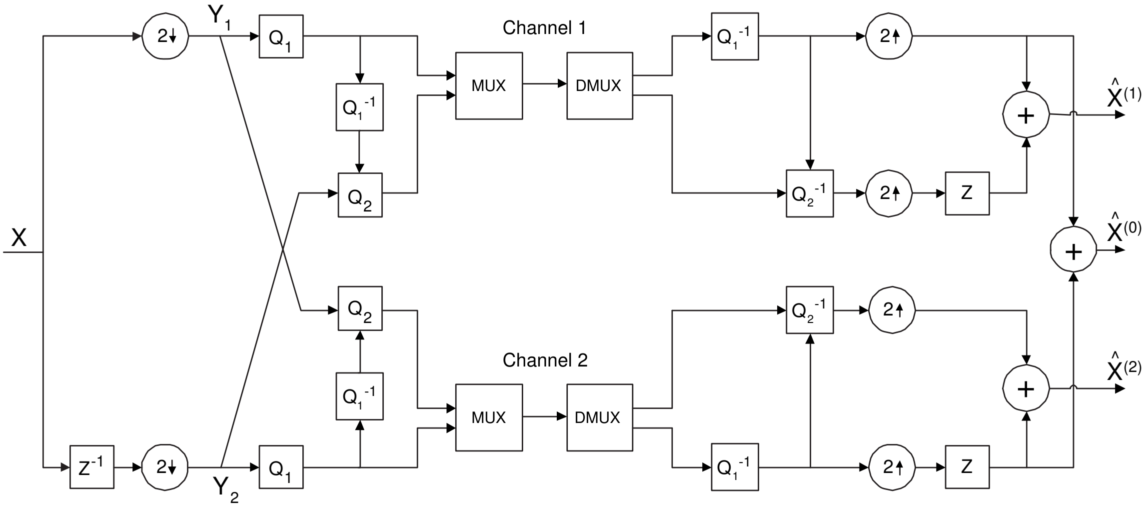

Figure 4.7: MD coding system based on polyphase transform that utilizes context-based quantization in coding the redundant information.

Although the polyphase components Y_1 and Y_2 in Figure 4.7 are both derived from the same input X (denoting the transform coefficients), it is not obvious *how* one of them can help to improve the quantization of the other. First, since the input X is typically the output of a decorrelating transform, the coefficients of the polyphase components Y_1 and Y_2 are approximately decorrelated. Thus, quantization schemes that rely on correlation, such as linear predictive quantization, are not expected to yield an improvement in the quantization of a polyphase component by using the (quantized) complementary polyphase component. Second, the structural similarities existing between the polyphase components Y_1 and Y_2 cannot be easily characterized and their utilization to improve the quantization efficiency of the quantizer Q_2 is therefore a complicated task. Moreover, since the polyphase decomposition generally breaks the natural structure of the transform decomposition X (e.g., wavelet decomposition), common statistical characterizations of the transform coefficients and their dependencies have to be adapted accordingly.

To overcome these difficulties, the proposed MD coder realizes an adaptive model-based quantization scheme that is aimed to utilize the special statistical properties of the wavelet decomposition. Adaptivity is achieved by means of context-based classification that also produces a controlled amount of side information, where the contexts are formed from finely quantized coefficients of the complementary polyphase component. The classification, performed on a coefficient-by-coefficient basis (as opposed to block-based classification [33]), enables the encoder to use a set of customized quantizers for improved coding efficiency, along with efficient parametric model-based quantization and bit allocation schemes.

Figure 4.8 shows a simplified block diagram of the proposed encoder. 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

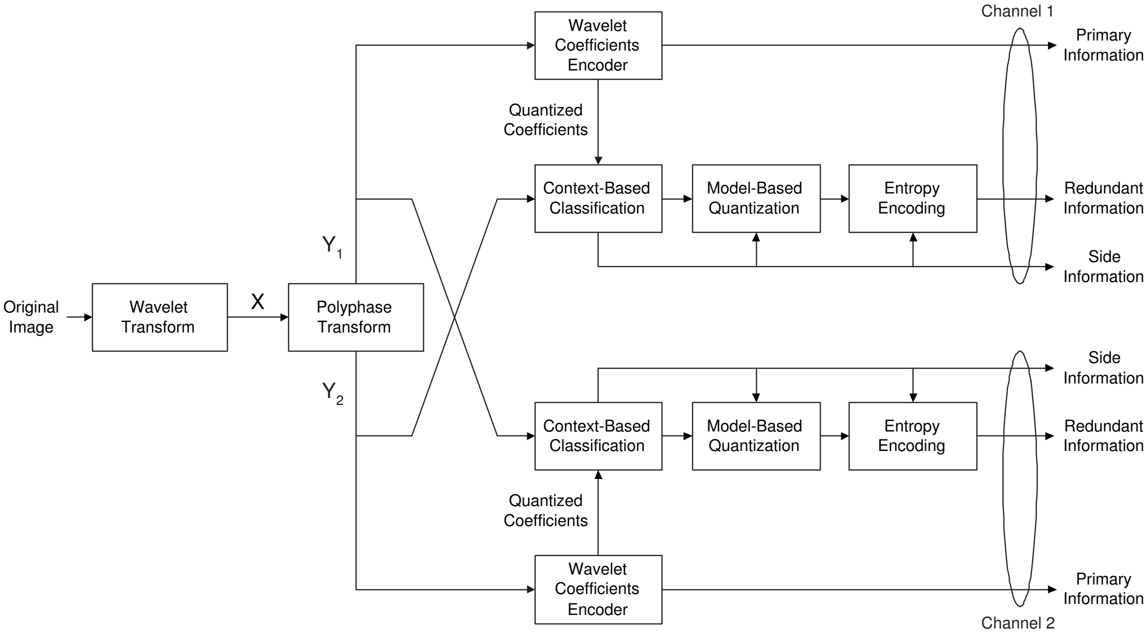

Figure 4.8: Simplified block diagram of the proposed MD encoder.

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 [45] 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 (see Section 4.1.2), 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 4.9. 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. To this end, entropy decoding is performed, followed by context-based classification and dequantization. The contexts for classification are formed from quantized coefficients of the polyphase component Y_i, and the auxiliary

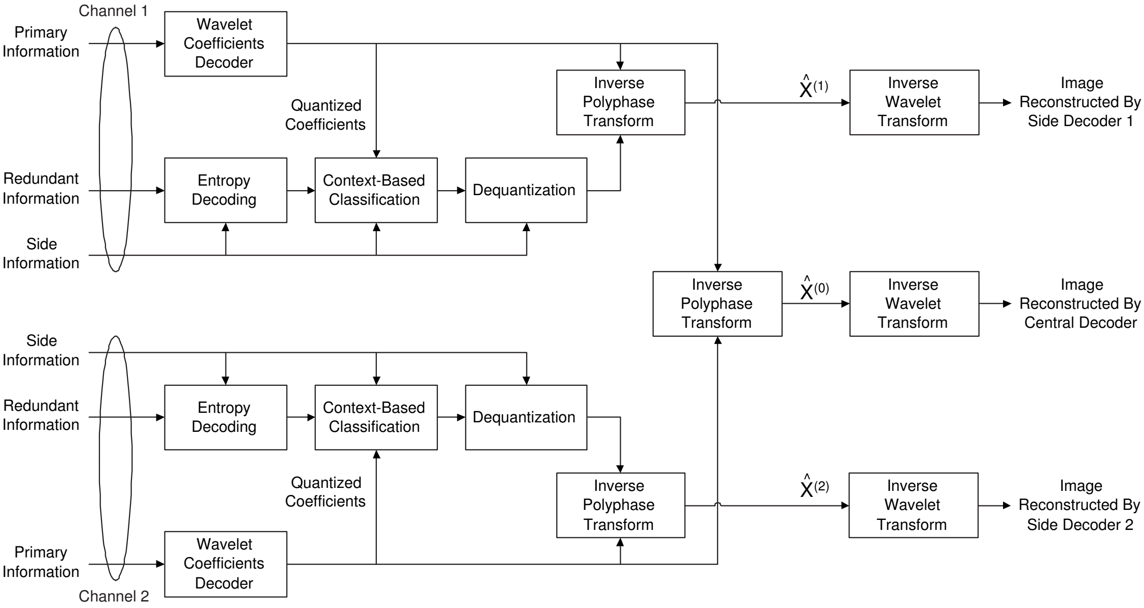

Figure 4.9: Simplified block diagram of the proposed MD decoder.

side information provides the required classification thresholds and source statistics of each class. Side decoder i then 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 next chapters provide a detailed description of the main building blocks of the proposed coder, which include the classification, quantization and bit allocation schemes. A careful treatment of some important implementation considerations, such as boundary extension and optimal bit allocation for biorthogonal wavelet filters, is also provided.

Chapter 5

Context-Based Multiple Description Wavelet Image Coding: Context-Based Classification

5.1 Introduction

Data classification is an effective technique used in various quantization applications. For an input source modeled as a mixture of sources with different distributions, classification 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, since the quantization can be adapted to the data based on the classification information [66].

As a simple demonstration of the potential benefit of classification in general subband image coding applications, consider Figure 4.4 (in Section 4.1.2), which shows the wavelet transform of the natural image shown in Figure 4.3. As can be seen, high-activity regions (corresponding to edges in the image, for example) tend to cluster at neighboring spatial locations, and are separated from surrounding low-activity regions. Identifying those high- and low-activity regions via classification would enable the utilization of separate quantizers, each suited to a different region.

A comprehensive study by Joshi et al. [33] demonstrates substantial coding gains due to classification in subband image coding applications. The considered classification techniques are all based on block-wise classification, where one of a finite number of classes is assigned to each block of coefficients within a subband. In the framework of [33], the class assigned to each block is explicitly stated as side information and transmitted to the decoder (this is known as *forward classification*). In order to limit the resulting overhead, the size of blocks used for classification must be large enough, and the number of potential classes must be kept small. Even then, however, the overhead rate may become excessive for low target bit rates.

Our 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. In order to perform well across a wide range of total bit rates (corresponding to the desired quality of reconstruction) and redundancy levels (derived from the probability of channel failure), the proposed coder avoids the penalty of forward classification. To this end, the classification is based on quantized coefficients of the complementary polyphase component, and utilizes the special statistical properties of the wavelet decomposition. Since the quantized coefficients of the complementary polyphase component are also available at the decoder, no transmission of side information is required. Nevertheless, a controlled amount of side information is still produced 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 distribution. It allows the classification process to adapt itself to the data, and enables the proposed coder to avoid the need for an explicit characterization of the statistical dependencies between neighboring wavelet coefficients. Compared to block-based classification techniques, our context-based classification procedure, inspired by that of [66], enables finer grain classification, down to the coefficient level (rather than a block level).

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 described in Section 4.1.2, 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 5.1 shows the context of a given wavelet coefficient $X_{i,j}$. For the ease of illustration, the primary and redundant polyphase components are interleaved in the figure. As shown, the classification of

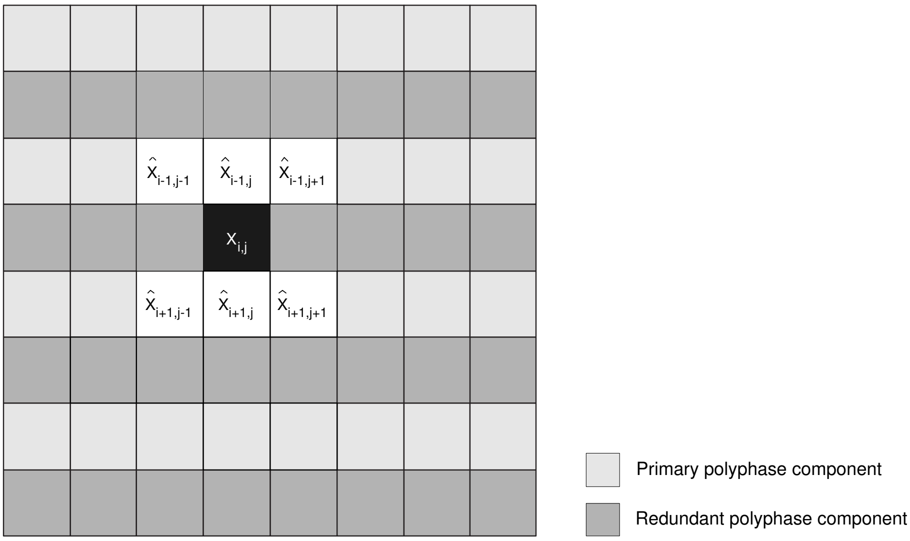

Figure 5.1: Context of a given wavelet coefficient $X_{i,j}$. The coefficient belongs to the redundant polyphase component, and is shown in black. The quantized coefficients that form the context belong to the primary polyphase component, and are shown in white.

the wavelet coefficient $X_{i,j}$ is based on the following context:

$$\mathcal{C}_{i,j} = \left\{ \hat{X}_{i-1,j-1},\ \hat{X}_{i-1,j},\ \hat{X}_{i-1,j+1},\ \hat{X}_{i+1,j-1},\ \hat{X}_{i+1,j},\ \hat{X}_{i+1,j+1} \right\} \tag{5.1}$$

where $\hat{X}$ denotes the value of X after quantization by the wavelet coefficients encoder (see Figure 4.8), and the coordinate system refers to the whole subband (i.e, prior to the polyphase transform).

5.2 Classification Rule

We now turn to describe the classification rule, which assigns one of a finite number of classes to a coefficient $X_{i,j}$, given its context $\mathcal{C}_{i,j}$. 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 $\mathcal{C}_{i,j}$. Define the *activity* $A_{i,j}$ of the coefficient $X_{i,j}$, predicted from its context $\mathcal{C}_{i,j}$, as

$$A_{i,j} = a_1|\hat{X}_{i-1,j-1}| + a_2|\hat{X}_{i-1,j}| + a_3|\hat{X}_{i-1,j+1}| + a_4|\hat{X}_{i+1,j-1}| + a_5|\hat{X}_{i+1,j}| + a_6|\hat{X}_{i+1,j+1}| \tag{5.2}$$

where the fixed relative weights $\{a_k\}$ satisfy $\sum_k a_k = 1$. One reasonable choice is to set the weights $\{a_k\}$ in accordance to the distances of the corresponding neighboring coefficients from the current coefficient $X_{i,j}$, such that closer coefficients will have greater contribution to the activity $A_{i,j}$ compared to more distant ones. In the proposed coder we thus set the weights $\{a_k\}$ to be inversely proportional to the Euclidean distances of the corresponding coefficients in $\mathcal{C}_{i,j}$ from the coefficient $X_{i,j}$.

In order to classify a coefficient $X_{i,j}$, its context $\mathcal{C}_{i,j}$ 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

$$T_0 = 0 < T_1 < T_2 < \cdots < T_{N-1} < \infty \,, \tag{5.3}$$

the coefficient $X_{i,j}$ is assigned to the class C_k if

$$A_{i,j} \in (T_{k-1},\; T_k] \tag{5.4}$$

where $T_{-1} = -\infty$ and $T_N = \infty$. This is shown in Figure 5.2.

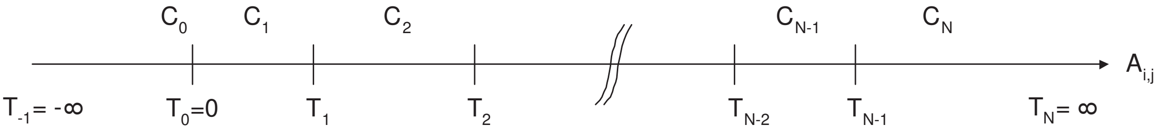

Figure 5.2: Classification of the coefficient $X_{i,j}$ based on its activity $A_{i,j}$.

Note that since the activity is a nonnegative quantity, a coefficient $X_{i,j}$ is assigned to the class C_0 if and only if $A_{i,j} = 0$. In other words, a coefficient is assigned to C_0 if and only if all the coefficients in its context are quantized to zero. This special case is more common at finer scales of the wavelet decomposition, and becomes more dominant at lower bit rates (with respect to the primary polyphase component, from which the context is extracted). In order to reliably model these coefficients having an all-zero context, a special class, namely C_0, is assigned only to them. This class is also called the *zero-context* class.

5.3 Classification Thresholds Design

This section describes in detail the classification thresholds design—the procedure 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. This procedure is performed independently for each of the subbands. In order to enable the decoder to perform the same classification as the encoder, the classification thresholds are also transmitted to the decoder as side information (in addition to the parametric source statistics of each class). Note that if we had also included in the context $\mathcal{C}_{i,j}$ already quantized coefficients of the redundant polyphase component, an interdependency between the classification and the choice of quantizers for encoding the redundant polyphase component would have been created. Namely, in such a case, the choice of quantizers for encoding the redundant polyphase component would have influenced the classification and vice versa. We therefore refrain from including in $\mathcal{C}_{i,j}$ already quantized coefficients of the redundant polyphase component, and so $\mathcal{C}_{i,j}$ is formed solely from quantized coefficients of the primary polyphase component. Thus, the process is naturally performed in a serial manner: the classification thresholds are determined first, then actual classification is performed, and only then the selection of quantizers is considered.

As will be detailed in Chapter 6, our class-adaptive quantization 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. Since we employ a predictive (DPCM-like) quantization scheme to encode the coefficients in the approximation subband (as will be described in Section 7.2.1), we assume a Laplacian distribution model (for the prediction errors) in the approximation subband as well.

The classification thresholds design involves estimation of the Laplacian distribution's parameter for each candidate class. We now derive the maximum likelihood estimator (MLE) [34, Ch. 7] for this case. Let $x_1, x_2, \ldots, x_n$ be n independent, identically distributed (i.i.d.) observations from a Laplacian source with an unknown parameter $\lambda\,(> 0)$. That is, independent samples drawn according to the probability density function (pdf)

$$f_\lambda(x) = \frac{\lambda}{2} e^{-\lambda |x|}\,. \tag{5.5}$$

The likelihood function of λ is given by

$$\mathcal{L}(\lambda) = \prod_{i=1}^{n} f_\lambda(x_i)\,. \tag{5.6}$$

Taking the natural logarithm to obtain the log-likelihood function yields

$$\mathcal{L}^*(\lambda) = \sum_{i=1}^{n} \log f_\lambda(x_i) = \sum_{i=1}^{n} \log\left(\frac{\lambda}{2} e^{-\lambda |x_i|}\right) = n \log\left(\frac{\lambda}{2}\right) - \lambda \sum_{i=1}^{n} |x_i|\,. \tag{5.7}$$

Maximization of (5.7) with respect to λ then yields the maximum likelihood estimator (MLE) of λ:

$$\hat{\lambda} = \frac{n}{\sum_{i=1}^{n} |x_i|} = \frac{1}{\frac{1}{n}\sum_{i=1}^{n} |x_i|}\,. \tag{5.8}$$

We now describe how to determine the classification thresholds $T_1, T_2, \ldots, T_{N-1}$ for a given subband. This design procedure is repeated independently for each of the subbands, and involves only the coefficients in the relevant subband. For brevity, we refer to the subband coefficients that belong to the redundant polyphase component as *redundant coefficients*, and to those that belong to the primary polyphase component as *primary coefficients*. Given the unquantized redundant coefficients (to be classified eventually using the obtained thresholds) and the quantized primary coefficients (that form the contexts for classification), we wish to determine the thresholds

$T_1, T_2, \ldots, T_{N-1}$, assuming a total of $N+1$ classes $C_0, C_1, \ldots, C_N$. Note that since $T_0 = 0$, there are only $N-1$ classification thresholds to determine.

As noted in Section 5.2, a redundant coefficient $X_{i,j}$ is assigned to the class C_0—the zero-context class—if and only if all the coefficients in its context are quantized to zero (i.e., $A_{i,j} = 0$). Thus, as a preliminary step, we first identify all the redundant coefficients having an activity of zero. These coefficients (and only them) will be assigned to the class C_0, and are therefore not considered in the classification thresholds design. The parameter λ of the Laplacian distribution for the class C_0 is also estimated at this stage using (5.8), and will be transmitted to the decoder. Note that we are now left with N "ordinary" potential classes (namely, $C_1, \ldots, C_N$) and need to determine $N-1$ classification thresholds (namely, $T_1, \ldots, T_{N-1}$).

The remainder of the design procedure follows that of [66] and is summarized below. We start by considering an oversized initial set of $N_0 - 1$ candidate thresholds (where $N_0 > N$)

$$\{T_k,\ 1 \le k \le N_0 - 1 \ :\ 0 < T_1 < \cdots < T_{N_0-1} < \infty\}\,. \tag{5.9}$$

This thresholds set defines N_0 initial classes $\{C_1, \ldots, C_{N_0}\}$ by the following classification rule: the redundant coefficient $X_{i,j}$ is assigned to the class C_k if

$$A_{i,j} \in (T_{k-1},\ T_k] \tag{5.10}$$

for $k = 1, \ldots, N_0$, and with $T_0 = 0$ and $T_{N_0} = \infty$. After the redundant coefficients are collected into the N_0 initial classes according to the classification rule (5.10), the maximum likelihood estimator (MLE) of the Laplacian parameter for each class is computed using (5.8). An iterative greedy algorithm is then applied, which removes one classification threshold (and one class) at each iteration, until the number of thresholds is reduced to $N-1$. At each iteration, the algorithm merges the pair of classes with the smallest classification gain[1] among all pairs of adjacent classes (i.e., classes corresponding to consecutive thresholds), where the classification gain is estimated by comparing the variances of the two adjacent classes with that of the union of the two classes (i.e., the merged class). The merging process is illustrated in Figure 5.3. The final $N-1$ classification thresholds are the output of the algorithm.

