

Nonlinear Transform Coding with Lossless Polar Coordinates

by

Demba Elimane Ba

Submitted to the Department of Electrical Engineering and Computer Science

in partial fulfillment of the requirements for the degree of

Master of Science in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2006

© Massachusetts Institute of Technology 2006. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
May 12, 2006

Certified by
Vivek K Goyal
Assistant Professor of Electrical Engineering
Thesis Supervisor

Accepted by
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Nonlinear Transform Coding with Lossless Polar Coordinates

by

Demba Elimane Ba

Submitted to the Department of Electrical Engineering and Computer Science
on May 12, 2006, in partial fulfillment of the
requirements for the degree of
Master of Science in Electrical Engineering

Abstract

In conventional transform coding, the importance of preserving desirable quantization partition cell shapes prevents one from considering the use of a nonlinear change of variables. If no linear transformation of a given source would yield independent components, this means having to encode it at a rate higher than its entropy, i.e. sub-optimally. This thesis proposes a new transform coding technique where the source samples are first uniformly scalar quantized and then transformed with an integer-to-integer approximation to a nonlinear transformation that would give independent components. In particular, we design a family of integer-to-integer approximations to the Cartesian-to-polar transformation and analyze its behavior for high rate transform coding. Among the benefits of such an approach is the ability to achieve redundancy reduction beyond decorrelation without limitation to orthogonal linear transformations of the original variables. A high resolution analysis is given, and for source models inspired by a sensor network application and by image compression, simulations show improvements over conventional transform coding. A comparison to state-of-the-art entropy-coded polar quantization techniques is also provided.

Thesis Supervisor: Vivek K Goyal

Title: Assistant Professor of Electrical Engineering

Acknowledgments

All praise is due to *Allah*, lord of the universe, the merciful, the most beneficent, who has given me the gift of an MIT education and without whom I would not have been able to rise to the challenge. In his great mercy and wisdom, He (SWT) has gathered around me the conditions and, most importantly, the people who have helped me to farewell in an endeavor very few are given the opportunity to undertake.

This is one of those moments when, before the completion of a stage in your life and the beginning of a new one, you sit back and realize that the drive and the hunger that animated you in times of hardship would have been absent had it not been for a select number of people who believed in you at times when you doubted your own self, and a few others who doubted you along the way and that you had to prove wrong. I would like to take a few moments to acknowledge the former group of people.

I would like to thank my family for their commitment and support. Yaye (Mom), your dedication and hard work have always amazed me and been a mystery to me; honestly, I don't know how you do it and I don't think I will ever find out because it's time to retire. Papa (Pop), it goes without saying that I feed off your rigueur and seriousness in everything I do. Pa (Grandpa), for the toughness I've inherited from you. Famani and Khalidou, I'm looking forward to attending both of your graduations. Baboye and Moctar, all three of us have master degrees now, I'm not sure who's oldest anymore, would that be me? Elimane "Le Zincou", Dakar, N'dioum, Paris and Nouakchott we've been together; next on the list in Bamako. Gogo, for so warmly welcoming me every time I visited Nouakchott. Tata Rougui and Tonton Hamat, for treating me both like a son and a younger brother. Tata Khady Hane, for spoiling me with amazing meals every time I visited Paris. Kawoum Bamoum (Tonton Moussa) and Rose, for so patiently taking care of my Grandpa, my Mom and my Dad. Oumou, for treating me like a younger brother. Mamadou, for treating me like an older brother. Tata Astou, Tata Aliatou, Mame Djeinaba, Mame Kadiata, Tata Marie-Claire, Tata Mame Penda, for always being there for me. To my niece Aicha and my nephews Abu-Bakr, Muhammad, Abdul Kareem and Moussa, all future

Beavers. Yacine and Awa, thanks for taking care of my niece and nephews. To my relatives in Boghe and N'dioum, thanks for always being there.

To my advisor Vivek Goyal. You've been there from day -1, helping me getting out of the undergrad mode and making sure my transition into MIT was not a painful one. Thanks for your guidance and rigueur, and for walking me through my baby steps in research.

I would like to thank members of the STIR Group, in particular Eric for helping me through the administrative hurdles and Lav for the fruitful discussions and for bearing with me while at DCC. Ruby, Adam and Julius for bearing with me as a suite mate.

I would like to thank my academic counselor Al Willsky for making sure I keep on track with the doctoral program. To Prof. Verghese who's been looking out for me since the day I first visited MIT for the open house. To Debb of MTL and Lourenco of the 002 lab mental and moral support.

To my friends in Maryland for believing in me and bearing with me all these years. Abdoul Ahad, Suze and the Diop family, Mamiko, Meriam, Malick Diop, Babacar.

To my friends in Boston. Thanks for making my life at MIT not just about work and study. Borjan, the first guy I met at MIT, we're in this together man, I got your back you got my front from the start to finish. Zahi and Akram, thanks for going out of your way to accommodate my dietary needs. We don't share the same blood and yet you're like family. Youssef and Val, thanks for being my friends. Jen, for convincing me that if one is not going to do anything to change things around them, one might as well shut up. Your drive and energy are contaminating. Alaa, Sourav and Tom of DSPG for the fruitful discussions and for hanging out with me at times when it was hard to do work. Hanaan, Obrad, Victor and BQ, thanks for being my friends.

Last but not least, to my friends from back home. Alioune, we've grown up together, thanks for always being there. Luson you're like family. Thierno Niang, Papis Konate, Said, Racine, Korsi, Loic, Stephan and those I used to hang out with back in Ouaga. Raby and Sally Ba, thanks for the moral support.

Contents

1	Introduction	11
2	Background	15
2.1	Integer-to-integer (i2i) Transform Coding	15
2.2	Gaussian Scale Mixtures (GSMs)	17
2.3	Parametric Audio Coding and Entropy-coded Polar Quantization . .	19
3	I2I Polar Coordinates	25
3.1	Designing the Transform	25
3.2	High Resolution Analysis	27
3.3	Results	31
3.3.1	Numerical Results for GSM Examples	32
3.3.2	Numerical Results for Sensor Networks Example	38
4	Generalizations of Integer-to-integer Polar Coordinates	43
4.1	Generalization to Higher Dimensions	43
4.2	Generalization to Sources Separable in <i>Extended</i> Polar coordinates . .	44
4.3	Lossless Coding with Integer-to-integer Polar Coordinates	46
5	Discussion	49

List of Figures

1-1	A conventional encoder with a continuous-domain transform $T : \mathbb{R}^N \rightarrow \mathbb{R}^N$. $Q(\cdot)$ is an unbounded uniform scalar quantizer with step size Δ .	12
1-2	A generic encoder with a continuous-domain transform $T_1 : \mathbb{R}^N \rightarrow \mathbb{R}^N$ and a discrete-domain transform $U : \Delta\mathbb{Z}^N \rightarrow \Delta\mathbb{Z}^N$. $Q(\cdot)$ is an unbounded uniform scalar quantizer with step size Δ .	13
2-1	For x jointly Gaussian, the i2i encoder matches the performance of that in Figure 1-1. $T_2 : \Delta\mathbb{Z}^N \rightarrow \Delta\mathbb{Z}^N$ is a discrete-domain transform. $Q(\cdot)$ is an unbounded uniform scalar quantizer with step size Δ .	17
2-2	Contours of equal probability for uncorrelated (left) and independent (right) Laplacian random variables	19
2-3	As reproduced from [16], partitions of the input space for ECRQ (left), ECSPQ (middle) and ECUPQ (right)	23
3-1	Visualizing i2i Polar Coordinates, (a) $M = 8$, (b) $M = 11$, (c) $M = 16$. The left panel shows cells with equal discrete radius P with the same color. The right panel shows regions of constant discrete angle A with constant color.	28
3-2	For high-resolution analysis of $H(P)$, a square cell with area Δ^2 (left panel) is replaced by a portion of a ring with area Δ^2 (right panel) as shown. The analysis finds the radius and thickness of the ring as a function of P .	29
3-3	Simulation results. $M = 8$ and a range of rates is obtained by varying Δ .	34

3-4	Comparing Monte Carlo simulation of i2i Polar transform coding ($M = 8$) to the performance predicted by the high-rate analysis	35
3-5	Analytical Comparison of i2i Polar transform coding to ECSPQ at high rate	36
3-6	Ratio of distortions at high rate as a function of λ	36
3-7	Simulation results. $M = 8$ and a range of rates is obtained by varying Δ	38
3-8	A depiction, reproduced from [13], of a highly non-Gaussian distribution that arises in a localization problem with ranging sensors. The distribution is very complicated, but is well-suited to a polar coordinate representation.	39
3-9	Simulation results. $M = 8$ and a range of rates is obtained by varying Δ	40
3-10	Comparing Monte Carlo simulation of i2i Polar transform coding ($M = 8$) to the performance predicted by the high-rate analysis	41
3-11	Analytical Comparison of i2i Polar transform coding to ECSPQ at high rate	41
4-1	An encoder for a polar-separable source with arbitrary covariance matrix with a discrete-domain transform $\hat{D} : \Delta\mathbb{Z}^N \rightarrow \Delta\mathbb{Z}^N$ and an integer-to-integer polar transform $U : \Delta\mathbb{Z}^N \rightarrow \Delta\mathbb{Z}^N$. $Q(\cdot)$ is an unbounded uniform scalar quantizer with step size Δ	45

Chapter 1

Introduction

Transform coding is the method of choice for lossy compression of multimedia content. It is central to compression standards for audio, image and video coding. Therefore, we are interested in understanding and, to whatever extent possible, addressing the limitations of conventional transform coding. From its introduction up to this day, the conventional transform coding paradigm has been limited to linear transformations for reasons pertaining mainly to the effect of nonlinear transformations on quantization partition cells. This thesis introduces a transform coding framework that accommodates nonlinear transformations of the original variables. Specifically, we consider the problem of high-rate transform coding of polar-separable sources using an integer-to-integer approximation to Cartesian-to-polar transformations. A detailed analysis and simulation results show that such an approach indeed provides an improvement over the conventional one, which limits itself to orthogonal transformation of the original variables. The hope is that the ideas developed in this thesis will set the stage for a more comprehensive theory of nonlinear transform coding that would accommodate *arbitrary* nonlinear transformations of the original variables.

The foundations of transform coding were laid down in 1963 by Huang and Schultheiss [11]: a source vector x undergoes a change of coordinates described by a linear transformation T , is scalar-quantized and subsequently encoded with a *fixed-rate* code. Here, we require that the output of the quantizer be scalar *entropy-coded* and refer to the resulting encoding structure, which is depicted in Figure 1-1, as




Figure 1-1: A conventional encoder with a continuous-domain transform $T : \mathbb{R}^N \rightarrow \mathbb{R}^N$. $Q(\cdot)$ is an unbounded uniform scalar quantizer with step size Δ .

conventional transform coding. By scalar entropy coding, we mean that the rate is measured by the sum of the entropies of the components of \hat{y} , which is larger than or equal to the entropy of the corresponding vector (with equality if and only if the components are independent).

The main features of conventional transform coding are the low computational complexity of uniform scalar quantization and scalar entropy coding. The 1963 paper by Huang and Schultheiss regarded T as an invertible linear transformation that produces uncorrelated transform coefficients. The contemporary view is that T should be optimized among orthogonal transforms. Intuitively, this is because nonorthogonal transforms deform the shape of the quantization partition (although to a lesser extent than nonlinear transforms do). In either of these cases, for jointly Gaussian sources and under some mild conditions, a Karhunen–Loève transform (KLT) is optimal [7]. Although limited to a specific class of sources, this result highlights some desirable properties of a transform code. Loosely, a well-behaved transform code should preserve the shape of the quantization partition cells and yield transform coefficients as close as possible to independent, so that scalar entropy coding can be used with no penalty in rate. The KLT happens to satisfy both of these conditions in the jointly Gaussian case. However, for non-Gaussian sources, these may be competing goals. Indeed, more often than not, any transformation that gives independent components is not orthogonal; in fact, it may even be nonlinear. Transform coding inherently suffers from a packing gain lost by scalar quantization. For non-Gaussian sources, the situation becomes even worse because one is having to separately encode random variables that are not independent, which we know, from basic information theory principles, is suboptimal.




Figure 1-2: A generic encoder with a continuous-domain transform $T_1 : \mathbb{R}^N \rightarrow \mathbb{R}^N$ and a discrete-domain transform $U : \Delta\mathbb{Z}^N \rightarrow \Delta\mathbb{Z}^N$. $Q(\cdot)$ is an unbounded uniform scalar quantizer with step size Δ .

The above discussion suggests the need of a more general approach to transform coding that will accommodate both linear and nonlinear transformations. This thesis proposes to assess the practicality of a theory for nonlinear transform coding by considering optimal redundancy reduction for sources that are better suited to polar rather than rectangular coordinate processing.

We start by taking an expansive view of transform coding in which the encoder performs some signal transformation, uniform scalar quantization, some further signal transformation, and then scalar entropy coding, as illustrated in Figure 1-2. This approach to transform coding is inspired by the integer-to-integer (i2i) framework introduced in [5] but goes much further by allowing $U(\cdot)$ to be an integer-to-integer approximation to a *continuous-domain* nonlinear transformation. Note that choosing $T_1(\cdot)$ to be a KLT and $U(\cdot)$ as the identity operator would lead to ordinary transform coding as in Figure 1-1.

We hypothesize that, for sources separable in polar coordinates, a system where $T_1(\cdot)$ is the identity operator and $U(\cdot)$ an i2i approximation to a Cartesian-to-polar transformation will outperform conventional transform coding at high-rate.

Chapter 2

Background

The aim of this chapter is to give an overview of the literature that is relevant to the proposed thesis. The connection between the generic encoder of Figure 1-2 and the integer-to-integer transform coding framework introduced in [5] is made precise. Moreover, we present previous work on ordinary transform coding of Gaussian Scale Mixtures, a class of sources amenable to our i2i polar framework. The class of i2i polar transformations developed in this thesis are also reminiscent of recent developments in entropy-coded polar quantization. We discuss such work to set the stage for a comparison with our scheme.

2.1 Integer-to-integer (i2i) Transform Coding

In the usual formulation of transform coding as depicted in Figure 1-1, performance is limited by two characteristics: *scalar* quantization and *scalar* entropy coding. With scalar quantization, to minimize the distortion at high rate the best one can hope for is a quantizer partition into hypercubes; and with scalar entropy coding, one desires independent transform coefficients to minimize the rate. Transform coding has a clean analytical theory for Gaussian sources because these two virtues can be achieved simultaneously. For other sources, ideal cell shape and independence of transform coefficients may be competing, and in fact the optimal transform will generally give neither.

For high-rate transform coding of a jointly Gaussian source, it was shown in [5] that the performance of conventional transform coding can be matched by a system that first quantizes the source and then implements an i2i approximation to a linear, non-orthogonal transformation. This encoding strategy is illustrated in Figure 2-1. For a k -dimensional jointly Gaussian source, the benefit of such a strategy is a k -fold reduction in memory requirements with no penalty in rate. Such a result is indeed remarkable. However, we would like to stress that this is not the feature of i2i transform coding that the proposed thesis hopes to exploit. The superiority of the i2i framework to the conventional approach stems from the additional degrees of freedom the former allows in the design of the transform $T_2(\cdot)$. Indeed, by allowing quantization to occur first, one is no longer restricted to orthogonal linear transformations as is the case in Figure 1-1. As long as it is invertible, $T_2(\cdot)$ may be non-orthogonal and even nonlinear. This is precisely what the proposed thesis hopes to build upon. Specifically, for sources separable in polar coordinates, we anticipate that choosing $U(\cdot)$ in Figure 1-2 as an i2i polar coordinate transformation will lead to the ideal situation of encoding independent random variables.

Additional intuition about i2i transform coding can be built by considering the per-component mean-squared error (MSE) distortion at high rate of the encoder in Figure 1-1 for an N -dimensional source. The said MSE is given by the following expression [6]:

$$D = \frac{1}{N} \underbrace{\frac{1}{12} \text{trace} (T^{-1}(T^{-1})')}_{\text{shape}} \underbrace{2^{\frac{2}{N} \sum_{i=1}^N h(y_i)}}_{\text{independence}} 2^{-2R}. \quad (2.1)$$

Assuming without loss of generality that T is a linear transform with determinant equal to 1, the equation above highlights two important facts:

1. Orthogonality is good because it minimizes the shape factor over unit-determinant matrices T .
2. Independence of transform coefficients is good because it minimizes $\sum_{i=1}^N h(y_i)$.




Figure 2-1: For x jointly Gaussian, the i2i encoder matches the performance of that in Figure 1-1. $T_2 : \Delta\mathbb{Z}^N \rightarrow \Delta\mathbb{Z}^N$ is a discrete-domain transform. $Q(\cdot)$ is an unbounded uniform scalar quantizer with step size Δ .