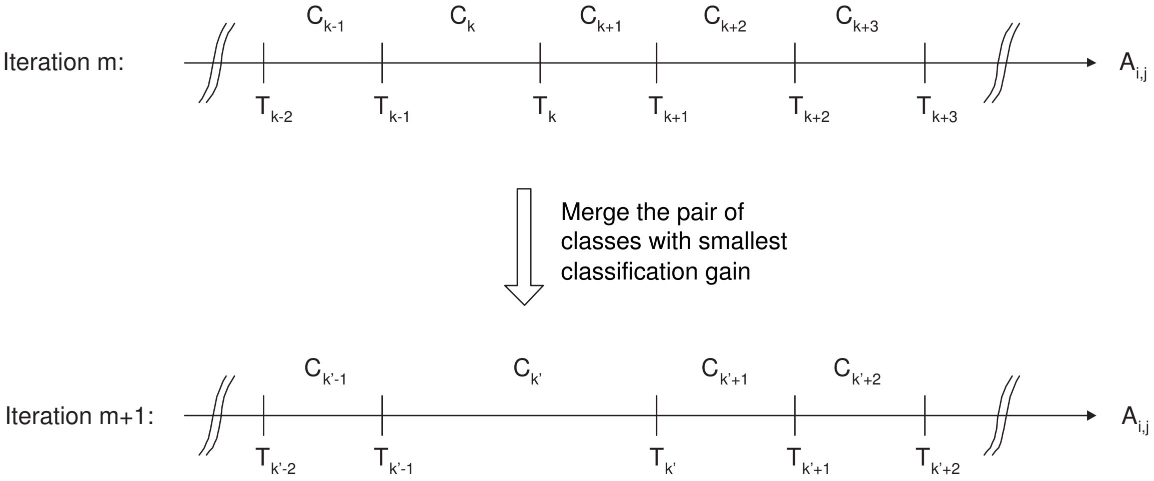

Figure 5.3: Merging process performed by the classification thresholds design algorithm. At each iteration, the algorithm merges the pair of classes with the smallest classification gain among all pairs of adjacent classes.

We now give a more detailed description of the various stages of the algorithm. We begin with the selection of the initial number of classes N_0, which also determines the size of the initial thresholds set ($N_0 - 1$). Since the final $N-1$ thresholds will be selected from the initial thresholds set, N_0 should be large enough so that the final thresholds can lie at almost any point within the activity range (i.e., between the minimum and maximum activity levels). Nevertheless, N_0 should not be too large,

---

[1] The reason for using the term "classification gain" in this context will be explained later in this section.

in order to enable each of the initial classes to contain enough coefficients, which will ensure a reliable estimation of the Laplacian parameter for each class. Regarding the selection of the initial thresholds $T_1, \ldots, T_{N_0-1}$ themselves, it is desirable that each of the resulting classes will contain approximately the same number of coefficients, in order to ensure equally reliable estimates for the different classes. For that reason, the trivial choice of uniformly spaced thresholds is not appropriate in general. A better choice is to select $N_0 - 1$ nonuniform thresholds, such that each of the initial classes will contain approximately the same number of coefficients.

As explained earlier, the greedy algorithm removes one classification threshold (and one class) at each iteration, until the number of thresholds is reduced to $N-1$. At each iteration, the algorithm merges the pair of classes with the smallest classification gain among all pairs of adjacent classes. Denote the number of coefficients in two adjacent classes C_k and C_{k+1} by n_k and n_{k+1}, respectively. Also, let σ_k^2 and σ_{k+1}^2 be the respective variances of the two classes, and let $\sigma_{k'}^2$ be the variance of the union of the two classes. As in [32], the classification gain is then estimated by

$$G = \frac{\sigma_{k'}^2}{\sigma_k^{2r_k} \sigma_{k+1}^{2r_{k+1}}} \tag{5.11}$$

where $r_k = n_k/(n_k + n_{k+1})$ and $r_{k+1} = 1 - r_k = n_{k+1}/(n_k + n_{k+1})$. We also note that under certain simplifying assumptions stated in [32], the potential gain in coding performance due to classification (rather than using one unified class) can be approximated by a ratio similar to (5.11), which is the reason for referring to G as the "classification gain" in this context.

Assuming that the Laplacian model is still valid for the union of the classes C_k and C_{k+1}, the maximum likelihood estimator (MLE) of the Laplacian parameter for the merged class can be computed as

$$\hat{\lambda}_{k'} = \frac{n_k + n_{k+1}}{n_k/\hat{\lambda}_k + n_{k+1}/\hat{\lambda}_{k+1}} \tag{5.12}$$

where $\hat{\lambda}_k$ and $\hat{\lambda}_{k+1}$ are the parameter estimates for C_k and C_{k+1}, respectively. This stems directly from (5.8). Under the Laplacian assumption, the variance is equal to $2/\lambda^2$ and thus the classification gain can be estimated by

$$G = G(\hat{\lambda}_k, \hat{\lambda}_{k+1}; n_k, n_{k+1}) = \frac{\hat{\lambda}_k^{2r_k} \hat{\lambda}_{k+1}^{2r_{k+1}}}{\hat{\lambda}_{k'}^2} \,. \tag{5.13}$$

As proved in Appendix A, since the classification gain G of (5.13) is virtually the ratio of a weighted arithmetic mean to a weighted geometric mean, it follows that $G \geq 1$, with equality if and only if $\hat{\lambda}_k = \hat{\lambda}_{k+1}$.

It should be noted that the utilization of the Laplacian model offers a considerable reduction in the complexity of the algorithm. Once the Laplacian parameter estimates for the initial classes are computed, there is no need to use the direct MLE computation of (5.8) at subsequent iterations. Instead, new parameter estimates can be obtained using (5.12), without affecting the other classes not involved in the merging, and the classification gains can be computed directly using (5.13).

The entire classification thresholds design procedure (for a given subband) under the Laplacian model assumption is summarized as follows.

Algorithm 5.1 (Classification Thresholds Design) *Given the unquantized redundant coefficients, the quantized primary coefficients, the desired total number of classes $N+1$, and $N_0 - 1$ initial thresholds (where $N_0 > N$) corresponding to the initial thresholds set (5.9):*

1. *Classify all the redundant coefficients with $A_{i,j} = 0$ to C_0 (the zero-context class). Compute the corresponding Laplacian parameter estimate $\hat{\lambda}_0$ using (5.8). These coefficients are not considered further by the algorithm.*

2. *Classify the remaining redundant coefficients according to (5.10). Compute the Laplacian parameter estimates* $\hat{\lambda}_1, \ldots, \hat{\lambda}_{N_0}$ *using (5.8). Set* $K = N_0$.

3. *Find* k^* *such that*

$$k^* = \arg \min_{1 \leq k < K} G(\hat{\lambda}_k, \hat{\lambda}_{k+1}; n_k, n_{k+1})$$

   where $G(\cdot)$ *is given by (5.13) and computed also using (5.12).*

4. *Merge the classes* C_{k^*} *and* C_{k^*+1} *and update accordingly:*
   - *For* $k = k^*$*:* $C_k = C_{k^*} \bigcup C_{k^*+1}$, $n_k = n_{k^*} + n_{k^*+1}$, *update* $\hat{\lambda}_k$ *according to (5.12).*
   - *For* $k > k^*$*:* $C_k = C_{k+1}$, $T_k = T_{k+1}$, $n_k = n_{k+1}$, $\hat{\lambda}_k = \hat{\lambda}_{k+1}$.

5. *Set* $K = K - 1$. *If* $K > N$, *go to step 3.*

6. *Finish: return the classification thresholds* $T_1, \ldots, T_{N-1}$ *and the Laplacian parameter estimates* $\hat{\lambda}_0, \hat{\lambda}_1, \ldots, \hat{\lambda}_N$.

The procedure is performed separately for each of the descriptions, and is repeated 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 in the next chapter.

Chapter 6

Context-Based Multiple Description Wavelet Image Coding: Quantization and Optimal Bit Allocation

6.1 Parametric Model-Based Uniform Threshold Quantization

As noted in Section 4.1.2, detail wavelet coefficients can be effectively modeled using a Laplacian distribution. Naturally, the more general option of using the generalized Gaussian distribution (GGD) to model the detail coefficients, will always provide a fit to the empirical distribution which is at least as good as that of the Laplacian (assuming that the model parameters can be accurately estimated). Nevertheless, preferring the Laplacian model offers several advantages: First, the Laplacian model makes the design of the classification thresholds, the design of the quantizers, and the bit allocation procedure, more tractable and computationally efficient. Additionally, as indicated in [6], the Laplacian model offers coding performance (in a rate versus distortion sense) which is almost the same as that offered by the GGD model. This slight difference in performance can be well compensated by the smaller modeling cost of the Laplacian model [66], in addition to the significant reduction in the complexity of the system.

The context-based classification enables the proposed coder to utilize a set of quantizers, each customized to an individual class. As demonstrated in [16], the optimum entropy-constrained scalar quantizer (ECSQ) for the Laplacian distribution (and more generally, the GGD), 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 (e.g., [54, 66]). 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 [25]. We focus on UTQs with odd number of levels ("mid-tread" UTQs) in order to be able to obtain low bit rates (i.e., low redundancy), if desired; as noted in [5] and proved in [16], 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. Figure 6.1 shows a UTQ with an odd number of levels.

A general entropy-constrained design algorithm of UTQ is given in [16]. For the Laplacian (and exponential) distribution, a non-iterative entropy-constrained design

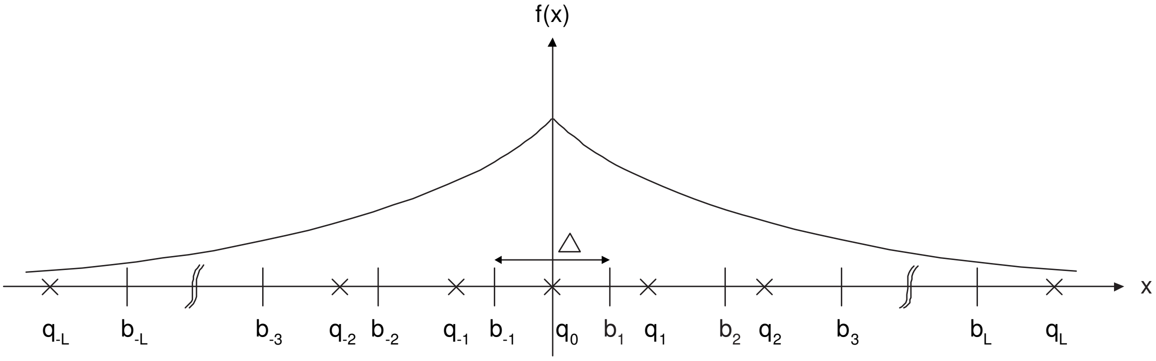

Figure 6.1: Uniform threshold quantizer (UTQ) with step size Δ and an odd number of levels $N = 2L + 1$. The bin boundaries are denoted by $\{b_j\}$ and the reconstruction levels by $\{q_j\}$. The Laplacian pdf is also shown for illustration.

algorithm of UTQ is provided by Sullivan [53]. 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.

We now derive the optimal uniform threshold quantizer (UTQ) Q with step size Δ for the Laplacian distribution, where the optimization is for minimum MSE distortion. Let X be a Laplacian random variable with parameter $\lambda\,(> 0)$, i.e., with pdf

$$f_X(x) = \frac{\lambda}{2} e^{-\lambda|x|} \,. \tag{6.1}$$

Let $N = 2L + 1$ be the number of levels of the quantizer (N is odd, as explained earlier in this section). Referring to Figure 6.1, denote the bin boundaries by

$$\{b_j,\, j = \pm 1, \ldots, \pm L\}\,, \tag{6.2}$$

the reconstruction levels by

$$\{q_j,\, j = 0, \pm 1, \ldots, \pm L\}\,, \tag{6.3}$$

and the corresponding bin probabilities by

$$\{p_j,\, j = 0, \pm 1, \ldots, \pm L\}\,. \tag{6.4}$$

The bins of the UTQ are of equal width Δ, and so the bin boundaries are given by

$$b_j = (j - \frac{1}{2})\Delta, \;\; j = 1, \ldots, L\,, \tag{6.5}$$

with

$$b_{-j} = -b_j, \;\; j = 1, \ldots, L\,. \tag{6.6}$$

Due to the symmetry of the pdf, the reconstruction levels are antisymmetric, i.e.,

$$q_{-j} = -q_j, \;\; j = 1, \ldots, L\,, \tag{6.7}$$

and the bin probabilities are symmetric, i.e.,

$$p_{-j} = p_j, \;\; j = 1, \ldots, L\,. \tag{6.8}$$

For the center bin of the UTQ:

$$p_0 = \Pr\left(X \in [b_{-1},\, b_1]\,\right) = 2 \cdot \Pr\left(X \in [0,\, b_1]\,\right) = 2 \int_0^{b_1} \frac{\lambda}{2} e^{-\lambda x}\, dx = 1 - e^{-\lambda b_1} \tag{6.9}$$

and

$$q_0 = E\left[X | X \in [b_{-1},\, b_1]\,\right] = 0 \tag{6.10}$$

where $E\left[\cdot\right]$ denotes expectation, and the last equality follows from symmetry. For the rightmost bin of the UTQ:

$$p_L = \Pr\left(X \in [b_L,\, \infty)\,\right) = \int_{b_L}^{\infty} \frac{\lambda}{2} e^{-\lambda x}\, dx = \frac{1}{2} e^{-\lambda b_L} \tag{6.11}$$

and

$$\begin{aligned} q_L &= E\left[X | X \in [b_L,\, \infty)\,\right] = \\ &= \frac{1}{p_L}\int_{b_L}^{\infty} x\frac{\lambda}{2} e^{-\lambda x}\, dx = \frac{1}{p_L} b_L \cdot \frac{1}{2} e^{-\lambda b_L} + \frac{1}{\lambda} = b_L + \frac{1}{\lambda}\,. \end{aligned} \tag{6.12}$$

For the rest of the bins on the positive half of the real line, i.e., for $j = 1, \ldots, L-1$:

$$p_j = \Pr\left(X \in [b_j,\, b_{j+1}]\,\right) = \int_{b_j}^{b_{j+1}} \frac{\lambda}{2} e^{-\lambda x}\, dx = \frac{1}{2}\left[e^{-\lambda b_j} - e^{-\lambda b_{j+1}}\right] \tag{6.13}$$

and

$$\begin{aligned} q_j &= E\left[X | X \in [b_j,\, b_{j+1}]\,\right] = \\ &= \frac{1}{p_j}\int_{b_j}^{b_{j+1}} x\frac{\lambda}{2} e^{-\lambda x}\, dx = \frac{1}{2p_j}\left[b_j e^{-\lambda b_j} - b_{j+1} e^{-\lambda b_{j+1}}\right] + \frac{1}{\lambda}\,. \end{aligned} \tag{6.14}$$

In Appendix B, 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

$$q_j = b_j + \delta, \quad j = 1, \ldots, L-1\,, \tag{6.15}$$

where

$$\delta = \frac{1}{\lambda} - \frac{\Delta}{e^{\lambda\Delta} - 1}\,. \tag{6.16}$$

This is a consequence of the following useful property of the Laplacian distribution: if X is a Laplacian random variable, and c is a fixed nonnegative real number, then given that $X \geq c$, the conditional distribution of $X - c$ is exponential. Refer to Appendix C for a proof of this property.

The reconstruction levels and bin probabilities for the bins on the negative half of the real line are obtained using (6.7) and (6.8), respectively. Obviously, the aforementioned property of constant reconstruction offsets also applies to the corresponding bins on the negative half of the real line, by symmetry.

Under the Laplacian assumption, the average (squared error) distortion D of the quantizer Q is given by

$$\begin{aligned} D &= E\left[(X - Q(X))^2\right] = \int_{-\infty}^{\infty} f_X(x)\,(x - Q(x))^2\,dx = \\ &= \int_{-b_1}^{b_1} f_X(x)(x-0)^2\,dx + \\ &\quad + 2\sum_{j=1}^{L-1}\int_{b_j}^{b_{j+1}} f_X(x)(x-q_j)^2\,dx + \\ &\quad + 2\int_{b_L}^{\infty} f_X(x)(x-q_L)^2\,dx\,. \end{aligned} \tag{6.17}$$

Based on the property of constant reconstruction offsets (6.15), we derive a closed form expression for the average distortion D, whose evaluation does not involve integration, in Appendix D. The obtained expression for the average distortion D is

$$\begin{aligned} D &= \frac{2}{\lambda^2} - e^{-\lambda\frac{\Delta}{2}}\left(\left(\frac{\Delta}{2}\right)^2 + \frac{\Delta}{\lambda} + \frac{2}{\lambda^2}\right) + \\ &\quad + \left(\sum_{j=1}^{L-1} e^{-\lambda q_j}\right)\cdot\left[e^{\lambda\delta}\left(\delta^2 - \frac{2\delta}{\lambda} + \frac{2}{\lambda^2}\right) - e^{\lambda(\delta-\Delta)}\left((\delta-\Delta)^2 - \frac{2(\delta-\Delta)}{\lambda} + \frac{2}{\lambda^2}\right)\right] + \\ &\quad + \frac{1}{\lambda^2}e^{1-\lambda q_L} \end{aligned} \tag{6.18}$$

where δ is given in (6.16).

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

$$H_Q = -\sum_{j=-L}^{j=L} p_j \log_2 p_j\,. \tag{6.19}$$

Due to the Laplacian model assumption, we have closed form expressions for the bin boundaries, reconstruction levels and bin probabilities of the UTQ, as well as for the distortion and rate estimates. Thus, compared to other distribution models, the Laplacian model also offers a simple and efficient UTQ design. By computing the average distortion D and the rate R for various closely spaced values of the step size Δ, we obtain the operational distortion rate function $\hat{D}(R)$ (also called quantizer function) of the UTQ. This function describes the average distortion D attained by the UTQ as a function of the rate R. Figure 6.2 shows the operational distortion rate function of the optimal UTQ for the unit-variance Laplacian distribution (with squared error distortion). It should be noted that the rate distortion function $R(D)$,

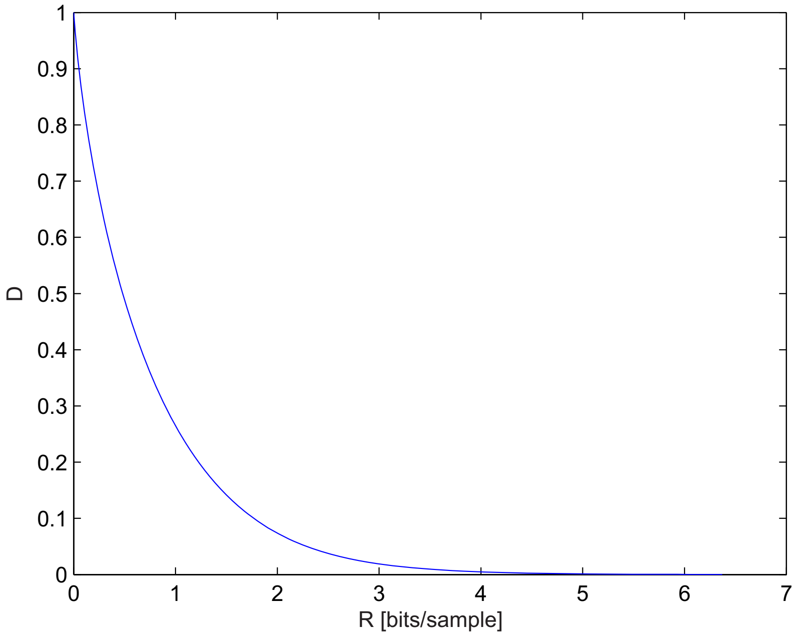

Figure 6.2: Operational distortion rate function of the optimal UTQ for the unit-variance Laplacian distribution with squared error distortion.

or the distortion rate function $D(R)$ [13], for a Laplacian source with squared error distortion, can be computed numerically using the Blahut-Arimoto algorithm [4, 7, 2], as done in [53]. For completeness, we also point out that for the case of absolute difference distortion, i.e., $d(x, \hat{x}) = |x - \hat{x}|$, there exists a closed form expression for the rate distortion function $D(R)$ for a Laplacian source [4, Ch. 4], [10].

6.2 Uniform Reconstruction with Unity Ratio Quantizers as an Alternative to Uniform Threshold Quantizers

As described in Section 6.1, the uniform threshold quantizer (UTQ) is used effectively by numerous image subband quantization schemes. Interestingly, another relatively simple and effective quantizer for the Laplacian distribution—the uniform reconstruction with unity ratio quantizer (URURQ) [53]—is hardly ever used in such applications. Although the UTQ well approximates the optimum entropy-constrained scalar quantizer (ECSQ) for the Laplacian distribution (assuming MSE distortion), the URURQ still offers a small improvement in rate distortion performance, compared to the UTQ [53]. For the proposed MD coder, we also examine the utilization of URURQs as an alternative to UTQs, to see how this slight performance improvement in theory (i.e., for simulated Laplacian data), translates into practice (i.e., in terms of the overall performance of the coder, with quantization of real wavelet data). In this section we provide a brief and concise description of the URURQ.

The URURQ is defined by its step size Δ—the distance between adjacent *reconstruction* levels. This is illustrated in Figure 6.3. As described in Section 6.1, we only focus on quantizers with odd number of levels ("mid-tread" quantizers) in order to be able to obtain low bit rates (i.e., low redundancy), if desired. For Laplacian

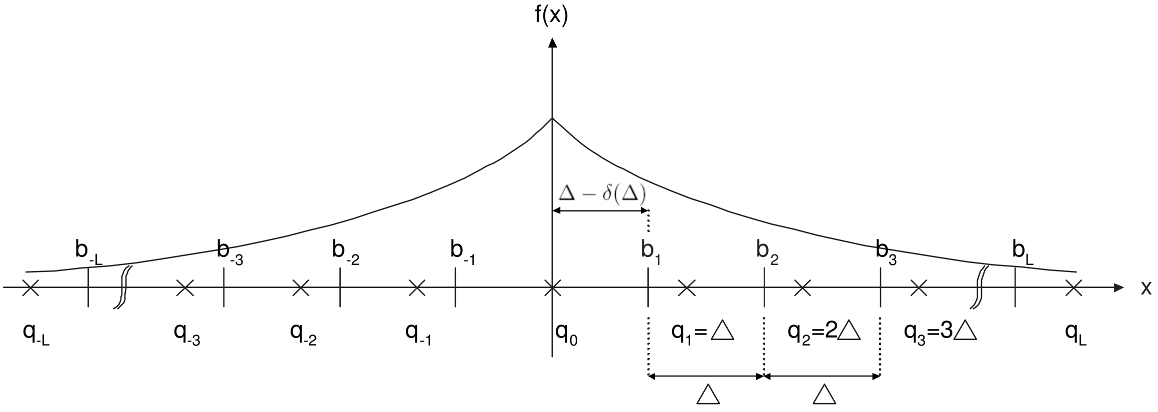

Figure 6.3: Uniform reconstruction with unity ratio quantizer (URURQ) with step size Δ and an odd number of levels $N = 2L + 1$. The bin boundaries are denoted by $\{b_j\}$ and the reconstruction levels by $\{q_j\}$. The Laplacian pdf is also shown for illustration.

distribution with parameter λ, the leftmost bin boundary on the positive half of the real line is

$$b_1 = \Delta - \delta(\Delta) \tag{6.20}$$

where

$$\delta(\Delta) = \arg\min_{\psi} \left\{ \int_0^{\Delta} (x - \psi)^2 e^{-\lambda x}\, dx \right\} . \tag{6.21}$$

Differentiating with respect to ψ and setting to zero to find the minimizer, yields

$$\delta(\Delta) = \frac{1}{\lambda} - \frac{\Delta}{e^{\lambda\Delta} - 1} \tag{6.22}$$

which also equals[1] δ of (6.16). The URURQ is symmetric, and the rest of the bin boundaries to the right of b_1 essentially construct a uniform threshold sub-quantizer

---

[1] This equality is not a coincidence, of course. It is a consequence of the useful property of the Laplacian distribution stated (and proved) in Appendix C.

(to the right of b_1) with step size Δ. Note that given the bin boundaries, the reconstruction levels of the URURQ satisfy the well-known centroid condition (i.e., conditional expectation), and therefore minimize the (MSE) distortion. Also, note that in [53], the dead-zone ratio z is defined as $z = b_1/(\Delta - \delta(\Delta))$, which equals one in the case of the URURQ, and reveals the origin of the term "unity ratio" in the quantizer's name.

Since the UTQ and the URURQ are very similar and also share many properties in common, we are able to straightforwardly adapt the proposed MD coder to operate with URURQs instead of UTQs. Specifically, we adapt all the coding procedures related to the UTQs also to the case of URURQs. In Appendix E we also derive a closed form expression for the average squared error distortion D of the URURQ, whose evaluation does not involve integration. In addition, some experimental results with URURQs are presented in Chapter 8.

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.

6.3 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 (using the wavelet coefficients encoder of Figure 4.8). 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. Note that here we use the convention that the averaging is with respect to the total number of image pixels, so that all rates from all descriptions can be summed to produce the total encoding rate for the image. 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. This efficient bit allocation scheme is based on variance scaling and on a pre-designed, indexed array of optimized uniform threshold quantizers (UTQs) of different step sizes. This array of UTQs is created only once, off-line, and can be regarded as a fixed resource of the coder, used by both the encoder and the decoder. It enables the proposed coder to avoid complex on-line bit allocation procedures, as well as to intelligently and instantly adapt the arithmetic entropy encoder to the varying coefficients statistics, by storing the pre-computed bin probabilities of each UTQ. It should be noted that this adaptation scheme enables the entropy encoder to fully exploit the higher level statistics captured by the Laplacian model-based classification and quantization algorithms, rather than using ad hoc updating rules, as done in many other entropy coding adaptation schemes.

The bit allocation, which is based on the Laplacian model, is performed in the general framework of Lagrangian optimization [40, 50, 55]. 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 (for the approximation subband, a predictive, DPCM-like, quantization scheme is employed, so the prediction errors are considered instead; this will be described in Section 7.2.1).

Given the desired redundancy rate for the description, we wish to find the optimal rate 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, as described in Section 5.3).

The following notation is used in the sequel:

D : Average squared error distortion (MSE) in the image domain, resulting from quantization of coefficients in the redundant polyphase component.

R : Average rate (in bits/pixel) for encoding the coefficients in the redundant polyphase component (i.e., the redundancy rate of the description).

R_T : Desired redundancy rate (in bits/pixel) for the description.

B : The total number of classes (from all subbands).

b : Index of a specific class ($1 \leq b \leq B$).

N_b : Number of coefficients in the class b.

N : Number of pixels in the image.

η_b : The ratio N_b/N, i.e., the ratio between the number of coefficients in the class b and the number of pixels in the image.

D_b : Average squared error distortion (MSE) in the transform domain in encoding the class b, where the averaging is with respect to the number of coefficients in the class (N_b).

R_b : Average rate (in bits/coefficient) for encoding the coefficients in the class b, where the averaging is with respect to the number of coefficients in the class (N_b).

G_b : The synthesis (inverse transform) energy gain factor associated with the subband to which the class b belongs. 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 $G_b = 1, \forall b$. A detailed description will be given in Section 7.1.2.

Thus, given the desired redundancy rate R_T, our objective is to find the optimal rates $\{R_b\}_{b=1}^{B}$ for the different classes, such that the distortion D is minimized, subject to the constraint $R \leq R_T$. We use Lagrangian optimization [40, 50, 55] to solve this constrained optimization problem, and instead seek to minimize the Lagrangian cost function

$$J(\xi) = D + \xi R \tag{6.23}$$

for a fixed Lagrange multiplier ξ. Later, we will describe how to determine the Lagrange multiplier ξ such that $R = R_T$ (or, more precisely, $R \in (R_T - \epsilon,\, R_T]$ for an arbitrary tolerance $\epsilon > 0$). This will also yield the solution to the original constrained optimization problem [50].

The average distortion and rate are given by

$$D = \sum_{b=1}^{B} \eta_b G_b D_b(R_b) \tag{6.24}$$

and

$$R = \sum_{b=1}^{B} \eta_b R_b\,, \tag{6.25}$$

respectively. Thus,

$$\begin{aligned}
\min_{R_1,\ldots,R_B} J(\xi) &= \min_{R_1,\ldots,R_B} \{D + \xi R\} = \\
&= \min_{R_1,\ldots,R_B} \left\{ \sum_{b=1}^{B} \eta_b G_b D_b(R_b) + \xi \sum_{b=1}^{B} \eta_b R_b \right\} = \\
&= \sum_{b=1}^{B} \min_{R_b} \{\eta_b G_b D_b(R_b) + \xi \eta_b R_b\} = \\
&= \sum_{b=1}^{B} \eta_b \min_{R_b} \{G_b D_b(R_b) + \xi R_b\} \qquad (6.26)
\end{aligned}$$

which suggests that the minimum can be computed independently for each class.

For a given class b, the optimal rate R_b which minimizes

$$G_b D_b(R_b) + \xi R_b \qquad (6.27)$$

can be found by differentiating (6.27) with respect to R_b and setting to zero. This yields

$$G_b D_b'(R_b) + \xi = 0 \,, \qquad (6.28)$$

or

$$D_b'(R_b) = -\frac{\xi}{G_b} \,, \qquad (6.29)$$

where $D_b'(R_b)$ is the derivative of the distortion D_b, evaluated at R_b bits/coefficient.

The rates $\{R_b\}_{b=1}^{B}$ are thus chosen to satisfy the following rate allocation equations

$$D_b'(R_b) = -\frac{\xi}{G_b}, \;\; b = 1, \ldots, B \,. \qquad (6.30)$$

For each class, the relevant equation specifies the required slope on the operational distortion rate function of the corresponding UTQ (and thus determines its step size). The rate allocation equations (6.30) realize the well-known “constant slope” principle [40] in Lagrangian optimization, adapted to the case of a non-orthonormal transform (i.e., arbitrary G_b’s).

For efficiency, our bit allocation scheme utilizes the following properties of the optimal UTQ for the Laplacian distribution, optimized for minimum MSE distortion:

1. For a given rate (output entropy), the optimal quantizer for an input with variance σ^2 is a scaled version, by σ, of the optimal quantizer for a unit-variance input (i.e., the bin boundaries and reconstruction levels of the unit-variance quantizer are to be multiplied by σ). The bin probabilities (and obviously the output entropy) of both quantizers are the same.

2. The distortion attained by the scaled quantizer above is larger by a factor of σ^2 compared to that attained by the unit-variance quantizer (each with its corresponding input).

These properties can be easily proved using the definitions and computations of Section 6.1. The important consequence of the aforementioned properties is that the operational distortion rate function $\hat{D}(R)$ of the optimal UTQ for an input with variance σ^2 can be obtained from the normalized distortion rate function (i.e., that for a unit-variance input; see Figure 6.2) by merely scaling the distortion axis by σ^2. Obviously, the slope (the derivative of the distortion with respect to the rate) is affected similarly (i.e., multiplied by a factor of σ^2).

Thus, in order to solve the rate allocation equations (6.30), we can 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 $\hat{\sigma}_b^2 = 2/\hat{\lambda}_b^2$). Instead, we only compute (off-line) 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 distortion rate function and the normalized slope $-\xi/(G_b\hat{\sigma}_b^2)$, as detailed next.

Off-line, we create an array of optimized UTQs for the unit-variance Laplacian distribution, where each quantizer has a different step size Δ. For each quantizer, we compute and store the bin boundaries, reconstruction levels, and bin probabilities, as well as the resulting rate and distortion, as detailed in Section 6.1. By using closely spaced values of the step size Δ, we also obtain the operational distortion rate function of the optimal UTQ for the unit-variance Laplacian distribution, and are able to well approximate its derivative (i.e., slope) at each operating point. For each quantizer in the array we also store the corresponding slope on the operational distortion rate function, which enables us to index the array by the slope (as required by the bit allocation algorithm).

During the encoding process itself, for a fixed Lagrangian multiplier ξ, and a given class b, we need to choose the appropriate quantizer (in order to quantize the coefficients of the class b). To do so, we first translate the (non-normalized) slope $-\xi/G_b$ from the relevant rate allocation equation of (6.30) to the corresponding normalized slope $-\xi/(G_b\hat{\sigma}_b^2)$ on the unit-variance operational distortion rate function, where $\hat{\sigma}_b^2 = 2/\hat{\lambda}_b^2$ is the variance estimate for the class b. Since the pre-computed array of optimized UTQs is indexed by the slope, the appropriate "normalized" quantizer (i.e., the one whose slope is closest to the required normalized slope) can be instantly obtained from the array. The actual quantizer for the class b is simply the scaled version, by $\hat{\sigma}_b$, of this "normalized" quantizer, as explained earlier in this section. Namely, the bin boundaries and reconstruction levels are multiplied by $\hat{\sigma}_b$, while the bin probabilities remain unchanged. The 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.

We now turn to describe how to determine the Lagrange multiplier ξ in the

Lagrangian cost function $J(\xi)$ of (6.23), such that $R = R_T$ (or, more precisely, $R \in (R_T - \epsilon, R_T]$ for an arbitrary tolerance $\epsilon > 0$). First, define the following function

$$f(\xi) = R(\xi) - R_T \tag{6.31}$$

where R_T is the desired redundancy rate, and $R(\xi)$ is the average redundancy rate resulting from minimization of the Lagrangian cost function of (6.23) (for the same Lagrange multiplier ξ). Finding a proper ξ such that $R(\xi) = R_T$ is then equivalent to finding a root of $f(\xi)$. Since $R(\xi)$ is monotonically non-increasing with ξ [50] (as ξ controls the penalty for an increased rate in the Lagrangian cost function of (6.23)), root finding of $f(\xi)$ can be easily performed using the bisection method [42, Ch. 9]. This involves bracketing a root (i.e., finding ξ_1 and ξ_2 such that $f(\xi_1) > 0$ and $f(\xi_2) < 0$), and then iterating until $R(\xi) \in (R_T - \epsilon, R_T]$, for an arbitrary tolerance $\epsilon > 0$ (e.g., $\epsilon = 0.005$).

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). For each class in each subband, the decoder can then match the same UTQ used by the encoder, by repeating the aforementioned variance scaling procedure performed by the encoder. The bin probabilities of the matched UTQs serve to properly adapt the arithmetic entropy decoder, and the classification and dequantization stages follow to correctly decode the redundant polyphase component of the description. Obviously, decoding this redundant part of information is needed only in the case that the other description is lost, as explained in Section 4.2.

Chapter 7

Context-Based Multiple Description Wavelet Image Coding: Implementation Considerations and Miscellaneous

7.1 Implementation Considerations

7.1.1 Boundary Extension for Wavelet Transform

Section 4.1.1 described the computation of the two-dimensional discrete wavelet transform of an image using a series of filtering (and downsampling) operations. Since the image and the filters are of finite length, special attention has to be paid in computing their convolution, at each stage of the decomposition process. This section addresses the problem of boundary extension for the wavelet transform, especially referring to coding applications. In order to simplify the discussion, some of the concepts in this section are illustrated using the one-dimensional wavelet transform, but hold equally well for the two-dimensional case (through the separable extension of the one-dimensional transform, as described in Section 4.1.1).

As detailed in [57], the standard orthogonal wavelet transform [51, Ch. 6] has a significant shortcoming that limits its usefulness in coding applications: it generally results in “coefficient expansion”, that is to say, the number of output wavelet coefficients is greater than the number of input samples. This increase in the number of coefficients generally degrades the performance of the coding system, since extra bits are required to represent the additional coefficients. This is the case, for example, if zero-padding boundary extension is used, which assumes that the image is zero outside its boundaries. Zero-padding extension may also result in unpleasant artifacts, usually in the form of dark areas, near the boundaries of the reconstructed image.

Coefficient expansion can be eliminated by using periodic-padding boundary extension (or circular convolution), which assumes that the image is periodic. This type of extension generally results in large discontinuities near the boundaries, however, and extra bits are required in order to code them. Moreover, annoying artifacts near the boundaries of the reconstructed image may appear due to quantization of the artificially introduced large wavelet coefficients at the locations of the discontinuities.

An effective way to avoid coefficient expansion and greatly reduce boundary artifacts is to use symmetric (or antisymmetric) wavelet filters along with symmetric boundary extension. By using symmetric boundary extension, the image is extended in a continuous way outside its boundaries, and thus artificial discontinuities are avoided. Due to the symmetry imposed by the extension, there is virtually no effective increase in the number of input samples, and coefficient expansion can be avoided if the symmetry is preserved at each stage of the decomposition process. The key underlying concept is that the convolution of a symmetric ("linear phase") input with a linear phase filter produces a linear phase output (this can be easily observed in the frequency domain as the linear exponents add and remain linear), and that the exact type of symmetric extension can be chosen such that the symmetry is preserved also by the downsampling operation.

Two types of symmetric extension should be considered in this context—*half-point* and *whole-point* symmetric extension—where the former repeats the outermost samples of the input, while the latter does not. This is illustrated in Figure 7.1. As shown, the extended input is periodic for either type of extension. For an input of length L, the period of the extended input is $2L$ in the case of half-point symmetric

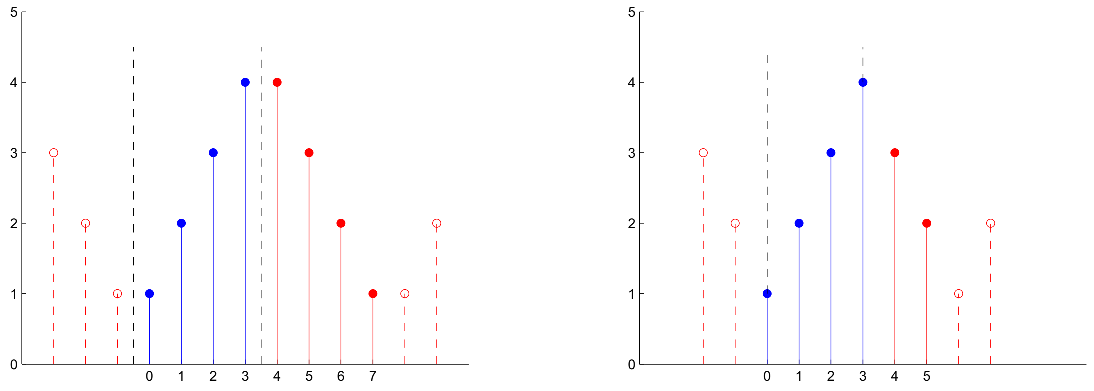

(a) (b)

Figure 7.1: Half-point and whole-point symmetric extensions. (a) Half-point symmetric extension. (b) Whole-point symmetric extension. The original (one-dimensional) input is shown in blue, and the extension is shown in red. The resulting axes of symmetry are shown in black, and are located at "half-point" locations in (a) and at "whole-point" locations in (b). One period of the extended input is shown in solid lines.

extension, and is $2(L-1)$ in the case of whole-point symmetric extension. In both cases, however, only L samples of the extended input are not redundant.

Two questions still have to be answered: First, are there any linear phase (i.e., symmetric or antisymmetric) wavelet filters? And second, can we guarantee that the downsampling operation will preserve the symmetry by choosing the proper type of symmetric extension?

Regarding the first question, unfortunately, for real valued, compactly supported orthogonal wavelets, the only set of linear phase filters is the trivial set of Haar filters [14]. Nevertheless, the more general form of biorthogonal wavelets [51, Ch. 6] permits the use of a much broader class of filters, which includes more sophisticated linear phase filters, and has been shown to provide an improved performance (compared to the use of orthogonal wavelets) in image coding applications [35]. It should be noted that unlike the orthonormal wavelet transform, the biorthogonal wavelet transform does not possess the useful property of being energy preserving, as will be discussed in Section 7.1.2.

We now give an answer to the second question, dealing with choosing the proper type of symmetric extension (half-point or whole-point) such that the symmetry can be preserved by the downsampling operation. The proper type of extension is related to the lengths of the wavelet filters used, and the following guideline should be followed:

- Use a half-point (repeat) symmetric extension for an even-length filter.
- Use a whole-point (no repeat) symmetric extension for an odd-length filter.

The various tricky details involved in deriving the above guideline, as well as detailed directions for proper downsampling of the symmetric output of the filters, are given in [51, Ch. 8].

In summary, coefficient expansion and unpleasant boundary artifacts can be avoided by using linear phase wavelet filters along with symmetric boundary extension. Depending on the lengths of the filters used, the proper type of extension (half-point or whole-point) should be chosen. The specific choice of wavelet filters and boundary extension for the proposed MD coder will be detailed in Chapter 8. The consequences of using the non-energy preserving biorthogonal wavelet transform, which enables the use of efficient linear phase filters, are addressed next.

7.1.2 Optimal Bit Allocation for Biorthogonal Wavelet Transform

As discussed in the previous section, biorthogonal wavelets permit the use of a much broader class of filters, compared to orthogonal wavelets, which includes non-trivial linear phase filters. Nevertheless, one of the main drawbacks of the biorthogonal wavelet transform is that, unlike the orthonormal wavelet transform, it does not possess the useful property of being energy preserving. Therefore, bit allocation algorithms which are optimal for energy preserving transforms (e.g., orthonormal wavelet transform), are no longer optimal in the case of biorthogonal wavelet transform. Yet, numerous coding systems that utilize biorthogonal wavelets, use bit allocation procedures that were designed for orthonormal wavelets, although these bit allocation procedures provide suboptimal performance, in terms of distortion versus rate, in the case of biorthogonal wavelets [8].

This section discusses optimal bit allocation for biorthogonal wavelet transform, and provides a procedure to adapt standard bit allocation algorithms to this special case. The adaptation is based on the fact that the biorthogonal transform essentially weights the MSE in each of the subbands to obtain the MSE in the reconstructed output, according to the energy expansion properties of the synthesis system (inverse transform), under the assumption that the quantization errors are uncorrelated, with zero mean [56]. Using the notation of Section 6.3, the MSE in the reconstructed image can therefore be presented as a weighted average of the MSEs in the different subbands, as in (6.24), to obtain

$$D = \sum_b \eta_b G_b D_b(R_b) \tag{7.1}$$

where the G_b's are the synthesis energy gain factors (weights) associated with the different subbands[1]. For an orthonormal transform $G_b = 1, \forall b$. It should be noted that although our discussion is generally oriented towards the two-dimensional case (i.e., towards image coding), the principles presented also apply to other cases.

The energy gain factor G_b of a given subband represents the amount of squared error in the reconstructed image introduced by a unit error in a transform coefficient of the subband. It can be shown that this energy gain factor is given by the squared l_2 norm (energy) of the synthesis basis vectors for this subband [55, Ch. 4].