The idea in i2i transform coding is that one can avoid the shape factor in (2.1) by substituting the system of Figure 1-1 with that of Figure 2-1 and setting T_2 equal to an i2i approximation of T obtained through the lifting scheme described in [5]. Essentially, what such a strategy achieves is independence of transform coefficients with the best possible MSE, while restricting ourselves to computationally easy operations, namely *scalar* quantization and *scalar* entropy coding.

2.2 Gaussian Scale Mixtures (GSMs)

A straightforward statistical description of any vector source could be in terms of the first and second-order moments of its components. In various signal processing tasks; however, such characterizations may not be sufficient to fully capture the dependencies between components. For instance, non-Gaussianity as well as statistical dependence despite uncorrelatedness are salient features of wavelet coefficients of natural images. In [17], the authors show that GSM models are well suited to fully capture the statistical intricacies of wavelet coefficients of natural images. More generally, GSMs have been used to model a variety of sources whose components mainly exhibit statistical correlations of order higher than two.

Formally, a random vector Y of size N is a GSM if it can be expressed as the product of a scalar non-negative random variable z and a zero-mean jointly Gaussian random vector X : $Y = zX$. The scalar random variable z is known as the multiplier, and z and X are independent.¹

¹From this point on, capital letters are used for vector random variables and lower case letters for scalar random variables

As a consequence of the above definition, any GSM has density of the form:

$$f_Y(Y) = \int_0^\infty \frac{1}{(2\pi)^{N/2} |z^2 K_X|^{1/2}} \exp\left(-\frac{Y^T K_X^{-1} Y}{2z^2}\right) f_z(z) dz, \quad (2.2)$$

where K_X is the covariance matrix of X .



In what follows, we focus on the case where $K_X = I$. In the two-dimensional case, the joint density specializes to:

$$f_Y(Y) = \int_0^\infty \frac{1}{2\pi z^2} \exp\left(-\frac{y_1^2 + y_2^2}{2z^2}\right) f_z(z) dz. \quad (2.3)$$

There are a number of well-known distributions that can be expressed as GSMs. In 1D for instance ($N = 1$), a Rayleigh prior on z (exponential prior on z^2) results in a random variable y which follows a Laplacian distribution [10]. This is also true in higher dimensions ($N > 1$): the components of a GSM vector with a Rayleigh prior on z are marginally Laplacian. However, they are not independent. This is shown in Figure 2-2 for $N = 2$, where the left panel shows contours of equal probability for a 2D GSM vector with a Rayleigh prior on z , while the right panel shows those of a 2D vector with *independent* identically-distributed (*iid*) Laplacian components. Note that the components of the 2D GSM vector are *uncorrelated* but *not* independent.

Consider the problem of encoding the sources of Figure 2-2 under the conventional transform coding framework. The wisdom of conventional transform coding tells us that they should be encoded the same way since they both have uncorrelated components. While not separable in rectangular coordinates, the 2D GSM vector, with probability density function as in Equation 2.3, is clearly separable in polar coordinates and therefore amenable to the i2i polar transform coding scheme proposed in this thesis.

Ordinary transform coding for GSMs was considered in [14]. Specifically, the authors showed that, among orthogonal transforms, a particular class of KLTs are optimal for high rate transform coding of a GSM. One can easily check that no linear transformation of a GSM will yield independent components. This means that, after transformation and quantization of a GSM as in Figure 1-1, some dependencies will



(a) Uncorrelated Laplacian Components (GSM) (b) Independent Laplacian Components

Figure 2-2: Contours of equal probability for uncorrelated (left) and independent (right) Laplacian random variables

remain. The latter can only be removed through nonlinear transformations, which we know are disliked in conventional transform coding. This thesis shows that the i2i transform coding formulation coupled with knowledge of the polar separability of GSMs can overcome such a difficulty. The challenge is to design an i2i approximation to a Cartesian-to-polar transformation that can be used as $U(\cdot)$ in the generic encoder of Figure 1-2.

2.3 Parametric Audio Coding and Entropy-coded Polar Quantization

In [4], parametric audio coding (PAC) is cast as an effective technique to address the non-scalability of conventional speech and transform coding based audio coders to low rates. The authors argue that transform coders are not optimal for low-rate coding of audio signals and that, although speech coders could be used instead, they are too specialized to be effective over a wide range of audio signals. PAC offers a more flexible alternative which scales to low rates, while accounting for the variety of audio signals that one may encounter in practice. In PAC, an audio signal is decomposed as

the sum of a transient component, a noise component and a sinusoidal components. The parameters relevant to each component are encoded and transmitted to a decoder whose task is to reconstruct a perceptually-faithful estimate of the input signal based on prior knowledge of the decomposition model. Efficient encoding and transmission of the model parameters is therefore essential to the performance of a parametric audio coder. The question then comes as to what the optimal encoding schemes are for each of the three components of an audio signal decomposed as in the PAC framework. In what follows, we focus on the problem of optimal quantization of the sinusoidal component.

When dealing with sinusoidal signals, two types of representations naturally come to mind. The first one is a polar representation where the signal is specified by its magnitude r and its phase θ . The second one is a rectangular representation where one specifies the coordinates x and y of the signal in Euclidean space. The transformation from polar to rectangular coordinates is as follows: $x = \text{Re}\{re^{j\theta}\}$ and $y = \text{Im}\{re^{j\theta}\}$. This naturally leads to two possible formulations of the quantizer design problem: one where the signal is quantized in polar coordinates and the other where the quantization is performed in rectangular coordinates.

The quantizer design problem can be formulated in two equivalent ways:

- Given a maximum tolerable distortion D , we want to know what minimum number of bits, R , is required to represent such a signal and whether this is achieved by polar or rectangular representation of the signal (or neither).
- Stated otherwise, we want to know what amount of distortion, D , we are bound to incur given that we are willing to spend at most R bits to represent the signal, and whether the distortion depends on the type of representation chosen (Cartesian, polar or other).

The rest of the discussion assumes that the second approach is the one adopted to tackle the quantizer design problem. This very approach was adopted by the authors in [16] who compare three different ways of quantizing and encoding signals that lend themselves to both Cartesian and polar representations. Assuming high rate (large

R), they derive three quantizer structures as well as their respective performances, which are given by expressions relating D to R . Note that the constraint on R is expressed as one on the entropy of the quantizer output. Below is a description of the said quantizer structures:

1. Entropy-Coded Rectangular Quantization (ECRQ): The signal is specified by its coordinates x and y in Euclidean space. The encoding process goes as follows: x and y are first scalar-quantized and subsequently scalar entropy-coded. Note that x and y are not necessarily independent or identically-distributed. Letting \hat{x} and \hat{y} be the quantized versions of x and y respectively, the quantizer design problem amounts to minimizing D subject to a constraint on $R = H(\hat{x}) + H(\hat{y})$. This results in the left-most quantizer structure of Figure 2-3. For such a quantizer structure, D as a function of R is given as follows:

$$D_{ECRQ} = \frac{2^{-(R-h(x)-h(y))}}{6}, \quad (2.4)$$

where $h(x)$ and $h(y)$ are the differential entropies of x and y respectively.

2. Entropy-Coded Strictly Polar Quantization (ECSPQ): Polar representation of the signal is assumed. The encoding process goes as follows: r and θ are scalar-quantized and subsequently *scalar* entropy-coded. Letting \hat{r} and $\hat{\theta}$ be the quantized versions of r and θ respectively, the quantizer design problem amounts to minimizing D subject to a constraint on $R = H(\hat{r}) + H(\hat{\theta})$. This results in the middle quantizer structure of Figure 2-3. The corresponding expression for D as a function of R (as derived in [16]) is:

$$D_{ECSPQ} = \sigma_r \frac{2^{-(R-h(r)-h(\theta))}}{6}, \quad (2.5)$$

where $h(r)$ and $h(\theta)$ are the differential entropies of r and θ respectively. Denoting the probability density r by $p(r)$, $\sigma_r = \sqrt{\int_0^\infty r^2 p(r) dr}$.

3. Entropy-Coded Unrestricted Polar Quantization (ECUPQ): Polar representation of the signal is assumed. The encoding structure is specified as follows: r and θ are first scalar-quantized and then *jointly* entropy-coded. The quantizer design problem amounts to minimizing D subject to a constraint on $R = H(\hat{r}, \hat{\theta})$. This results in the right-most quantizer structure of Figure 2-3. The corresponding expression for D as a function of R (the derivation of which can be found in [16]) is:

$$D_{ECUPQ} = \frac{2^{-(R-h(r)-h(\theta)-b(r))}}{6}, \quad (2.6)$$

where $b(r) = \int_0^\infty p(r) \log_2(r) dr$.