Computation of the energy gain factor of a subband given the synthesis filters can be performed by computing the squared l_2 norm of the single equivalent filter which takes the subband coefficients directly to the reconstructed output [57]. Specifically, in the case of a two-dimensional separable wavelet transform, the energy gain factor of a subband is the product of its "vertical weight" and its "horizontal weight". The vertical (horizontal) weight of a subband is related to the synthesis filters applied in the vertical (horizontal) dimension during the synthesis from the subband to the

---

[1] In the notation of Section 6.3, a single class in each subband is assumed here.

reconstructed image. As an example, coefficients belonging to the horizontal detail subband at level three of the wavelet transform (d_3^H), undergo one highpass filtering operation followed by two lowpass filtering operations in the vertical dimension during the synthesis, and three lowpass filtering operations in the horizontal dimension. Denote the impulse responses of the one-dimensional lowpass and highpass synthesis filters by l_r and h_r, respectively. The vertical weight is computed by first upsampling h_r by a factor of two (i.e., adding one zero between each coefficient in h_r) and convolving the result with l_r. The result of the convolution is then upsampled by a factor of two and convolved with l_r. The vertical weight is given by the squared l_2 norm (i.e., the sum of squares) of the final convolution result. The horizontal weight is computed in a similar manner, using l_r three times. The product of the vertical and horizontal weights then gives the energy gain factor of the subband under discussion (d_3^H).

As an example, consider the Cohen-Daubechies-Feauveau (CDF) biorthogonal wavelet transform using 9/7-tap wavelet filters [11, 1], extensively utilized in image compression applications. Using the above procedure, we computed the energy gain factors of the various subbands of the corresponding three-level two-dimensional wavelet decomposition. The obtained energy gain factors are shown in Table 7.1. As can be seen, the energy gain factors are not equal to one, suggesting that the transform is not energy preserving. Nevertheless, the energy gain factors for this particular transform deviate only by a few percent from one, implying that the transform is reasonably "close" to being orthogonal. Other biorthogonal transforms may generally present much higher deviations from orthogonality, which must be taken into consideration by the bit allocation algorithm of a coding system utilizing the biorthogonal wavelet transform, in order to achieve optimal performance for the system.

Level	Subband	Energy Gain Factor
1	d_1^D (Diag. details)	1.08250699475332
	d_1^V (Vert. details)	1.02270033578441
	d_1^H (Horiz. details)	1.02270033578441
2	d_2^D (Diag. details)	0.93550641543575
	d_2^V (Vert. details)	0.99681499726207
	d_2^H (Horiz. details)	0.99681499726207
3	d_3^D (Diag. details)	1.08082593648928
	d_3^V (Vert. details)	1.09378514093722
	d_3^H (Horiz. details)	1.09378514093722
	a_3 (Approximation)	1.10689972746311

Table 7.1: Energy gain factors for the three-level two-dimensional biorthogonal wavelet transform using Cohen-Daubechies-Feauveau (CDF) 9/7-tap wavelet filters.

7.2 Miscellaneous

7.2.1 Prediction Scheme for Approximation Coefficients

In the proposed context-based MD encoder (see Figure 4.8 in Section 4.2), the various components in the encoding path of the redundant polyphase component of each description model the detail wavelet coefficients (i.e., the coefficients from all subbands excluding the approximation) using a (zero-mean) Laplacian distribution. For the approximation subband, a prediction scheme is employed first, based on quantized approximation coefficients of the complementary polyphase component, and the resulting prediction errors are then encoded. Since the prediction errors are modeled using the same Laplacian distribution model used for the detail coefficients, the same encoding scheme can be utilized for all subbands (where for the approximation subband the prediction errors are encoded).

The prediction scheme employed for the approximation subband predicts the value of each approximation coefficient in the redundant polyphase component based on quantized neighboring coefficients of the complementary polyphase component. This is shown in Figure 7.2. (Notice the resemblance to contexts formed for context-based classification.) The predicted value for the approximation coefficient $X_{i,j}$ is given by

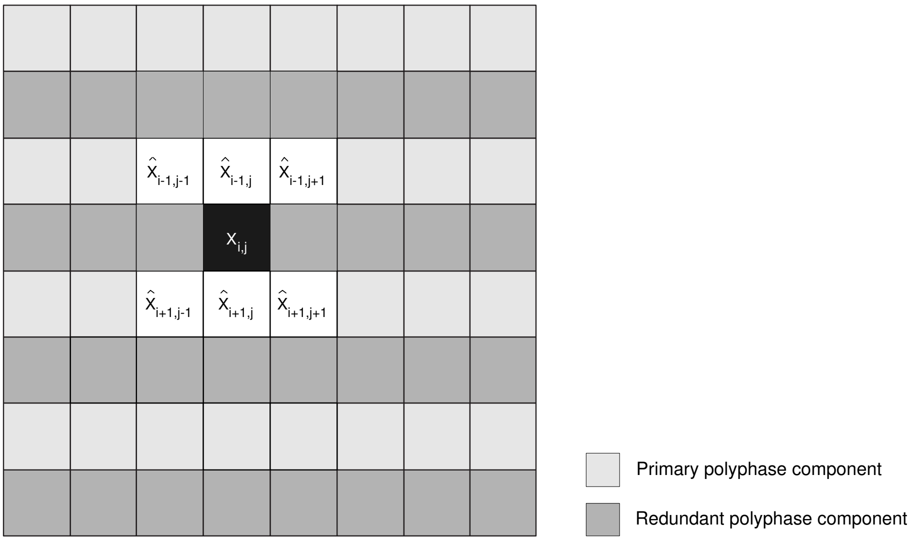

Figure 7.2: Prediction of an approximation coefficient $X_{i,j}$ based on its quantized neighbors. The approximation coefficient $X_{i,j}$ belongs to the redundant polyphase component, and is shown in black. The quantized coefficients used for the prediction belong to the primary polyphase component, and are shown in white.

a weighted average of its quantized neighbors shown. The proposed coder uses the following prediction mask

$$\frac{1}{2\left(1+\sqrt{2}\right)} \cdot \begin{array}{|c|c|c|} \hline \frac{1}{\sqrt{2}} & 1 & \frac{1}{\sqrt{2}} \\ \hline & \cdot & \\ \hline \frac{1}{\sqrt{2}} & 1 & \frac{1}{\sqrt{2}} \\ \hline \end{array} \tag{7.2}$$

so that the weights are inversely proportional to the Euclidean distances of the corresponding quantized neighbors from the coefficient $X_{i,j}$, and sum to one. The prediction error is then given by the difference between $X_{i,j}$ and its predicted value.

The context-based classification of the prediction errors (corresponding to approximation coefficients of the redundant polyphase component) is performed in a very similar manner to that performed for detail coefficients. The distinction is that, instead of being based on magnitudes of neighboring quantized coefficients from the primary polyphase component of the description, the classification procedure in the approximation subband is based on magnitudes of generated "prediction errors", computed using quantized approximation coefficients of the primary polyphase component. Note that these generated prediction errors are neither encoded nor transmitted in any way, but only serve to provide contextual information for the classification of the actual prediction errors (corresponding to approximation coefficients of the redundant polyphase component). The underlying assumption is that large prediction errors in the approximation subband tend to cluster together at relatively active areas (e.g., edges), and that these dependencies can be utilized to efficiently encode the prediction errors to be transmitted. Thus, for the approximation coefficients of the redundant polyphase component, contexts are formed from generated prediction errors, computed using quantized approximation coefficients of the primary polyphase component. Specifically, instead of using the context coefficients shown in Figure 5.1 themselves, each coefficient in the context is replaced by its corresponding generated prediction error. This generated prediction error is calculated as the difference between the quantized context coefficient (belonging to the primary polyphase component) and a weighted average of three of its quantized neighbors (which also belong to the primary polyphase component). For illustration, see Figure 7.3, which resembles Figure 5.1, but also marks with asterisks the three neighboring coefficients used to “predict” the value of the quantized approximation coefficient $\hat{X}_{i-1,j-1}$. In the formation of the context for the classification of the actual prediction error of $X_{i,j}$, the context coefficient $\hat{X}_{i-1,j-1}$ is replaced by the difference between its (quantized) value and a weighted average of its three quantized neighbors marked with asterisks. The weights in the above weighted average can be chosen, for example, to correspond to the following weighting mask [19, Ch. 8]

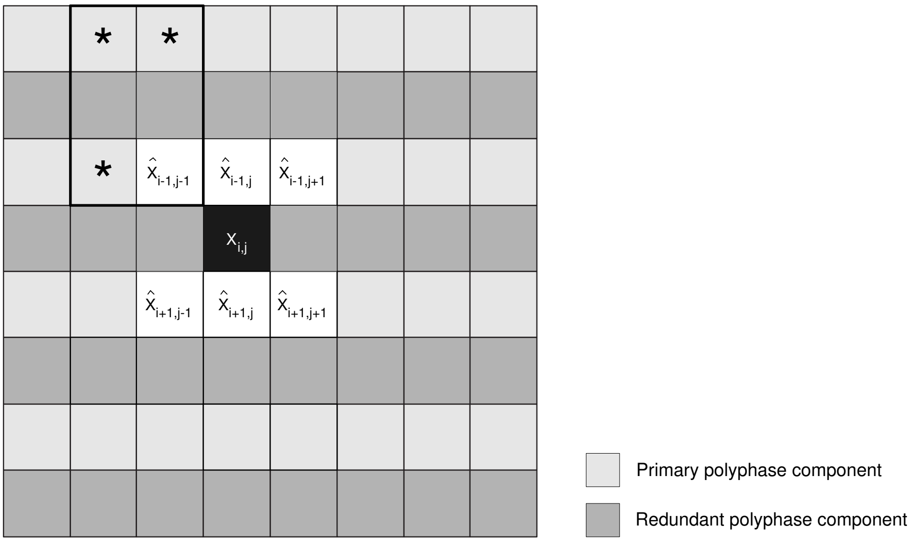

Figure 7.3: Quantized coefficients from the primary polyphase component (marked with asterisks) used to “predict” the value of the quantized approximation coefficient $\hat{X}_{i-1,j-1}$.

$$\begin{array}{|c|c|}\hline -0.5 & 0.75 \\ \hline & \\ \hline 0.75 & \cdot \\ \hline \end{array} \tag{7.3}$$

whose elements sum to one. The rest of the context is formed in a similar manner, with the required spatial adjustments. The context-based classification procedure for the prediction errors corresponding to approximation coefficients of the redundant polyphase component, given the contexts formed from generated “prediction errors”, is identical to that used for the detail coefficients (described in Chapter 5).

Finally, we note that some standard (single description) wavelet coders employ a predictive (DPCM-like) quantization scheme for the approximation wavelet coefficients, and encode the resulting quantized prediction errors. If such a coder is utilized in the proposed MD coding scheme (as the “wavelet coefficients encoder” in Figure 4.8) to encode the primary polyphase component of each description, these quantized prediction errors can directly replace the aforementioned generated “prediction errors”, and provide the contextual information for the classification of the actual prediction errors (corresponding to approximation coefficients of the redundant polyphase component). Obviously, in such a case, there is no need to explicitly generate the aforementioned “prediction errors” in order to provide the required contextual information.

Chapter 8

Experimental Results

This chapter provides experimental results obtained by the proposed context-based MD wavelet image coding system. We begin with a detailed description of the actual configuration of the system used to produce the results. Demonstrations related to the quality of the context-based classification follow, as well as both objective and subjective performance results, which also include a comparison to other relevant MD image coding systems. In addition, the effect of various choices made by the system on the resulting performance is presented and analyzed. We then demonstrate the improvement in coding efficiency due to the effective utilization of contextual information, which we refer to as the *context gain*. Finally, we describe how to determine the optimal operating point for the proposed MD coder, based on the properties of the communication channel.

8.1 System Configuration

A simplified block diagram of the proposed context-based MD encoder was presented in Figure 4.8 (in Section 4.2). Referring to the block diagram shown, this section provides a detailed description of the actual configuration of the various components of the system, as used to produce the experimental results presented in the chapter.

Regarding the choice of wavelet transform, we avoid coefficient expansion and unpleasant boundary artifacts by utilizing linear phase wavelet filters along with symmetric boundary extension, as discussed in Section 7.1.1. Specifically, a two-dimensional biorthogonal wavelet transform using Cohen-Daubechies-Feauveau (CDF) 9/7-tap wavelet filters [11, 1] is employed, a choice which has been shown experimentally to provide excellent performance for image compression [64]. Since the wavelet filters are of odd length, a whole-point (no repeat) symmetric boundary extension is used, and coefficient expansion is avoided (see Section 7.1.1). The number of levels in the decomposition is set to three, and the resulting number of subbands is therefore 10.

The polyphase transform used by the proposed coder is the plain polyphase transform, which, for each column (or row) in each subband, simply groups all the odd-numbered coefficients into one polyphase component and all the even-numbered coefficients into the other component. This is shown in Figure 8.1. The orientation of the polyphase transform is adapted to the orientation of the subband on a subband-by-subband basis, in order to enable the context-based classification procedure to more effectively utilize of the contextual information extracted from the complementary polyphase component. Thus, for all subbands except those of horizontal details the polyphase transform shown in Figure 8.1 is used, while for subbands of horizontal

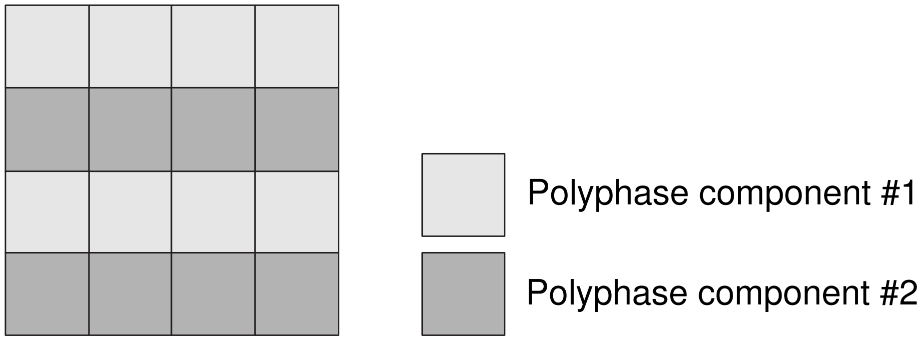

Figure 8.1: Plain polyphase transform used by the proposed MD coder.

details (at all scales) a transposed version of the transform shown is used (i.e., odd-indexed columns constitute one polyphase components, and even-indexed columns constitute the other component). Of course, the orientation of the contexts formed for context-based classification of wavelet coefficients belonging to subbands of horizontal details, is changed accordingly (i.e., a transposed version of the context shown in Figure 5.1).

As described in Section 4.2, the primary polyphase component of each description is encoded independently by a wavelet coefficients encoder, which can be a standard wavelet encoder, operating directly on the wavelet coefficients. The single description wavelet encoder of Yoo et al. [66] (with no adaptive parameter estimation) is used by the proposed MD coder for this task. The resulting quantized coefficients also serve to provide contextual information for efficient encoding of the redundant polyphase component of the corresponding description, as shown in Figure 4.8.

The context-based classification procedure uses a total of four classes per subband, unless stated otherwise, and the initial number of classes N_0 in Algorithm 5.1 (the classification thresholds design algorithm; see Section 5.3) is chosen such that each class initially contains approximately 16 coefficients. This choice for the initial number of classes offers a relatively dense set of initial thresholds (from which the final thresholds are selected), and yet enables a reliable estimation of the Laplacian parameter for each class. We experimentally found this choice to yield very good classification results, compared to other choices.

For each description, the resulting classification thresholds and Laplacian parameter estimates (for each of the subbands) are transmitted to the decoder as side information. The side information for each description also includes the value of the Lagrange multiplier ξ, which specifies the bit allocation (see Section 6.3). Each of the side information parameters is represented using 16 bits (in fixed-point representation). Note that the resulting overhead due to the side information is generally minimal. For example, for the various choices stated above, the overhead due to side information translates to an excess rate of less than 0.004 bpp (bits per pixel), per description, in the case of a grayscale input image of size 512×512 pixels.

Finally, we note that we consider in our experiments the balanced MD coding scenario, in which the generated descriptions are of statistically equal rate and importance. This balanced case is generally the one considered in MD coding applications, as it is well matched to communication systems with no priority mechanisms for data delivery (e.g., the Internet) [30].

8.2 Quality of Classification

This section demonstrates the subjective quality of the context-based classification, which enables the proposed MD coder to efficiently encode the redundant polyphase component of each description, based on quantized coefficients of the complementary polyphase component.

As an example, consider the original (grayscale) image Lena, of size 512 × 512 pixels, shown in Figure 8.2(a). The wavelet transform (in absolute value) of the image Lena, with approximation coefficients replaced by their corresponding prediction errors (see Section 7.2.1), is shown in Figure 8.3(a). Note that all subbands have been enhanced to make their underlying structure more visible. Figure 8.3(b) shows the corresponding classification map, where the classification of each coefficient is based on quantized coefficients from the complementary polyphase component (in this example the quantization of the primary polyphase component of each description is at an average rate of 0.48 bpp, per description). This classification map has only four gray levels, corresponding to the four possible classes in each subband (coefficients

(a) (b)

Figure 8.2: Original images (grayscale) of size 512 × 512 pixels.
(a) Lena. (b) Goldhill.

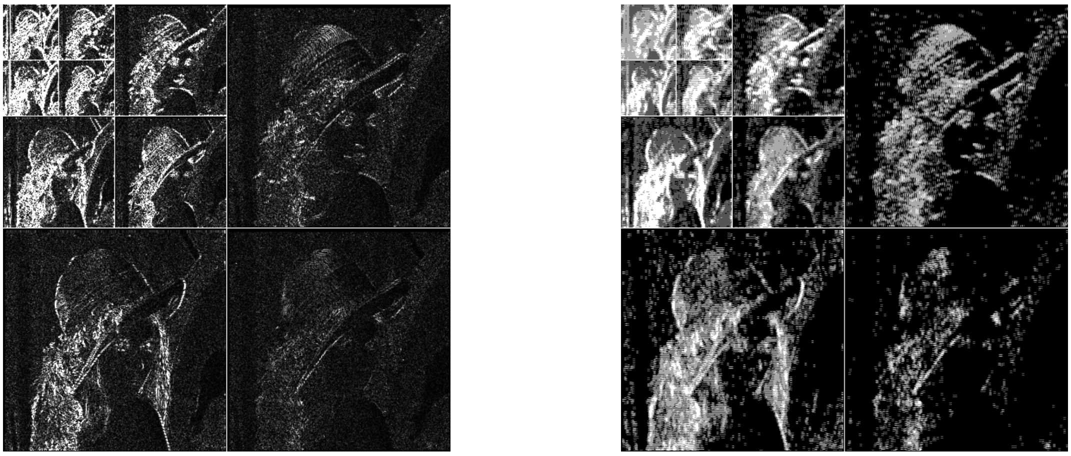

(a)

(b)

Figure 8.3: (a) Wavelet transform (in absolute value) of the original image Lena, with approximation coefficients replaced by their corresponding prediction errors. (b) Corresponding classification map (with four gray levels).

assigned to the least active class C_0 are shown in black, and those assigned to the most active class C_3 are shown in white). As shown, despite the fact that the classification of each coefficient is based solely on quantized neighboring coefficients from the complementary polyphase component, areas with different energy levels are successfully identified by the context-based classification procedure, and effective classification is achieved.

Figure 8.4 shows the histograms corresponding to the four classes assigned to the coefficients in subband d_1^H of the redundant polyphase component of description 1. Each of the histograms shows the actual distribution of the coefficients assigned to its corresponding class, and also includes a plot of the fitted Laplacian pdf for the class (using maximum likelihood estimation of the Laplacian parameter, as described in Section 5.3). As shown, the classified coefficients are successfully separated into four different classes according to their activity, where classes corresponding to a higher activity indeed contain coefficients with a larger variance. In addition, the Laplacian pdf matched to the coefficients in each class fits their actual distribution reasonably well.

8.3 Performance of the System

This section provides both objective and subjective performance results for the proposed MD coder, including a comparison to other relevant MD coding systems. It also presents and analyzes the effect of various choices made by the system on the resulting performance.

The PSNR (Peak Signal-to-Noise Ratio) values corresponding to the central and

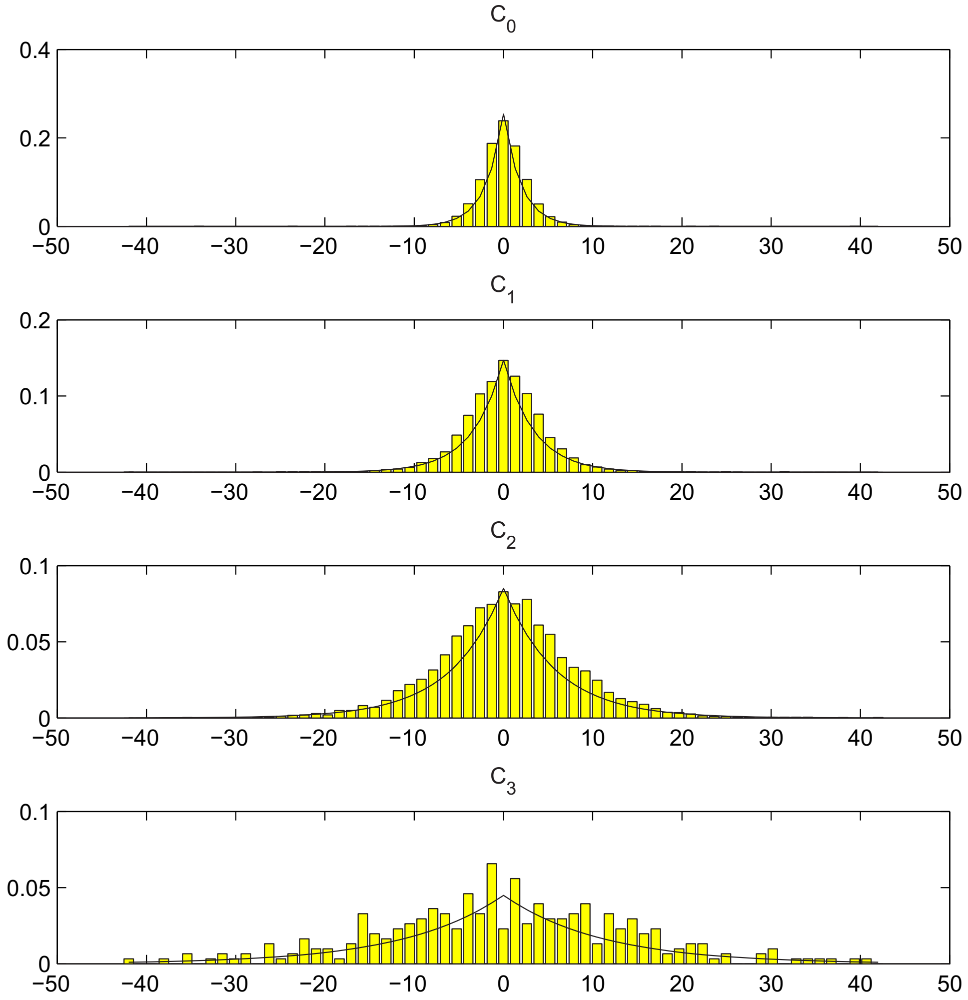

Figure 8.4: Histograms corresponding to the four classes assigned to the coefficients in subband d_1^H of the redundant polyphase component of description 1. Each histogram also includes a plot of the fitted Laplacian pdf for the class.

(average) side distortions attained for different amounts of redundancy for the images Lena and Goldhill (shown in Figure 8.2) are shown in Figure 8.5 and Figure 8.6, respectively. The total rate, which takes into account both descriptions altogether, including the overhead due to side information (i.e., every single bit transmitted) is fixed, and is set to 1 bpp in this example. Each performance curve is virtually a parametric curve, where different points on the curve correspond to different amounts

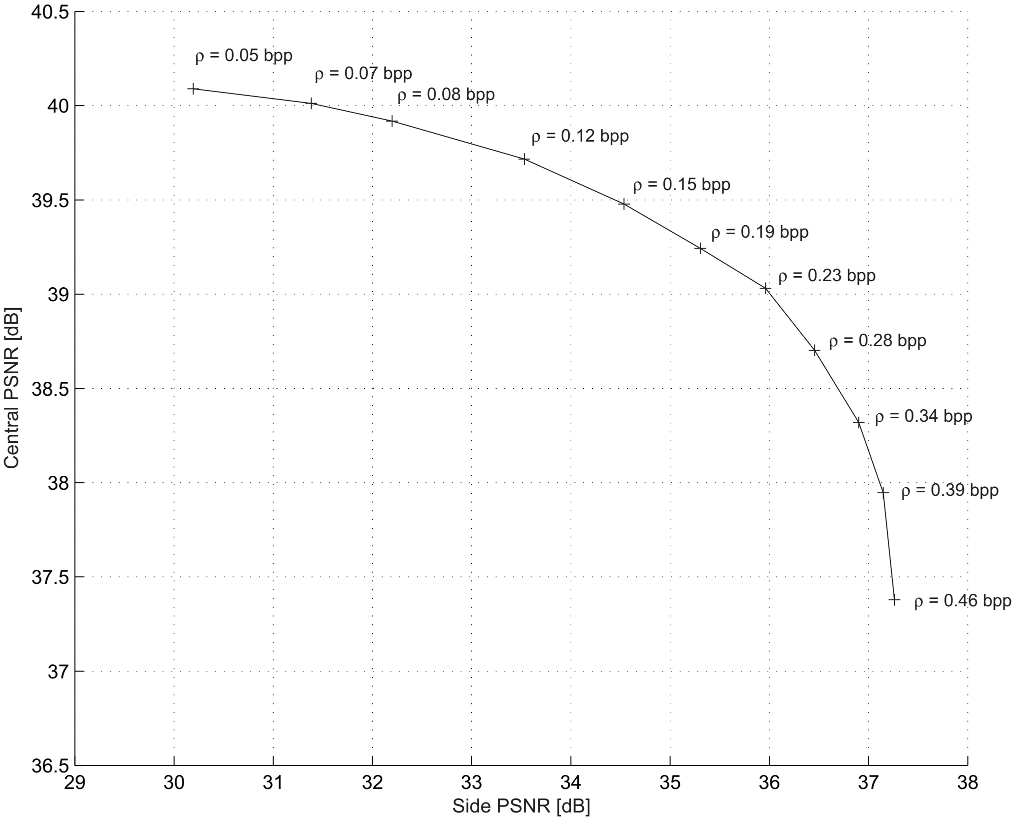

Figure 8.5: Performance of the proposed MD coder for the image Lena (total rate of 1 bpp). Also shown is the redundancy rate ρ corresponding to various points on the performance curve.

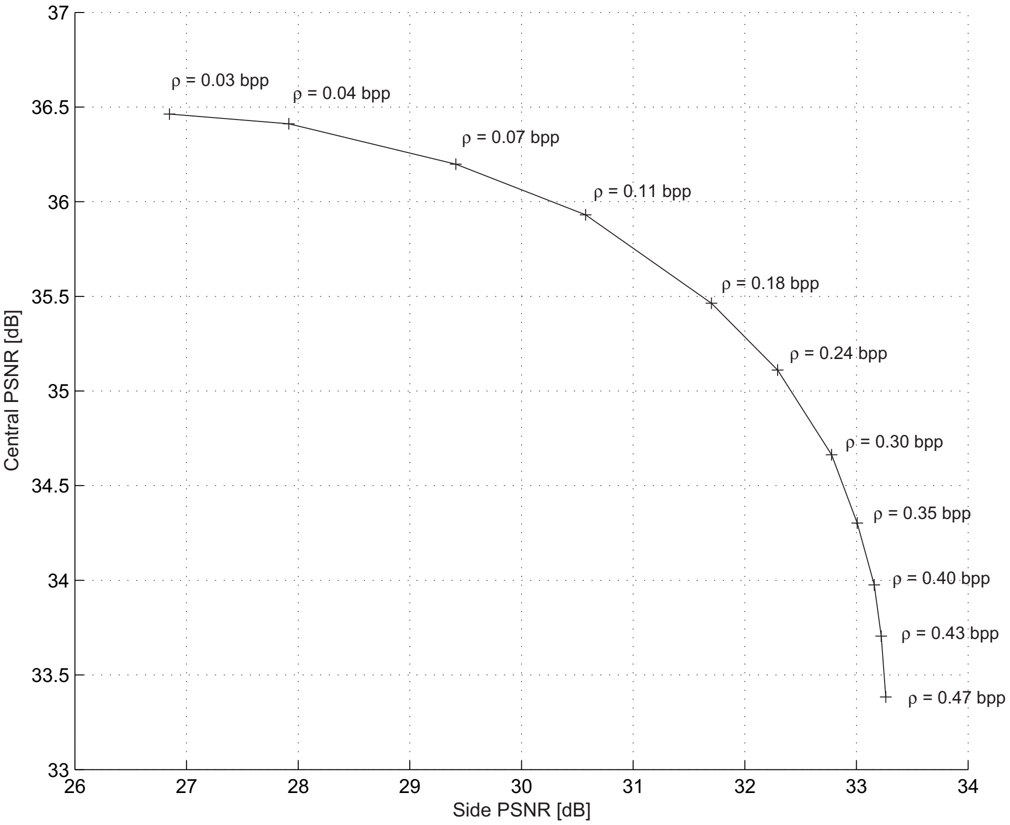

Figure 8.6: Performance of the proposed MD coder for the image Goldhill (total rate of 1 bpp). Also shown is the redundancy rate ρ corresponding to various points on the performance curve.

of redundancy (i.e., the rate dedicated for encoding the redundant polyphase components in both descriptions), for a fixed total rate (which also includes the redundancy rate, obviously). Each figure also shows the redundancy rate ρ corresponding to various points on the performance curve presented. As shown, for a fixed total rate, increasing the redundancy rate significantly improves the quality of the side reconstructions (i.e., yields higher side PSNR, or lower side distortion), at the expense of an increasingly degraded central reconstruction (i.e., lower central PSNR, or higher central distortion). Obviously, one also has to consider the relative importance of the side reconstructions, compared to that of the central reconstruction, which is generally related to the probability of channel failure (i.e., lost description). Thus, determining the optimal operating point for the system, namely, the optimal amount of redundancy for a fixed total rate, should be based on the properties of the communication channel. This will be described in Section 8.5.

Recall that our proposed MD coder aims to improve the performance of the original polyphase transform-based coding scheme of Jiang and Ortega [30], by utilizing the special statistical properties of the wavelet decomposition to improve the coding efficiency of the redundant polyphase component of each description (see Section 4.2). In other words, for a given quality of central reconstruction, the proposed coder aims to provide improved side reconstructions. As an example, Figure 8.7 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 (see Section 3.2), for the image Lena 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 will see in Section 8.4 that this improvement is due to the

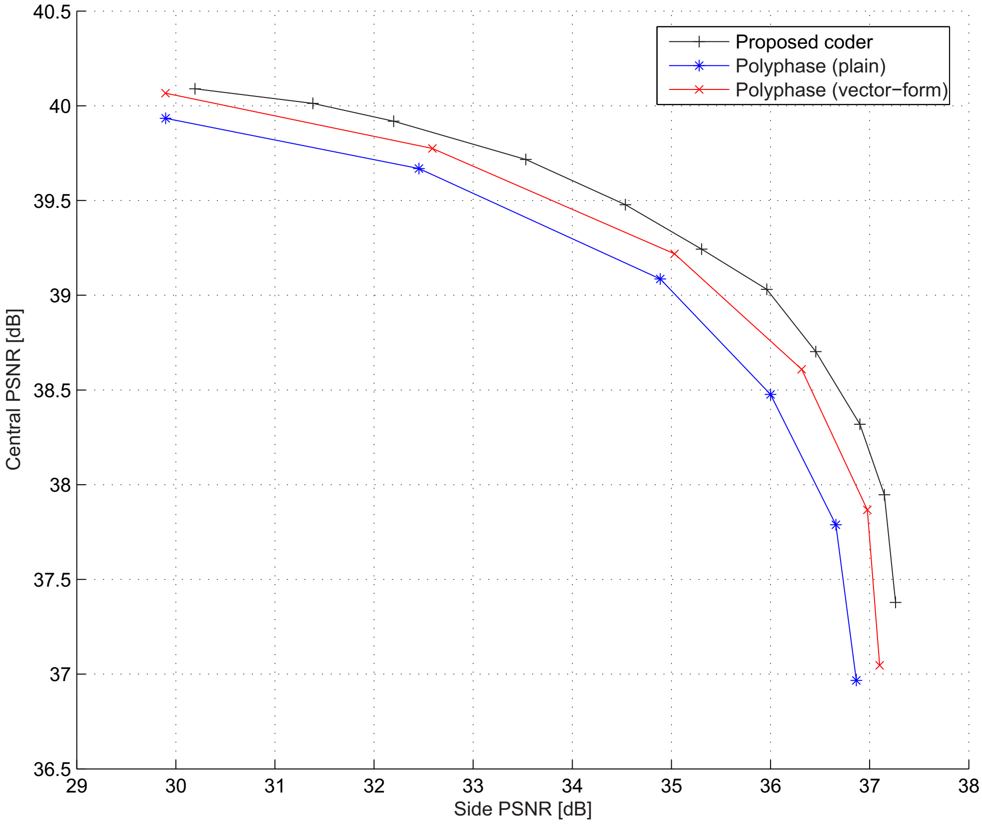

Figure 8.7: Performance of the proposed MD coder (in black), compared to that of the original polyphase transform-based MD coder [30], utilizing either the plain polyphase transform or the vector-form polyphase transform (for the image Lena and a total rate of 1 bpp).

effective utilization of contextual information, obtained from quantized wavelet coefficients of the primary polyphase component of each description, for improving the coding efficiency of the corresponding redundant polyphase component.

Figures 8.8 and 8.9 illustrate the subjective performance results obtained by the proposed MD coder, and also demonstrate the possible tradeoff offered by MD coding, compared to single description (SD) coding. Figure 8.8(a) shows again the original image Goldhill of size 512 × 512 pixels, which can be compressed, for example, to a rate of 1 bpp by a JPEG2000 coder[1] [55]—a state of the art SD wavelet coder—to yield a relatively high quality reconstructed image (PSNR of 36.27 dB), as shown in Figure 8.8(b). However, if two separate bitstreams ("descriptions") need to be created and transmitted, with the possibility that one of them will be lost, this SD coder does not offer a straightforward way to tradeoff the quality of the reconstruction from both bitstreams with that of the reconstruction from either bitstream alone. Obviously, as an MD coder, the proposed coder does offer such a tradeoff. For example, for the same total rate of 1 bpp, with a redundancy rate of 0.47 bpp (and two balanced descriptions), the proposed MD coder yields a good quality reconstruction from both descriptions (central PSNR of 33.38 dB), as shown in Figure 8.8(c), as well as from a single description (side PSNR of 33.26 dB), as shown in Figure 8.8(d).

Figure 8.9 shows a similar illustration for the image Lena of size 512 × 512 pixels (shown again in Figure 8.9(a)). In this case, for an encoding rate of 1 bpp, the JPEG2000 coder attains a PSNR of 40.03 dB (see Figure 8.9(b)). For the same total rate of 1 bpp, with a redundancy rate of 0.46 bpp, for example, the proposed MD coder attains a central PSNR of 37.38 dB (see Figure 8.9(c)), and an average side

---

[1] JPEG2000 compression results obtained using Adobe® Photoshop® CS2, with Adobe® JPEG2000 Plug-in v1.6.

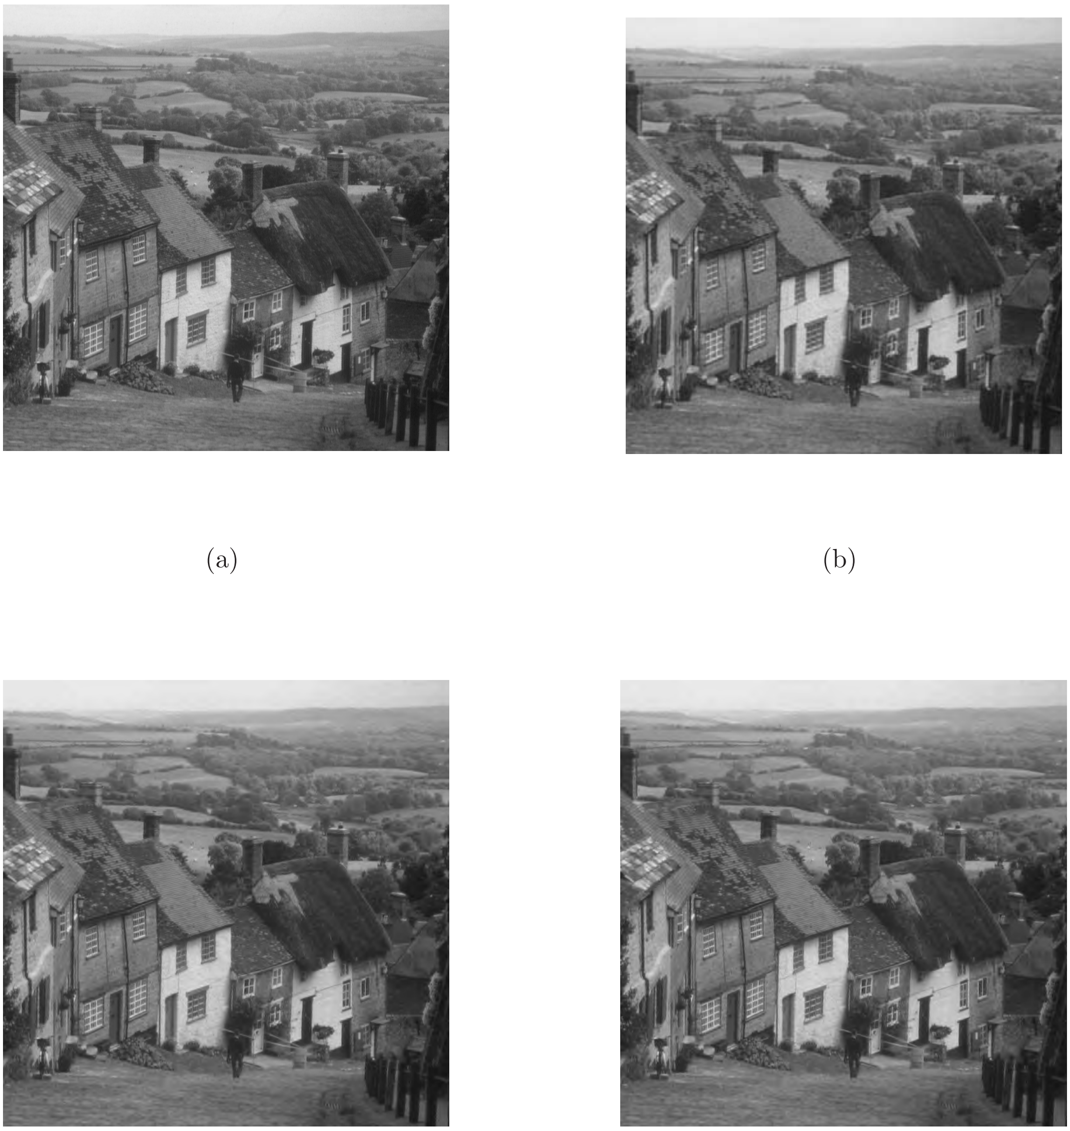

(a) (b) (c) (d)

Figure 8.8: Subjective performance results for the proposed MD coder, also in comparison to SD coding (for the image Goldhill). (a) Original image Goldhill. (b) Reconstructed using JPEG2000 (SD) coder (rate 1 bpp). (c) Central reconstruction by the proposed coder (total rate 1 bpp, redundancy rate 0.47 bpp). (d) Side reconstruction by the proposed coder (total rate 1 bpp, redundancy rate 0.47 bpp).

(a) (b)

(c) (d)

Figure 8.9: Subjective performance results for the proposed MD coder, also in comparison to SD coding (for the image Lena). (a) Original image Lena. (b) Reconstructed using JPEG2000 (SD) coder (rate 1 bpp). (c) Central reconstruction by the proposed coder (total rate 1 bpp, redundancy rate 0.46 bpp). (d) Side reconstruction by the proposed coder (total rate 1 bpp, redundancy rate 0.46 bpp).

PSNR of 37.26 dB (see Figure 8.9(d)).

In order to demonstrate the effect of the number of classes used by the context-based classification procedure on the performance of the proposed coder, Figure 8.10 shows the performance obtained by the coder for various choices of the total number of classes per subband (for the image Lena and a total rate of 1 bpp). As shown, the best overall performance is achieved using four classes per subband, a choice which we found to be robust also for other rates and other test images. Note that as the

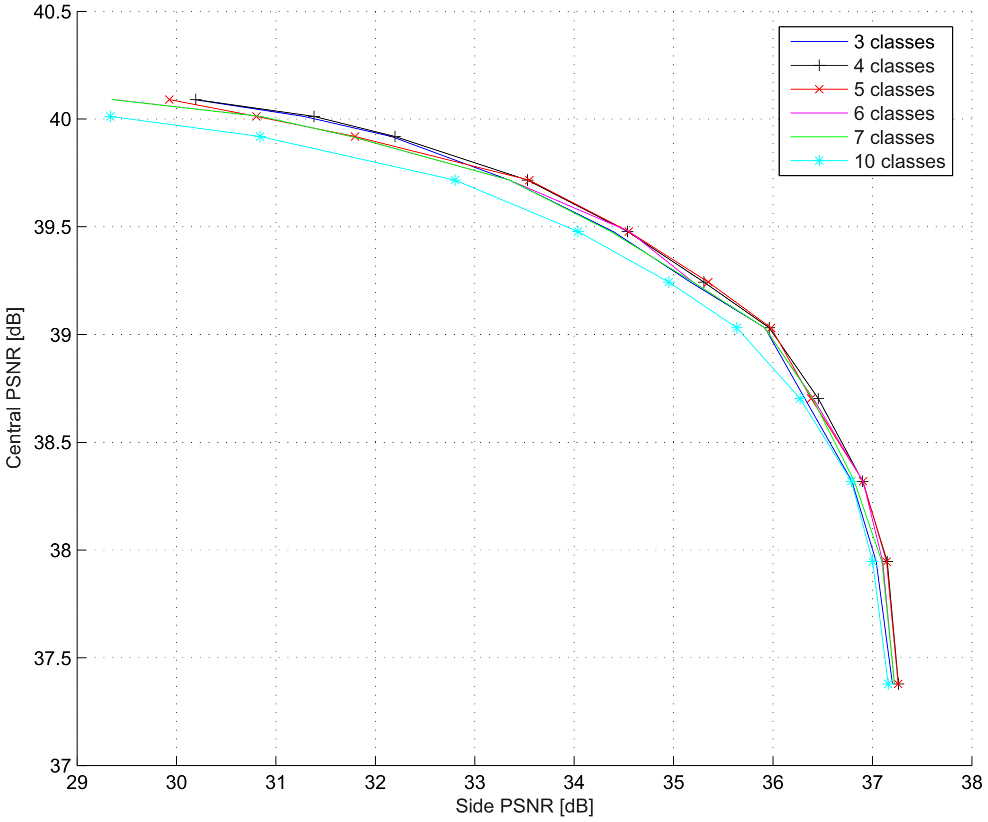

Figure 8.10: Performance of the proposed MD coder for various choices of the total number of classes per subband (for the image Lena and a total rate of 1 bpp).

number of classes is increased, so does the overhead due to side information (as well as the complexity of system), and thus using an excessively large number of classes results in a degradation in the performance of the system.