Important facts to note about all three quantization schemes are:

- Each of the three quantizer structures can be expressed as the product of two scalar quantizers. This means that they have low computational complexity, which makes them attractive when compared to a vector quantizer.
- It is a well-known result that uniform scalar quantizers are optimal for high rate entropy-coded quantization of a scalar random variable [8]. This latter fact is reflected in Figure 2-3.
- The authors of [16] show that ECUPQ and ECRQ achieve the same performance for high rate quantization of a two-dimensional jointly Gaussian source with independent, identically-distributed components. ECSPQ achieves the worst performance of all three quantization schemes.

By allowing phase resolution to increase with magnitude, ECUPQ is able to maintain *close-to-square* quantization cells, which, if we restrict our attention to outermost cells, have lower second moments than their counterparts in ECSPQ. Since distortion is proportional to the second moment of quantization cells, this translates into lower distortion at a given rate for ECUPQ. The i2i polar transforms developed in this thesis are able to achieve *strictly* square quantization cells while allowing polar




Figure 2-3: As reproduced from [16], partitions of the input space for ECRQ (left), ECSPQ (middle) and ECUPQ (right)

representation of the data. It would therefore be interesting to compare the i2i polar transform coding scheme to ECUPQ in terms of their coding performance.

Chapter 3

I2I Polar Coordinates

In an information-theoretic sense, it is suboptimal to separately encode the components of a source whose components are not independent. If no linear transformation of the source would yield independent components, one may need to consider using a nonlinear change of variables. However, as developed in Chapter 1, nonlinear transformations are disliked because of their effect on the quantization partition. In that chapter, we also made the case for i2i transform coding by highlighting its ability to overcome the undesirable effect of nonlinear transforms on quantization partition cells in the conventional transform coding framework. In this chapter, we develop an i2i approximation to the Cartesian-to-polar transformation and show its superiority to the conventional approach for transform coding of sources separable in polar coordinates.

3.1 Designing the Transform

Consider a two-dimensional vector X with components x_1 and x_2 —these may correspond to pairs of proximate wavelet coefficients modeled by a GSM as in Section 2.2.¹ If X has independent or nearly independent radial and angular components, we hypothesized earlier that scalar quantization followed by an i2i approximation to the

¹note that restricting K_X in (2.2) to be the identity matrix is without loss of generality as we shall see in Chapter 4.

Cartesian-to-polar transformation (as in Figure 1-2) could do better at coding such a source than the conventional approach (Figure 1-1). Note that in Figure 1-2, $T_1(\cdot)$ can be omitted because the components of the source considered are already uncorrelated. Here, we test our hypothesis by considering the design of a family of i2i approximations to the Cartesian-to-polar transformation.

Let $(\hat{x}_1, \hat{x}_2) = ([x_1]_\Delta, [x_2]_\Delta)$, where Δ is the step size of the uniform scalar quantizer used in both dimensions and $[\cdot]_\Delta$ denotes rounding to the nearest multiple of Δ . Note that $(\hat{x}_1, \hat{x}_2) \in \Delta\mathbb{Z}^2$, which is the same as saying that $(\hat{x}_1/\Delta, \hat{x}_2/\Delta) \in \mathbb{Z}^2$. The design of our i2i polar transform therefore amounts to a relabeling of the integer coordinate pairs $(\hat{x}_1/\Delta, \hat{x}_2/\Delta)$ into integer pairs (P, A) , where P and A represent our discrete “pseudoradius” and angle variables respectively. In order to avoid loss of information, the mapping from $(\hat{x}_1/\Delta, \hat{x}_2/\Delta)$ to (P, A) should be invertible.

Note that simply quantizing a polar coordinate representation will not work. Let $r = \sqrt{(\hat{x}_1/\Delta)^2 + (\hat{x}_2/\Delta)^2}$ and $\theta = \tan^{-1}[(\hat{x}_2/\Delta)/(\hat{x}_1/\Delta)]$. Further assume $A \in \{0, 1, \dots, M-1\}$ with $M \in \mathbb{Z}^+$. For illustration, one can easily check that

$$(P, A) = ([r]_1, \left\lfloor \frac{M}{\pi} [\tan^{-1}[(\hat{x}_2/\Delta)/(\hat{x}_1/\Delta)] \bmod \pi] \right\rfloor)$$

is not invertible. Instead of trying to find an approximation to (r, θ) that is invertible on $\Delta\mathbb{Z}^2$, we adopt a sorting approach as follows:

- Divide \mathbb{Z}^2 into M disjoint discrete angle regions by choosing

$$A = \left\lfloor \frac{M}{\pi} [\tan^{-1}(\hat{x}_2/\hat{x}_1) \bmod \pi] \right\rfloor.$$

- Within each angle region (except the zeroth), the $(\hat{x}_1/\Delta, \hat{x}_2/\Delta)$ pair with smallest value of $r = \sqrt{(\hat{x}_1/\Delta)^2 + (\hat{x}_2/\Delta)^2}$ is assigned “pseudoradius” $P = 1$. In the case of the zeroth region, the origin is the point with smallest value of r ; its corresponding “pseudoradius” is $P = 0$. The remaining values of P are assigned by going through each point in increasing order of r and incrementing the value of P at each step; ties are broken by any deterministic rule.

The above construction is illustrated in Figure 3-1 for different values of M . By construction, such a mapping is invertible. The main drawback of such an approach is the lack of a simple expression for the mapping from $(\hat{x}_1/\Delta, \hat{x}_2/\Delta)$ to (P, A) .

3.2 High Resolution Analysis

The key feature of the i2i approach is that any invertible transformation of \hat{y} can be used without affecting the distortion. Thus to consider our performance analytically, we only need to approximate the rate, i.e., the entropies $H(P)$ and $H(A)$.

Approximating $H(A)$ is not difficult or interesting. In our examples, A is uniformly distributed on $\{0, 1, \dots, M-1\}$ so $H(A) = \log_2 M$.

To evaluate $H(P)$ directly, we would need the probability mass function of P , but this is difficult to relate to f_{x_1, x_2} precisely because of the complicated mapping from (\hat{x}_1, \hat{x}_2) to (P, A) . Instead, we relate $H(P)$ to some relevant differential entropy. This is summarized in the following theorem:

Theorem: Let X be a two-dimensional vector with components x_1 and x_2 in rectangular coordinates. If the polar coordinate representation (r, θ) has independent components, then


$$H(P) \approx h(r^2) + \log_2(\pi/M) - 2 \log_2 \Delta, \quad (3.1)$$

with vanishing relative error as $\Delta \rightarrow 0$.²


Proof: We obtain (3.1) by relating P to the uniform quantization (with unit step size) of a continuous-valued quantity $\pi r^2/(M\Delta^2)$.

Consider the cell of area Δ^2 centered at (\hat{x}_1, \hat{x}_2) as shown in the left panel of Figure 3-2, and denote the “pseudoradius” label of the cell by $P = k$. We associate it with a section of a ring, also with area Δ^2 , as shown in the right panel.


²This seems to be a good approximation even if r and θ are not independent; a precise statement would require M to grow as Δ shrinks.



(a) $M = 8$



(b) $M = 11$



(c) $M = 16$

Figure 3-1: Visualizing i2i Polar Coordinates, (a) $M = 8$, (b) $M = 11$, (c) $M = 16$. The left panel shows cells with equal discrete radius P with the same color. The right panel shows regions of constant discrete angle A with constant color.




Figure 3-2: For high-resolution analysis of $H(P)$, a square cell with area Δ^2 (left panel) is replaced by a portion of a ring with area Δ^2 (right panel) as shown. The analysis finds the radius and thickness of the ring as a function of P .

When (a) Δ is small; (b) r has a smooth density; and (c) r and θ are independent, then the probabilities of the two regions are equal.³ We obtained the desired approximation by equating two estimates of the area of the wedge bounded by an arc going through (\hat{x}_1, \hat{x}_2) . First, because of the point density $1/\Delta^2$, the area should be approximately $k\Delta^2$; we have simply counted points that are assigned smaller P values. On the other hand, the area is $\pi r_k^2/M$, where r_k is the distance from the origin to the point of interest. Equating the areas,

$$k \approx \frac{\pi r_k^2}{M\Delta^2}.$$

Thus, the (integral) P values are approximately the same as what would be obtained by rounding $\pi r^2/(M\Delta^2)$ to the nearest integer. This completes the proof. We validate the above proof by considering the case where x_1 and x_2 are i.i.d., zero-mean Gaussian random variables with variance σ^2 .