As an alternative to uniform threshold quantizers (UTQs), we also experimented with uniform reconstruction with unity ratio quantizers (URURQs) [53] for the Laplacian distribution. As described in Section 6.2, the rate distortion performance of the URURQ is very close to that of the optimum entropy-constrained scalar quantizer (ECSQ) for the Laplacian distribution, assuming MSE distortion. Compared to the UTQ, the URURQ does offer some improvement in performance, but this improvement is quite small. Nevertheless, it is interesting to examine how this slight improvement in theory (i.e., for simulated Laplacian data), translates into practice (i.e., in terms of overall performance, with quantization of real wavelet data). Figure 8.10 shows the performance of the proposed MD coder, with utilization of URURQs for quantization of the wavelet coefficients of the redundant polyphase component of each description, compared to utilization of UTQs (for the image Lena and a total rate of 1 bpp). As shown, the overall performance of the proposed coder is practically the same for both types of quantizers. Thus, no significant gain is achieved by the utilization of URURQs instead of UTQs. Nevertheless, since both types of quantizers are relatively simple to design and operate, the rarely used URURQ should also be considered as a viable option for quantization of wavelet coefficients, modeled using the Laplacian distribution, in various coding applications.

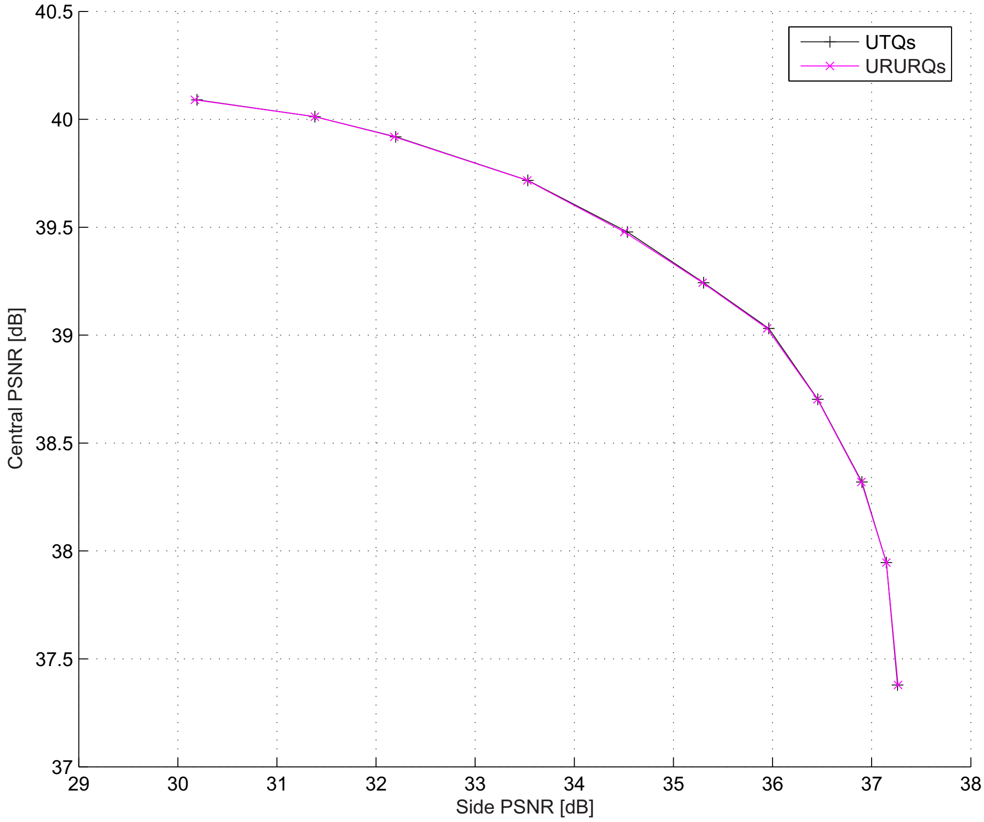

Figure 8.11: Performance of the proposed MD coder, with utilization of URURQs, compared to utilization of UTQs (for the image Lena and a total rate of 1 bpp).

8.4 Context Gain

The proposed MD coder aims to improve the performance of the original polyphase transform-based coding scheme of Jiang and Ortega [30], by expoliting 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.

As demonstrated earlier in this chapter, the proposed coder indeed attains improved side reconstructions for a given quality of central reconstruction (across the entire redundancy range). Since the encoding path of the redundant polyphase component of each description is comprised of various coding components (as described above), it is interesting to examine whether the achieved improvement in coding efficiency can indeed be attributed primarily to the effective utilization of contextual information. In the sequel, 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 default configuration (see Section 8.1), 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. Other than the fact that the "no-context coder" uses a single class for each subband, all other coding procedures remain the same as in the proposed coder. Specifically, the "no-context coder" still performs optimal bit allocation (between the various subbands).

Figure 8.12 shows the performance of the "no-context coder", compared to that of the proposed coder in its default 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

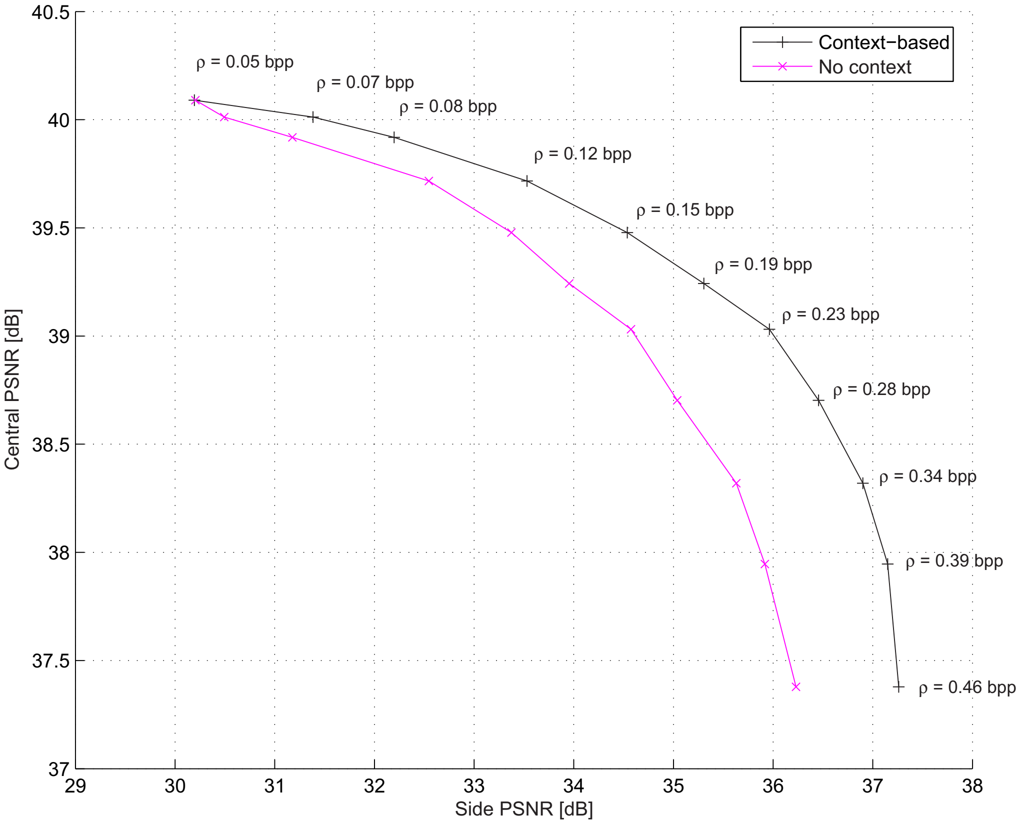

Figure 8.12: Performance of the "no-context coder", compared to that of the proposed context-based coder in its default configuration (for the image Lena and a total rate of 1 bpp). Also shown is the redundancy rate ρ corresponding to various points on the performance curves.

“no-context coder”. As expected, at very low redundancies, the context gain is small. The reason for this is that a very low redundancy means that only few bits are used to encode the redundant polyphase component of each description. Thus, although the contextual information does exist, it cannot be utilized effectively, since the representation of the redundant polyphase components is too coarse. As the redundancy increases, and more bits are available for the encoding of the redundant polyphase component of each description, the context gain rapidly becomes significant (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.

8.5 Determining the Optimal Operating Point Based on Channel Properties

This section describes how to determine the optimal operating point for the proposed MD coder, based on the properties of the communication channel. Namely, given the channel model, we wish to determine the optimal amount of redundancy such that the average distortion attained is minimized, subject to a total rate constraint.

We consider the following commonly used model for the communication channel:

- Descriptions are sent over two independent channels.
- Each channel fails with probability p.

Thus, four different situations at the receiver are possible, each with its corresponding probability:

1. Both descriptions are received—with probability $(1-p)^2$.
2. Only description 1 is received—with probability $p(1-p)$.
3. Only description 2 is received—with probability $p(1-p)$.
4. None of the descriptions is received—with probability p^2.

For the balanced case (descriptions of statistically equal rate and importance), instead of considering each of the side distortions individually, we can consider the average side distortion D_s. Thus, the average distortion at the receiver is given by

$$\bar{D} = (1-p)^2 D_c + 2p(1-p) D_s + p^2 D_{\text{none}} \tag{8.1}$$

where D_c is the central distortion, and D_{none} is the distortion attained if none of the descriptions is received.

Our problem is to determine the optimal amount of redundancy such that the average distortion $\bar{D}$ is minimized, subject to a total rate constraint. Since D_{none} is not affected by the amount of redundancy, our problem reduces to minimization, subject to a total rate constraint, of

$$D_c + \alpha D_s \tag{8.2}$$

where

$$\alpha = \frac{2p(1-p)}{(1-p)^2} = \frac{2p}{(1-p)}\ . \tag{8.3}$$

Note that α is fixed (given the channel failure probability p).

As an example, consider Figure 8.13, which shows the performance of the proposed MD coder (for the image Lena and a total rate of 1 bpp), this time with the vertical and horizontal axes corresponding to the central distortion and average side distortion,

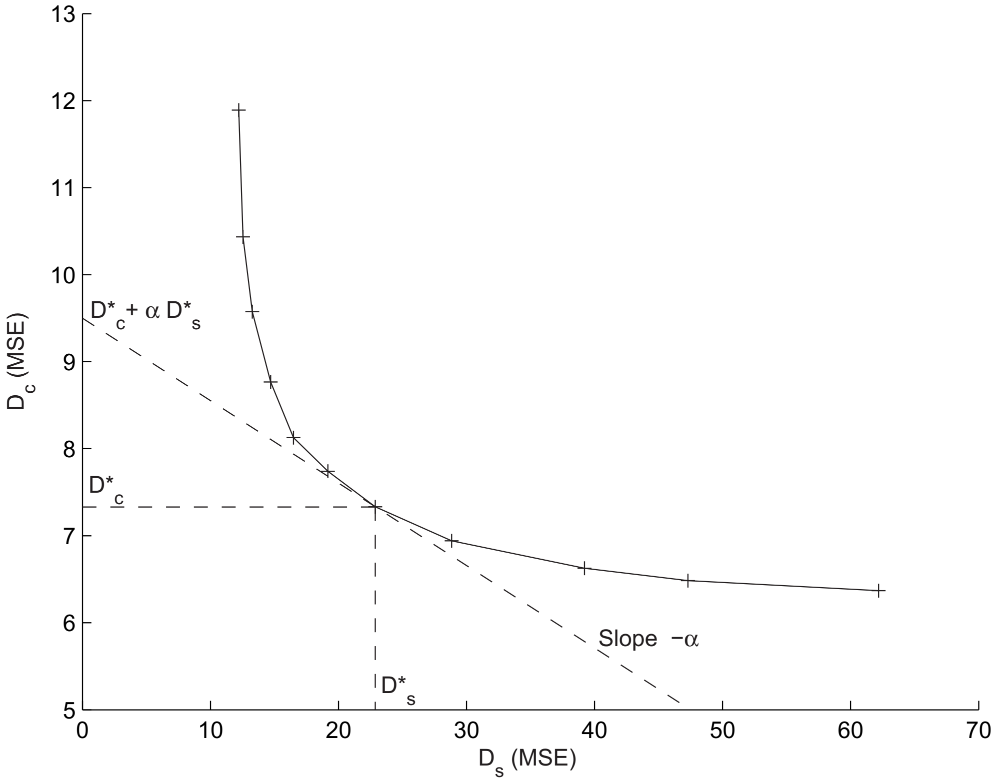

Figure 8.13: Performance of the proposed MD coder, with the vertical and horizontal axes corresponding to the central distortion and average side distortion, respectively (for the image Lena and a total rate of 1 bpp).

respectively (instead of the corresponding PSNR, as before). Since every point on the performance curve shown satisfies the total rate constraint (1 bpp in this example), it is left to find the point on the curve that minimizes $D_c + \alpha D_s$. Due to the convexity of the curve, the point that minimizes $D_c + \alpha D_s$ is the point on the curve that is "hit" first by a "plane wave" of slope α (in absolute value), as illustrated graphically in Figure 8.13. Note that $p \to 0^+$ (i.e., a description is lost with very low probability) yields $\alpha \to 0$, which corresponds to minimization of the central distortion, while $p \to 1^-$ (i.e., a description is lost with very high probability) yields $\alpha \to \infty$, which corresponds to minimization of the side distortions.

Practically, since the proposed coder is able to achieve any desired total rate with various amounts of redundancy, finding the optimal operating point that minimizes $D_c + \alpha D_s$ is essentially a one-dimensional (1-D) minimization problem (where the redundancy is the independent variable). Therefore, various methods for minimization of a one-dimensional function can be utilized to find the optimal operating point, such as the golden section search [42, Ch. 10].

Chapter 9

Conclusion

9.1 Summary

Among prior works, MDs for image coding were generated via the utilization of a polyphase transform and selective quantization, performed in the wavelet domain. Our research work explored 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, 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 have used an effective context-based classification procedure. 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 still produced 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 have examined two types of quantizers: the uniform threshold quantizer (UTQ) and the uniform reconstruction with unity ratio quantier (URURQ). Both of these quantizers well approximate the optimum entropy-constrained scalar quantizer (ECSQ) for the Laplacian distribution, assuming 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 have proposed 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 have also derived closed form expressions for the distortions attained by both the UTQ and the URURQ.

For bit allocation between the various classes in the different subbands, we have developed 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 online bit allocation procedures, as well as to intelligently and instantly adapt the arithmetic entropy encoder to the varying coefficients statistics. This adaptation scheme enables the entropy encoder to fully exploit the higher level statistics captured by the Laplacian model-based classification and quantization algorithms, rather than using ad hoc updating rules, as done in many other entropy coding adaptation schemes.

We provided various experimental results that clearly demonstrate the advantages of the proposed MD coder. Specifically, it was also shown that the proposed coder outperforms its predecessor—the polyphase transform-based coder of Jiang and Ortega—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. We also provided a way to optimally determine, in practice, the operating point for the proposed MD coder, based on the properties of the communication channel.

Finally, we express our hope that parts of this thesis will also be able to serve as a unified and clear reference for different aspects related to efficient implementation of wavelet coders. To this end, we comprehensively addressed various relevant issues in this thesis, such as how to avoid coefficient expansion and unpleasant boundary artifacts by properly choosing the wavelet filters, along with the appropriate type of boundary extension, as well as how to generally perform optimal bit allocation in the case of a biorthogonal wavelet transform.

9.2 Main Contributions

In this section we list the main contributions of this research work.

- **Providing a novel context-based coding scheme that enables effective utilization of the special statistical properties of the wavelet decomposition for improved MD coding efficiency.** Our experimental results also demonstrate the superiority of the resulting context-based coder over its predecessor, and prove that the improvement in coding efficiency can indeed be attributed primarily to the effective utilization of contextual information.

- **Developing an efficient design strategy for quantization, using either UTQs or URURQs, based on a pre-designed array of corresponding MSE-optimized quantizers of different step sizes for the Laplacian distribution.** Particularly, we also derive closed form expressions for the distortions attained by both the UTQ and the URURQ, in order to further reduce the complexity of the proposed design strategy.

- **Developing an optimal and efficient model-based bit allocation scheme, in the general framework of Lagrangian optimization.** Specifically, the proposed bit allocation scheme also takes into account the non-energy preserving nature of the biorthogonal wavelet transform, if such a transform is used.

9.3 Future Directions

In this section we briefly describe several possible directions for future research.

Reverse Context-Based MD Coding System (with Combined Contexts)

Recall that in the original polyphase transform-based MD coding system [30], as well as in the proposed MD coding system, the purpose of the redundant information in each description is to enable an acceptable reconstruction of the lost polyphase component, in the case of a channel failure. However, if both descriptions arrive, this redundant information is of no use, unfortunately.

In order to utilize this redundant information in any case (i.e., whether descriptions are lost or not), and thus obtain an additional coding gain, we suggest to use a "reverse" context-based MD coding system, as illustrated in Figure 9.1. Compared to Figure 4.7, the coding scheme presented here utilizes context-based quantization in coding the *primary* information. Its main advantage is that if both descriptions arrive, the redundant information is still utilized, by providing contextual information

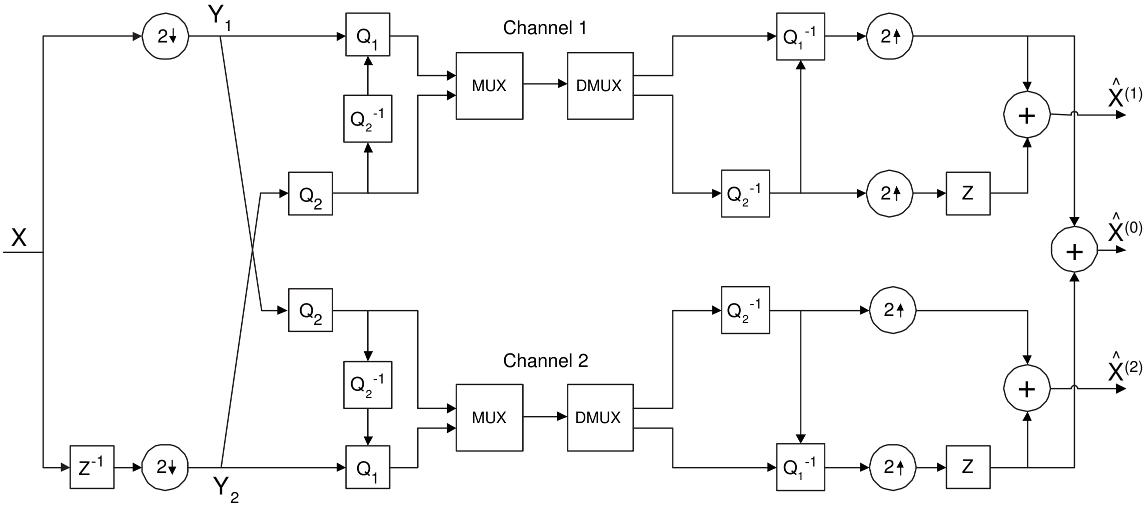

Figure 9.1: Polyphase transform-based MD coding system that utilizes context-based quantization in coding the primary information (reverse context-based MD coding system).

to improve the efficiency in encoding the primary information.

In order to perform well across the entire redundancy range, we additionally suggest the use of a weighted *combined* context for Q_1, which is formed from quantized coefficients of the corresponding redundant polyphase component (via Q_2^{-1}), as well as from already quantized coefficients of the primary polyphase component. To optimize the performance, the weights corresponding to the two sets should be determined according to the relative resolutions of Q_2 and Q_1, or, in other words, based on the relative amount of redundancy.

Utilization of Across-Scale Dependencies

Compared to early wavelet image coders, modern wavelet coders usually utilize the multiresolution nature of the wavelet decomposition in order to achieve improved coding efficiency (e.g., [49, 45]). As also described in Section 4.1.2, experiments demonstrate a significant statistical dependence between a wavelet coefficient and the coefficient corresponding to the same spatial location at the next coarser scale (i.e., its parent). Other scale-to-scale dependencies in the wavelet decomposition also exist (even if to a lesser extent).

Consequently, in the framework of our context-based MD coding system, it would be highly beneficial to use polyphase transforms that enable the formation of contexts utilizing both intra- and inter-subband statistical dependencies. Such polyphase transforms and corresponding contexts should therefore be sought after.

Extensions to More Than Two Descriptions

Many MD coding schemes focus on the two-description case, mainly due to its relative tractability, and so does our coder. Nevertheless, it is possible to extend the proposed context-based MD coding scheme to the case of more than two descriptions.

Note that the original polyphase transform-based MD coding system [30] can be adapted to the case of more than two descriptions relatively easily, due to its explicit separation between description generation and redundancy addition, as detailed in [30]. For example, to protect an arbitrary number of descriptions, a polyphase transform with a suitable number of polyphase components can be used, and the redundancy between the descriptions can be carried in a sequential (cyclic) way (i.e., description 1 carries redundancy to protect description 2, while description 2 carries redundancy to protect description 3, and so on). Naturally, for the proposed context-based MD coder, the extraction of contextual information should be adapted accordingly.

For context-based MD coding, a systematic and effective co-design of the polyphase transform and the context formation procedure, offers an additional potential coding gain, especially in the case of an arbitrary number of descriptions. Efficient utilization of contextual information, using the approaches presented in this work, among others, is expected to yield improved coding performance, also in the case of more than two descriptions.

Appendix A

Proof that the Classification Gain G Satisfies G $\geq$ 1

In this appendix we prove that the classification gain G of (5.13), given by

$$G = \frac{\hat{\lambda}_k^{2r_k} \hat{\lambda}_{k+1}^{2r_{k+1}}}{\hat{\lambda}_{k'}^2} , \tag{A.1}$$

satisfies $G \geq 1$, with equality if and only if $\hat{\lambda}_k = \hat{\lambda}_{k+1}$. Recall from Section 5.3 that $r_k = n_k/(n_k + n_{k+1})$ and $r_{k+1} = 1 - r_k = n_{k+1}/(n_k + n_{k+1})$, where n_k and n_{k+1} are the number of coefficients in two adjacent classes C_k and C_{k+1}. Also recall that $\hat{\lambda}_k$, $\hat{\lambda}_{k+1}$ and $\hat{\lambda}_{k'}$ are the maximum likelihood estimators of the Laplacian parameters of the classes C_k, C_{k+1} and the union of the two classes, respectively.

We first prove the following inequality of weighted arithmetic and weighted geometric means.

Lemma A.1 *Let x_1 and x_2 be positive real numbers. Also, let the weights w_1 and w_2 be positive real numbers, and set $w = w_1 + w_2$. Then the following inequality holds*

$$\frac{w_1}{w} x_1 + \frac{w_2}{w} x_2 \geq \sqrt[w]{x_1^{w_1} x_2^{w_2}} , \tag{A.2}$$

with equality if and only if $x_1 = x_2$*.*

Proof First, if $x_1 = x_2$ then both sides of the inequality reduce to x_1, and equality holds.

Now assume that $x_1 \neq x_2$. Since the natural logarithm is a strictly concave function, it follows that

$$\log\left(\frac{w_1}{w}x_1 + \frac{w_2}{w}x_2\right) > \frac{w_1}{w}\log x_1 + \frac{w_2}{w}\log x_2 \,. \tag{A.3}$$

We therefore get

$$\log\left(\frac{w_1}{w}x_1 + \frac{w_2}{w}x_2\right) > \log\left(\sqrt[w]{x_1^{w_1}x_2^{w_2}}\right) \tag{A.4}$$

and since the natural logarithm is a strictly increasing function, it follows that

$$\frac{w_1}{w}x_1 + \frac{w_2}{w}x_2 > \sqrt[w]{x_1^{w_1}x_2^{w_2}} \,. \tag{A.5}$$

□

We now prove that the classification gain G satisfies $G \geq 1$, with equality if and only if $\hat{\lambda}_k = \hat{\lambda}_{k+1}$. From (5.12)

$$\frac{1}{\hat{\lambda}_{k'}} = \frac{n_k}{n_k + n_{k+1}} \cdot \frac{1}{\hat{\lambda}_k} + \frac{n_{k+1}}{n_k + n_{k+1}} \cdot \frac{1}{\hat{\lambda}_{k+1}} = r_k\frac{1}{\hat{\lambda}_k} + r_{k+1}\frac{1}{\hat{\lambda}_{k+1}} \tag{A.6}$$

which is the weighted arithmetic mean of $1/\hat{\lambda}_k$ and $1/\hat{\lambda}_{k+1}$. Using Lemma A.1 we therefore get

$$\frac{1}{\hat{\lambda}_{k'}} \geq \frac{1}{\hat{\lambda}_k^{r_k}} \cdot \frac{1}{\hat{\lambda}_{k+1}^{r_{k+1}}} \tag{A.7}$$

with equality if and only if $\hat{\lambda}_k = \hat{\lambda}_{k+1}$. Rearranging (A.7) and squaring both sides of the inequality yields

$$\frac{\hat{\lambda}_k^{2r_k}\hat{\lambda}_{k+1}^{2r_{k+1}}}{\hat{\lambda}_{k'}^2} \geq 1 \tag{A.8}$$

with equality if and only if $\hat{\lambda}_k = \hat{\lambda}_{k+1}$. From (A.1) it then follows that the classification gain G satisfies $G \geq 1$, with equality if and only if $\hat{\lambda}_k = \hat{\lambda}_{k+1}$.

Appendix B

Reconstruction Offsets of the Optimal Uniform Threshold Quantizer for the Laplacian Distribution

An expression for the reconstruction levels of a uniform threshold quantizer (UTQ) for the Laplacian distribution, optimized for minimum MSE distortion, was derived in Section 6.1. In this appendix we prove that some of the reconstruction offsets of this UTQ are equal.

Using the same notation as in Section 6.1, let $\lambda\,(> 0)$ be the parameter of the Laplacian distribution, let Δ be the step size of the UTQ, and denote the number of levels of the UTQ by $N = 2L + 1$. The bin boundaries $b_j,\ j = 1, \ldots, L$, of the UTQ are given in (6.5), and the bin probabilities $p_j,\ j = 1, \ldots, L-1$, are given in (6.13).

From (6.14), we then get for the reconstruction levels q_j, $j = 1, \ldots, L-1$:

$$\begin{aligned}
q_j &= \frac{1}{2p_j}\left[b_j e^{-\lambda b_j} - b_{j+1} e^{-\lambda b_{j+1}}\right] + \frac{1}{\lambda} = \\
&= \frac{1}{e^{-\lambda b_j} - e^{-\lambda b_{j+1}}}\left[(j-\frac{1}{2})\Delta e^{-\lambda b_j} - (j - \frac{1}{2} + 1)\Delta e^{-\lambda b_{j+1}}\right] + \frac{1}{\lambda} = \\
&= \frac{1}{e^{-\lambda b_j} - e^{-\lambda b_{j+1}}}\left[(j-\frac{1}{2})\Delta\left(e^{-\lambda b_j} - e^{-\lambda b_{j+1}}\right) - \Delta e^{-\lambda b_{j+1}}\right] + \frac{1}{\lambda} = \\
&= (j-\frac{1}{2})\Delta - \Delta\frac{e^{-\lambda b_{j+1}}}{e^{-\lambda b_j} - e^{-\lambda b_{j+1}}} + \frac{1}{\lambda} = \\
&= (j-\frac{1}{2})\Delta - \Delta\frac{e^{-\lambda\left(j+\frac{1}{2}\right)\Delta}}{e^{-\lambda\left(j-\frac{1}{2}\right)\Delta} - e^{-\lambda\left(j+\frac{1}{2}\right)\Delta}} + \frac{1}{\lambda} = \\
&= (j-\frac{1}{2})\Delta - \frac{\Delta}{e^{-\lambda\left(j-\frac{1}{2}\right)\Delta + \lambda\left(j+\frac{1}{2}\right)\Delta} - 1} + \frac{1}{\lambda} = \\
&= (j-\frac{1}{2})\Delta - \frac{\Delta}{e^{\lambda\Delta} - 1} + \frac{1}{\lambda} = \\
&= b_j + \frac{1}{\lambda} - \frac{\Delta}{e^{\lambda\Delta} - 1}\,. \qquad \text{(B.1)}
\end{aligned}$$

Thus, the reconstruction offsets $q_j - b_j$, for $j = 1, \ldots, L-1$, do not depend on j. Namely,

$$q_j = b_j + \delta, \quad j = 1, \ldots, L-1\,, \qquad \text{(B.2)}$$

where

$$\delta = \frac{1}{\lambda} - \frac{\Delta}{e^{\lambda\Delta} - 1}\,. \qquad \text{(B.3)}$$

Appendix C

Proof of a Useful Property of the Laplacian Distribution

The following useful property of the Laplacian distribution is proved in this appendix.

Property C.1 *Let X be a Laplacian random variable with parameter $\lambda\,(> 0)$, i.e.,*

$$f_X(x) = \frac{\lambda}{2} e^{-\lambda|x|}, \;\; -\infty < x < \infty \tag{C.1}$$

and let c be a fixed nonnegative real number. Then given that $X \geq c$, the conditional distribution of $Y = X - c$ is exponential with parameter λ, i.e.,

$$f_{Y|X\geq c}(y|X \geq c) = \lambda e^{-\lambda y}, \;\; y \geq 0\,. \tag{C.2}$$

Proof The conditional density function of X, conditioned on $X \geq c$, is given by

$$\begin{aligned}
f_{X|X\geq c}(x|X \geq c) &= \begin{cases} f_X(x)/\int_c^\infty f_X(\alpha)\,d\alpha, & x \geq c \\ 0, & \text{else} \end{cases} = \\
&= \begin{cases} \frac{\lambda}{2}e^{-\lambda x}/\int_c^\infty \frac{\lambda}{2}e^{-\lambda\alpha}\,d\alpha, & x \geq c \\ 0, & \text{else} \end{cases} = \\
&= \begin{cases} e^{-\lambda x}/\left(\frac{1}{\lambda}e^{-\lambda c}\right), & x \geq c \\ 0, & \text{else} \end{cases} = \\
&= \begin{cases} \lambda e^{-\lambda(x-c)}, & x \geq c \\ 0, & \text{else} \end{cases} . \qquad \text{(C.3)}
\end{aligned}$$

Defining $Y = X - c$, we get

$$f_{Y|X\geq c}(y|X \geq c) = \lambda e^{-\lambda y},\ \ y \geq 0\,. \qquad \text{(C.4)}$$

That is, given that $X \geq c$, the conditional distribution of $Y = X - c$ is exponential with parameter λ. □

Appendix D

Squared Error Distortion of the Optimal Uniform Threshold Quantizer for the Laplacian Distribution

An expression for the average squared error distortion of a uniform threshold quantizer (UTQ) for the Laplacian distribution, optimized for minimum distortion, was derived in Section 6.1. In this appendix we further develop this expression to obtain a closed form expression for the average distortion D, whose evaluation does not involve integration.

As in Section 6.1, let Δ be the step size of the UTQ. The bin boundaries $\{b_j\}$ of the UTQ are given in (6.5), and the reconstruction levels $\{q_j\}$ are given in (6.12) and

(6.15). According to (6.17), the average distortion of the UTQ is then given by

$$\begin{aligned}
D &= \int_{-b_1}^{b_1} f_X(x)(x-0)^2\,dx + 2\sum_{j=1}^{L-1}\int_{b_j}^{b_{j+1}} f_X(x)(x-q_j)^2\,dx + 2\int_{b_L}^{\infty} f_X(x)(x-q_L)^2\,dx = \\
&= 2\left[\int_0^{b_1} x^2 f_X(x)\,dx + \sum_{j=1}^{L-1}\int_{b_j}^{b_{j+1}}(x-q_j)^2 f_X(x)\,dx + \int_{b_L}^{\infty}(x-q_L)^2 f_X(x)\,dx\right] = \\
&= 2\left[\int_0^{b_1} x^2\frac{\lambda}{2}e^{-\lambda x}\,dx + \sum_{j=1}^{L-1}\int_{b_j}^{b_{j+1}}(x-q_j)^2\frac{\lambda}{2}e^{-\lambda x}\,dx + \int_{b_L}^{\infty}(x-q_L)^2\frac{\lambda}{2}e^{-\lambda x}\,dx\right] = \\
&= \lambda\left[\int_0^{b_1} x^2 e^{-\lambda x}\,dx + \sum_{j=1}^{L-1}\int_{b_j}^{b_{j+1}}(x-q_j)^2 e^{-\lambda x}\,dx + \int_{b_L}^{\infty}(x-q_L)^2 e^{-\lambda x}\,dx\right] \qquad \text{(D.1)}
\end{aligned}$$

where $\lambda\,(>0)$ is the parameter of the Laplacian distribution, and the number of levels of the quantizer is $N = 2L+1$.

To calculate the above integrals, we use the following indefinite integral, obtained directly using integration by parts:

$$\int x^2 e^{-\lambda x}\,dx = -\frac{1}{\lambda}e^{-\lambda x}\left(x^2+\frac{2x}{\lambda}+\frac{2}{\lambda^2}\right)+C \qquad \text{(D.2)}$$

where C is an arbitrary constant of integration. The first term in the brackets of (D.1) then becomes

$$\begin{aligned}
\int_0^{b_1} x^2 e^{-\lambda x}\,dx &= \frac{1}{\lambda}\left[e^{-\lambda x}\left(x^2+\frac{2x}{\lambda}+\frac{2}{\lambda^2}\right)\right]_{b_1}^{0} = \\
&= \frac{1}{\lambda}\left[\frac{2}{\lambda^2}-e^{-\lambda b_1}\left(b_1^2+\frac{2b_1}{\lambda}+\frac{2}{\lambda^2}\right)\right] = \\
&= \frac{1}{\lambda}\left[\frac{2}{\lambda^2}-e^{-\lambda\frac{\Delta}{2}}\left(\left(\frac{\Delta}{2}\right)^2+\frac{\Delta}{\lambda}+\frac{2}{\lambda^2}\right)\right]. \qquad \text{(D.3)}
\end{aligned}$$

The second term in the brackets of (D.1) is calculated in a similar manner, along with the substitution of variables $t = x - q_j$ (translation), as follows. From (6.15) we get

$$q_j - b_j = \delta, \quad j = 1,\ldots,L-1\,, \qquad \text{(D.4)}$$

where

$$\delta = \frac{1}{\lambda} - \frac{\Delta}{e^{\lambda\Delta} - 1} \,. \tag{D.5}$$

We thus obtain for the second term

$$\begin{aligned}
\sum_{j=1}^{L-1} \int_{b_j}^{b_{j+1}} (x-q_j)^2 e^{-\lambda x}\, dx &= \sum_{j=1}^{L-1} \int_{b_j - q_j}^{b_{j+1}-q_j} t^2 e^{-\lambda(t+q_j)}\, dt = \\
&= \sum_{j=1}^{L-1} e^{-\lambda q_j} \int_{-\delta}^{-\delta+\Delta} t^2 e^{-\lambda t}\, dt = \\
&= \sum_{j=1}^{L-1} e^{-\lambda q_j} \frac{1}{\lambda} \left[e^{-\lambda t} \left(t^2 + \frac{2t}{\lambda} + \frac{2}{\lambda^2} \right) \right]_{-\delta+\Delta}^{-\delta} = \\
&= \frac{1}{\lambda} \left(\sum_{j=1}^{L-1} e^{-\lambda q_j} \right) \cdot \left[e^{\lambda\delta} \left(\delta^2 - \frac{2\delta}{\lambda} + \frac{2}{\lambda^2} \right) + \right. \\
&\quad \left. - e^{\lambda(\delta-\Delta)} \left((\delta-\Delta)^2 - \frac{2(\delta-\Delta)}{\lambda} + \frac{2}{\lambda^2} \right) \right] .
\end{aligned} \tag{D.6}$$

Similarly, for the third term in the brackets of (D.1), using (6.12) and the substitution of variables $t = x - q_L$, we get

$$\begin{aligned}
\int_{b_L}^{\infty} (x-q_L)^2 e^{-\lambda x}\, dx &= \int_{b_L - q_L}^{\infty} t^2 e^{-\lambda(t+q_L)}\, dt = e^{-\lambda q_L} \int_{-\frac{1}{\lambda}}^{\infty} t^2 e^{-\lambda t}\, dt = \\
&= e^{-\lambda q_L} \frac{1}{\lambda} \left[e^{-\lambda t} \left(t^2 + \frac{2t}{\lambda} + \frac{2}{\lambda^2} \right) \right]_{\infty}^{-\frac{1}{\lambda}} = \\
&= \frac{1}{\lambda} e^{-\lambda q_L} \cdot e \cdot \left(\frac{1}{\lambda^2} - \frac{2}{\lambda^2} + \frac{2}{\lambda^2} \right) = \\
&= \frac{1}{\lambda^3} e^{1-\lambda q_L} \,.
\end{aligned} \tag{D.7}$$

From (D.1), (D.3), (D.6) and (D.7), we finally obtain a closed form expression for the average distortion of the UTQ:

$$\begin{aligned} D &= \frac{2}{\lambda^2} - e^{-\lambda\frac{\Delta}{2}}\left(\left(\frac{\Delta}{2}\right)^2 + \frac{\Delta}{\lambda} + \frac{2}{\lambda^2}\right) + \\ &\quad + \left(\sum_{j=1}^{L-1} e^{-\lambda q_j}\right) \cdot \left[e^{\lambda\delta}\left(\delta^2 - \frac{2\delta}{\lambda} + \frac{2}{\lambda^2}\right) - e^{\lambda(\delta-\Delta)}\left((\delta-\Delta)^2 - \frac{2(\delta-\Delta)}{\lambda} + \frac{2}{\lambda^2}\right)\right] + \\ &\quad + \frac{1}{\lambda^2} e^{1-\lambda q_L} \end{aligned} \tag{D.8}$$

where δ is given in (D.5), and where

$$q_j = b_j + \delta = (j - \frac{1}{2})\Delta + \delta, \quad j = 1, \ldots, L-1 \tag{D.9}$$

and

$$q_L = b_L + \frac{1}{\lambda} = (L - \frac{1}{2})\Delta + \frac{1}{\lambda}\,. \tag{D.10}$$

Appendix E

Squared Error Distortion of the Uniform Reconstruction with Unity Ratio Quantizer for the Laplacian Distribution

The uniform reconstruction with unity ratio quantizer (URURQ) for the Laplacian distribution was presented in Section 6.2. In this appendix we derive a closed form expression for the average squared error distortion D of the URURQ, whose evaluation does not involve integration.

Using the same notation as in Section 6.2, let $\lambda\,(> 0)$ be the parameter of the Laplacian distribution, let Δ be the step size of the URURQ, and denote the number of levels of the URURQ by $N = 2L+1$. Following Appendix D, the average distortion of the URURQ is given by

$$D = \lambda \left[\int_0^{b_1} x^2 e^{-\lambda x}\, dx + \sum_{j=1}^{L-1} \int_{b_j}^{b_{j+1}} (x - q_j)^2 e^{-\lambda x}\, dx + \int_{b_L}^{\infty} (x - q_L)^2 e^{-\lambda x}\, dx \right] \quad \text{(E.1)}$$

where the bin boundaries $\{b_j\}$ and the reconstruction levels $\{q_j\}$ of the URURQ are defined in Section 6.2. Note that since the number of levels is finite, we optimize the outermost reconstruction levels (i.e., q_L and q_{-L}) for minimum MSE distortion, according to the well-known centroid condition. This yields

$$q_L = E\left[X | X > b_L\right] = b_L + E\left[X - b_L | X > b_L\right] = b_L + \frac{1}{\lambda} \quad \text{(E.2)}$$

where the last equality follows from the fact that given that $X > b_L$, the conditional distribution of $X - b_L$ is exponential with parameter λ (see Appendix C), and thus has a mean of $1/\lambda$.

Using the indefinite integral (D.2), the first term in the brackets of (E.1) is given by

$$\begin{aligned} \int_0^{b_1} x^2 e^{-\lambda x}\, dx &= \frac{1}{\lambda} \left[e^{-\lambda x} \left(x^2 + \frac{2x}{\lambda} + \frac{2}{\lambda^2} \right) \right]_{b_1}^{0} = \\ &= \frac{1}{\lambda} \left[\frac{2}{\lambda^2} - e^{-\lambda b_1} \left(b_1^2 + \frac{2b_1}{\lambda} + \frac{2}{\lambda^2} \right) \right] = \\ &= \frac{1}{\lambda} \left[\frac{2}{\lambda^2} - e^{\lambda(\delta - \Delta)} \left((\delta - \Delta)^2 - \frac{2(\delta - \Delta)}{\lambda} + \frac{2}{\lambda^2} \right) \right] \quad \text{(E.3)} \end{aligned}$$

where

$$\delta = \delta(\Delta) = \frac{1}{\lambda} - \frac{\Delta}{e^{\lambda \Delta} - 1}\,. \quad \text{(E.4)}$$

Since $q_j - b_j = \delta$, for $j = 1, \ldots, L-1$, we obtain, following Appendix D, for the second and the third terms in the brackets of (E.1)

$$\begin{aligned} \sum_{j=1}^{L-1} \int_{b_j}^{b_{j+1}} (x - q_j)^2 e^{-\lambda x}\, dx &= \frac{1}{\lambda} \left(\sum_{j=1}^{L-1} e^{-\lambda q_j} \right) \cdot \left[e^{\lambda \delta} \left(\delta^2 - \frac{2\delta}{\lambda} + \frac{2}{\lambda^2} \right) + \right. \\ &\quad \left. - e^{\lambda(\delta - \Delta)} \left((\delta - \Delta)^2 - \frac{2(\delta - \Delta)}{\lambda} + \frac{2}{\lambda^2} \right) \right] \text{, (E.5)} \end{aligned}$$

and

$$\int_{b_L}^{\infty} (x - q_L)^2 e^{-\lambda x}\, dx = \frac{1}{\lambda^3} e^{1-\lambda q_L} \,, \tag{E.6}$$

respectively.

From (E.1), (E.3), (E.5) and (E.6), we finally obtain a closed form expression for the average distortion of the URURQ:

$$\begin{aligned} D \;=\; & \frac{2}{\lambda^2} - e^{\lambda(\delta-\Delta)}\left((\delta-\Delta)^2 - \frac{2(\delta-\Delta)}{\lambda} + \frac{2}{\lambda^2}\right) + \\ & + \left(\sum_{j=1}^{L-1} e^{-\lambda q_j}\right) \cdot \left[e^{\lambda\delta}\left(\delta^2 - \frac{2\delta}{\lambda} + \frac{2}{\lambda^2}\right) - e^{\lambda(\delta-\Delta)}\left((\delta-\Delta)^2 - \frac{2(\delta-\Delta)}{\lambda} + \frac{2}{\lambda^2}\right)\right] + \\ & + \frac{1}{\lambda^2} e^{1-\lambda q_L} \end{aligned} \tag{E.7}$$

where δ is given in (E.4), and where

$$q_j = j\Delta \,, \quad j = 1, \ldots, L-1 \tag{E.8}$$

and

$$q_L = b_L + \frac{1}{\lambda} = L\Delta - \delta + \frac{1}{\lambda} \,. \tag{E.9}$$

References

[1] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, “Image coding using wavelet transform,” *IEEE Trans. Image Process.*, vol. 1, no. 2, pp. 205–220, Apr 1992.