Example: Suppose x_1 and x_2 are i.i.d., zero-mean Gaussian with variance σ^2 . In polar coordinates we have r^2 exponentially distributed with mean $2\sigma^2$. Using standard

³Requirements (b) and (c) could be replaced with smoothness of the joint density of r and θ if the angular extent were small.

differential entropy calculations [3], one can evaluate as follows:

$$\begin{aligned} h(x_1) + h(x_2) &= \frac{1}{2} \log_2(2\pi e\sigma^2) + \frac{1}{2} \log_2(2\pi e\sigma^2) \text{ bits} \\ [h(r^2) + \log_2(\pi/M)] + H(A) &= [\log_2(2e\sigma^2) + \log_2(\pi/M)] + \log_2 M \text{ bits.} \end{aligned}$$

Our high-resolution analysis thus implies

$$H(\hat{x}_1) + H(\hat{x}_2) \approx H(P) + H(A),$$

as we would expect.

For sources that are independent in polar coordinates but not in Cartesian coordinates, we will obtain $H(P) + H(A) < H(\hat{x}_1) + H(\hat{x}_2)$.

In what follows, we use the result from the above proof to compute the distortion-rate performance of the i2i polar scheme. Let $R = H(P) + H(A)$ be the entropy constraint on the encoding and $D = E[(x_1 - \hat{x}_1)^2 + (x_2 - \hat{x}_2)^2] = \frac{\Delta^2}{6}$ be the mean-squared error (MSE) resulting from scalar quantization of x_1 and x_2 (and consequently from the i2i polar scheme since the transformation from (\hat{x}_1, \hat{x}_2) to (P, A) is lossless). Then for large P ,

$$\begin{aligned} R &= H(P) + H(A) \\ &= h(r^2) + \log_2(\pi/M) - 2 \log_2 \Delta + h(\theta) - \log_2(2\pi/M) \\ &= h(r^2) + h(\theta) - 2 \log_2 \Delta - 1 \\ &= h(r) + E \left[\log_2 \left| \frac{d}{dr}(r^2) \right| \right] + h(\theta) - \log_2 \Delta^2 - 1 \\ &= h(r) + 1 + E[\log_2 r] + h(\theta) - \log_2 \Delta^2 - 1 \\ &= h(r) + E[\log_2 r] + h(\theta) - \log_2 \Delta^2, \end{aligned}$$

where the fourth equality follows from well-known properties of the differential entropy of functions of a random variable [3].

Rearranging the above equation gives the following expression for D as a function

of R :

$$D = \frac{2^{-(R-h(r)-h(\theta)-b(r))}}{6}, \quad (3.2)$$

where $b(r) = \int_0^\infty p(r) \log_2(r) dr$.

Note that the high-rate performance of the i2i polar scheme does not depend on the number of angle regions M , as can be seen from (3.2).

The analysis, in layman's terms: An alternative way of deriving the high-resolution performance of the proposed i2i polar scheme is by considering the area covered by the event $\{P \leq k\}$. We use Figure 3-2 as a visual aid to estimate the said area. Loosely, the calculation amounts to finding the aggregate probability of M wedges with roughly k quantization cells of area Δ^2 per wedge and equating it to the area of a circle with radius r_k :

$$M \times k \Delta^2 = \pi r_k^2.$$

Rearranging the above, we arrive at the same conclusion as in the more sophisticated high-rate analysis,

$$k \approx \frac{\pi r_k^2}{M \Delta^2}.$$

This completes the analysis of the i2i polar scheme. Next, we present results of numerical simulations that validate the analysis above.

3.3 Results

In this section, we validate the i2i polar transform coding approach on an idealization of a class of sources inspired from sensor networks as well as on examples of Gaussian scale mixtures. We show the superiority of the i2i polar scheme to the conventional approach through Monte Carlo simulation. We also verify that the simulation results are in agreement with the performance predicted by the high-rate analysis. Last but not least, we compare our approach to entropy-coded unrestricted polar quantiza-

tion (ECUPQ) and entropy-coded strictly polar quantization (ECSPQ) using results from [16]. The i2i polar scheme achieves the same performance as ECUPQ both for GSMs and the source inspired from sensor networks. It achieves the same performance as ECSPQ on the sensor network example but significantly outperforms it for GSMs.

3.3.1 Numerical Results for GSM Examples

We present results for two examples of GSMs, one with a Rayleigh prior on z (exponential prior on z^2) and the other with an exponential prior on z .

Exponential prior on z^2 (Rayleigh prior on z)

If z^2 is an exponential random variable with parameter λ , it can be easily shown that z follows a Rayleigh distribution with parameter $\frac{1}{\sqrt{2\lambda}}$, denoted as $R(\frac{1}{\sqrt{2\lambda}})$. Assume a 2D GSM vector Y with multiplier $z \sim R(\frac{1}{\sqrt{2\lambda}})$. Then each of its components, y_1 and y_2 , follows a Laplacian distribution with parameter $\sqrt{2\lambda}$, which we denote as $L(\sqrt{2\lambda})$ [10]. It is straightforward to see, from the definition of a GSM, that y_1 and y_2 are uncorrelated but not independent. It can also be easily verified that their joint pdf is separable in polar coordinates (2.3). In particular, the phase, θ , is uniformly distributed between 0 and 2π (denoted as $U([0,2\pi])$) and the pdf of the radius, r , can be obtained in closed-form from (2.3):

$$f_r(r) = 2\lambda r K_0(\sqrt{2\lambda}r),$$

where $K_0(\sqrt{2\lambda}r)$ is a modified Bessel function of the second kind.

In what follows, we assume, without loss of generality, that $\lambda = 1$. The above information on the marginal pdfs allows us to derive $h(y_1)$, $h(y_2)$ and $h(\theta)$ in closed-form. Using numerical integration when needed, we were able to compute or approximate

$h(r)$, $b(r)$ and σ_r . Below is a summary of the values for these parameters:

$$\begin{aligned} h(y_1) &= \frac{1}{2} \log_2(2e^2) \\ h(y_2) &= \frac{1}{2} \log_2(2e^2) \\ h(\theta) &= \log_2 2\pi \\ h(r) &\approx 1.505 \\ b(r) &\approx -0.338 \\ \sigma_r &\approx \sqrt{2}. \end{aligned}$$

Putting all of the above together yields the following expressions for the performance of various entropy-coding techniques applied to a GSM with an E(1) prior on z^2 :

$$\begin{aligned} D_{ECI2I}(R) &\approx 14.11 \frac{2^{-R}}{6} \\ D_{ECUPQ}(R) &\approx 14.11 \frac{2^{-R}}{6} \\ D_{ECRQ}(R) &\approx 14.78 \frac{2^{-R}}{6} \\ D_{ECSPQ}(R) &\approx 25.22 \frac{2^{-R}}{6}. \end{aligned}$$

Simulation results for the source considered in this sub-section are reported in Figure 3-3. The figure shows that, at high rate, both ECI2I and ECUPQ (red curve) outperform the conventional transform coding approach denoted as ECRQ (black curve) and are very close to optimal (green curve), namely to the performance that could be obtained by scalar quantization followed by joint entropy coding. At low rate, the $(P, A) = (0, 0)$ point has high probability and thus merits special attention. To improve the low-rate performance, we do not encode A when $P = 0$. The rate of this encoding strategy is $H(P) + P(P \neq 0)H(A | P \neq 0)$. Regardless of whether this adjustment is made, the i2i polar approach outperforms the conventional one at high rate. While the improvement is relatively small, it should be noted that it is for a source for which no improvement was thought possible with transform coding [14].




Figure 3-3: Simulation results. $M = 8$ and a range of rates is obtained by varying Δ .

In Figure 3-4, we compare simulation results to the performance predicted by the high-rate analysis. As can be seen, the simulation results agree well with theory.

Figure 3-5 shows a comparison, from analytical results, between ECI2I and ECSPQ. The latter significantly outperforms the former at high rate. In fact, the ratio of distortions at high rate is given by:

$$\frac{D_{ECSPQ}}{D_{ECI2I}} = \frac{\sigma_r}{2^{b(r)}} \approx 1.79.$$

Both ECI2I and ECUPQ aim at maintaining quantization partition cells that are as close as possible to squares (in ECI2I, the quantization partition cells are exact squares), while the cells that result from ECSPQ become further away from squares as one gets further away from the origin. Therefore, the outermost cells resulting from ECSPQ have higher second-order moments than those induced by ECI2I and ECUPQ, which translates into higher distortion (at the same rate).