[2] S. Arimoto, “An algorithm for computing the capacity of arbitrary discrete memoryless channels,” *IEEE Trans. Inf. Theory*, vol. 18, no. 1, pp. 14–20, Jan 1972.

[3] J.-C. Batllo and V. A. Vaishampayan, “Asymptotic performance of multiple description transform codes,” *IEEE Trans. Inf. Theory*, vol. 43, no. 2, pp. 703–707, Mar 1997.

[4] T. Berger, *Rate Distortion Theory: A Mathematical Basis for Data Compression.* Englewood Cliffs, NJ: Prentice-Hall, 1971.

[5] ——, “Minimum entropy quantizers and permutation codes,” *IEEE Trans. Inf. Theory*, vol. 28, no. 2, pp. 149–157, Mar 1982.

[6] K. A. Birney and T. R. Fischer, “On the modeling of DCT and subband image data for compression,” *IEEE Trans. Image Process.*, vol. 4, no. 2, pp. 186–193, Feb 1995.

[7] R. Blahut, "Computation of channel capacity and rate-distortion functions," *IEEE Trans. Inf. Theory*, vol. 18, no. 4, pp. 460–473, Jul 1972.

[8] J. N. Bradley and C. M. Brislawn, "The wavelet/scalar quantization compression standard for digital fingerprint images," in *Proc. IEEE Int. Symposium on Circuits and Systems*, May–Jun 1994, vol. 3, pp. 205–208.

[9] 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.

[10] T. Chiang and Y.-Q. Zhang, "A new rate control scheme using quadratic rate distortion model," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 7, no. 1, pp. 246–250, Feb 1997.

[11] A. Cohen, I. Daubechies, and J.-C. Feauveau, "Biorthogonal bases of compactly supported wavelets," *Commun. Pure and Appl. Math.*, vol. 45, pp. 485–560, 1992.

[12] J. H. Conway and N. J. A. Sloane, "Fast quantizing and decoding algorithms for lattice quantizers and codes," *IEEE Trans. Inf. Theory*, vol. 28, no. 2, pp. 227–232, Mar 1982.

[13] T. M. Cover and J. A. Thomas, *Elements of Information Theory.* New York: Wiley, 1991.

[14] I. Daubechies, "Orthonormal bases of compactly supported wavelets," *Commun. Pure and Appl. Math.*, vol. XLI, pp. 909–996, 1988.

[15] A. A. El Gamal and T. M. Cover, "Achievable rates for multiple descriptions," *IEEE Trans. Inf. Theory*, vol. 28, no. 6, pp. 851–857, Nov 1982.

[16] 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.

[17] M. Fleming and M. Effros, "Generalized multiple description vector quantization," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar 1999, pp. 3–12.

[18] Y. Frank-Dayan and R. Zamir, "Universal lattice-based quantizers for multiple descriptions," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar 2000, pp. 500–509.

[19] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 3rd ed. Prentice Hall, 2008.

[20] V. K. Goyal, "Single and multiple description transform coding with bases and frames," Philadelphia, PA: SIAM, 2001.

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

[22] ——, "Theoretical foundations of transform coding," *IEEE Signal Process. Mag.*, vol. 18, no. 5, pp. 9–21, Sep 2001.

[23] V. K. Goyal and J. Kovačević, "Optimal multiple description transform coding of Gaussian vectors," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar–Apr 1998, pp. 388–397.

[24] ——, “Generalized multiple description coding with correlating transforms,” *IEEE Trans. Inf. Theory*, vol. 47, no. 6, pp. 2199–2224, Sep 2001.

[25] R. M. Gray and D. L. Neuhoff, “Quantization,” *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2325–2383, Oct 1998.

[26] H. Jafarkhani and V. Tarokh, “Multiple description trellis-coded quantization,” *IEEE Trans. Commun.*, vol. 47, no. 6, pp. 799–803, Jun 1999.

[27] A. K. Jain, *Fundamentals of Digital Image Processing.* Englewood Cliffs, NJ: Prentice-Hall, 1989.

[28] N. S. Janyant, “Subsampling of a DPCM speech channel to provide two ‘self-contained’ half-rate channels,” *Bell Syst. Tech. J.*, vol. 60, no. 4, pp. 501–509, Apr 1981.

[29] N. S. Jayant and S. W. Christensen, “Effects of packet losses in waveform coded speech and improvements due to an odd-even sample-interpolation procedure,” *IEEE Trans. Commun.*, vol. 29, no. 2, pp. 101–109, Feb 1981.

[30] 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.

[31] R. L. Joshi, V. J. Crump, and T. R. Fischer, “Image subband coding using arithmetic coded trellis coded quantization,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 5, no. 6, pp. 515–523, Dec 1995.

[32] R. L. Joshi, T. R. Fischer, and R. H. Bamberger, “Optimum classification in subband coding of images,” in *Proc. IEEE Int. Conf. Image Processing*, Austin, Texas, Nov 1994, vol. 2, pp. 883–887.

[33] R. L. Joshi, H. Jafarkhani, J. H. Kasner, T. R. Fischer, N. Farvardin, M. W. Marcellin, and R. H. Bamberger, “Comparison of different methods of classification in subband coding of images,” *IEEE Trans. Image Process.*, vol. 6, no. 11, pp. 1473–1486, Nov 1997.

[34] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory.* Englewood Cliffs, NJ: Prentice-Hall, 1993.

[35] M. Lightstone, E. Majani, and S. K. Mitra, “Low bit-rate design considerations for wavelet-based image coding,” *Multidimensional Syst. Signal Processing*, vol. 8, pp. 111–128, Jan 1997.

[36] S. M. LoPresto, K. Ramchandran, and M. T. Orchard, “Image coding based on mixture modeling of wavelet coefficients and a fast estimation-quantization framework,” in *Proc. Data Compression Conf.*, Mar 1997, pp. 221–230.

[37] S. G. Mallat, “A theory for multiresolution signal decomposition: the wavelet representation,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, no. 7, pp. 674–693, Jul 1989.

[38] A. E. Mohr, E. A. Riskin, and R. E. Ladner, “Unequal loss protection: graceful degradation of image quality over packet erasure channels through forward error correction,” *IEEE J. Select. Areas Commun.*, vol. 18, no. 6, pp. 819–828, Jun 2000.

[39] M. T. Orchard, Y. Wang, V. A. Vaishampayan, and A. R. Reibman, "Redundancy rate-distortion analysis of multiple description coding using pairwise correlating transforms," in *Proc. IEEE Int. Conf. Image Processing*, Santa Barbara, CA, Oct 1997, vol. 1, pp. 608–611.

[40] 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.

[41] L. Ozarow, "On a source-coding problem with two channels and three receivers," *Bell Syst. Tech. J.*, vol. 59, no. 10, pp. 1909–1921, Dec 1980.

[42] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. Cambridge University Press, 2007.

[43] R. Puri and K. Ramchandran, "Multiple description source coding using forward error correction codes," in *Conf. Rec. 33rd Asilomar Conf. Sig., Sys. & Computers*, Pacific Grove, CA, Oct 1999, vol. 1, pp. 342–346.

[44] D. O. Reudink, "The channel splitting problem with interpolative coders," Bell Labs, Tech. Rep. TM80-134-1, Oct 1980.

[45] 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.

[46] S. D. Servetto, K. Ramchandran, V. A. Vaishampayan, and K. Nahrstedt, "Multiple-description wavelet based image coding," in *Proc. IEEE Int. Conf. Image Processing*, Oct 1998, vol. 1, pp. 659–663.

[47] S. D. Servetto, V. A. Vaishampayan, and N. J. A. Sloane, “Multiple description lattice vector quantization,” in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar 1999, pp. 13–22.

[48] C. E. Shannon, “Coding theorems for a discrete source with a fidelity criterion,” *IRE Nat. Conv. Rec.*, Mar 1959, reprinted in *Key Papers in the Development of Information Theory.* New York: IEEE.

[49] J. M. Shapiro, “Embedded image coding using zerotrees of wavelet coefficients,” *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3445–3462, Dec 1993.

[50] Y. Shoham and A. Gersho, “Efficient bit allocation for an arbitrary set of quantizers,” *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, no. 9, pp. 1445–1453, Sep 1988.

[51] G. Strang and T. Nguyen, *Wavelets and Filter Banks*, 2nd ed. Wellesley-Cambridge Press, 1996.

[52] P. Subrahmanya and T. Berger, “Multiple descriptions encoding of images,” in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar 1997, p. 470.

[53] 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.

[54] N. Tanabe and N. Farvardin, “Subband image coding using entropy-coded quantization over noisy channels,” *IEEE J. Sel. Areas Commun.*, vol. 10, no. 5, pp. 926–943, Jun 1992.

[55] D. S. Taubman and M. W. Marcellin, *JPEG2000: Image Compression Fundamentals, Standards and Practice.* Springer, 2002.

[56] B. E. Usevitch, “Optimal bit allocation for biorthogonal wavelet coding,” in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar–Apr 1996, pp. 387–395.

[57] ——, “A tutorial on modern lossy wavelet image compression: foundations of JPEG 2000,” *IEEE Signal Process. Mag.*, vol. 18, no. 5, pp. 22–35, Sep 2001.

[58] B. E. Usevitch and M. T. Orchard, “Smooth wavelets, transform coding, and Markov-1 processes,” *IEEE Trans. Signal Process.*, vol. 43, no. 11, pp. 2561–2569, Nov 1995.

[59] V. A. Vaishampayan, “Design of multiple description scalar quantizers,” *IEEE Trans. Inf. Theory*, vol. 39, no. 3, pp. 821–834, May 1993.

[60] V. A. Vaishampayan and J.-C. Batllo, “Asymptotic analysis of multiple description quantizers,” *IEEE Trans. Inf. Theory*, vol. 44, no. 1, pp. 278–284, Jan 1998.

[61] V. A. Vaishampayan and J. Domaszewicz, “Design of entropy-constrained multiple-description scalar quantizers,” *IEEE Trans. Inf. Theory*, vol. 40, no. 1, pp. 245–250, Jan 1994.

[62] V. A. Vaishampayan, N. J. A. Sloane, and S. D. Servetto, “Multiple-description vector quantization with lattice codebooks: design and analysis,” *IEEE Trans. Inf. Theory*, vol. 47, no. 5, pp. 1718–1734, Jul 2001.

[63] R. Venkataramani, G. Kramer, and V. K. Goyal, “Multiple description coding with many channels,” *IEEE Trans. Inf. Theory*, vol. 49, no. 9, pp. 2106–2114, Sep 2003.

[64] J. D. Villasenor, B. Belzer, and J. Liao, “Wavelet filter evaluation for image compression,” *IEEE Trans. Image Process.*, vol. 4, no. 8, pp. 1053–1060, Aug 1995.

[65] Y. Wang, M. T. Orchard, and A. R. Reibman, “Multiple description image coding for noisy channels by pairing transform coefficients,” in *Proc. IEEE Workshop on Multimedia Signal Processing*, Princeton, NJ, Jun 1997, pp. 419–424.

[66] 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.

[67] R. Zamir, “Gaussian codes and shannon bounds for multiple descriptions,” *IEEE Trans. Inf. Theory*, vol. 45, no. 7, pp. 2629–2636, Nov 1999.

[68] Z. Zhang and T. Berger, “New results in binary multiple descriptions,” *IEEE Trans. Inf. Theory*, vol. 33, no. 4, pp. 502–521, Jul 1987.

קידוד תמונות מרובה תיאורים מבוסס
קונטקסט בתחום ה-Wavelet

דרור פורת

קידוד תמונות מרובה תיאורים מבוסס קונטקסט בתחום ה-Wavelet

חיבור על מחקר

לשם מילוי חלקי של הדרישות לקבלת תואר
מגיסטר למדעים בהנדסת חשמל

דרור פורת

הוגש לסנט הטכניון — מכון טכנולוגי לישראל
תשרי תש"ע חיפה אוקטובר 2009

המחקר נעשה בהנחיית פרופ’ דוד מלאך
בפקולטה להנדסת חשמל

הכרת תודה

ברצוני להביע את תודתי העמוקה והכנה למנחה שלי, פרופ’ דוד מלאך, על הדרכתו המקצועית המעולה, תמיכתו האיתנה ויחסו החם. ברצוני גם להודות לנמרוד פלג וליתר אנשי צוות המעבדה לעיבוד אותות ותמונות (SIPL) על הסיוע הטכני הרציף. תודה מעומק לבי להוריי, שמואל וזהבה, על תמיכתם האינסופית ועידודם במהלך הדרך. לבסוף, ברצוני להודות ליקירתי שירה על אהבתה ותמיכתה בי ללא תנאי.

אני מודה לטכניון על התמיכה הכספית הנדיבה בהשתלמותי

תוכן עניינים

רשימת איורים

רשימת טבלאות

תקציר

קידוד מרובה תיאורים (Multiple description—MD) הינו טכניקת קידוד המפיקה מספר תיאורים של מקור אינפורמציה יחיד (למשל, תמונה), באופן כזה שאיכויות שחזור שונות מתקבלות מתתי-קבוצות שונות של התיאורים. המטרה של קידוד מרובה תיאורים הינה לספק עמידות כנגד שגיאות עבור אינפורמציה המשודרת ברשתות שבהן עלול להתרחש אובדן של חבילות מידע. היות שקידוד מרובה תיאורים אינו יוצר היררכיה בין התיאורים השונים, ייצוגים הנוצרים על-ידי קידוד מסוג זה מביאים לכך שכל התיאורים המתקבלים הינם מועילים. זאת בניגוד, לדוגמה, לקידוד בשכבות (Layered coding), שבו אובדן של שכבה אחת עשוי להפוך גם שכבות אחרות לחסרות תועלת. על כן, קידוד מרובה תיאורים מתאים במיוחד לרשתות שאין בהן מנגנונים להעדפת חבילות מידע מסוימות על-פני אחרות, כמו למשל רשת האינטרנט. מובן שעל מנת לאפשר עמידות כנגד אובדן אפשרי של תיאורים, נדרש להקריב חלק מיעילות הקידוד (כלומר, הייצוג הינו יתיר). זו גם הסיבה לכך שקידוד מרובה תיאורים הוא כדאי לשימוש רק במקרה שבו אכן ישנו צורך בהגנה שכזו כנגד אובדן אפשרי של תיאורים. סכימות רבות לקידוד מרובה תיאורים מתמקדות במקרה של שני תיאורים, בעיקר עקב הנוחות היחסית בטיפול בו, וכך גם עבודה זו.

בעבודה קודמת, תיאורים מרובים לקידוד תמונה נוצרו באמצעות שימוש בפירוק למרכיבים דמויי-פוליפאזה (התמרת פוליפאזה) וקוונטיזציה סלקטיבית, המבוצעים בתחום ה-Wavelet. בניגוד לטכניקות רבות אחרות לקידוד מרובה תיאורים, כמו למשל קוונטיזציה סקלרית לריבוי תיאורים (MD scalar quantization—MDSQ) והתמרות מצמדות לריבוי

תיאורים (MD correlating transforms), טכניקת הקידוד באמצעות התמרת פוליפאזה וקוונטיזציה סלקטיבית מפרידה באופן מפורש בין יצירת התיאורים והוספת היתירות, דבר המביא להפחתה משמעותית במורכבות התכן והמימוש של המערכת. טכניקה זו גם מאפשרת יצירה נוחה של תיאורים בעלי קצב זהה (מבחינה סטטיסטית) וחשיבות זהה, תכונה המתאימה במיוחד לרשתות החסרות מנגנונים להעדפת חבילות מידע מסוימות על-פני אחרות.

ליצירת התיאורים, במקרה של שני תיאורים הנדון כאן, טכניקה זו מפעילה התמרת פוליפאזה בתחום ה-Wavelet, וכל אחד משני מרכיבי הפוליפאזה הנוצרים מקודד באופן בלתי-תלוי בקצב "קידוד מקור" ומהווה את חלק האינפורמציה העיקרי עבור התיאור המתאים לו. על מנת להוסיף יתירות באופן מפורש לכל תיאור, מרכיב הפוליפאזה האחר מקודד אז בקצב "קידוד יתירות" (הנמוך יותר, לרוב) באמצעות קוונטיזציה סלקטיבית ומצורף לתיאור זה. במקרה של אובדן תיאור, יתירות זו מאפשרת שחזור סביר של המרכיב האבוד.

בעבודת מחקר זו, אנו מפתחים ובוחנים דרך יעילה לנצל את התכונות הסטטיסטיות המיוחדות של פירוק ה-Wavelet על מנת להשיג יעילות קידוד משופרת, במסגרת הכללית של קידוד מרובה תיאורים מבוסס התמרת פוליפאזה. אנו מציעים סכימת קידוד חדשה, אשר מנצלת ביעילות אינפורמציה קונטקסטואלית המחולצת ממרכיב הפוליפאזה העיקרי של כל תיאור, על מנת לשפר את יעילות הקידוד של מרכיב הפוליפאזה היתיר המתאים, ובכך מאפשרת למקודד ה-MD המוצע להשיג ביצועים כוללים משופרים. הדרך להשיג זאת הינה באמצעות פרוצדורות קידוד שונות, כגון סיווג מבוסס קונטקסט של מקדמי ה-Wavelet, קוונטיזציה מסתגלת מבוססת מודל פרמטרי, הקצאת סיביות אופטימלית יעילה וקידוד אנטרופיה מסתגל.

על מנת לנצל באופן יעיל את התלויות הסטטיסטיות בין מקדמי Wavelet סמוכים ולהימנע מהצורך באפיון מפורש של תלויות אלה, אנו משתמשים בפרוצדורה יעילה לסיווג מבוסס קונטקסט. הסיווג מבוסס על קונטקסטים הנוצרים ממקדמים מקוונטטים של

מרכיב הפוליפאזה העיקרי בתיאור, והזמינים גם במפענח, ובכך נמנע הצורך בשידור של מידע צד. אף על פי כן, כמות מבוקרת של מידע צד עדיין מופקת ומשודרת, על מנת לשפר את ביצועי המערכת. מידע צד זה כולל את ספי הסיווג, המאפשרים לשייך מקדם למחלקה המתאימה בהינתן הקונטקסט שלו, כאשר כל מחלקה ממודלת באמצעות פילוג פרמטרי לפלסיאני. המידול הפרמטרי מנוצל גם על-ידי סכימת הקצאת הסיביות ופרוצדורות הקוונטיזציה וקידוד האנטרופיה העוקבות.

פרוצדורת הקוונטיזציה מבוססת הקונטקסט מאפשרת למקודד המוצע להשתמש באוסף של קוונטייזרים, שכל אחד מהם מותאם למחלקה אחרת. למטרה זו, אנו בוחנים שני סוגים של קוונטייזרים: קוונטייזר מסוג UTQ (Uniform threshold quantizer) וקוונטייזר מסוג URURQ (Uniform reconstruction with unity ratio quantizer). שני קוונטייזרים אלה מקרבים היטב את הקוונטייזר הסקלרי מאולץ-האנטרופיה (Entropy-constrained scalar quantizer—ECSQ) האופטימלי עבור הפילוג הלפלסיאני, בהנחת עיוות ריבועי ממוצע, והם פשוטים למדי לתכנון ולהפעלה. על מנת להימנע מהסיבוכיות הגבוהה של אלגוריתמי תכן מאולץ-אנטרופיה לקוונטייזרים, אנו מציעים אסטרטגיית תכן יעילה, המבוססת על מערך ממופתח (Indexed) המתוכנן מראש (Off-line) של קוונטייזרים אופטימליים (במובן שגיאה ריבועית ממוצעת) לפילוג הלפלסיאני בעלי גודל צעד שונה. על מנת להפחית עוד יותר את הסיבוכיות, אנו מפתחים ביטויים סגורים לעיוות המושג הן על-ידי ה-UTQ והן על-ידי ה-URURQ.

עבור הקצאת סיביות בין המחלקות השונות בפסי-הצד (Subbands) השונים, אנו מפתחים סכימת הקצאת סיביות אופטימלית ויעילה מבוססת-מודל, תחת המסגרת של אופטימיזציה לגרנג'יאנית, אשר מביאה בחשבון גם את העובדה שהתמרת ה-Wavelet הבי-אורתוגונלית אינה משמרת אנרגיה. סכימת הקצאת הסיביות, המבוססת על התאמת שונות (Variance scaling) ועל מערך הקוונטייזרים הממופתח המוזכר לעיל, מאפשרת למקודד המוצע להשיג סיבוכיות מופחתת במהלך הקידוד עצמו (On-line), וכן לבצע הסתגלות מושכלת ומיידית למקודד האנטרופיה האריתמטי בהתאם לסטטיסטיקת המקדמים המשתנה.

אנו מציגים תוצאות נסיוניות מגוונות המדגימות באופן ברור את יתרונותיו של מקודד ה-MD מבוסס הקונטקסט המוצע. באופן ספציפי, אנו מראים גם כי המקודד המוצע עולה בביצועיו על המקודד המקורי מבוסס התמרת הפוליפאזה, המספק את המסגרת למקודד המוצע, וזאת על-פני תחום היתירות כולו. כמו כן אנו מראים כי השיפור ביעילות הקידוד אכן יכול להיות מיוחס בראש ובראשונה לניצול היעיל של אינפורמציה קונטקסטואלית ושל התכונות הסטטיסטיות המיוחדות של פירוק ה-Wavelet. בנוסף, אנו מציגים דרך לקבוע, הלכה למעשה, את נקודת העבודה האופטימלית של מקודד ה-MD המוצע, בהתבסס על מאפייני ערוץ התקשורת.