It is also interesting to note that the performance gain does not depend on the value of the parameter λ of the GSM multiplier z . This is portrayed in Figure 3-6, which shows that the ratios D_{ECSPQ}/D_{ECI2I} and D_{ECRQ}/D_{ECI2I} are constant as a

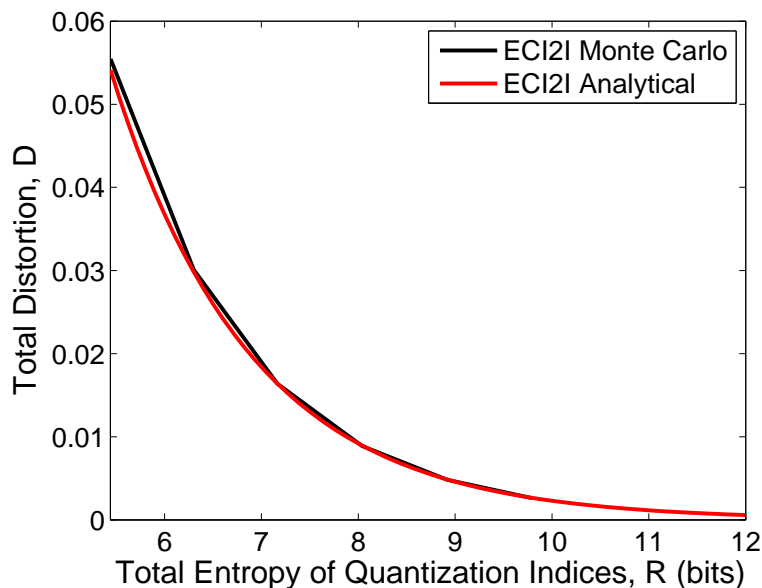


Figure 3-4: Comparing Monte Carlo simulation of i2i Polar transform coding ($M = 8$) to the performance predicted by the high-rate analysis

function of λ . This is relevant to the image coding motivating example where the multiplier z of the GSM would depend on the subbands from which parent and child pairs of wavelet coefficients are extracted. Figure 3-6 essentially tells us that, at high rate, the ratio of the distortions considered above would be the same regardless of the subbands from which parent and child are extracted.

Exponential prior on z

Calculations of differential entropies and of the parameters required to compute the analytical distortion-rate performance were not as tractable for this GSM as they were for the one considered in the previous subsection. To the best of our knowledge, the marginal pdfs of y_1 and y_2 do not admit a closed-form, which means that we were not able to compute $h(y_1)$ and $h(y_2)$ or even estimate them. Although the pdf of the radius, r , has a closed-form, not all of the parameters that depend on it are easy to obtain, even through numerical integration. Indeed, we were able to estimate $b(r)$ and σ_r but not $h(r)$. Below is a summary of the marginal pdfs and parameters we

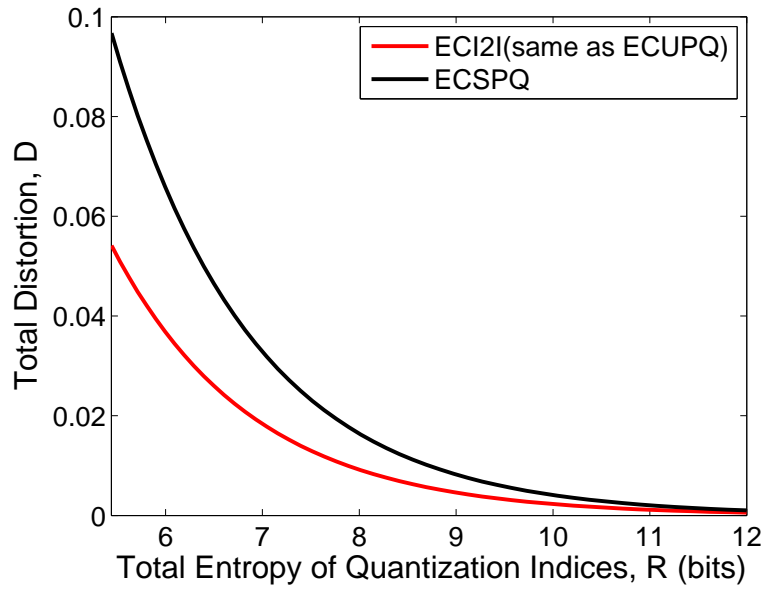


Figure 3-5: Analytical Comparison of i2i Polar transform coding to ECSPQ at high rate

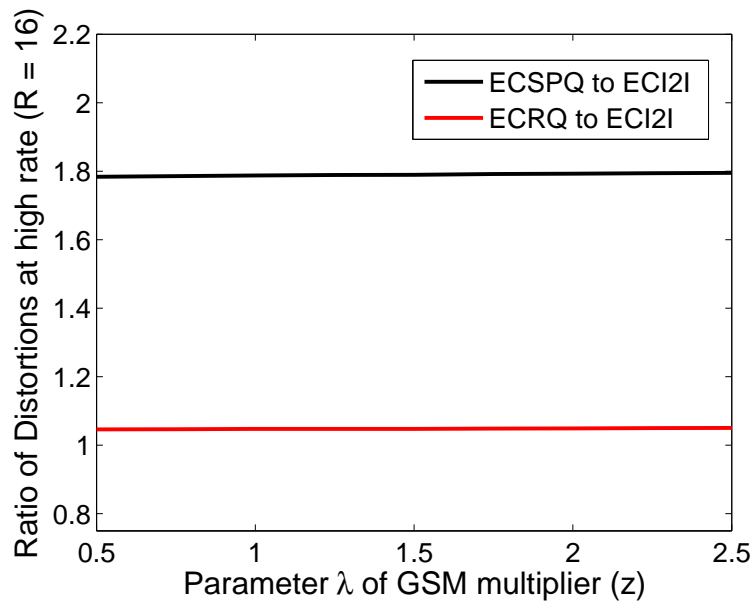


Figure 3-6: Ratio of distortions at high rate as a function of λ

were able to obtain for this GSM:

$$\begin{aligned}
f_\theta(\theta) &= \frac{1}{2\pi} \\
h(\theta) &= \log_2 2\pi \\
f_r(r) &= G_{0,3}^{3,0} \left(\frac{\lambda^2 r^2}{8} \middle| \begin{array}{c} - \\ 1/2, 1/2, 0 \end{array} \right) \\
b(r) &\approx -0.749 \\
\sigma_r &\approx 2,
\end{aligned}$$

where $G_{0,3}^{3,0} \left(\frac{\lambda^2 r^2}{8} \middle| \begin{array}{c} - \\ 1/2, 1/2, 0 \end{array} \right)$ is an instance of the Meijer G-function, which is a very general function that reduces to simpler well-known functions in many cases [15].

Considering all of the above, our results in this section are in the form of Monte Carlo simulation as well as the ratio of distortions of ECI2I and ECSPQ.

Simulation results are reported in Figure 3-7. The figure shows that, at high rate, both ECEI2I and ECUPQ (red curve) outperform the conventional transform coding approach denoted as ECRQ (black curve) and are very close to optimal (green curve), namely to the performance that could be obtained by scalar quantization followed by joint entropy coding. We adopted the same strategy as in the previous subsection to improve the low-rate performance. Note that for this particular GSM, the gain from ECI2I over ECRQ is much larger than in the previous subsection. This is also reflected in the ratio of distortions of ECI2I and ECSPQ:

$$\frac{D_{ECSPQ}}{D_{ECI2I}} \approx 3.36.$$

The intuition behind such a significant gain follows an argument along the lines of the one put forth in the previous subsection: outermost quantization cells in ECSPQ are further away from squares than in ECI2I and ECUPQ, which translates into higher contribution to overall distortion from such cells. The gain is larger in this case (3.36 vs. 1.79) because the pdf of the radius decays slower, which means that outermost




Figure 3-7: Simulation results. $M = 8$ and a range of rates is obtained by varying Δ .

ECSPQ quantization cells have more probability mass associated with them than in the example of the previous subsection.

3.3.2 Numerical Results for Sensor Networks Example

Consider a network of nodes that track one or more objects using a range-only sensing modality. If the nodes share their (low-quality, range-only) measurements, object locations can be estimated precisely by multilateration. Suppose a central processing unit performs the multilateration and sends (high-quality) location estimates back to non-central nodes. What is the nature of the source coding problem at the central node (transmitter)?

The most efficient communication is achieved when the transmitter uses the receiver's conditional object location density, where the conditioning is on the receiver's own measurement. Because range measurements are taken, these conditional densities are highly non-Gaussian and much more amenable to polar coordinates than Cartesian coordinates. An example, reproduced from [13], is shown in Figure 3-8. (This source coding problem is not considered in [13].)




Figure 3-8: A depiction, reproduced from [13], of a highly non-Gaussian distribution that arises in a localization problem with ranging sensors. The distribution is very complicated, but is well-suited to a polar coordinate representation.

Motivated by a localization problem described above, we consider a source with density similar to that observed in Figure 3-8. Specifically, we chose θ uniformly distributed on $[0, \pi]$ and $r = |L(5) + 4|$, where $L(5)$ denotes a Laplacian random variable with parameter equal to 5. The motivation behind such a choice stems from the fact that we want the joint density to be peaky around some particular radius; a shifted Laplacian random variable seemed like a reasonable choice. Taking the absolute value enforces the fact that r must be non-negative.

Below is a summary of pdfs and relevant parameters that we are able to obtain, either in closed-form or through numerical integration, for the source described above:

$$\begin{aligned}
 f_{\theta}(\theta) &= \frac{1}{\pi} \\
 h(\theta) &= \log_2 \pi \\
 f_r(r) &= \frac{5}{2}(e^{-5|r-4|} + e^{-5|x+4|}) \\
 h(r) &\approx -0.121 \\
 b(r) &\approx 2 \\
 \sigma_r &\approx 4.
 \end{aligned}$$

Simulation results are reported in Figure 3-9. The figure shows that, at high rate, both ECEI2I and ECUPQ (red curve) outperform the conventional transform coding approach denoted as ECRQ (black curve) and are very close to optimal (green curve), namely to the performance that could be obtained by scalar quantization followed by joint entropy coding. We adopted the same strategy as in the previous subsection to

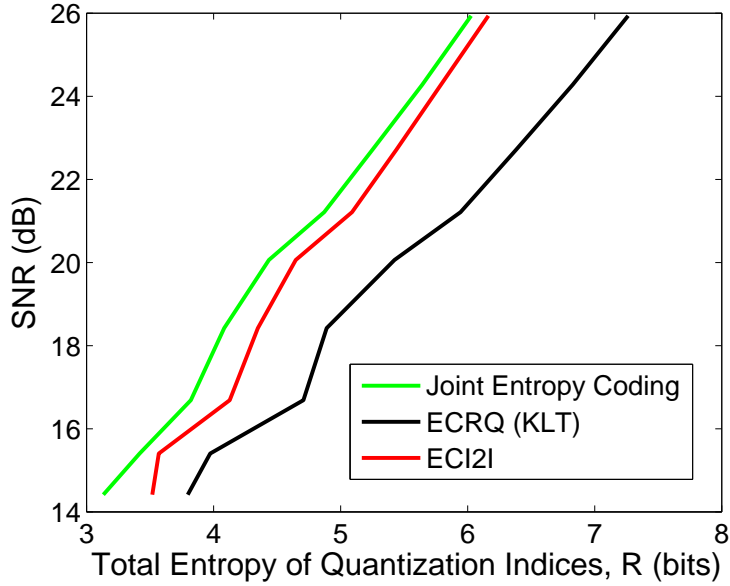


Figure 3-9: Simulation results. $M = 8$ and a range of rates is obtained by varying Δ .

improve the low-rate performance. Note that for this particular source, the gain from ECI2I over ECRQ is much more important than in the case of GSMs.

In Figure 3-10, we compare simulation results to the performance predicted by the high-rate analysis. As can be seen, the simulation results agree well with theory.

Figure 3-11 shows a comparison, from analytical results, between ECI2I and EC-SQPQ. As can be seen, the high-rate performance of both schemes is the same. In fact, the ratio of distortions at high rate is approximately 1:

$$\frac{D_{ECSPQ}}{D_{ECI2I}} \approx 1.$$

The model we chose as representative of the distribution shown in Figure 3-8 is characterized by a peak at $r = 4$ as well as the low variance of its radial component. This means that it does not matter whether phase quantization depends on magnitude or not, which explains why ECI2I/ECUPQ and ECSPQ have approximately the same performance at high rate. Intuitively, outermost quantization cells of ECSPQ have very low probability mass associated with them, which means that their not being close to squares does not affect overall distortion as much as it did for GSMs.




Figure 3-10: Comparing Monte Carlo simulation of i2i Polar transform coding ($M = 8$) to the performance predicted by the high-rate analysis




Figure 3-11: Analytical Comparison of i2i Polar transform coding to ECSPQ at high rate

Chapter 4

Generalizations of Integer-to-integer Polar Coordinates

Thus far, we have limited our treatment of i2i polar transform coding to two dimensions mainly due to the simplicity of the manipulations involved at that level. We have also assumed that the sources we wish to encode have uncorrelated components. In this section, we touch on how i2i polar transform coding may generalize to higher dimensions and show that the restriction to sources with identity covariance matrix is without loss of generality. Moreover, the ability to perform lossless polar coding is presented as a direct application of i2i polar transform coding.

4.1 Generalization to Higher Dimensions

In the general n -dimensional case, a source is said to be spherically symmetric if its probability density function can be expressed as the product of a term that depends on the magnitude r only and $n - 1$ angle terms (a latitude term θ_1 and $n - 2$ longitude terms θ_i , $i = 2, 3, \dots, n - 1$). Mathematically, this is expressed as follows:

$$p(r, \theta_1, \theta_2, \dots, \theta_{k-1}) = p(r)p(\theta_1)p(\theta_2)\dots p(\theta_{k-1}),$$

where

$$r \in [0, \infty)$$

$$\theta_1 \in [0, 2\pi]$$

$$\theta_i \in [0, \pi], i = 2, 3, \dots, n - 1$$

When going from 2D to higher dimensions, the only difference is that there are new angle variables that all take values in $[0, \pi]$.

A 3D Example: Let x_1, x_2 and x_3 be three independent, jointly Gaussian random variables, each with zero mean and unit variance. Using well-known rules for computing the joint probability density function of transformations of random variables, it can be shown that, in spherical coordinates, r is a Maxwell random variable with parameter 1, θ_1 is uniformly distributed in $[0, 2\pi]$ and $f_{\theta_2}(\theta_2) = \sin(\theta_2)/2$. Moreover, r, θ_1 and θ_2 are independent.

The idea behind i2i polar coordinates in higher dimensions is to break \mathbb{Z}^n into M disjoint wedges, each corresponding to a discrete angle region as in Figure 3-2 (but in a higher dimensional space), and repeat the procedure outlined in section 3.1 for the design of the i2i polar transformation.

4.2 Generalization to Sources Separable in *Extended* Polar coordinates

Definition: Assume a two-dimensional source vector X whose components x_1 and x_2 are uncorrelated with respective variances σ_1^2 and σ_2^2 . Further assume a 2×2 diagonal matrix D with elements $1/\sigma_1$ and $1/\sigma_2$. X is said to be separable in extended polar coordinates if the polar representation of $Y = DX$ has independent components.

Example: Let X in the definition above be jointly Gaussian and (r, θ) be its polar representation. The Jacobian of the transformation from X to polar coordinates is




Figure 4-1: An encoder for a polar-separable source with arbitrary covariance matrix with a discrete-domain transform $\hat{D} : \Delta\mathbb{Z}^N \rightarrow \Delta\mathbb{Z}^N$ and an integer-to-integer polar transform $U : \Delta\mathbb{Z}^N \rightarrow \Delta\mathbb{Z}^N$. $Q(\cdot)$ is an unbounded uniform scalar quantizer with step size Δ .

given by $J_X = r$. It can be easily checked that the polar representation of X does not have independent components. The polar representation of $Y = DX$, on the other hand, does have independent components. Indeed, the Jacobian of the latter transformation is given by $J_X = \sigma_1\sigma_2r$, from which it is straightforward to verify that X is separable in extended polar coordinates.

Let X be an two-dimensional source vector belonging to the class of sources separable in extended polar coordinates and let K_X denote its covariance matrix. To encode such a source under the i2i polar transform coding framework, we reduce it, as mathematicians would say, to a previously solved problem (Section 3.1). Let $Y = DX$, where D is a transformation of X such that $K_X = I$, the identity operator. D is a 2×2 matrix which, in general, is nonorthogonal. Just like nonlinear transforms, nonorthogonal transforms are also disliked in conventional transform coding because they tend to deform the quantization partition (although to a smaller extent). To overcome such an issue, we invoke the linear i2i transform coding framework developed in [5], which is based on the result that any 2×2 matrix with determinant 1 can be implemented such that it is invertible on $\Delta\mathbb{Z}^2$ through the lifting scheme. A more extensive treatment of matrix factorizations that are possible through the lifting scheme can be found in [9].

In summary, a source X separable in extended polar coordinates should first be quantized and then transformed through an i2i approximation, \hat{D} , to D , which reduces our problem to that previously solved for polar-separable sources with identity covariance matrices (Figure 4-1).

4.3 Lossless Coding with Integer-to-integer Polar Coordinates

The problem of compressing two-dimensional discrete-domain data is one that has long been of interest to the source coding community. In their pioneering work [1], Bruekers and van den Eden proposed a framework that allows perfect inversion and perfect reconstruction in filter banks. Later work by Calderbank et al. [2] developed integer-to-integer versions of wavelet transforms and showed promising applications to image coding. More recently, a general matrix factorization theory for reversible integer mapping of invertible linear transforms was developed in [9]. Such a theory forms the basis of the integer-to-integer transform coding framework introduced in [5], although not available then in the form presented in [9].

This brief, and by no means exhaustive, review of the integer-transforms literature shows that reversible mappings that map integers to integers play an important role in transform coding. To our knowledge, prior work on lossless transform coding has focused on developing integer-to-integer versions of underlying continuous-domain *linear* transformations. This is mainly because continuous-domain *nonlinear* transformations are not of much interest in conventional transform coding, let alone integer-to-integer versions of such transformations.

We have previously shown that ECI2I and ECUPQ have the same distortion rate performance for coding of two-dimensional *continuous-domain* data that is amenable to polar coordinates. Consider now the problem of efficiently encoding two-dimensional *discrete-domain* data that is amenable to polar coordinates. The standard approach would be to interpolate the data in polar coordinates, thereby introducing error/loss, and then use an appropriate entropy-coded polar quantization technique (such as ECUPQ or ECSPQ). An alternative approach is to employ the family of i2i polar transformations introduced in this thesis, followed by scalar entropy coding of “pseudoradius” and angle. The latter approach has the advantage of being lossless, while exploiting the suitability of the data to polar-coordinate processing. Indeed, the

saving in bits would still be:

$$H(P) + H(A) - H(\hat{x}_1) - H(\hat{x}_2) \approx h(r) + E[\log_2 r] + h(\theta) - h(x_1) - h(x_2),$$

as predicted by the high-rate analysis. However, the total distortion is now zero, as opposed to what is given in (3.2). Note that the bit saving is given in terms of continuous-domain data because we can assume that the discrete-domain data was obtained by quantization (with step-size $\Delta = 1$) of some underlying continuous source specified by the pairs (x_1, x_2) in Euclidean space and (r, θ) in polar coordinates.

Chapter 5

Discussion

In conventional transform coding, the desire to preserve the shape of quantization partition cells as well as have the luxury of encoding independent quantities are competing goals. Integer-to-integer transform coding is a method that aims at simultaneously achieving both of these goals by allowing quantization to occur first and then implementing an invertible integer-to-integer approximation to a possibly nonlinear transformation that will yield independent components. Inspired by the integer-to-integer (i2i) framework [5], this thesis introduces the generic encoder of Figure 1-2 for nonlinear transform coding of polar separable sources. Specifically, we design an integer-to-integer approximation to the Cartesian-to-polar transformation and analytically show its superiority to the conventional approach (shown in Figure 1-1) for high-rate transform coding of polar-separable sources. The claim is that i2i polar transform coding is optimal for high-rate transform coding of sources separable in polar coordinates, where optimality is measured by the performance that could be achieved by *uniform* scalar quantization (in rectangular coordinates) followed by *joint* entropy coding. Our analysis is supported by simulation results on an example of sources borrowed from sensor networks applications, as well as Gaussian scale mixtures (GSMs), a class of sources for which no improvement was previously thought possible under the conventional transform coding paradigm [14]. Moreover, comparisons to entropy-coded polar quantization techniques show that the high-rate performance is the same as that of entropy-coded unrestricted polar quantization

(ECUPQ) at arguably lower complexity [16]. We also show that the i2i polar scheme outperforms entropy-coded strictly polar quantization (ECSPQ) on two examples of GSMs and does equally-well on the source inspired from sensor networks applications.

The main drawback of our approach is that the transformation is designed through a sorting procedure rather than being given by a simple expression.

Our results, however, may set the stage for a more comprehensive theory of nonlinear transform coding. Indeed, a question that is implicitly asked in our treatment is the following: given two independent random variables $y_1 = f(x_1, x_2)$ and $y_2 = g(x_1, x_2)$, how can one approximate the mapping $(x_1, x_2) \mapsto (y_1, y_2)$ such that it is invertible on $\Delta\mathbb{Z}^2$? A satisfactory answer to this question will facilitate nonlinear transform coding. Results from nonlinear independent component analysis (ICA) suggest a way to approach the general problem posed above. If the space of functions f and g is not limited, the authors in [12] show that there exists infinitely many functions f and g to choose from. However, under mild conditions on f and g , they also show that the following mapping will always result in random variables y_1 and y_2 that are jointly uniformly distributed on $[0, 1]^2$:

$$\begin{aligned} Y_1 &= F_{x_1}(X_1) \\ Y_2 &= F_{x_2|X_1}(X_2|X_1) \end{aligned}$$

where X_1 , X_2 , Y_1 and Y_2 are realizations of x_1 , x_2 , y_1 and y_2 respectively.

Now that we have reduced the number of candidate functions f and g to the above two, the nonlinear transform coding problem becomes that of finding a generic way of approximating the above mapping so that it is invertible on $\Delta\mathbb{Z}^2$.

Bibliography

- [1] F. A. M. L. Bruekers and A. W. M. van den Enden. New networks for perfect inversion and perfect reconstruction. *IEEE J. Sel. Areas Comm.*, 10:130–137, January 1992.
- [2] R. Calderbank, I. Daubechies, W. Sweldens, and B.-L. Yeo. Wavelet transforms that map integers to integers. *Appl. Computational Harmonic Anal.*, 5(3):332–369, July 1998.
- [3] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, New York, 1991.
- [4] B. Edler and H. Purnhagen. Parametric audio coding. In *Proc. 5th Int. Conf. Signal Processing, Beijing, China*, pages 614–617, August 2000.
- [5] V. K Goyal. Transform coding with integer-to-integer transforms. *IEEE Trans. Inform. Theory*, 46(2):465–473, March 2000.
- [6] V. K Goyal. A toy problem in transform coding, or: How I learned to stop worrying and love nonorthogonality. École Polytechnique Fédérale de Lausanne, Computer and Communication Sciences Summer Research Institute, July 2003.
- [7] V. K Goyal, J. Zhuang, and M. Vetterli. Transform coding with backward adaptive updates. *IEEE Trans. Inform. Theory*, 46(4):1623–1633, July 2000.
- [8] R. M. Gray and D. L. Neuhoff. Quantization. *IEEE Trans. Inform. Theory*, 44(6):2325–2383, October 1998.

- [9] P. Hao and M. Q. Shi. Matrix Factorizations for Reversible Integer Mapping. *IEEE Trans. Signal Proc.*, 49(10):2314–2324, October 2001.
- [10] A. Hjørungnes, J. Lervik, and T. Ramstad. Entropy coding of composite sources modeled by infinite Gaussian mixture distributions. In *Proc. IEEE Digital Signal Processing Workshop*, pages 235–238, September 1996.
- [11] J. J. Y. Huang and P. M. Schultheiss. Block quantization of correlated Gaussian random variables. *IEEE Trans. Commun. Syst.*, 11:289–296, September 1963.
- [12] A. Hyvarinen and P. Pajunen. Nonlinear independent component analysis: Existence and uniqueness results. *Neural Networks*, 12:429–439, 1999.
- [13] A. T. Ihler, J. W. Fisher III, R. L. Moses, and A. S. Willsky. Nonparametric belief propagation for self-localization of sensor networks. *IEEE J. Sel. Areas Comm.*, 23(4):809–819, April 2005.
- [14] S. Jana and P. Moulin. Optimality of KLT for encoding Gaussian vector-scale mixtures: Application to reconstruction, estimation and classification. *IEEE Trans. Inform. Theory*, 2004. submitted.
- [15] A. M. Mathai. *A Handbook of Generalized Special Functions for Statistical and Physical Sciences*. Oxford University Press, New York, 1993.
- [16] R. Vafin and W. B. Klejin. Entropy-constrained polar quantization and its application to audio coding. *IEEE Trans. Speech Audio Proc.*, 13(2):220–232, March 2005.
- [17] M. J. Wainwright and E. P. Simoncelli. Scale mixtures of Gaussians and the statistics of natural images. In *Advances in Neural Information Processing Systems*, pages 855–861. MIT Press, 2000.